SHARP APL *for UNIX*

Documentation Suite

USING THIS DOCUMENTATION
MASTER INDEX
HANDBOOK
AUXILIARY PROCESSORS MANUAL
FILE SYSTEM MANUAL
INTRINSIC FUNCTIONS MANUAL
LANGUAGE GUIDE
SVP MANUAL
SYSTEM GUIDE

SOLITON ASSOCIATES

Using This Documentation

All Soliton documentation delivered in Portable Document Format (PDF) is designed for optimal printing as well as for on-screen viewing.

Colors. Navigation screens, divider pages, and cross-references appear in various *colors* (to indicate hypertext links); however, the subject documentation exhibits a more traditional black-on-white "paper" look.

Navigation. Although you can turn pages one at a time using the PAGE UP and PAGE DOWN keys, Acrobat also includes tools to help you find text, locate specific pages, and return easily to previous page views. Consult your Acrobat documentation for instructions.

Magnification. It may also be helpful to increase the magnification of some detailed figures and examples using the zoom-in tool in the tool bar, or the magnification box at the bottom of the Acrobat window.

Cross-References. *Colored* hypertext links appear throughout the PDF wherever one expression cross-references another. They also appear in the divider pages, the Table of Contents, and the Index. Your mouse pointer changes into an index finger when it is positioned over a link — just click to activate the link.

Bookmarks. Bookmarks are special links that serve as an instant table of contents for the whole PDF. Click Window > Show Bookmarks to display the list of bookmarks on the left side of the document.

The list is presented in a hierarchy where subordinate Bookmarks appear indented below the parent. When a parent Bookmark is collapsed, all its subordinates are hidden and it shows a plus sign (Windows) or a triangle (Mac OS) next to it. Click the plus sign or triangle to display hidden subordinates.

S
A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

SHARP APL *for UNIX*

Master Index

CLICK ON THE LETTER BAR TO PROCEED

PAGE NUMBERS ARE HYPERTEXT LINKED

PAGE NUMBER LEGEND

A	AUXILIARY PROCESSORS MANUAL
F	FILE SYSTEM MANUAL
H	HANDBOOK
IF	INTRINSIC FUNCTIONS MANUAL
L	LANGUAGE GUIDE
S	SYSTEM GUIDE
SV	SVP MANUAL

© 2000 by Soliton Associates Limited
44 Victoria Street, Suite 2100, Toronto, Ontario, Canada M5C 1Y2

Pubcode: UW-000-3001 (0007)

), recall previous line *S6-4*
)*binds*, display bound intrinsic functions IF4-1, *S6-5*
)*clear*, clear active workspace S5-57, S5-64, *S6-5*
)*copy*, copy objects into workspace S5-64, *S6-7*
)*drop*, delete a workspace *S6-8*
)*edit*, full-screen editor H3-3, L1-13, S6-8, *S7-14*
)*erase*, erase global objects *S6-8*
)*fns*, display function names S6-9
)*ifdef*, display IF definition/redefine IF IF4-2, *S6-9*
)*iferase*, remove IF definitions IF4-3, *S6-10*
)*iflist*, list IF groups and their contents IF4-4, *S6-11*
)*lib*, display saved workspace names S6-12
)*load*, activate a workspace F1-4, S5-57, S5-64, *S6-13*
)*off*, terminate APL session H2-7, *S6-14*
)*pcopy*, protected copy *S6-7*
)*save*, save active workspace S5-64, *S6-15*, S6-20
)*sesm*, session manager conditions H3-4, H4-6, *S6-16*
)*si*, state indicator *S6-18*
)*sic*, state indicator clear *S6-18*, S7-9
)*sinl*, state indicator with name list *S6-18*
)*symbols*, set symbol table size *S6-19*
)*vars*, display variable names *S6-20*
)*wsid*, set workspace name *S6-20*
)*xload*, activate a workspace S5-64, S6-13
⋆ (star/asterisk); ***See*** exponential, power L4-7
+ (plus sign); ***See*** conjugate, identity, plus, addition L4-6
, (comma); ***See*** ravel, catenate L4-61
− (minus sign); ***See*** minus, negate, subtraction L4-6
< (less-than sign); ***See*** less than, enclose L4-27
= (equal sign); ***See*** equals, nubin L4-25, S5-17
> (greater-than sign); ***See*** greater than, disclose, open L4-27, S5-17
? (question mark, query); ***See*** roll, deal L4-19, S5-35
∧ (and sign, caret); ***See*** and, least common multiple L4-17, S5-17
~ (tilde,not sign); ***See*** not, less L4-19, S5-17
⌹ (domino); ***See*** matrix inverse, matrix divide L4-20
∨ (or sign, inverted caret); ***See*** or, greatest common divisor L4-17, S5-17
≠ (not equal sign); ***See*** not equals nubsieve L4-27
≥ (greater-than-or-equal sign); ***See*** greater than or equal L4-27, S5-17
⍟ (log sign/circle-star); ***See*** logarithm L4-8
⍱ (nor sign); ***See*** nor L4-18
≡ (equal-underbar); ***See*** match L4-28, S5-17
⍪ (comma-bar); ***See*** table, catenate-down L4-61
⊣ (left tack, lev); ***See*** stop, left L4-74
→ (right arrow); ***See*** branch L4-94
× (times sign); ***See*** times, mutilplication, signum L4-6
⌽ (circle-stile); ***See*** rotate, reverse L4-66, S5-35
≤ (less-than-or-equal sign); ***See*** less than or equal L4-27, S5-17
⍲ (nand sign); ***See*** nand L4-18
⊖ (circle-bar); ***See*** rotate-down, reverse-down L4-66, S5-35
∈ (epsilon); ***See*** member L4-28, S5-17
⍷ (epsilon-underbar); ***See*** in L4-30
⍎ (pawn/tack), ***See*** execute L4-91, S5-35, S5-63, S5-74, S5-76
⌊ (down stile); ***See*** floor, minimum L4-12, S5-17
⊥ (up tack); ***See*** decode, base-10 value L4-24
∇ (del editor) H2-6, L4-92, S5-21
⍒ (del-stile); ***See*** grade L4-38, S5-35
⍳ (iota); ***See*** count, index of L4-35, S5-17, S5-35
⍸ (iota-underbar); ***See*** index L4-37
⍋ (delta-stile); ***See*** grade L4-38, S5-35
⍕ (thorn); ***See*** format L4-78
⍞, character input/output H3-6, L3-7, S5-1 *to* S5-2, *S5-4 to S5-9*, *S5-53*, S5-56
⎕, evaluated input/output H3-6, L3-12, S5-2, *S5-4 to S5-9*, S5-35, S5-74
⎕*ai*, accounting information *S5-10*

□*append*/□*appendr*, append variable to tied file F2-16, ***S5-10***
□*arbin*/□*arbout*, arbitrary input/output H3-6, H4-8, ***S5-11***
□*av*, atomic vector H3-3, L3-9, L3-17, S5-11 *to* ***S5-12***
□*avail*, file system availability F2-8, ***S5-13***
□*avm*, mainframe atomic vector H3-3, ***S5-12***
□*bind*, intrinsic function bind IF1-1, IF3-1, IF3-6, ***S5-13*** *to* ***S5-14***
□*bounce*, terminate task H2-2, H2-7, H3-6, S5-4, ***S5-14***
□*cr*, canonical representation S5-2, ***S5-15***, S7-11
□*create*, create a file F2-10, ***S5-15***
□*ct*, comparison tolerance L4-12, L4-37 *to* L4-38, ***S5-15***
□*dl*, delay execution ***S5-18***
□*drop*, drop components F2-18, F3-9, ***S5-18***
□*ec*, evironment condition S5-18, S9-1, ***S9-6***
□*EDIT*, system editor S5-19, S7-14
□*er*, event report S5-19, ***S9-5***
□*erase*, erase a tied file F2-4, F2-11, ***S5-19***
□*ex*, expunge objects S5-4, ***S5-19***
□*fc*, format control ***S5-20***
□*fd*, function definition ***S5-20*** *to* ***S5-23***, S5-33, S5-38, S5-66, S7-11
□*fhold*, hold tied files F1-4 *to* F1-5, ***F2-24*** *to* ***F2-27***, S5-23
□*fi*, input format conversion ***S5-23***
□*fmt*, format output L3-16 *to* L3-21, ***S5-24*** *to* ***S5-31***
□*fsm*, full-screen manager ***S5-32***
□*fx*, fix a definition ***S5-32***, S5-38, S5-66
□*hold*, hold tied files F1-4, ***F2-24*** *to* ***F2-27***, S5-23
□*ht*, horizontal tab ***S5-33***
□*io*, index origin L4-19, L4-22, L4-31 *to* L4-45, L4-65, L4-73, ***S5-34***
□*lc*, line counter ***S5-35***
□*lib*, file library F2-8, ***S5-35***
□*load*, load a workspace ***S5-36***, S5-64
□*lx*, latent expression S5-1, ***S5-37***, S5-62, S5-64, S6-13
□*MF*, monitoring facility S5-1, S5-33, ***S5-37*** *to* ***S5-39***, S5-49
□*names*, library names of tied files F2-9, ***S5-41***
□*nc*, name class of identifiers S5-4, ***S5-41***, S5-51
□*nl*, name list ***S5-42***, S5-73, S6-20
□*nums*, tie numbers of tied files F2-9, ***S5-42***
□*out*, record session I/O to file ***S5-43*** *to* ***S5-47***
□*pack*, build a package L3-11, ***S5-47***
□*PARSE*, pre-parse user-defined function ***S5-48***
□*paths*, pathnames of tied files F2-10, ***S5-49***
□*pdef*, package define L3-11, ***S5-49***
□*pex*, package expunge L3-11, ***S5-50***
□*pins*, package insert L3-11, ***S5-50***
□*plock*, package lock L3-11, ***S5-50***
□*pnames*, package names L3-11, ***S5-51***
□*pnc*, package name class L3-11, ***S5-51***
□*pp*, print precision L4-78, ***S5-52***
□*ppdef*, protective package design L3-11, ***S5-53***
□*PR*, prompt replacement S5-1, ***S5-53***
□*ps*, position and spacing in display L3-15, L4-56, L4-78, ***S5-54***
□*psel*, package select L3-11, ***S5-56***
□*pval*, package value L3-11, ***S5-56***
□*pw*, printing width H4-8, L3-16 *to* L3-21, S5-11, ***S5-56***, S6-4 *to* S6-5, S7-2
□*qload*, quiet load ***S5-36***
□*rdac*, read access matrix F2-23, ***S5-57***
□*rdci*, read component information F2-19, ***S5-57***
□*rdfi*, read file information F2-20, ***S5-57***
□*read*, read component of a tied file F2-16, ***S5-58***
□*rename*, rename a tied file F2-12, ***S5-58***
□*replace*, replace variable stored in a component F2-17, F3-9, ***S5-58***
□*resize*, reset file size limit F2-3, F2-21, ***S5-58***
□*rl*, random link L4-19, L4-22, ***S5-58***
□*run*, start N-task H2-2, F2-4, S1-3, ***S5-59*** *to* ***S5-62***
□*runs*, report running tasks ***S5-63***

□*sc*, shared variable state change SV2-5, ***SV3-22*** *to* ***SV3-24***, S4-2, S5-63
□*signal*, signal event S5-63, S7-9, S9-1, ***S9-3***, ***S9-12***
□*size*, size of tied file F2-22, ***S5-64***
□*sp*, session parameter H3-3, S5-36, S5-61, ***S5-64***, S6-5
□*stac*, set access matrix F2-23, ***S5-65***
□*stie*, share-tie F2-4, F2-13 *to* F2-15, ***S5-65***
□*stop*, stop execution of user-defined function S5-33, S5-38, S5-49, ***S5-65***
□*svc*, shared variable access control SV2-4, ***SV3-16*** *to* ***SV3-22***, S4-2, S5-66
□*svn*, shared variable clone ID ***SV3-7*** *to* ***SV3-12***, S4-2, S5-66
NSVP fuctionality SV4-4
□*svo*, shared variable offer SV2-4, ***SV3-1*** *to* ***SV3-5***, S4-2, S5-2, S5-67
NSVP fuctionality SV4-5
□*svq*, shared variable query ***SV3-12*** *to* ***SV3-13***, S4-2, S5-67
NSVP fuctionality SV4-6
□*svr*, shared variable retraction ***SV3-5*** *to* ***SV3-6***, S4-2, S5-67
□*svs*, shared variable state ***SV3-14*** *to* ***SV3-16***, S5-67
□*SYSVER*, system version ***S5-67***
□*tie*, exclusive-tie F2-4, F2-13 *to* F2-15, ***S5-68***
□*trace*, trace execution of a user-defined function S5-33, S5-38, S5-49, ***S5-68***
□*trap*, event trap S5-69, S9-3, ***S9-6*** *to* ***S9-11***
□*ts*, time stamp ***S5-69***
□*twsid*, termination workspace identification H2-7, S3-6, ***S5-70*** *to* ***S5-71***
□*ul*, user load ***S5-72***
□*untie*, untie tied files F2-4, F2-15, ***S5-72***
□*USERID*, UNIX user name ***S5-72***
□*vi*, verify input S5-24, ***S5-72***
□*wa*, work area available ***S5-72***
□*ws*, workspace information ***S5-73*** *to* ***S5-76***, S6-5, S6-20

○ (circle); ***See*** Pi times, circle functions L4-13
⊤ (down tack); ***See*** encode, representation L4-23
| (stile); ***See*** magnitude, absolute value, residue, remainder L4-9, S5-17
⍉ (circle-back-slash); ***See*** transpose L4-69, S5-17
↓ (down arrow); ***See*** drop, raze L4-52
⍴ (rho); ***See*** shape, reshape L4-46
⊃ (right shoe); ***See*** enclose, link L4-58
⌈ (up stile); ***See*** ceiling, maximum L4-12, S5-17
÷ (division sign); ***See*** reciprocal, divide L4-7
⍙ (delta-underbar) F2-3, L2-1
⊢ (right tack, dex); ***See*** right, pass L4-73
↑ (up arrow); ***See*** take, nub L4-49
! (shriek, exclamation point, bang); ***See*** factorial, out-of, combinations L4-10
& (ampersand); ***See*** select, merge L5-6
' quote mark L6-3
. (dot); ***See*** alternat, inner-product, ply, tie L5-15
: (colon); ***See*** label L6-9, S7-8
@ (at); ***See*** from, all L4-31, L4-78
_ (underbar) F2-3, L2-1
{ (left brace); ***See*** deprecated primitives
} (right brace); ***See*** deprecated primitives
⍵ (omega) L2-1
⍺ (alpha) L2-1
¨ (diersis); ***See*** with L5-15
/ (slash); ***See*** f-reduce, copy L5-6
⌿ (slash-bar); ***See*** f-reduce-down, copy-down L5-6
⍝ (lamp); ***See*** comment
∆ (delta) F2-3, L2-1
∘ (jot), exclude none L4-35, L4-51
⍤ (paw); ***See*** on, rank, cut L5-15
⍥ (hoof); ***See*** upon, under L5-15
⊂ (left shoe); ***See*** swap L5-6
⍀ (back-slash-bar); ***See*** f-scan-down, expand-down L5-6
← (left arrow); ***See*** assignment
`.saxif` IF description file IF2-3
`.sf` UNIX suffix for component files F1-4

.sw UNIX suffix for workspace files F1-4
\ (back-slash); ***See*** expand, f-scan
¯ (high minus); ***See*** negative number
◇ (diamond separator) L6-2
1 *ap124*, public workspace A4-3, A4-20, ***A4-29***, S8-1
1 *del*, public workspace S8-1
1 *hostap*, public workspace ***A3-2***, S8-1
1 *hsp*, public workspace S8-1
1 *if*, public workspace ***IF1-2***, S8-1
1 *pc108*, public workspace A3-5, S8-1
1 *socket*, public workspace ***IF5-11***, S8-1
1 *ssql*, public workspace S8-1
1 *toolkit*, public workspace S8-2
1 *tools*, public workspace A2-9, A3-2, S8-2
filehelper function F2-6, F2-24
timen function F2-20
1 *unix*, public workspace ***IF5-2***, ***IF5-10***, S8-2
1 *wssearch*, public workspace S8-2
1 *wstransfer*, public workspace S8-2
1 *xfer*, public workspace S8-2

A

absolute value (primitive function |) ***L4-9***
access matrix F2-23 *to* F2-24, S5-57, S5-65
passnumber F2-3 *to* F2-4
access table S3-8
accounting information (system function □*ai*) ***S5-10***
account number F1-1 *to* F1-3, F2-4, S1-2, S3-3 *to* S3-8
activate a workspace
(system command)*load*) S5-57, S5-64, S6-13
(system command)*xload*) S6-13
active workspace F1-1 *to* F1-3, L1-7, S3-3 *to* S3-8
addition (primitive function +) ***L4-6***
adminstrator; UNIX system administrator H1-2, F2-8, F3-9
⍫ (del-tilde); ***See*** lock definition
□*ai*, accounting information ***S5-10***
all (primitive function @) L1-17, ***L4-78***
alpha (⍺) L2-1
alternat (dyadic operator .) ***L5-20***
ampersand (monadic operator &); ***See*** select, merge
and (primitive function ∧) ***L4-17***
AP1, APL (S-task) AP H2-7, ***A2-1*** *to* ***A2-10***, S1-6, S2-1, S4-3
prefix codes A2-3
AP11, Host AP H4-7, ***A3-1*** *to* ***A3-11***, S1-6, S3-4, S4-3
commands for closed files A3-4
commands on opened files A3-6
return codes A3-3
system facility commands A3-10
AP124, full-screen AP H4-7, ***A4-1*** *to* ***A4-29***, S1-6, S4-3
commands A4-20 *to* A4-28
communication A4-7 *to* A4-12
differences under UNIX and OS/390 A4-28
format table A4-16
keyboard A4-5
keys H4-4
mapping fields A4-12
reading and writing fields A4-5
return codes A4-10
session manager control A4-2 *to* A4-7
AP daemon A1-1
APL H2-7
APL/SV SV2-1
APL and UNIX naming conventions F1-3 *to* F1-4, S1-4
array L1-8
character support H4-2 *to* H4-6, L3-8, S7-4
postscript printing H5-1 *to* H5-8
X Window installation H4-7
component files F1-1 *to* F1-3, F2-1 *to* F2-28, S1-4, S3-1 *to* S3-9
concepts F2-2 *to* F2-6
See also components
utilities F1-4, F3-1 *to* F3-19

Dictionary of APL L2-5
files ***overview*** F2-1
file system F1-1 *to* F1-5, F2-1 *to* F2-28, S1-4, S3-3 *to* S3-8
functions for access control F2-23 *to* F2-24
functions for concurrency F2-24 *to* F2-27
functions for description and limits F2-18 *to* F2-22
functions for existence F2-10 *to* F2-13
functions for inquiry F2-8 *to* F2-10
functions for linkage F2-13 *to* F2-15
functions for reading and writing F2-15 *to* F2-18
font
See APL character support S7-4
language interpreters H3-5 *to* H3-7, L1-1, S1-1
language ***overview*** L1-1 *to* L1-17
primitive functions, keyword substitutes H4-6
primitives ***L1-2***
rules of syntax ***L1-10***
sharing with non-APL programs SV2-2
statement L1-1 *to* L1-17
storage facility S3-1 *to* S3-9
symbols or tokens L1-1 *to* L1-17
table of APL symbols L1-3
table of APL symbols(by description) L1-4
table of APL symbols(syntactic classes) L1-6
tasks H2-2, S1-3, S2-1
See also task
terminate APL session (system command `)off`) H2-7, S6-14
workspaces ***overview*** S1-3
See also workspace
See also SHARP APL
`apl-8` language interpreter H1-3, H2-5, H3-5 *to* H3-7, H4-6, S1-1, S5-56, S6-16
APL expression ***L1-5***
`APLFONT` UNIX environment variable H3-2
`apl` language interpreter H1-3, H2-5, H3-5 *to* H3-7, S1-1
`APLPARMS` UNIX environment variable H3-2, H3-5, S3-5
`APLPATH` UNIX environment variable H3-2, H3-5 *to* H3-7
`APLSIZE_COLS` UNIX environment variable H3-2
`APLSIZE_ROWS` UNIX environment variable H3-2
`apl` utility function A2-9
`□append`/*`□appendr`*, append variable to tied file F2-16, ***S5-10***
application H2-3 *to* H2-5
definition L1-12
`□arbin`/*`□arbout`*, arbitrary input/output H3-6, H4-8, ***S5-11***
arbitrary input/output (system functions *`□arbin`*/*`□arbout`*) H3-6, H4-8, S5-11
argument, definition ***L1-7***
array L1-8, L3-1 *to* L3-21
boxed; ***See*** enclosed array
definition ***L1-1*** *to* ***L1-8***
elements, types of L3-5 *to* L3-12
empty array L3-21
enclosed array L1-8, L3-9, L4-56
functions for array manipulation L1-8
heterogeneous L3-10
higher-rank L3-14
homogeneous L3-10
nested array L1-8
position and spacing in an array L3-15
processing L1-17
rank and shape of L3-1 *to* L3-12
simple, open or unenclosed arrays L3-9
See also scalar, vector, matrix
arrow
down arrow (primitive function ↓); ***See*** drop, raze
left arrow (syntax symbol ←); ***See*** assignment L1-7
right arrow (syntax symbol →); ***See*** branch
up arrow (primitive function ↑); ***See*** take, nub
assignment (syntax symbol ←) L1-7, L2-1 *to* L2-5, L6-6
indexed assignment L2-4, ***L4-44***
indexed assignment, obsolescence L2-5

indirect assignment L2-3
multiple assignment L2-3
naming variables L2-1 *to* L2-5
reassigning L2-4
result L2-2
associated processors S1-6
□*NA* S1-1
asterisk (primitive function *); ***See*** exponential, power
at (primitive function @); ***See*** from, all
atomic vector (system variable □*av*) H3-3, L3-9, L3-17, S5-11 *to* ***S5-12***
attention/interrupt S7-11
ATTN key
See interrupt
auxiliary processor ***A1-1*** *to* ***A1-4***, S1-6
AP1 H2-7, ***A2-1*** *to* ***A2-10***, S1-6, S4-3
prefix codes A2-3
AP11 ***A3-1*** *to* ***A3-11***, S1-6, S3-4, S4-3
AP124 A4-1 *to* A4-29, S1-6, S4-3
AP daemon A1-1, S4-3
clone ID SV3-9
See also AP1, AP11, and AP124
□*av*, atomic vector H3-3, S5-11 *to* S5-12
□*avail*, file system availability F2-8, ***S5-13***
behaviour with/without Filserver F4-13
□*avm*, mainframe atomic vector H3-3, ***S5-12***
axis
bracket notation L5-7, L6-6
frame L4-22
specifier L5-7
trailing L4-22

B

back-slash (monadic operator \); ***See*** expand, f-scan
back-slash-bar (monadic operator ⍀); ***See*** f-scan-down
bang (primitive function !); ***See*** factorial, out-of, combinations
base-10 value (primitive function ⊥) ***L4-24***
□*bind*, intrinsic function bind IF1-1, IF3-1, IF3-6, ***S5-13*** *to* ***S5-14***
)*binds*, display bound intrinsic functions IF4-1, ***S6-5***
Boolean
arguments L4-25
numbers L3-6
partition L5-29
□*bounce*, terminate task H2-2, H2-7, H3-6, S5-4, S5-14
box (⊃); ***See*** enclose
boxes; ***See*** enclosed arrays
braces
left brace (primitive function {); ***See*** deprecated primitives
right brace (primitive operator }); ***See*** deprecated primitives
bracket
axis bracket notation L5-7, L6-6
bracket indexing (syntax symbols []) ***L4-42***, L6-6
branch (syntax symbol →) L6-2, S7-9
as grouping device L6-8
within execute L4-94
break AP11 command A3-10
BREAK key
See interrupt
BUFFERS Fileserver configuration parameter F4-10
build a package (system function □*pack*) ***S5-47***

C

canonical representation (system function □*cr*) S5-2, ***S5-15***, S7-11
caret (primitive function ∧); ***See*** and, least common multiple
cartesian coordinates L3-1

cases of dyadic operators L5-15
catenate (primitive function ,) ***L4-62***
  with axis specifier L5-7
catenate-down (primitive function ⍪) ***L4-62***
  with axis specifier L5-7
ceiling (primitive function ⌈) ***L4-12***
  fuzzy ceiling ***L4-13***
cell L1-12, L1-14, L3-2 *to* L3-5
  major cells L3-4
  repetition of cells in a frame L1-16
`cfcheck` component file utility F3-13 *to* F3-19
character
  APL character support H4-2 *to* H4-6, L3-8, S7-4
    X Window installation H4-7
  character grade down (primitive function ⍒) ***L4-39***
  character grade up (primitive function ⍋) ***L4-39***
  character input/output (system variable ⎕) H3-6, L3-7, S5-1 *to* S5-2, ***S5-4*** *to* ***S5-9***, S5-53, S5-56
  character set; ***See*** APL character support, APL symbols or tokens L3-8
  characters grouped into tokens L1-6
  character support
    backspace key S7-6
    escape key S7-5
    insert/replace modes S7-4
    overstrike mode S7-5
  National Use characters H4-5, HA-9 *to* HA-11, S7-4
  quotes delimit character constant L6-3
  representing a number L3-12
  text-delimiting characters S5-28
  weights L4-40
  See also data L3-12
circle-back-slash (primitive function ⍉); ***See*** transpose
circle-bar (primitive function ⊖); ***See*** rotate-down, reverse-down
circle-star (primitive function ⍟); ***See*** logarithm
circle-stile (primitive function ⌽); ***See*** rotate, reverse
`)clear`, clear active workspace S5-57, S5-64, S6-5
`clear ws` H2-3
clone ID SV3-3, SV3-7, S5-66
  of an auxiliary processor SV3-9
`close` AP11 command A3-7
colon (syntax symbol `:`) L6-9, S7-8
comma-bar (primitive function ⍪); ***See*** table, catenate-down
command
  command processor S2-1
  summary of system commands S6-2
  system commands S1-6, S6-1 *to* S6-21
comment (syntax symbol ⍝) L6-9
communication
  discipline of SV2-1 *to* SV2-5
comparison tolerance (system variable ⎕`ct`) L4-12, L4-37 *to* L4-38, ***S5-15***
complete separators L6-7
component files F1-1 *to* F1-3, F2-1 *to* F2-28, S3-1 *to* S3-9
  `.sf` UNIX suffix for component files F1-4, S3-7
  checking F3-13 *to* F3-19
  compacting F3-9
  concepts F2-2 *to* F2-6
  sorting F3-10 *to* F3-13
  tie numbers F2-4
  utilities F1-4, F3-1 *to* F3-19
components F2-1, F2-19 *to* F2-28, S1-4, S5-57 *to* S5-58
  as variables F2-2
  component number F2-2
`COMPONENTS` Fileserver configuration parameter F4-11
composition operators L5-21
compress (monadic operator ⌿); ***See*** copy
compress-down (monadic operator ⍀); ***See*** copy-down
concealed suffix F1-4, S3-7
conditional enclose (⊃); ***See*** enclose
configuration
  Fileserver installation and configuration F4-8 *to* F4-12
  SHARP APL configuration H3-1 *to* H3-9
conjugate (primitive function +) ***L4-6***

constant
quotes delimit character constant L6-3
representation of L3-12
side-by-side constants L5-5
continue workspace H2-7
copula (assignment symbol ←); ***See*** assignment
)*copy*, copy objects into workspace S5-64, ***S6-7***
copy (monadic operator /) ***L5-11***, L5-14
copy-down (monadic operator ⌿) ***L5-11***, L5-14
count (primitive function ⍳) ***L4-35***
⎕*cr*, canonical representation S5-2, ***S5-15***, S7-11
⎕*create*, create a file F2-10, ***S5-15***
⎕*ct*, comparison tolerance L4-12, L4-37 *to* L4-38, ***S5-15***
cube; ***See*** power
cut (dyadic operator ⍤) L5-28 *to* L5-32

D

daemon
AP daemon A1-1
NSVP daemon SV4-2
data H2-3, L3-1 *to* L3-21
composition of L1-6
differentiating character from numeric L3-13
displaying L3-12 *to* L3-21
file is collection of data F2-1
input/output L3-12 *to* L3-21
See also variable
deal (primitive function ?) ***L4-22***
decode (primitive function ⊥) ***L4-24***
decorators S5-30
defining functions and variables S7-8
definition mode S7-7 *to* S7-11
del (∇) editor H2-6, L4-92, S5-15, S5-21, S5-41, S7-11 *to* S7-13
comments L6-10
delay execution (system function ⎕*dl*) ***S5-18***
delete a workspace (system command)*drop*) S6-8
delta (name symbol ∆) F2-3, L2-1
delta-underbar (name symbol ⍙) F2-3, L2-1
deprecated primitives: braces L1-17
derived function L1-7, L1-11
`detachsax` script H3-8
`detach` utility (detach APL session) H3-8
dex (primitive function ⊢); ***See*** right, pass
diamond separator (syntax symbol ⋄) L6-2
Dictionary of APL L2-5
diersis (dyadic operator ⍤); ***See*** with
directory S1-4, S3-1, S3-4 *to* S3-9, S6-13
format S3-6, S5-73
home directory F2-10, F2-14, ***S3-3***, S5-36
`pwd`, Show Working Directory (UNIX command) S3-5
root directory F2-8, F2-10, S3-1
subdirectory S3-1
disclose (primitive function >) L1-9, ***L4-59***
display
display bound intrinsic functions (system command)*binds*) IF4-1, S6-5
display IF definition (system command)*ifdef*) IF4-2, S6-9
displaying data L3-12 *to* L3-21
display library numbers and directory paths (system command)*libs*) S3-4, ***S6-13***
display saved workspace names (system command)*lib*) ***S6-12***
`DISPLAY` UNIX command line option H2-1
display variable names (system command)*vars*) ***S6-20***
folding wide displays L3-19 *to* L3-20
of variables L3-15
position and spacing in display (system L4-56, L4-78
position and spacing in display (system variable ⎕*ps*) L3-15, ***S5-54***

scissoring wide displays H3-4, L3-16 *to* L3-19, S5-56, S6-16
distinguished name ***S1-2***, S5-1, S7-8
divide (primitive function ÷) ***L4-7***
`⎕dl`, delay execution ***S5-18***
`domain error` L4-7, L4-20, L4-26, L4-91, S5-20 *to* S5-21, S5-38, S5-49
domino (primitive function ⌹); ***See*** matrix inverse, matrix divide
dot (dyadic operator `.`); ***See*** alternat, inner-product, ply, tie
`)drop`, delete a workspace ***S6-8***
`⎕drop`, drop components F2-18, F3-9, ***S5-18***
drop (primitive function ↓) ***L4-52***, L5-14
dyadic L1-11
cases of dyadic operators L5-15
function S5-4
operator L1-7, L1-9, L5-1, ***L5-15*** *to* ***L5-32***, S5-2
See also niladic and monadic S5-1
table of dyadic operators L5-2

E

`⎕ec`, environment condition S5-18, S9-1, ***S9-6***
editor L1-13, S7-1
(∇) editor H2-6, S5-21, S5-41, ***S7-11*** *to* ***S7-13***
comments L6-10
execute function limitation ***L4-92***
full-screen editor S5-19, S5-38, S5-66, ***S7-13*** *to* ***S7-22***
`)edit` H3-3, L1-13, S6-8, ***S7-14***
activation S7-14
APL session versus editor screen S7-15
commands S7-15 *to* S7-22
errors S7-20
`⎕EDIT`, system editor S5-19, S7-14
lock definition (line editor symbol ∀) S7-13, S9-5
`EDITSIZE` UNIX environment variable H3-2
e exponential symbol L3-5
element
definition ***L1-8***
enclosed L3-9
identity element L4-2
numeric elements L3-5 *to* L3-7
types of array elements L3-5 *to* L3-12
See also array, matrix
`emacs` H3-7
emergency exit H2-8
empty frame L4-4
emulators H4-1
enclose
(primitive function <) L3-9, ***L4-56***
conditional enclose (primitive function ⊃) ***L4-58***
enclosed array L3-9, L4-56
encode(primitive function ⊤) ***L4-23***
ending a session H2-7
environment condition (system variable `⎕ec`) S5-18, ***S9-6***
environment variables H3-1 *to* H3-9
epsilon (primitive function ∈); ***See*** member
epsilon-underbar (primitive function ⍷); ***See*** in
equal
equals (primitive function =) ***L4-25***
equal-underbar (primitive function ≡); ***See*** match
not equals (primitive function ≠) ***L4-27***
`⎕er`, event report S5-19, ***S9-5***
`)erase`, erase global objects ***S6-8***
`⎕erase`, erase a tied file F2-4, F2-11, ***S5-19***
error
interpreter errors S9-2
signalled errors S9-3
See `domain error`
See `file index error`
See `rank error`
See `value error`
escape key S7-5
Ethernet H4-10

euro symbol L3-8
evaluated input/output (system variable □) H3-6, L3-12, S5-2, ***S5-4** to **S5-9***, S5-35, S5-74
event ***S9-2***
  event report (system variable □*er*) S5-19, S9-5
  handling S1-8, ***S9-1** to **S9-13***
  numbers S9-1
    listed S9-3
  signal event S5-63, S7-9, S9-3, S9-12
  trapping S1-8, S5-7, S5-18, S5-66
  trapping (system variable □*trap*) S5-69, ***S9-6** to **S9-11***
    untrappable events S9-3
  unreportable or untrappable S9-3
□*ex*, expunge objects S5-4, ***S5-19***
eXceed H4-8
exclamation point (primitive function !); ***See*** factorial, out-of, combinations
excluded index selection L4-34
exclusive-tie (system function □*tie*) F2-13 *to* F2-15, ***S5-68***
execute (primitive function ⍎) ***L4-91***, S5-35, S5-63, S5-74, S5-76
  branch within ***L4-94***
  error conditions ***L4-94***
  evaluation of compound statement ***L4-93***
  occurrence in the state indicator ***L4-93***
execution
  order of L1-10
  stack S2-6
exit
  emergency exit H2-8
expand (monadic operator \) ***L5-11***
expand-down (monadic operator ⍀) ***L5-11***
experimental system names S5-1
explicit output S5-7
exponential
  (primitive function *) ***L4-7***
  exponential numbers (numeric symbol *e*) L3-5
expression ***L1-5***
expunge
  names S5-22
  package expunge (system function □*pex*) ***S5-50***

F

factorial (primitive function !) ***L4-10***
□*fc*, format control L4-90, ***S5-20***
`fcntl()` F1-4
fcntl AP11 command A3-9
`fcreatesw` startup option H3-4, S3-8
□*fd*, function definition ***S5-20** to **S5-23***, S5-33, S5-38, S5-66, S7-11
□*fhold*, hold tied files F1-4 *to* F1-5, ***F2-24** to **F2-27***, S5-23
□*fi*, input format conversion ***S5-23***
file
  access control F2-3 *to* F2-4
  accessing non-APL files F1-3
  APL component files F2-1 *to* F2-28, S1-4
    concepts F2-2 *to* F2-6
    utilities F1-4, F3-1 *to* F3-19
  APL file system F1-1 *to* F1-5, F2-1 *to* F2-28, S3-3 *to* S3-8
    functions for access control F2-23 *to* F2-24
    functions for concurrency F2-24 *to* F2-27
    functions for description and limits F2-18 *to* F2-22
    functions for existence F2-10 *to* F2-13
    functions for inquiry F2-8 *to* F2-10
    functions for linkage F2-13 *to* F2-15
    functions for reading and writing F2-15 *to* F2-18
  compaction F2-28
  file locking semaphore information F3-3
  file system S1-4
  file system availability (system function □*avail*) F2-8, ***S5-13***

File System Server; ***See*** Fileserver
interrupt S9-2
library (system function □*lib*) F2-8, ***S5-35***
names F2-3, F2-11
naming conventions H1-4, S3-6
read file information F2-20, ***S5-57***
size limit F2-3, F2-11
tieing a file F2-4
transfer H4-11
UNIX file permissions S3-8
See also Fileserver
file access error F2-6
filehelper function F2-6, F2-24
file index error F2-17 *to* F2-18
FILENAMES Fileserver configuration parameter F4-12
Fileserver F1-5, F4-1 *to* F4-16
overview F4-1
abort F4-5
activating the Fileserver F4-5
clearing the transition semaphore F4-6
configuration parameters F4-10
connecting the interpreter F4-7
fsios server processes F4-3
fsutil command F4-3
installation and configuration F4-8 *to* F4-12
list of advantages over conventional system F1-5
modes (Online, Offline, Abort, Inconsistent) F4-6
operations F4-3 *to* F4-8
restart actions F4-7
saxfs.cnf configuration file F4-3, F4-8
saxfs.install install script F4-8
shutting down F4-5
UNIX characteristics F4-13 *to* F4-16
user privileges F4-14
fix
a definition (system function □*fx*) ***S5-32***, S5-38, S5-66
fixed-point fraction S5-29
floating-point numbers L3-6
flockers component file utility F1-4, F3-7 *to* F3-9
behaviour with/without Filserver F4-16
floor (primitive function ⌊) ***L4-12***
fuzzy floor ***L4-13***
□*fmt*, format output L3-16 *to* L3-21, L6-8, ***S5-24*** *to* ***S5-31***
)*fns*, display function names S6-9
folding wide displays L3-19 *to* L3-20
fonts
See APL character support
Terminal Emulators package H4-11 *to* H4-12
format
control (system variable □*fc*) ***S5-20***
dyadic format (primitive function ⍕) ***L4-84***
format control (□*fc*) L4-90
format output (system function □*fmt*) L3-16 *to* L3-21, L6-8
input format conversion (system function □*fi*) S5-23
monadic format (primitive function ⍕) L3-15 *to* L3-16, ***L4-78***
output (system function □*fmt*) ***S5-24*** *to* ***S5-31***
frame L1-12, L1-14, L3-2 *to* L3-5
extending the empty frame L4-4
frame axes L4-22
f-reduce (monadic operator ⌿) ***L5-8***
f-reduce-down (monadic operator ⍀) ***L5-8***
from (primitive function @) L1-17, ***L4-31***, L5-14
f-scan (monadic operator \) ***L5-8***
f-scan-down (monadic operator ⍀) ***L5-8***
FS in transition or being queried Fileserver message F4-6
fsios Filerserver processes F4-3
□*fsm*, full-screen manager ***S5-32***
fsutil Filerserver command F4-3
fsync startup option H3-4
full-screen
editor H2-6, H3-3, L1-13, S5-19, S5-38, S5-66, S6-8,

S7-13 to S7-22
activation S7-14
APL session versus editor screen S7-15
commands S7-15 *to* S7-22
errors S7-20
manager (system function ⎕*fsm*) ***S5-32***
function H2-3, F1-2, L1-7, L1-9, L4-1 *to* L4-94, L5-25, S2-1, S2-5
access control system functions F2-23 *to* F2-24
arithmetic functions, non-scalar L4-20 *to* L4-25
arithmetic functions, scalar L4-5 *to* L4-20
categories of primitive functions L4-1 *to* L4-5
categories of system functions S5-76
defining functions S7-8
definition (system function ⎕*fd*) S5-20 *to* S5-23, S5-33, S5-38, S5-66, S7-11
derived from composition operators L5-23
derived functions L1-7, L1-11, L5-1 *to* L5-32
display function names (system command)*fns*) S6-9
hybrid functions S1-6, S5-39 *to* S5-41
indexing functions L4-31 *to* L4-45
inquiry system functions F2-8 *to* F2-10
linkage system functions F2-13 *to* F2-15
miscellaneous functions L4-73 *to* L4-94
modified by axis bracket notation L5-7
modified functions L1-14
monadic selection functions L5-14
primitives L1-2, L4-1 *to* L4-94, S5-2, S5-12, S7-11
primitives dependent on comparison tolerance S5-17
reading and writing system functions F2-15 *to* F2-18
relational functions L4-25 *to* L4-31
scalar L4-4 *to* L4-5
See also intrinsic function
structural functions L4-45 *to* L4-73
system functions for concurrency F2-24 *to* F2-27
system functions for description and limits F2-18 *to* F2-22
system functions for existence F2-10 *to* F2-13
system variables and functions S1-2, S5-1 *to* S5-77
table of system functions S5-2
used with axis specifier L5-7
user-defined functions H2-4, L1-12 *to* L1-13, L1-17, L2-5, L6-2, S1-6, S5-20 *to* S5-23, S5-33, S5-38 *to* S5-41, S5-48, S5-65 *to* S5-66, S5-68, S6-8 *to* S6-9, S7-8, S7-11
utility functions S8-1 *to* S8-2
See also intrinsic functions
See also primitive functions and variables
Fuzzy
floor and ceiling L4-13
GCD and LCM L4-18
⎕*fx*, fix a definition ***S5-32***, S5-38, S5-66

G

GCD; ***See*** greatest common divisor
get AP11 command A3-4
getenv AP11 command A3-9
`gethostid ()` system call F3-4
glyphs
See APL symbols or tokens
grade
character grade down (primitive function ⍒) ***L4-39***
character grade up (primitive function ⍋) ***L4-39***
numeric grade down (primitive function ⍒) ***L4-38***
numeric grade up (primitive function ⍋) ***L4-38***
greater than (primitive function >) ***L4-27***
greater than or equal (primitive function ≥) ***L4-27***
greatest common divisor (primitive function ∨) ***L4-17***
fuzzy GCD ***L4-18***

H

HDS108 H4-10
□*hold*, hold tied files F1-4, ***F2-24*** *to* ***F2-27***, S5-23
`holders` component file utility F1-4, F3-5 *to* F3-7
`HOLDS` Fileserver configuration parameter F4-11
hold tied files (system function □*fhold*) F1-4 *to* F1-5, ***F2-24*** *to* ***F2-27***, S5-23
home
  home directory F2-10, F2-14, ***S3-3***, S5-36
  home library ***S3-4*** *to* ***S3-5***
hoof (dyadic operator ⍤); ***See*** upon, under
horizontal tab (system variable □*ht*) ***S5-33***
host utility function A3-2 *to* A3-3, A3-5 *to* A3-6
`HQUEUES` Fileserver configuration parameter F4-12
□*ht*, horizontal tab ***S5-33***
Hummingbird H4-8
hybrid functions S1-6, S5-39 *to* S5-41
  See also user-defined functions

I

identification in use SV3-9
identifiers, name class of S5-4, S5-41, S5-51
identity (primitive function +) ***L4-6***
IF ; ***See*** intrinsic function
IFbind IF3-6
`ifcc` IF utility IF2-8 *to* IF2-10
IFdef IF3-9
)*ifdef*, display IF definition/redefine IF IF4-2, ***S6-9***
IFerase IF3-10
)*iferase*, remove IF definitions IF4-3, ***S6-10***
`ifgen` IF utility IF2-8 *to* IF2-10
IFlist IF3-11
)*iflist*, list IF groups and their contents IF4-4, ***S6-11***
IFrep IF3-8
immediate-execution mode H1-3, H2-3, L1-1, L6-1, S1-2 *to* S1-3, S6-1, S7-7 *to* S7-11
implications of network sharing SV3-11
implicit iteration L1-16
in (primitive function ∈) ***L4-30***
independence bit (AP1) A2-7
index
  bracket indexing (syntax symbols `[]`) ***L4-42***, L6-6
  excluded index selection L4-34
  index (primitive function ⍳) ***L4-37***
  indexed assignment L2-4, L4-44
    obsolescence L2-5
  indexing functions L4-31 *to* L4-45
  index of (primitive function ⍳) ***L4-36***
  index origin (system variable □*io*) L4-19, L4-22, L4-31 *to* L4-45, L4-65, L4-73, ***S5-34***
  indicies; enclosed L4-32
  sign-on index SV3-4, SV3-9
  index error L4-32
indirect assignment L2-3
infix notation L1-11
`INIT` UNIX environment variable H3-2
inner-product (dyadic operator `.`) ***L5-18***
input/output L3-12 *to* L3-21, S7-2 *to* S7-7
  character input/output, ⍞ (quote-quad) H3-6, L3-7, S5-1 *to* S5-2, ***S5-4*** *to* ***S5-9***, S5-53, S5-56
  evaluated input/output, □ (quad) H3-6, L3-12, S5-2, ***S5-4*** *to* ***S5-9***, S5-35, S5-74
input format conversion (system function □*fi*) ***S5-23***
input interrupt S5-7, S7-11, S9-2
insert/replace S7-4
installation
  Fileserver installation and configuration F4-8 *to* F4-12
  X Window APL font installation H4-7
integer L3-5
interlocked system variable S5-63
international characters H4-5, HA-9 *to* HA-11

interpreter
connection to Fileserver F4-7
errors S9-2
See, apl, apl-8, rtapl, rtapl-8 H3-5
See also APL language interpreters
interrupt S9-2
attention S7-11
current line S7-11
hard and soft interrupts S7-11
input interrupt S5-7, S7-11, S9-2
shared variable reference (*sv interrupt*) S7-11
signalling an interrupt(ATTN or BREAK key) S7-6
weak interrupt L6-3
intrinsic function ***IF1-1 to IF1-3***
building intrinsic functions IF2-8 *to* IF2-10
display bound intrinsic functions (system command)*binds*) ***IF4-1***, S6-5
display IF definition (system command)*ifdef*) ***IF4-2***, S6-9
IF system commands IF1-3, ***IF4-1 to IF4-4***
included in package IF2-13
list IF groups and their contents (system command)*iflist*) ***IF4-4***, S6-11
message queue intrinsic functions IF5-2 *to* IF5-4
redefine IF (system command)*ifdef*) ***IF4-2***, S6-9
remove IF definition (system command)*iferase*) ***IF4-3***, S6-10
semaphore intrinsic functions IF5-4 *to* IF5-7
shared memory intrinsic functions IF5-7 *to* IF5-10
socket intrinsic functions IF5-11
supplied intrinsic functions IF1-3, ***IF5-1 to IF5-11***
system intrinsic functions IF1-2, ***IF3-1 to IF3-15***
UNIX system intrinsic functions IF5-1
user-defined intrinsic functions ***IF2-1 to IF2-14***
intrinsic function bind (system function □*bind*) IF1-1, IF3-1, IF3-6, ***S5-13 to S5-14***
□*io*, index origin L4-19, L4-22, L4-31 *to* L4-45, L4-65, L4-73, ***S5-34***
ioctl AP11 command A3-8
IOSERVERS Fileserver configuration parameter F4-11
iota (primitive function ⍳); ***See*** count, index of
□*run* S2-2, S2-4
is assigned (syntax symbol ←); ***See*** assignment
ISO standard APL L5-7
item; ***See*** scalar

J

jot (∘), exclude none L4-35, L4-51

K

Kermit H1-5, H4-10
KEYBASE Fileserver configuration parameter F4-12
keyboard
AP124 keys H4-4, ***A4-5***
APL escape key S7-5
APL overstrike mode S7-5
backspace key S7-6
interrupt S7-11, S9-2
mapped for X Window system H4-3
Meta key H4-2
overstrike combinations H4-3
See APL character support
session manager keys H4-4
union keyboard H4-2
keyword substitutes H4-6
kill, Signal a Process (UNIX command) H2-7, S5-14, S9-3

L

label (syntax symbol :) L6-9, S7-8
lamp (syntax symbol ⍝) L6-9
language; ***See*** APL language S1-1
latent expression (system variable ⎕*lx*) S5-1, ***S5-37***, S5-62, S5-64, S6-13
⎕*lc*, line counter ***S5-35***
LCM; ***See*** least common multiple
least common multiple (primitive function ∧) ***L4-17***
  fuzzy LCM ***L4-18***
left (primitive function ⊣) ***L4-74***
length error L4-46
length of axis L3-2
less (primitive function ~) ***L4-77***
less than (primitive function <) ***L4-56***
  See also enclose
less than or equal (primitive function ≤) ***L4-27***
lev (primitive function ⊣); ***See*** stop, left
)*lib*, display saved workspace names ***S6-12***
⎕*lib*, file library F2-8, ***S5-35***
library H1-4, F1-1 *to* F1-3, L1-12, S1-4, S3-1, S3-3 *to* S3-9
  display library numbers and directory paths (system command)*libs*) S3-4, S6-13
  file library (system function ⎕*lib*) F2-8, ***S5-35***
  format S3-6, S5-73
  home library ***S3-4*** *to* ***S3-5***
  library 1 S8-1 *to* S8-2
  mapping S3-5
  names of tied files (system function ⎕*names*) F2-9, S5-41
  numbers S3-4, S6-13
  of workspaces H2-3
  private libraries S3-4, S6-13
  public libraries S8-1 *to* S8-2
  workspace libraries S1-4
)*libs*, display library numbers and directory paths S3-4, ***S6-13***
limits, hardware and operating system S2-5
line, the unit of execution L1-7, L6-1 *to* L6-3
line counter (system function ⎕*lc*) ***S5-35***
link (primitive function ⊃) L3-14, ***L4-58***
list; ***See*** vector
list IF groups and their contents (system command)*iflist*) IF4-4, S6-11
)*load*, activate a workspace F1-4, S5-57, S5-64, ***S6-13***
⎕*load*, load a workspace ***S5-36***, S5-64
localization of names L2-5 *to* L2-6
lock
  functions S5-23
  interlocked system variable S5-63
  package lock S5-50
lock AP11 command A3-10
lock definition (line editor symbol ⍒) S7-13, S9-5
  See also del (∇) editor
logarithm (primitive function ⍟)
  base-⍺ ***L4-8***
  natural ***L4-8***
look-ahead L5-5
lseek AP11 command A3-8
⎕*lx*, latent expression S5-1, ***S5-37***, S5-62, S5-64, S6-13

M

magnitude (primitive function |) ***L4-9***
mainframe atomic vector (system variable ⎕*avm*) H3-3, ***S5-12***
maintenance
  NSVP maintenance SV5-1
  SVP maintenance SV5-2, SV5-4
match (primitive function ≡) ***L4-28***
mathematical notation
  APL resembles mathematical notation L1-1
matrix L3-1 *to* L3-12

access matrix F2-4, F2-23 *to* F2-24, S5-57
access table S3-8
divide (primitive function ⌹) ***L4-20***
inverse (primitive function ⌹) ***L4-20***
See also scalar, vector, array
`MAXFILE` Fileserver configuration parameter F4-10
maximum (primitive function ⌈) ***L4-13***
`MAXPROC` Fileserver configuration parameter F4-10
member (primitive function ∈) ***L4-28***
merge (monadic operator ⍏) ***L1-17***, ***L5-14***
message queue intrinsic functions IF5-2 *to* IF5-4
Meta key H4-2
⎕*MF*, monitoring facility S5-1, S5-33, ***S5-37*** *to* ***S5-39***, S5-49
Microsoft; ***See*** MS Windows, MS-DOS, Windows NT H4-9
minimum (primitive function ⌊) ***L4-12***
minus
high minus (numeric symbol ¯) L3-5, L3-13
minus (primitive function −) L3-13, ***L4-6***
See also negative number
miscellaneous functions L4-73 *to* L4-94
mode
APL overstrike mode S5-9
definition mode S7-7 *to* S7-11
immediate execution mode L6-1 *to* L6-3
immediate-execution mode S1-2, S6-1, S7-7 *to* S7-11
insert mode S7-4
modes of operation H1-3
replace mode S7-4
modified functions L1-14
monadic L1-11
function S5-4
operator S5-2
operators L1-7, L1-9, L5-1, L5-6 *to* L5-15
See also dyadic and niladic S5-1
selection functions L5-14
table of monadic operators L5-2
monitoring facility (system function ⎕*MF*) S5-1, S5-33, ***S5-37*** *to* ***S5-39***, S5-49
MS-DOS H4-10
MS Windows H4-9
multiple assignment L2-3
multiplication (primitive function ×) ***L4-7***

N

⎕*NA*, name association ***S1-6***, ***S5-39*** *to* ***S5-41***
name S5-1
naming SHARP APL user objects L2-1 *to* L2-6
APL and UNIX naming conventions F1-3 *to* F1-4, S1-4
assigned to line (label) L6-9, S7-8
assigning names to variables L2-1 *to* L2-5
association (system function ⎕*NA*) ***S1-6***, ***S5-39*** *to* ***S5-41***
defined functions denoted by arbitrary names L1-14
display function names (system command)*fns*) S6-9
distinguished name ***S1-2***
distinguished names S7-8
expunge names S5-22
file and workspace naming conventions S3-6
is a token L1-6
library names of tied files (system function ⎕*names*) F2-9, ***S5-41***
localization of names L2-5 *to* L2-6
multiple names for same value L2-4
name class by state indicator level S5-76
name class of identifiers (system function ⎕*nc*) S5-4, ***S5-41***, S5-51
name list (system function ⎕*nl*) ***S5-42***, S5-73, S6-20
name symbols L2-1
naming conventions for files and workspaces H1-4
package name class (system function ⎕*pnc*) ***S5-51***

package names (system function □*pnames*) ***S5-51***
pathnames of tied files F2-10, S5-49
permitted file names F2-3, F2-11
permitted workspace names F2-3
reassigning L2-4
See also shared name
shadowed L2-6
sharing SV2-1
termination workspace name S5-70
within a package L1-9, L3-11
workspace S6-2
See also name symbols: delta, delta-underbar, alpha, omega, underbar
□*names*, library names of tied files F2-9, ***S5-41***
nand (primitive function ⍲) ***L4-18***
National Use characters H4-5, HA-9 *to* HA-11, S7-4
See APL character support
□*nc*, name class of identifiers S5-4, ***S5-41***, S5-51
negate (primitive function −) ***L4-6***
negative number L3-4 *to* L3-5
network shared variable processor (NSVP) A1-2, SV1-1 *to* SV1-2, ***SV4-1*** *to* ***SV4-6***, S1-5, S4-1 *to* S4-4
daemon SV4-2
□*svo*, □*svq*, and □*svn* function extensions SV4-4
maintenance SV5-1
newline character L6-1
niladic
function S5-4
See also dyadic and monadic S5-1
□*nl*, name list ***S5-42***, S5-73, S6-20
non-scalar function L4-20 *to* L4-25
nor (primitive function ⍱) ***L4-18***
no shares effecting programs via UNIX scripts H3-9
not (primitive function ~) ***L4-19***
notation
APL resembles mathematical notation L1-1
axis bracket notation L5-7
infix notation L1-11
not equals (primitive function ≠) ***L4-27***
See also nubsieve
NSVP (network shared variable processor) A1-2, SV1-1 *to* SV1-2, SV3-11, ***SV4-1*** *to* ***SV4-6***, S1-5, S4-1 *to* S4-4
daemon SV4-2
□*svo*, □*svq*, and □*svn* function extensions SV4-4
maintenance SV5-1
N-task H2-2, H2-7, A2-8, F2-4, S1-3, S2-2 *to* S2-3, S5-59 *to* S5-62, S5-64
nub (primitive function ⍬) ***L4-76***
nubin (primitive function =) ***L4-76***
nubsieve (primitive function ≠) ***L4-75***
□*nums*, tie numbers of tied files F2-9, ***S5-42***
number
account number F2-4
Boolean numbers L3-6
component number F2-2
floating-point numbers L3-6
integer L3-5
internal representation of numbers L3-6
is token L1-6
negative number L3-4 *to* L3-5
real number L3-5
represented by a character L3-12
storage of numbers L3-6
types of numbers L3-5 *to* L3-7

O

)*off*, terminate APL session H2-7, ***S6-14***
omega (⍵) L2-1
on (dyadic operator ⍠) ***L5-21*** *to* L5-22
open (primitive function >) ***L4-59***
open AP11 command A3-7
operator L1-7, L5-1 *to* L5-32
composition operators L5-21
consecutive operators L5-4

dyadic operators L5-15 *to* L5-32
modifying functions L1-14
monadic operators L5-6 *to* L5-15
precedence of operators L5-3
primitives L1-2, L5-1 *to* L5-32
syntactic classes of operators L5-3
table of dyadic operators L5-2
treatment of L1-11
valence L5-3
options
session startup options H3-2
"suit yourself" startup options H3-2, H3-4
terminal startup options H3-7
or (primitive function ∨) ***L4-17***
order of execution L1-10
OS/390
AP124 differences under UNIX and OS/390 A4-28
⎕*out*, record session I/O to file ***S5-43*** *to* ***S5-47***
out (primitive function !) ***L4-11***
outer-product; ***See*** tie
output vector L6-8
overstrike
character combinations H4-3
mode S7-5
owner access control; ***See*** file access control, access matrix

P

package L1-9, L3-11 *to* L3-12
build a package (system function ⎕*pack*) L3-11, ***S5-47***, ***S5-53***
data may also be L1-6
package define (system function ⎕*pdef*) L3-11, ***S5-49***
package expunge (system function ⎕*pex*) L3-11, ***S5-50***
package insert (system function ⎕*pins*) L3-11, ***S5-50***
package lock (system function ⎕*plock*) L3-11, ***S5-50***
package name class (system function ⎕*pnc*) L3-11, ***S5-51***
package names (system function ⎕*pnames*) L3-11, ***S5-51***
package select (system function ⎕*psel*) L3-11, ***S5-56***
package value (system function ⎕*pval*) L3-11, ***S5-56***
protective package define (system function ⎕*ppdef*) L3-11, ***S5-53***
parentheses L1-7, L6-4
recall previous line (system command)) S6-4
redundant parentheses L6-5
⎕*PARSE*, pre-parse user-defined function ***S5-48***
pass (primitive function ⊢) ***L4-73***
passnumber F2-3 *to* F2-4
path H1-4, A3-10, F1-3 *to* F1-4, S1-4, S3-1
mapping S3-5
pathname F2-10, S3-1
⎕*paths*, pathnames of tied files F2-10, ***S5-49***
paw (dyadic operator ⍤); ***See*** on, rank, cut
pawn (primitive function ⍎); ***See*** execute S5-35, S5-63, S5-74, S5-76
PC108 H1-5, H4-10
`1 pc108`, public workspace H4-11
)*pcopy*, protected copy ***S6-7***
⎕*pdef*, package define L3-11, ***S5-49***
permission
file system functions with permissions F2-7
permission table; ***See***access matrix
UNIX file permissions S3-8
See file access control, access matrix
permission code F2-4
⎕*pex*, package expunge ***S5-50***
⎕*pins*, package insert L3-11, ***S5-50***

Pi times (primitive function ○) ***L4-13** to **L4-14***
□*lock*, package lock L3-11, ***S5-50***
plus (primitive function +) ***L4-6***
ply (dyadic operator .) ***L5-18***
□*pnames*, package names L3-11, ***S5-51***
□*pnc*, package name class L3-11, ***S5-51***
position and spacing in display (system variable □*ps*) L3-15, L4-56, L4-78, ***S5-54***
postscript support H5-1 *to* H5-8
power (primitive function ⋆) ***L4-8***
□*pp*, print precision L4-78, ***S5-52***
□*ppdef*, protective package define L3-11, ***S5-53***
□*PR*, prompt replacement S5-1, ***S5-53***
precedence rules; ***See*** APL rules of syntax
primitive ***L1-2***
  deprecated functions and variables L1-17
  functions L1-2, L4-1 *to* L4-94, S1-2, S5-2, S7-11
  operators L1-2, L5-1 *to* L5-32
  table of APL symbols L1-3
  table of APL symbols(by description) L1-4
  table of APL symbols (syntactic classes) L1-6
  variables S1-2, S7-11
printing utility H5-1 *to* H5-8
printing width (system variable □*pw*) H4-8, L3-16 *to* L3-21, S5-11, ***S5-56***, S6-4 *to* S6-5, S7-2
print precision (system variable □*pp*) L4-78, ***S5-52***
processor S4-1 *to* S4-4
  associated processor S1-6
  associated processor, auxiliary processor, and shared variable processor S1-1
  auxiliary processor (AP) ***A1-1** to **A1-4***, S1-6
  ID SV3-2
  network shar S1-5
  NSVP ***SV4-1** to **SV4-6***
  number SV3-3
  SVP A1-2, SV1-1 *to* SV1-2, S1-5, S4-1 *to* S4-4
profile H1-4, H3-2, S3-5, S7-2
program,***See*** function H2-3, L1-7
prompt
  discarding the prompt S5-9
  input on same line as character prompt S5-8
  replacement (system variable □*PR*) S5-1, ***S5-53***
protected copy (system command)*pcopy*) ***S6-7***
protective package define (system function □*ppdef*) ***S5-53***
□*ps*, position and spacing in display L3-15, L4-56, L4-78, ***S5-54***
□*psel*, package select L3-11, ***S5-56***
pseudorandom number L4-23
public libraries S8-1 *to* S8-2
public workspaces
  1 *ap124* A4-3, S8-1
    utility functions A4-20, ***A4-29***
  1 *del* S8-1
  1 *hostap* ***A3-2***, S8-1
  1 *hsp* S8-1
  1 *if* ***IF1-2***, S8-1
  1 *pc108* A3-5, S8-1
  1 *socket* ***IF5-11***, S8-1
  1 *ssql* S8-1
  1 *toolkit* S8-2
  1 *tools* A2-9, A3-2, S8-2
  1 *unix* ***IF5-2***, IF5-10, S8-2
  1 *wssearch* S8-2
  1 *wstransfer* S8-2
  1 *xfer* S8-2
punctuation; ***See*** syntax
put AP11 command A3-5
putenv AP11 command A3-9
putu utility function A3-5
□*pval*, package value L3-11, ***S5-56***
□*pw*, printing width H4-8, L3-16 *to* L3-21, S5-11, S5-56, S6-4 *to* S6-5, S7-2
pwd, Show Working Directory (UNIX command) S3-5

Q

□*qload*, quiet load ***S5-36***
quad (□) H3-6, L3-12, S5-2, S5-4 *to* S5-9, S5-35, S5-74
qualifying suffix F1-4, S3-7
query (primitive function *?*); ***See*** roll, deal
question mark (primitive function *?*); ***See*** roll, deal
quiet load (system function □*qload*) ***S5-36***
`quit` character H2-8
quote mark (syntax symbol ') L6-3
quote-quad (⍞) H3-6, L3-7, S5-1 *to* S5-2, ***S5-4*** *to* ***S5-9***

R

random link (system variable □*rl*) L4-19, L4-22, ***S5-58***
rank L1-14, L3-1 *to* L3-12
  argument ranks of functions L4-1, S5-4
  infinite L3-5, L4-2
  of derived function L5-21
  rank (dyadic operator ⍤) L4-2, ***L5-25***, S5-4
  result rank L4-4
  undecided L4-2
rank error L4-20, L4-38, SV3-11
ravel (primitive function ,) ***L4-61***
raze (primitive function ↓) ***L4-53***
□*rdac*, read access matrix F2-23, ***S5-57***
□*rdci*, read component information F2-19, ***S5-57***
□*rdfi*, read file information F2-20, ***S5-57***
read
  access matrix S5-57
  read component information (system function □*rdci*) F2-19, ***S5-57***
  read component of a tied file (system function □*read*) F2-16, ***S5-58***
  read file information (system function □*rdfi*) F2-20, ***S5-57***
read AP11 command A3-8
real number L3-5
recall previous line (system command]) ***S6-4***
reciprocal (primitive function ÷) ***L4-7***
recommended setting H2-5
record session I/O to file (system function □*out*) ***S5-43*** *to* ***S5-47***
recursion L1-14
redefine IF (system command]*ifdef*) IF4-2, S6-9
reduce; ***See*** f-reduce, f-reduce-down
relational functions L4-25 *to* L4-31
remainder (primitive function |) ***L4-9***
remote no shares SV3-11
remove IF definition (system command]*iferase*) IF4-3, S6-10
□*rename*, rename a tied file F2-12, ***S5-58***
□*replace*, replace variable stored in a component F2-17, F3-9, ***S5-58***
report
  event report (system variable □*er*) S5-19, ***S9-5***
  report running tasks (system function □*runs*) ***S5-63***
report shape and rank; shape (primitive function ⍴) L3-2
representation (primitive function ⍕) ***L4-23***
reset file size limit (system function □*resize*) F2-3, F2-21, ***S5-58***
reshape (primitive function ⍴) ***L4-46***
`RESIDENT` Fileserver configuration parameter F4-12
residue (primitive function |) ***L4-9***
□*resize*, reset file size limit F2-3, F2-21, ***S5-58***
result rank L4-4
reverse
  reverse (primitive function ⌽) ***L4-66***, L5-7
  reverse-down (primitive function ⊖) ***L4-66***, L5-7
  reverse-down with axis specifier L4-66
review of system requirements H1-1
rho (primitive function ⍴); ***See*** shape, reshape

right (primitive function ⊢) ***L4-74***
right shoe (⊃); ***See*** enclose, link
⎕*rl*, random link L4-19, L4-22, S5-58
roll (primitive function ?) ***L4-19***
root
  directory F2-8, F2-10, S3-1
  expression S5-68
  `root` (superuser) F4-3, SV5-4
  See also power
rotate
  rotate (primitive function ⌽) ***L4-67***, L5-7
  rotate-down (primitive function ⊖) ***L4-68***, L5-7
  with axis specifier L4-67
row-major order
  See also ravel
`rtapl-8` language interpreter H3-5 *to* H3-7, S1-1
`rtapl` language interpreter H3-5 *to* H3-7, S1-1
rules of APL syntax L1-10
⎕*run*, start N-task H2-2, F2-4, S1-3, S2-2, S2-4, ***S5-59*** *to* ***S5-62***
⎕*runs*, report running tasks ***S5-63***

S

SAMI SV4-2
sapl.retcode system intrinsic function IF3-1
sapl.singleref system intrinsic function IF3-2
sapl.translate system intrinsic function IF3-3
sapl.type system intrinsic function IF3-4
sapl.version system intrinsic function IF3-5
)*save*, save active workspace S5-64, ***S6-15***, S6-20
save
  saved workspace F1-1 *to* F1-3, L1-7, ***S1-3***
  saving the execution stack S2-6
  saving workspace containing shared variables SV2-5
`sax` script H1-2 *to* H1-4, H2-1, H3-1 *to* H3-9
`SAXBRACELOG` UNIX environment variable H3-2, ***L1-18***
`SAXCNF` UNIX environment variable H3-2
`SAXDIR` UNIX environment variable H3-2
`saxfs.cnf` Filerserver configuration file F4-3, F4-8
`saxfs.install` Fileserver install script F4-8
`SAXGROUP` Fileserver configuration parameter F4-11
`.saxif` IF description file IF2-3
`saxrc` script H1-5
`SAXRC` UNIX environment variable H3-2
`saxscript` H3-6
`SAXUSER` Fileserver configuration parameter F4-11
⎕*sc*, shared variable state change SV2-5, ***SV3-22*** *to* ***SV3-24***, S4-2, S5-63
scalar L3-1 *to* L3-12
  drop applied to L4-53
  scalar function L4-4 *to* L4-20
  singleton treated as L4-5
  take applied to L4-52
  See also array, vector, matrix
scan; ***See*** f-scan, f-scan-down
scissoring wide displays H3-4, L3-16 *to* L3-19, S5-56, S6-16
`scissor` startup option H3-4
scripts
  `detachsax` H3-8
  `nsvp` SV4-2
  `sax` H1-2 *to* H1-4, H2-1, H3-1 *to* H3-9
  `saxrc` H1-5
  `saxscript` H3-6
  shell script H1-4
scrolling the session log S7-3
select (monadic operator &) ***L1-17***, ***L5-14***
semaphore
  clearing the transition semaphore F4-6
  file locking semaphore information F3-3
semaphore intrinsic functions IF5-4 *to* IF5-7
semicolon (syntax symbol `;`) L6-7
  in vector L6-8

sentence; ***See*** statement
separator L6-3
  complete separators L6-7
  diamond (syntax symbol ◇) L6-2
)*sesm*, session manager conditions H3-4, H4-6, ***S6-16***
`sesmtrans` startup option H3-4
session
  detached SHARP APL session H3-8
  ending a session H2-7
  □*pw* as session variable S5-57
  log H1-3, H2-6, S7-2
    reusing a line S7-3
  manager H1-3, H2-5, H4-4, S5-56, S7-2 *to* S7-7
    control by AP124 A4-2
  session manager conditions (system command)*sesm*) H3-4, H4-6, ***S6-16***
  session parameter (system variable □*sp*) H3-3, S5-36, S5-61, ***S5-64***, S6-5
  srolling the session log S7-3
  starting a session H1-2 *to* H1-3, H2-1
  startup options H3-2
  terminate APL session (system command)*off*) H2-7, S6-14
set access matrix (system function □*stac*) F2-23, ***S5-65***
set symbol table size (system command)*symbols*) S6-19
set workspace name (system command)*wsid*) ***S6-20***
shape (primitive function ⍴) L1-15, L3-1 *to* L3-12, ***L4-44***, ***L4-46***
shared memory intrinsic functions IF5-7 *to* IF5-10
shared name; ***See*** shared variable
shared variable
  access control (system function □*svc*) SV2-4, ***SV3-16*** *to* ***SV3-22***, S4-2, S5-66
  clone ID (system function □*svn*) ***SV3-7*** *to* ***SV3-12***, S4-2, S5-66
    NSVP fuctionality SV4-4
  communication SV1-1 *to* SV1-2, S1-5
  discipline of communication SV2-1 *to* SV2-5
  duration of sharing SV3-5
  interrupt S9-2
  network shared variable processor (NSVP) S1-5, S4-1 *to* S4-4
  network sharing SV3-11
  offer (system function □*svo*) SV2-4, ***SV3-1*** *to* ***SV3-5***, S4-2, S5-2, S5-67
    NSVP fuctionality SV4-5
  query (system function □*svq*) ***SV3-12*** *to* ***SV3-13***, S4-2, S5-67
    NSVP fuctionality SV4-6
  retraction (system function □*svr*) ***SV3-5*** *to* ***SV3-6***, S4-2, S5-67
  shared names (origin of) SV2-1
  shared variable processor (SVP) A1-2, SV1-1 *to* SV1-2, S1-5, S4-1 *to* S4-4
    maintenance SV5-2, SV5-4
    variables and functions SV3-1 *to* SV3-24
  shared variables vs shared names SV2-1
  sharing with non-APL programs SV1-2
  state (system function □*svs*) ***SV3-14*** *to* ***SV3-16***, S5-67
  state change (system variable □*sc*) SV2-5, ***SV3-22*** *to* ***SV3-24***, S4-2, S5-63
  used to control AP124 screen A4-8
share-tie (system function □*stie*) F2-4, F2-13 *to* F2-15, ***S5-65***
sharing with non-APL programs SV2-2
SHARP APL
  language overview L1-1 *to* L1-17
  as UNIX filter H3-6, S5-56
  configuration H3-1 *to* H3-9
  profile H3-2
  system overview S1-1
  system requirements H1-1
  using SHARP APL H2-1 *to* H2-8
  without a terminal H3-8
  See APL

shell AP11 command A3-11
shell script H1-4
SHMADDR Fileserver configuration parameter F4-12
shoe
  left shoe (monadic operator ⊂); ***See*** swap
  right shoe (⊃); ***See*** enclose, link
shriek (primitive function !); ***See*** factorial, out-of, combinations
)*si*, state indicator ***S6-18***
)*sinl*, state indicator with name list ***S6-18***
)*sic*, state indicator clear ***S6-18***, S7-9
side-by-side constants L5-5
sift component file utility F1-4, F2-28, F3-9
  behaviour with/without Filserver F4-16
□*signal*, signal event S5-63, S7-9, S9-1, ***S9-3***, ***S9-12***
signalled errors S9-3
signalling an interrupt (ATTN or BREAK key) S7-6
sign-on index SV3-9
signum (primitive function ×) ***L4-6***
singleton L4-5
  See also scalar
size
  active workspace size S2-5
  file size limit F2-11
  □*size*, size of tied file F2-22, ***S5-64***
  size AP11 command A3-6
slash (monadic operator /); ***See*** f-reduce, copy
slash-bar (monadic operator ⌿); ***See*** f-reduce-down
socket facility (IFs) IF5-11
sortio component file utility F1-4, F3-10 *to* F3-13
□*sp*, session parameter H3-3, S5-36, S5-61, ***S5-64***, S6-5
specifying rank L5-25
square root; ***See*** power
□*stac*, set access matrix F2-23, ***S5-65***
star (primitive function ⋆); ***See*** exponential, power
start
  start N-task (system function □*run*) H2-2, F2-4, S1-3, ***S5-59*** *to* S5-62
STARTAPS UNIX environment variable H3-2
starting a session H1-2 *to* H1-3, H2-1
startup options H3-2, H3-7
S-task ***A2-1*** *to* ***A2-10***, S1-3, S2-1, S5-53, S5-64
  duration A2-2
  freestanding A2-7
state
  indicator S5-74, S5-76
  shared variable state SV3-14 *to* SV3-16, S5-67
  state indicator (system command)*si*) ***S6-18***
  state indicator clear (system command)*sic*) S6-18, S7-9
  state indicator with name list (system command)*sinl*) ***S6-18***
statement L1-1 *to* L1-17, L6-2
  grouping and separators L6-3
□*stie*, share-tie F2-4, F2-13 *to* F2-15, ***S5-65***
stile
  (primitive function |); ***See*** magnitude, absolute value, residue, remainder
  del-stile (primitive function ⍒); ***See*** grade
  delta-stile (primitive function ⍋); ***See*** grade
  down stile (primitive function ⌊); ***See*** floor, minimum
  up stile (primitive function ⌈); ***See*** ceiling, maximum
□*stop*, stop execution of user-defined function S5-33, S5-38, S5-49, ***S5-65***
stop (primitive function ⍱) ***L4-74***
strong interrupt S9-2
structural functions L4-45 *to* L4-73
subtraction (primitive function −) ***L4-6***
suffixes
  .sf UNIX suffix for component files F1-4, S3-7
  .sw UNIX suffix for workspace files F1-4, S3-7
"suit yourself" startup options H3-4
summary of system variables and functions S5-76
superuser (root)
  See also UNIX system administrator
□*svc*, shared variable access control SV2-4, ***SV3-16*** *to*

SV3-22, S4-2, S5-66
sv interrupt S7-11
□*svn*, shared variable clone ID ***SV3-7** to **SV3-12***, S4-2, S5-66
NSVP fuctionality SV4-4
□*svo*, shared variable offer SV2-4, ***SV3-1** to **SV3-5***, S4-2, S5-2, S5-67
NSVP fuctionality SV4-5
SVP (shared variable processor) A1-2, SV1-1 *to* SV1-2, S1-5, S4-1 *to* S4-4
maintenance SV5-2, SV5-4
variables and functions SV3-1 *to* SV3-24
`svpdump` command SV5-2
`svpmaint` command SV5-4
□*svq*, shared variable query ***SV3-12** to **SV3-13***, S4-2, S5-67
NSVP fuctionality SV4-6
□*svr*, shared variable retraction ***SV3-5** to **SV3-6***, S4-2, S5-67
□*svs*, shared variable state ***SV3-14** to **SV3-16***, S5-67
`.sw` UNIX suffix
See workspace S6-2
swap (monadic operator ⊂) ***L5-13***
symbol table full S5-62
)*symbols*, set symbol table size ***S6-19***
symbols
See APL character support
table of APL symbols L1-3
table of APL symbols(by description) L1-4
table of APL symbols(syntactic classes) L1-6
syntax L1-6, L6-1 *to* L6-10
rules of APL syntax L1-10
syntax error L3-14
system
administrator F2-8
commands S1-6, S6-1 *to* S6-21
limitation S6-1
editor (system function □*EDIT*) S5-19, S7-14
facilities S7-1
overview S1-2
intrinsic functions IF1-2, IF3-1 *to* IF3-15
modes of operation H1-3
names S5-1
overview S1-1
profile S3-5, S7-2
summary of system commands S6-2
system version (system function □*SYSVER*) ***S5-67***
table of system functions S5-2
table of system variables S5-2
variables and functions S1-2, S5-1 *to* S5-77
X Window systems H4-7 *to* H4-9
system.bindrep system intrinsic function IF3-8
system.bind system intrinsic function IF3-6
system.def system intrinsic function IF3-9
system.erase system intrinsic function IF3-10
system.list system intrinsic function IF3-11
system administrator H1-2, F3-9
□*SYSVER*, system version ***S5-67***

T

table
AP124 format table A4-16
See also matrix
set symbol table size (system command)*symbols*) ***S6-19***
table (primitive function ⍪) ***L4-61***
tack
(primitive function ⍎) , execute S5-35, S5-63, S5-74, S5-76
(primitive function ⍎), execute ***L4-91***
down tack (⊤); ***See*** encode, representation
left tack (primitive function ⊣); ***See*** stop, left
right tack (primitive function ⊢); ***See*** right, pass
up tack (⊥); ***See*** decode, base-10 value
take (primitive function ↑) ***L4-49***, L5-14

talk utility function A2-9
task H2-2, S1-1, S1-3, S2-1
  N-task A2-8, S1-3
  S-task ***A2-1** to **A2-10***, S1-3
    duration A2-2
    freestanding A2-7
  terminate task (system function ⎕*bounce*) H2-2, H2-7, H3-6, S5-4, S5-14
  T-task A2-1, S1-3
TCP/IP H4-10
TEMPPATH UNIX environment variable H3-2
TERMCAP UNIX environment variable H3-2
terminal
  initialization H3-7
  startup options H3-7
  support H1-5, H3-7, H4-1 *to* H4-14
  terminal emulators H4-10
  Terminal Emulators package H4-2
  types H4-1
terminate APL session (system command)*off*) H2-7, ***S6-14***
terminate task (system function ⎕*bounce*) H2-2, H2-7, H3-6, S5-4, ***S5-14***
termination workspace identification (system function ⎕*twsid*) H2-7, S3-6, S5-70 *to* S5-71
TERMINFO UNIX environment variable H3-2
tessellation ***L5-28** to **L5-32***
text-delimiting characters S5-25, S5-28
thorn (primitive function ⍭); ***See*** format
tie
  exclusive-tie (system function ⎕*tie*) F2-4, F2-13 *to* F2-15, ***S5-68***
  number F2-4
  share-tie (system function ⎕*stie*) F2-4, F2-13 *to* F2-15, ***S5-65***
  tie (dyadic operator .) L5-15
  tieing a file F2-4
  tieing file with a passnumber F2-5
  untie tied files (system function ⎕*untie*) F2-4, F2-15, ***S5-72***
tiers component file utility F1-4, F3-2 *to* F3-5
timen function F2-20
times (primitive function ×) ***L4-7***
time stamp (system function ⎕*ts*) ***S5-69***
Token Ring H4-10
tokens L1-6 *to* L1-7
  six syntactic classes L1-6
⎕*trace*, trace execution of a user-defined function S5-33, S5-38, S5-49, ***S5-68***
trailing axis L4-22
transliteration H3-4, H4-6, S6-16
transpose
  (primitive function ⍉) L4-55
  dyadic transpose (primitive function ⍉) ***L4-71***
  monadic transpose (primitive function ⍉) ***L4-69***
⎕*trap*, event trap S5-69, ***S9-6** to **S9-11***
  untrappable events S9-3
⎕*ts*, time stamp ***S5-69***
T-task A2-1, S1-3, S2-1, S5-14, S5-53, S5-64
⎕*twsid*, termination workspace identification H2-7, S3-6, ***S5-70** to **S5-71***

U

⎕*ul*, user load ***S5-72***
umask, Mask Default File Permissions (UNIX command) H2-5, A3-5, S3-8
under (dyadic operator ⍢) ***L5-21** to* L5-23
underbar (name symbol_) F2-3, ***L2-1***
underline: ***See*** underbar (name symbol_) L2-1
underscore: ***See*** underbar (name symbol_) L2-1
UNIX
  AP124 differences under UNIX and OS/390 A4-28
  APL and UNIX naming conventions F1-3 *to* F1-4, S1-4
  environment variable S3-5

file S3-1
Fileserver characteristics under UNIX F4-13 *to* F4-16
filter H3-6, S5-56
`kill` command H2-7, S5-14, S9-3
path A3-10, F1-3 *to* F1-4, S1-4, S3-1
profile H1-4
read/write access F2-5
root directory F2-8
shell script H1-4
system administrator H1-2, F2-8, F3-9
system function facility (IFs) IF5-1
unix.errno system intrinsic function IF3-11
unix.filesync system intrinsic function IF3-12
⎕*untie*, untie tied files F2-4, F2-15, ***S5-72***
upon (dyadic operator ⍤) ***L5-21***
`usefilelocks` startup option H3-4
user
access control; ***See*** file access control, access matrix
Fileserver user privileges F4-14
naming SHARP APL user objects L2-1 *to* L2-6
user-defined functions L1-12 *to* L1-13, L1-17, L2-5, L6-2, S1-6, S5-20 *to* S5-23, S5-33, S5-38, S5-48, S5-65 *to* S5-66, ***S5-68***, S6-8 *to* S6-9, S7-8
See also hybrid functions
user-defined intrinsic functions ***IF2-1 to IF2-14***
user load (system function *ul*) ***S5-72***
user name (system function ⎕*USERID*) ***S5-72***
⎕*USERID*, UNIX user name ***S5-72***
utilities
APL file utilities F3-1 *to* F3-19

V

valence L1-11
of monadic operators and dyadic operators L5-3
value error S9-5
variable H2-3, F1-2, L3-1 *to* L3-21, S2-1, S2-5
assigning names to variables L2-1 *to* L2-5
boxed variable; ***See*** enclosed array
categories of system variables S5-76
components as variables F2-2
constants L3-12
defining variables S7-8
displaying variables L3-15
display variable names (system command)*vars*) ***S6-20***
interlocked system variable S5-63
⎕*pw* as session variable S5-57
mixed data L3-14
See also shared variables
shape of an array L1-15
system variables and functions S1-2, S5-1 *to* S5-77
table of system variables S5-2
UNIX environment variable (`APLPARMS`) S3-5
See also primitive functions and variables
)*vars*, display variable names ***S6-20***
vector L3-1 *to* L3-12
output vector L6-8
use of semicolons in vectors L6-8
See also scalar, array, matrix
verify input (system function ⎕*vi*) S5-24, ***S5-72***
⎕*vi*, verify input S5-24, ***S5-72***
VT100/VT320 H4-10

W

⎕*wa*, work area available ***S5-72***
weak interrupt L6-3, S9-2
window; ***See*** X Window system, MS Windows H4-8
with (dyadic operator ⍠) ***L5-32***
workspace H2-3, L1-7, S1-3, S2-5, S3-1 *to* S3-9
)*save*, save active workspace S5-64, ***S6-15***, S6-20
`.sw` UNIX suffix for workspace files F1-4, S3-7

See also qualifying suffix S6-2
activate a workspace (system command)*load*) F1-4, ***S6-13***
activate a workspace, (system command)*load*) S5-57, S5-64
activate a workspace, (system command)*xload*) S5-64
active workspace H2-3, H2-7, F1-1 *to* F1-3, S1-3, S2-5, S3-3 *to* S3-8
clear active workspace (system command)*clear*) S5-57, S5-64, ***S6-5***
clear ws H2-3
continue workspace H2-7
copy (protected) objects into workspace (system command)*pcopy*) ***S6-7***
copy objects into workspace (system command)*copy*) S5-64, ***S6-7***
delete a workspace (system command)*drop*) S6-8
display saved workspace names (system command)*lib*) ***S6-12***
environment S5-75
identification S5-73
information S5-75
library of workspaces H2-3
load a workspace (system function ⎕*load*) ***S5-36***, S5-64
names F2-3, S6-2
naming conventions S3-6
public workspaces S8-1 *to* S8-2
saved workspace H2-3, F1-1 *to* F1-3, S1-3, S2-6, S3-3 *to* S3-8
  containing shared variables SV2-5
set workspace name (system command)*wsid*) ***S6-20***
size of active workspace S2-5
termination H2-7, S3-6
termination workspace identification (system function ⎕*twsid*) S5-70 *to* S5-71
termination workspace name S5-70
  workspace libraries S1-4
write AP11 command A3-8
⎕*ws*, workspace information ***S5-73*** *to* ***S5-76***, S6-5, S6-20
ws full S5-20, S5-62
)*wsid*, set workspace name ***S6-20***
WSSIZE UNIX environment variable H3-2
wys utility H5-1 *to* H5-8

X

)*xload*, activate a workspace S5-64, S6-13
xvalchk startup option H3-4
X Window APL font
X Window system H2-1, H4-7 *to* H4-9
  APL font H4-7
  emulator H4-8
  keyboard map for APL characters H4-3
  window resizing H4-8
  X server/X client H4-7 *to* H4-9
  xterm H4-7 *to* H4-9

SHARP APL *for UNIX*

Handbook

JUMP TO ...

CHAPTER 1. INTRODUCTION
CONTENTS
PREFACE
MASTER INDEX
USING THIS DOCUMENTATION

HANDBOOK
AUXILIARY PROCESSORS MANUAL
FILE SYSTEM MANUAL
INTRINSIC FUNCTIONS MANUAL
LANGUAGE GUIDE
SVP MANUAL
SYSTEM GUIDE

© 2000 by Soliton Associates Limited
44 Victoria Street, Suite 2100, Toronto, Ontario, Canada M5C 1Y2

Pubcode: UW-000-0401 (0007)

SHARP APL *for UNIX*

USER LIBRARY

Handbook

SOLITON ASSOCIATES

© 2000 by Soliton Associates Limited–Printed in Canada
44 Victoria Street, Suite 2100, Toronto, Ontario, Canada M5C 1Y2

All rights reserved. No part of this work may be reproduced in any form or by any means without prior written permission of the copyright owner.

Soliton Associates Limited shall not be liable for errors contained herein or for incidental consequential damages in connection with the use of this material. The information contained within this document is subject to change without notice.

Acknowledgments

SHARP APL is a registered trademark of Soliton Associates Limited.

SVP and NSVP are trademarks of Soliton Associates Limited.

IBM and AIX are registered trademarks of International Business Machines Corporation.

RISC System/6000 is a trademark of International Business Machines Corporation.

UNIX is a registered trademark of UNIX System Laboratories, Inc.

Sun Workstation is a registered trademark of Sun Microsystems, Inc.

Sun, Solaris, SunOS, Open Windows, and SunView are trademarks of Sun Microsystems, Inc.

All other products referred to in this document are trademarks or registered trademarks of their respective trademark holders.

Publication code: UW-000-0401 (0009)

Contents

Tables and Figures vii

Preface

Introduction ix
Chapter Outlines ix
Conventions x
Documentation Summary xi
Contacting Soliton Associates xi

1. Introduction

What Is SHARP APL? 1-1
Accessing SHARP APL 1-1
The `sax` Script 1-2
Modes of Operation 1-3
File Naming Conventions 1-4
Setting Up Your Environment 1-4
The `sax` Script and Your Profile 1-4
Terminal Support 1-5
Print Utility 1-5

2. Using SHARP APL

Starting an APL Session 2-1
The APL Task 2-2
The Interpreter and Immediate Execution 2-3
The Workspace 2-3
Running an Application 2-4

Loading an Application from the SHARP APL Prompt 2-4
Environment Issues 2-5
Setting Default File Access for Shared Files 2-5
Basic System Facilities 2-5
Ending an APL Session 2-7
Communications Failure 2-7
Terminating a Task 2-7
Emergency Exit 2-8

3. SHARP APL Environment
Environment Variables 3-1
Profile for SHARP APL 3-2
Session Startup Options 3-2
Language Interpreters 3-5
`apl` 3-5
`apl-8` 3-5
SHARP APL as a UNIX Filter 3-6
`rtapl` 3-7
Terminal Startup Options 3-7
Terminal Initialization 3-7
SHARP APL Without a Terminal 3-8
One Application of `detach` 3-9
SHARP APL for UNIX Logging Facility 3-9

4. Terminal Support
Introduction 4-1
Terminal Types 4-1
SHARP APL Character Support 4-2
Session Manager Keys 4-4
AP124 Keys 4-4
National Use Characters 4-5
Keyword Substitutes for APL Primitives 4-6
X Window Systems 4-7
X Window APL Font 4-7

Window Resizing 4-8
X Window Emulator 4-8
Using Terminal Emulators 4-10
Asynchronous Serial Line Connection 4-10
TCP/IP Connection 4-10
File Transfer 4-11
`egaset` APL Font 4-11
Support for Color/Highlighting in AP124 4-12
Terminal Emulators Installation and Configuration 4-12

5. Printing
APL PostScript Printing 5-1
The `wys` Utility 5-1
Installed Files 5-1
Installing a Postscript Error Handler 5-2
Accessing the `wys` Script 5-3
Operation 5-3
Environment Variables 5-4
Print Formats 5-7
`wys` Examples 5-8

Appendix A. National Use Characters A-1

Index IX-1

Tables and Figures

Figure 2.1. Introductory messages. 2-2
Table 3.1. Environment variables for the `sax` script. 3-2
Table 3.2. Session startup options. 3-3
Table 3.3. “Suit Yourself” options. 3-4
Table 3.4. Terminal startup options. 3-7
Figure 4.1. SHARP APL keyboard map for X Window system. 4-3
Table 4.1. APL overstrike combinations. 4-3
Table 4.2. Transliteration keywords. 4-6
Table 5.1. `wys` files. 5-2
Table 5.3. `wys` print formats. 5-7
Table A.1. Ordinary characters. A-2
Table A.2. National Use characters. A-3

Preface

Introduction

This document contains the operating instructions for SHARP APL for UNIX. It covers basic topics intended for novice users and, when necessary, provides expanded details for those who are more familiar with the system (i.e., system setup, terminal support, and printer support). Some of the information in this manual is contingent on your terminal/workstation or UNIX host system. Consult your system administrator if you are unsure about your hardware or operating system.

This manual does not supply the operating instructions for applications that run under SHARP APL. Rather than reproduce existing material, references to related SHARP APL publications are supplied where applicable.

Chapter Outlines

The SHARP APL for UNIX ***Handbook*** is organized into the chapters described below.

Chapter 1, "Introduction," provides a brief description of SHARP APL, lists the system requirements, then discusses system access and modes of operation.

Chapter 2, "Using SHARP APL," gets into more detail about starting and ending a session, loading and running applications, and using system facilities such as the session manager and full-screen editor.

Chapter 3, "Configuring SHARP APL," describes UNIX environment variables and startup options used to specify operational characteristics for your session.

Chapter 4, “Terminal Support,” provides instructions for setting up SHARP APL terminal emulation and character support.

Chapter 5, “Printing,” discusses PostScript printer support for APL characters using the `wys` utility.

Conventions

The following conventions are used throughout this documentation:

`⎕io←0`	Although the default value for `⎕io` in a clear workspace is *one*, all examples in this manual assume that the index origin is *zero*.
α and ω	Two APL symbols are used in syntax descriptions of monadic and dyadic functions; *alpha* (α) identifies the left argument and *omega* (ω) identifies the right argument.
`constant width`	Examples of non-APL input and output are represented in a `constant width` typeface.
default (%) prompt	Examples in this book assume you are using the C shell. The default prompt for the C shell is the *percent sign* (%).
passno	Argument variables appear in *sans-serif italic* type.
[ ]	Optional arguments are shown between square brackets. For example, in the following APL statement, the argument *passno* is optional:

```
'filename' ⎕stie tieno [,passno]
```

These default environment variables represent frequently used pathnames in SHARP APL for UNIX documentation and scripts:

```
$SAXLIB    →    /usr/sax
$SAXDIR    →    /usr/sax/rel
$SAXCNF    →    /usr/sax/local
$HOME      →    home directory of the current user.
```

Documentation Summary

The SHARP APL for UNIX reference documentation is published in several volumes. One or more of the following may be referenced in this document.

SHARP APL for UNIX,

- *Language Guide,* publication code UW-000-0802
- *System Guide,* publication code UW-000-0902
- *SVP Manual,* publication code UW-001-0501
- *File System Manual,* publication code UW-037-0501
- *Auxiliary Processors Manual,* publication code UW-033-0501
- *Intrinsic Functions Manual,* publication code UW-007-0501

For a complete list of SHARP APL publications, please refer to the ***Support*** link on the Soliton Associates website: `http://www.soliton.com.`

Contacting Soliton Associates

Problems or questions regarding your SHARP APL system or this guide should be directed to the Soliton Associates Technical Support group:

`support@soliton.com`

Requests for additional technical documentation should be directed to the Soliton Associates Sales group:

`sales@soliton.com`

1
Introduction

What Is SHARP APL?

SHARP APL is a rich and versatile computing environment comprising an APL language interpreter, a file system, and a suite of facilities for programming and implementing powerful software applications. It is used extensively in such diverse applications as database management, actuarial science, financial and statistical analysis, simulation, modeling, and countless other commercial and academic enterprises.

All computation is expressed in APL, the theoretical notation described in Kenneth E. Iverson's book *A Programming Language.* SHARP APL functions and operators are drawn from logic, matrix theory, linear algebra, numerical analysis and trigonometry.

SHARP APL extends the high-level capabilities of the APL language. It incorporates an integrated file system for storing and manipulating massive amounts of data. Shared variable technology and intrinsic functions enable SHARP APL to interface with a variety of environments and applications. Currently, SHARP APL is supported across OS/390 and UNIX platforms.

For complete reference information on SHARP APL for UNIX functions and commands refer to the *Language Guide* or the *System Guide*.

Accessing SHARP APL

Before you get started, review what you'll need to use SHARP APL for UNIX:

- SHARP APL, fully installed and configured on a UNIX workstation and most likely set up for multi-user network access. See "System Configuration" later in this chapter for more information.

- An account number for the UNIX host (i.e., where SHARP APL resides).
- A UNIX workstation or a PC terminal emulator that is connected to SHARP APL's UNIX host. Review "Chapter 4. Terminal Support", for more information.
- If you are going to write or display programs written in APL, your terminal should also be equipped for APL character support. Please review "Chapter 4. Terminal Support", for more information.
- A PostScript printer (optional). Please review "Chapter 5. Printing", for more information.
- Access to related SHARP APL documentation. Please review the "Documentation Summary" in the Preface.

Starting a SHARP APL Session

Once you are enrolled as a user on the UNIX host, and the necessary environmental parameters are set, you should be able to start a SHARP APL session by entering the following at the UNIX prompt:

```
% sax
```

This means that the `sax` ***shell script*** has already been set up by your system administrator(probably when SHARP APL was installed on your system).

The `sax` Script

A script is a collection of commands that are grouped into one executable UNIX file. Since the SHARP APL system files can be installed in almost any location on your system, it is convenient to have the command text, along with associated startup parameters, in a single script referenced in the system search path. This allows the user to invoke a SHARP APL session as a one-word command (i.e., `sax`).

Depending on your system, your system administrator can choose either of the following standard procedures to make the `sax` script available to the user:

- Link the script into a standard path (i.e., an entry in the `PATH` environmental variable)

  ```
  ln –s $SAXDIR/bin/sax /usr/bin/sax
  ```

- If there is a global user shell login file, you can define the following environmental values

```
setenv SAXDIR /usr/sax/rel
setenv PATH $SAXDIR/bin/:$PATH
```

  Substitute the path for `sax` on your system in the appropriate places..

If neither of these methods is implemented on your system, you can also define `SAXDIR` and `PATH` in your shell login script in your home directory.

Ending a SHARP APL Session

The SHARP APL session remains active until you exit using the following system command:

`)off`

The characteristics of a SHARP APL session are explained in greater detail in "Chapter 2. Using SHARP APL".

Modes of Operation

The default mode in which the system operates is called ***immediate execution;*** whenever you enter an instruction, the system immediately interprets and executes it. Instructions to the system may be in the form of system commands, system functions and variables, APL statements, or user-defined functions. Each active user is assigned an amount of storage called a ***workspace,*** which serves both as the place in which calculation occurs during the session and the environment in which names (functions, variables, etc.) are understood.

There are other input modes, such as those for editing rather than executing definitions, or for supplying input to programs; however, immediate execution is the default mode and the mode to which the system returns when work is completed.

Depending on your requirements, the system can be configured to start different versions of the language interpreter: `apl` (default) or `apl-8`. Other facilities are provided by the system to integrate the interpreter with its environment, both within the active workspace and beyond it. These include a ***session manager,*** a ***session log,*** and two types of editors.

More information on SHARP APL interpreters and system facilities is provided in "Chapter 2. Using SHARP APL".

File Naming Conventions

SHARP APL for UNIX offers two different ways to name a file or workspace. One is consistent with the naming conventions of earlier APL systems, and the other is consistent with current UNIX practice:

APL library style	The name is written in the form `123 data`
UNIX path style	The name is written in the form `/home/sue/data`

The APL system maps its numeric account and library information to the UNIX system's directory of names. Refer to the *System Guide, Chapter 3* for a complete description of APL workspaces, files, and libraries.

Setting Up Your Environment

The system adminstrator is responsible for the installation, configuration, and maintenance of the base SHARP APL system. However, various options related to terminal support, printing, and system utilities enable users to customize their individual envionments (depending on the system and any restrictions imposed by the system adminstrator).

The `sax` Script and Your Profile

By entering `sax` at the UNIX prompt, you automatically:

1. define environment variables (establishing the size of the initial clear workspace, the maximum size of a requested workspace, and the size of the file devoted to the session log)
2. start the AP-Daemon
3. invoke one of several SHARP APL interpreters and,
4. clean up any processes started during your SHARP APL session after you sign off.

This script also provides for a private file, your *profile*, in which you can store a UNIX shell script that customizes SHARP APL for your use. The name of the file in which you keep your customized settings is, by default, stored in the environment variable `SAXRC`, for which the default is `$HOME/.saxrc`.

More information on the startup script and profile is provided in "Chapter 3. SHARP APL Environment".

Terminal Support

SHARP APL for UNIX supports two basic types of terminals: terminals using (or emulating) the X Window system and PCs using PC108 or the Kermit VT100/VT320 terminal emulator. Although it is recommended that you use a terminal that is configured for ***APL character support***, several other terminal types can be used (e.g., Hyperterminal or `telnet`).

Terminal and APL character support options are described in detail in "Chapter 4. Terminal Support".

Print Utility

SHARP APL comes equipped with a PostScript Type 1 font and the tools to install it on a PostScript printer. The UNIX command command `wys` allows printing of files in a variety of fonts (including APL), point sizes, and formats. The `wys` command and PostScript printing are explained in "Chapter 5. Printing".

2
Using SHARP APL

When SHARP APL is running, it interposes itself between you and UNIX and controls all work. Although SHARP APL may also be used for automated (batch) processing, the fundamental style of work is interactive. All dialogue between you and an APL session is maintained by the SHARP APL session manager, interpreting entries from your keyboard (possibly as characters in the APL character set) and displaying output on your terminal screen.

This chapter describes a the characteristics of a typical SHARP APL session and explains some of the facilities associated with SHARP APL. For a detailed reference on the SHARP APL language (i.e.,primitive functions), please consult the *Language Guide*. System functions and facilities are documented in the *System Guide*.

Starting an APL Session

As described in Chapter 1, you begin a session in SHARP APL by entering

```
% $SAXDIR/bin/sax
```

or

```
% sax
```

from the UNIX prompt. The `sax` script invokes specific actions depending on how your system is set up and on the kind of terminal you are using,

- in X Window mode (with the `DISPLAY` environment variable set) a new window appears displaying the Soliton Associates copyright banner while the APL interpreter loads from disk; or,

- in other modes, the screen clears (the UNIX prompt disappears) then displays the Soliton Associates copyright banner while the APL interpreter loads from disk.

Refer to "Chapter 4. Terminal Support" for more information on the various modes for accessing SHARP APL for UNIX.

If you are granted access, your SHARP APL session opens with a few introductory messages: i.e., the task number of your APL session (UNIX process number), a date/time stamp, your user name, and the current SHARP APL for UNIX version (see Figure 2.1). Once you start an APL session it remains active until you decide to terminate it.

```
------------------------------------------------------------------------
   Copyright 2000, Soliton Associates Limited. All rights reserved
------------------------------------------------------------------------
45380) 2000-06-10 09:59:47 akhodge
SHARP APL for UNIX    Version 6.0.0
clear ws
     _
```

Figure 2.1. Introductory messages.

The task number (`45380` in the example above) is what you'll need to distinguish this APL task from another; for example, when you use AP1 or □*run* to start multiple APL tasks. It is also required to identify this task should you decide to terminate it using □*bounce* (see "Ending an APL Session" at the end of this chapter).

The APL Task

Any work carried out during an APL session is referred to as an APL ***task.*** As a user of the SHARP APL system, you may have several different tasks running at the same time. Usually only one of them is directly linked to the terminal. The different types of tasks are discussed in detail in the *System Guide, Chapter 2*.

The Interpreter and Immediate Execution

After session initiation, the cursor appears on a new line resting six blank spaces from the left margin (below the last message, `clear ws`). This is the SHARP APL system prompt indicating that a new interactive task has started and is now awaiting your input.

SHARP APL provides an ***interpreter*** for the instructions (APL ***statements***) you enter at the system prompt. The interpreter proceeds directly from APL statements to execution; there are no intermediate stages of compilation, assembly, or linkage. One of four versions of the language interpreter may be installed during SHARP APL startup: `apl`, `apl-8`, `rtapl`, and `rtapl-8` (see "Chapter 3. SHARP APL Environment", for more information).

The APL task is, by default, in ***immediate execution*** mode so entering a simple statement at the system prompt should cause the interpreter to evaluate it at once. Note that the system's response is ***not*** indented. This makes it easy to tell which line was your statement and which lines were the system's response (on the screen or in the session log).

Example:

```
      2÷3
0.6666666667
```

Because the system interprets each statement as you present it, there is a continual alternation between two phases: you submit a statement, and the system executes it. When execution is complete, the system displays the result (if any) and awaits your next statement. The execution environment is known as the ***active workspace*** (explained below).

The Workspace

In APL, a workspace serves as the environment in which definitions are understood and the working storage area in which calculations take place. All APL functions (programs) or variables (data) involved in a calculation must be present in the active workspace at the moment the calculation is performed.

The contents of the active workspace can be copied into a ***saved workspace*** for later use. When an APL application (a set of user-defined functions and variables) is developed, the workspace must then be saved so that it is available for later use. It is also possible to build a new application by copying functions from one or more

workspaces into a new saved workspace. A set of saved workspaces is called a ***library*** of workspaces.

Note: The message *clear ws* that is displayed when you begin your session indicates that no saved workspace is currently present in the execution environment. *ws* is an abbreviation for "workspace."

Running an Application

As mentioned in the previous section, an application is a set of related user-defined functions and variables that are saved together in a workspace ready for execution. Ordinarily, you run an application by selecting from a library the workspace you wish to use—however, some systems are configured to load a SHARP APL application directly from the UNIX prompt, or via a graphical user interface.

Note: Because this procedure varies from application to application, you will need to check with your SHARP APL system adminstrator for exact login details.

The following operating instructions are relevant to most SHARP APL for UNIX applications. User instructions for a particular application would be found in separate documentation written specifically for that application.

Loading an Application from the SHARP APL Prompt

You start a SHARP APL application by loading its saved workspace. This is done by typing *)load*, followed by the workspace identification:

)load reports ***or*** *)load* 123 *accounts* ***or*** *)load /usr/application/ws*

When you press ENTER, the system should display *saved* followed by the time and date the workspace was last saved. If you receive *ws not found* or *incorrect command*, check to see that you typed the command correctly, and try again. If you still get *ws not found*, it probably means that the library or workspace name is incorrect or no longer exists.

Once the workspace is loaded, the application contained in the workspace may be executed. If the application is designed to start automatically, follow the instructions issued by the program to begin to use it — this is usually the case for menu-driven applications such as the Soliton Associates' data management and reporting application, VIEWPOINT. If the application is not designed to start automatically, you will have to select and enter each function you wish to execute.

Environment Issues

If your application does not perform as expected, or the keyboard does not behave according to design, you might be running the wrong SHARP APL environment for that application. The appropriate UNIX terminal type and communications software combinations will be listed in the application's documentation. Further information can be found in "Chapter 4. Terminal Support".

Setting Default File Access for Shared Files

Some applications are designed to allow users to create and own files that they can then choose to share or not share with other users. It is important to understand that the default UNIX file access mode ***overrides*** any file sharing constraints specified by the user in an APL session. The recommended `umask` setting of `022` allows the creator read and write access and restricts access for other users to read-only. (Also note, there should be no spaces when you reference a `umask` setting.)

Refer to your UNIX documentation for more information on `umask`. The UNIX file permission can be set in one of the session startup options explained in "Chapter 3. SHARP APL Environment".

Basic System Facilities

SHARP APL is equipped with some basic facilities which, while not strictly speaking primitive functions of APL, are nevertheless supported as part of the environment: the session manager and session log, and two types of editors. The following system facilities are discussed in detail in the *System Guide, Chapter 7*.

The Session Manager

The session manager is responsible for the management of an interactive APL task. It is responsible for transferring output from the APL system to the device driver, for receiving and interpreting input from the device drivers, and for providing basic services to the user at a terminal (such as scrolling, editing, copying, and moving text). SHARP APL provides a built-in session manager with its default interpreter `apl`. It also gives you the option of using a UNIX session manager with the interpreter `apl-8`.

The Session Log

The session manager maintains a log in which all session input and output are recorded. This is a record of each side of the dialogue between you and the APL system, with a newline character at the end of each line. A useful feature of the session log is the ability to scan the log and reuse earlier input or output as part of your new input.

The ∇-editor

This traditional APL editor is used for entry and revision of user-defined functions. You invoke it by typing the symbol ∇ followed by the name of the function to be edited. The ∇-editor was designed for the sole purpose of creating definitions of user-defined functions and is ill-suited for other kinds of editing.

The Full-Screen Editor

The session manager provides a full-screen editor for your APL session, and also makes it available for editing user-defined functions and variables of rank `1` or `2` character type. The full-screen editor, invoked by `⎕EDIT` or `)edit`, allows you to perform several editing tasks:

- create new functions or variables, or modify existing ones
- edit several objects at the same time
- edit an image of your APL session and save lines of text from it into user functions or variables
- use extended cursor movement, scrolling, and paging
- copy, insert, move, and delete blocks of text within an edit screen
- move blocks of text between edit screens
- search for specific character sequences, and optionally replace them.

For more information on the full-screen editor see the *System Guide, Chapter 7.*

Ending an APL Session

Typing the command `)off` ends your APL session and returns you to the prompt of the UNIX shell. Ending the APL session does not end work altogether; it simply terminates the APL process as far as UNIX is concerned:

- All files that had been opened during the APL session are closed
- The contents of the active workspace are discarded
- The keyboard (usually) reverts to its behavior before APL was invoked.

Although the APL task keeps an internal record of its elapsed time and cpu time (and you can use that for calibrating performance of APL programs), SHARP APL has no mechanism for passing that information to UNIX, and does not return an accounting statement when you terminate an APL task as it does in SHARP APL for OS/390.

Communications Failure

When you are using SHARP APL from a remote terminal, a communications failure ("line drop") terminates your UNIX session as well as your APL session. In the UNIX environment, such terminations may cause the system to save an APL `continue` workspace if you have properly conditioned the termination workspace facility (see `⎕twsid` in the *System Guide, Chapter 5*).

Terminating a Task

The system function `⎕bounce` can be used to terminate (bounce) any APL process(es) identified by task number. A task number is assigned whenever a task is started. The UNIX system closes any files that the terminated tasks had opened, but does not otherwise save their work. However, an APL task with `⎕twsid` set does save the active workspace.

Note: Because APL task IDs and UNIX process IDs are equivalent, executing `⎕bounce` *n* is the same as using the UNIX `kill` command for process ID *n*.

An S-task started by another APL task (using the auxiliary processor AP1) can be terminated by sending it the command `)off`, or by using `⎕bounce`. Under some conditions, such a task may terminate automatically when you cease to share the name that controls it, or when it finishes what it is doing. An N-task is terminated using `⎕bounce`, by a return to immediate execution, or if it solicits input.

For more information on tasks, □*bounce*, and other systems functions and commands, refer to the *System Guide*.

Emergency Exit

In the event that your APL session becomes unresponsive and you are unable to terminate your APL task, you may still be able to signal UNIX by typing the QUIT key, which is usually CTRL-\.

Note: This causes an abnormal termination. No trace of your session is retained and data may be lost.

3
SHARP APL Environment

This chapter describes the various executable files, both scripts and binaries, which help to establish and maintain an efficient working environment for SHARP APL under UNIX. It also provides information on the startup options that allow you to customize your SHARP APL environment.

Note: The degree to which you are able to customize your environment depends on your level of access. Consult your system adminstrator for more information.

Environment Variables

Each environment variable used in the `sax` script has a default value, which is set only if that variable is not already defined. This gives you the opportunity to override the defaults if you wish. Some of the variables are used to control the SHARP APL interpreter (which recognizes a wide range of startup parameters); others determine general aspects of the environment, provide for auxiliary processors, and name the system configuration files that control printing and sharing of data with other machines on your network. Table 3.1 lists the variables for the `sax` script.

The `sax` script assigns a default value to an environment variable only when one of the named variables has no value. You can assign a value directly in UNIX; for example, if you want to use the version of the interpreter that does not have a built-in session manager, you could set the environmental variable `APLPATH` to the value `$SAXDIR/bin/apl-8` and the `sax` script would use the value you supply.

In `.saxrc` (see "Profile for SHARP APL" below), you can do it without exporting the value to your working environment by entering `APLPATH=$SAXDIR/apl-8`.

Table 3.1. Environment variables for the `sax` *script.*

Name	*Description*	*Default*
`APLFONT`	APL font for X Window.	`saxmedium`
`APLPARMS`	Startup parameters.	*None*
`APLPATH`	Interpreter file.	`$SAXDIR/bin/apl`
`APLSIZE_ROWS`	Rows on APL screen.	`34`
`APLSIZE_COLS`	Columns on APL screen.	`80`
`APLTRACE`	Generate log files.	`$SAXDIR/trc`
`EDITSIZE`	Size of storage for editor.	`65536` *bytes*
`INIT`	Use non-default terminal type.	`$SAXDIR/lib/term/$TERM.init`
`SAXBRACELOG`	Log use of brace characters.	`0` (*off*)
`SAXCNF`	Configuration directory.	`/usr/sax/local`
`SAXDIR`	SHARP APL directory.	`/usr/sax/rel`
`SAXRC`	User-specific profile.	`$HOME/.saxrc`
`STARTAPS`	Start the AP daemon.	`yes`
`TEMPPATH`	Path for temp files.	`/var/tmp`
`TERMCAP`	Terminal database (obsolete).	`$SAXDIR/etc/atermcap`
`TERMINFO`	Terminal database for AP124.	`$SAXDIR/lib/term/terminfo`
`WSSIZE`	Workspace size.	`500000` *bytes*

Profile for SHARP APL

The `sax` script provides for a private file in which you can store a fragment of a UNIX shell script that you find appropriate to your use of SHARP APL. The name of the file in which you keep your private settings is, by default, stored in the environment variable `SAXRC`, for which the default is `$HOME/.saxrc`.

Session Startup Options

All SHARP APL interpreters (see "Language Interpreters") accept startup options: each option begins with a dash prefix (`-`) followed by an option letter and, if additional parameters are required, followed by certain values. For example, the `S` option specifies the number of screens that the session manager maintains for your session log; for 10 screens, the expression you supply is `-S10`. Some options, however, have an effect only on `apl`, the interpreter with a built-in session manager.

The session startup options are outlined in Tables 3.2 and 3.3, below. In addition to these, there are also startup options related to terminal support. These are outlined in the section "Terminal Startup Options" on page 3-7.

Table 3.2. Session startup options.

Option	*Description*	*Default*	*Default Script Settting*
`-B`	Turn session heading ***off***	Not turned off	
`-D`*txt*	Set S-Tasks ***local***: `-Dl` Set S-Tasks ***global***: `-Dg`	`-Dl`	
`-E`*txt*	Set editor command character. (7-bit version)	`-D'\t'` (TAB)	
`-G`	Generate core dump	Not generated.	Generated.
`-I`	Turn status line display ***off***	Not turned off.	
`-J`*txt*	Set number of spaces for APL prompt (8-bit version)	`-J6`	

`-L`*txt*	Set library/path mapping	No mapping	`-L1 $SAXDIR/lib/wss`
			`-L6 $SAXDIR/lib/wss6`
			`-L8 $SAXDIR/lib/wss8`

`-M`*txt*	Set editor memory in bytes for *)edit* (7-bit version)	`-M16384`	`-M65536`
`-N`*txt*	Set terminal configuration file	No configuration	`-N$SAXDIR/lib/term/$TERM.init`
`-O`*txt*	Output translation table		`-O$SAXDIR/lib/term/aplotab`
`-S`*txt*	Session parameter (⎕*sp*)	***empty*** (⎕*sp*←'')	
`-S`*txt*	Number of session screens (7-bit version)	`-S4`	`-S10`
`-T`*txt*	Alternate HOME	`-T$HOME`	
`-W`*txt*	Autostart WSID	*clear ws*	
`-X`*txt*	Select ⎕*avm* translate table. See note below.	Mainframe ⎕*av*	
`-Y`*txt*	"Suit Yourself" options	Miscellaneous options of the form *option=number*. Options are separated by commas, no blanks. See Table 3.3, below.	
number	Workspace size in bytes;	`16384`	`500000`

Note: If specified, `-X` must be followed immediately by the name of a file containing exactly 256 bytes. This file must be readable by the user. It may be specified either relative to the current working directory (e.g., `-Xfooxlate`) or as an absolute path name (e.g., `-X/usr/sax/local/ourxlate`) . The file is read once when the interpreter starts, and the contents are returned whenever ⎕*avm* is used.

Table 3.3. "Suit Yourself" options.

Option	*Description*	*Default*
`fcreatesw=0|1`	File access matrix implementation If `0`, use `chmod` to alter UNIX permissions If `1`, use `umask` to set UNIX permissions Refer to the *System Guide*, *Chapter 3*	`1`
`fsync=0|1|2`	For synching during file writes. If `0`, no synching If `1`, uses `synch()` If `2`, uses `fsynch()` Consult your UNIX documentation for more information.	`2`
`scissor=0|1`	If `0`, "folds" wide displays If `1`, "scissors" wide displays Refer to the *System Guide*, *Chapter 5*.	`1`
`sesmtrans=0|1`	If `0`, `)sesm` translation is off If `1`, `)sesm` translation keyword `#` (only 8 bit) Refer to the *System Guide*, *Chapter 6* .	`0`
`usefilelocks=0|1|2|3`	If `0`, APL does not perform file locking during file operations. If `1`, APL uses `fcntl()` to lock files during file operations. If `2`, APL uses semaphores to lock files during file operations. If `3`, APL uses the fileserver to perform file operations. Refer to the *File System Manual*. ***Note:*** Subtasks are not started with the selected value. If this is required, set `APLPARMS`.	`1`

There are two ways to override the default settings:

1. You can use arguments directly as options to the `sax` script. For example,

```
% sax -W"1 tools" 900000
```

   requests that the workspace to load at the start of the session be `1 tools` (using the `-W` option), and that a 900-kb workspace size override the 500-kb default.

2. You can use the environment variable APLPARMS. In this case, you also have two choices: set it directly in UNIX, or indirectly in your .saxrc file. In this example, the current contents of APLPARMS are catenated to your new value in your .saxrc file.

```
APLPARMS=$APLPARMS" -L2/usr2/coryc/fmt"
```

   If you have not set a value in $APLPARMS by either method, no default value for the variable is assumed.

Language Interpreters

Depending on your requirements, you can choose one of four different APL language interpreters, apl, apl-8, rtapl, and rtapl-8, by setting the APLPATH environment variable in your UNIX environment prior to starting SHARP APL:

Example:

```
APLPATH=/usr/sax/rel/apl-8
```

This overrides the default interpreter (apl) and loads apl-8 instead. Each interpreter offers functionality that is suited for a particular mode of operation. The different interpreters are explained in the sections that follow.

apl

This version of the interpreter, the default, has a built-in session manager which includes a full-screen editor. When you use apl as your interpreter, you must indicate how SHARP APL is to support your particular terminal.

apl-8

Many users prefer a UNIX session manager to one that is restricted only to SHARP APL sessions; they want that session manager, not SHARP APL, to provide their editing and scrolling facilities. The version of the SHARP APL interpreter ideally suited to this mode of operation supports the 8-bit characters required to represent the 256 positions in the APL character set.

The interpreter called apl-8 takes its input from stdin and places its output on

stdout. If you have an environment that permits interactive use of apl-8, you may often be able to run apl-8 without a cover script. If you require APs, however, the simplest approach is to use the sax script as usual with APLPATH set to $SAXDIR/bin/apl-8 in your UNIX environment.

When apl-8 encounters end-of-file (EOF) while reading from the standard input device (stdin), the action taken depends on the input mode of apl-8 at the time.

The possible input modes, and actions taken, are as follows:

Immediate Execution and Function Definition	If any characters are present before EOF, those characters are processed normally, and the EOF is ignored. If there are no characters before the EOF (a naked EOF), the apl-8 process terminates as if it had received a □*bounce*.
□ and ⍞ Input	If any characters are present before the EOF, those characters are processed normally, and the EOF is ignored. No characters before the EOF (a naked EOF) cause event 1004, input interrupt, to be signalled.
□*arbin*	If any characters are present before the EOF, those characters are processed normally, and the EOF is ignored. If there are no characters before the EOF (a naked EOF), □*arbin* returns an empty list.

Input buffer overflow of stdin in apl-8 is ignored by all input modes except □*arbin*. All modes but □*arbin* treat the characters received before the buffer overflow as a complete input line. □*arbin* reacts to overflow of the input buffer by signalling event 1004, input interrupt.

SHARP APL as a UNIX Filter

You can have apl-8 accept all its input from a UNIX file, and possibly redirect its output to a UNIX file. This lets you prepare test scripts or other sequences of expressions to produce output you desire. The script provided for this purpose is called saxscript. Its first argument must be the name of a UNIX file which contains your input lines. Any additional arguments you provide are the same as those appropriate to the sax script.

```
% saxscript /usr2/joe/wsin 500000
```

If you want to redirect output, use standard UNIX syntax:

```
% saxscript /usr2/joe/wsin >/usr2/joe/wsout
```

rtapl

The run-time interpreters `rtapl` and `rtapl-8` support the execution of closed applications and do not allow the user into APL immediate execution.

Terminal Startup Options

SHARP APL for UNIX supports users at terminals in one of two ways: If your interpreter is `apl` (with built-in session manager), then files in a terminal database provide support for your terminal. If, on the other hand, you use a session manager available in UNIX (such as `emacs`), no terminal database is required and the interpreter you use is `apl-8` (with no built-in session manager).

The APL session manager supports a set of startup options to allow customization for specific terminals. These options are described in Table 3.4.

Table 3.4. Terminal startup options.

Option	*Meaning*	*Default*
`-a`	Send to terminal at start of APL session	CTRL-n
`-A`	Send to terminal at end of APL session	CTRL-o
`-f`	Refresh screen	CTRL-l
`-k`	Undo (discard input)	CTRL-b
`-o`	Overstrike	CTRL-p
`-z`	O-U-T	CTRL-g
`-Z`	Next character in APL font	CTRL-a

Terminal Initialization

With the default interpreter `apl`, the UNIX environment variable `TERM` is used as the basis for terminal support. For example, if you use a VT100, the value in `TERM` is `vt100`. The terminal database in `$SAXDIR/lib/term` must contain a file called `$TERM.init`, which is used by the interpreter to initialize values global to the APL

session manager in such a way that they are appropriate to the characteristics of your terminal.

To understand the connection of SHARP APL to terminals, let us inspect the initialization file that supports a VT100. The UNIX file to initialize this terminal (which the command `apl` refers to using the `-N` option) is called `vt100.init`:

```
# Initialization file for SAX for DEC VT100 and
# other terminals with ANSI Standard keyboard.
#
# Initialization of terminal:
#
# ANSI mode:                            \033<
# Reset cursor key application mode:    \033[?1l
# PFkey native programs:                \033[2u
-a\033<\033[?1l
#
# Function              Keystroke
# ========              =========
#  overstrike           ctrl-P
-o^P
#  APL escape           ctrl-A
-Z^A
#
# choose output translate table for typewriter paired
# apl/ascii terminal
#
-O/usr/sax/rel/lib/term/aplrtab
```

To inspect initialization files, see the directory `$SAXDIR/lib/term`. See "Chapter 4. Terminal Support" for more complete documentation.

SHARP APL Without a Terminal

Some application systems require an active SHARP APL workspace but without a terminal attached. The `detach` utility starts a process with a new process group ID thereby detaching the new process from the controlling terminal.

The single argument required for `detach` is the name of a startup script. Other arguments can be passed to the script if it is written to accept arguments. An example, `detachsax`, is provided which sets the default for the interpreter to

`apl-8` (only `apl-8` or `rtapl-8` should be used without a controlling terminal) and redirects standard input and output to `/dev/null`. You can also use the normal `sax` script as the one that `detach` invokes. For example, an application, `odaserver`, which receives input via shared variables, and is to remain running indefinitely might use the following expression to start a detached SHARP APL session:

```
% detach detachsax -W/home/oda/odaserver
```

No & is needed to force background usage: this is automatically handled for you by `detach`.

One Application of `detach`

It is generally a good idea to use `detach` to run any `apl-8` process without a controlling terminal. For example, `detach` can be used to resolve a *`no shares`* problem that affects programs run via UNIX boot scripts.

When a SHARP APL program is run from a UNIX boot script, its process group ID becomes `0`. A process with a group ID of `0` that wishes to share a variable is denied use of the SVP for security reasons. However, if it is ***detached*** from the controlling terminal, the program's process group ID becomes non-zero and the SVP becomes accessible.

SHARP APL for UNIX Logging Facility

A system logging facility can be implemented to monitor elements of APL, APD, the Fileserver, and the auxiliary processors (this logging does not include the SAMI daemon). However, if your SHARP APL system is running properly, few entries are made to the log file.

The default log file name is determined by the `SAXLOG` environmental variable. The default name, if `SAXLOG` is not set, is `/var/tmp/sax.log`. To disable this logging mechanism, set `SAXLOG` to an empty string.

The format of entries in the logging file is as follows:

```
ptype msgid uid pid dd/mm/yy hh:mm:ss text
```

`ptype` Process type in the form of a 3-character acronym (APL, APD, etc.,).

`msgid` Message number, unique within `ptype` messages, `1 – 9999`. If ≤ `0`, neither `ptype` and `msgid` are output.

`uid` User ID

`pid` Process ID

No log messages are written on normal start and termination of processes. The message classes that correspond to error numbers are listed in Table 2.1.

Table 2.1. `SAXLOG` *message numbers*

Error Number	*Class of Messages*
`1`	***Cannot open log file***
`2`	***Cannot allocate ring memory***
`3`	***Successful initialization***
`4`	***Normal termination***
`100`	***Error detected***
`104`	***Abnormal termination***
`110`	***Initialization failed messages***
`120`	***Allocation failure messages***
`130`	***UNIX system call failure***
`200`	***SVP error***
`300`	***File System error***
`400`	***SHARP APL interpreter error***
`500`	***File Server error***
`600`	***File Server*** `fsios` ***error***

To find the most recent error messages in the logging file for a specific user (*userid*) type

`grep` *userid* `$SAXLOG | tail`

at the UNIX prompt.

4
Terminal Support

Introduction

This chapter explores your options for terminal support once SHARP APL and the APL font are installed on a host UNIX system.

The following features are provided for the supported terminal types:

- A line-oriented session manager, with a built-in `vi`-style function editor, for conducting a normal interactive SHARP APL session.
- AP124, which allows SHARP APL applications to drive the screen as a 3270-style full-screen device, including the use of color.
- APL character display, both by the session manager and AP124 applications.
- The SHARP APL "union keyboard," in which APL characters are entered using the ALT key as a shift key. The entry of APL characters is supported by both the session manager and AP124 applications.

Terminal Types

The following terminals are supported by SHARP APL for UNIX:

- Any UNIX workstation or PC running X server software.
- A PC running Microsoft DOS or Windows 3.1/95/98/NT that is connected to the UNIX host by Ethernet or asynchronous serial line. For such a terminal, access is achieved using either of two terminal emulators: PC108, which emulates an HDS 108 terminal, or the VT100/VT320 terminal emulator with MS-DOS Kermit.

Terminal Emulators Package

For APL support, three features are necessary in a PC terminal emulator: keyboard remap, extended font support, and translation of output. The ***Terminal Emulators*** package distributed with the SHARP APL for UNIX distribution CD provides display fonts, some proprietary and shareware emulators, and the requisite document files describing how to implement APL character support.

SHARP APL Character Support

If you are using an X Window terminal supported by SHARP APL for UNIX, your keyboard is mapped for entering APL characters, both when using the line-oriented session manager, and when using AP124. If you are unable to display APL characters, ensure that the X Window APL font is installed (See "X Window APL Font" on page 4-7)

APL character support is based on the conventional *union keyboard* where APL characters are produced using APL shifting-key combinations. For most terminals, ALT serves as the APL shifting key, although a Meta key may be configured to perform the function of ALT in some situations. The following discussions assume ALT as the APL shifting key for key sequence examples. Figure 4.1 presents the complete APL character map available for the X Window system.

Most APL characters are entered using a single alphabetic key while holding down the ALT key; for example, ALT-a produces the character ⍺ and ALT-/ produces the character ⍀. Some additional APL characters can be produced using a combination involving the SHIFT key as well; for example, SHIFT-ALT-m produces the character ⌶. Figure 4.1 describes the combinations available for producing APL and non-APL characters: *unshifted*, SHIFT, ALT, and SHIFT-ALT.

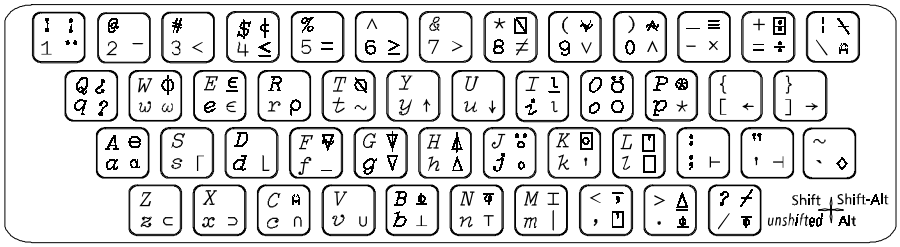

Figure 4.1. SHARP APL keyboard map for X Window system

Overstrike Combinations

Any APL character not available using ALT or SHIFT-ALT can be produced using a combination of two available APL characters. These combinations consist of two APL characters separated by an *overstrike* key (normally CTRL-p).

For example, ⊡ is a combination of the APL characters □ and ∘ which is produced by the sequence ALT-j + CTRL-p+ ALT-l. Table 4.1 lists the characters that can be produced using overstrike combinations.

Table 4.1. APL overstrike combinations.

APL Character	*Overstrike Pair*	*APL Character*	*Overstrike Pair*	*APL Character*	*Overstrike Pair*
⍨	− ,	⊡	□ ÷	≠	− /
⍲	∧ ~	⍱	∨ ~	⊡	□ ∘
⍂	□ \	≡	= _	⍀	− \
⍪	⊥ ⊤	⊖	○ −	⊆	∈ _
⍫	▽ ~	⍒	▽ |	⍋	△ |
⍳	⍳ _	⍤	∘ ¨	⍥	○ ¨
⍋	⊥ ∘	⍒	⊤ ∘	⍞	□ '
⍟	○ *	⍂	○ \	⍵	| ○
⍎	| ⊂	⍙	△ _		

Although CTRL-p is the default, you can define your own overstrike key as a startup option; see "Chapter 3. SHARP APL Environment" for more information on startup options.

Session Manager Keys

In addition to keys for entering APL characters, certain other keys are mapped to provide useful functions when the line-oriented session manager is in control. *System Guide, Chapter 7* provides more details on the keys outlined below:

INSERT	Toggles insert/typeover state.
DELETE	Deletes character at cursor position.
BACKSPACE	Deletes one character left of cursor position.
CTRL-e	Erases to end of line.
ENTER	Submits current line to APL interpreter.
CTRL-g	Causes input interrupt or O-U-T.
CTRL-c	Sends attention to SHARP APL interpreter.
CTRL-l	Refreshes current display.
arrow keys **PG UP/PG DOWN** **HOME and END**	Provide cursor movement within session manager log.

AP124 Keys

In addition to keys for entering APL characters, certain other keys are mapped to provide useful and/or necessary functions when an application using AP124 is in control, including keys equivalent to 3270 keys such as the PA keys and Clear key. The *Auxiliary Processors Manual, Chapter 4* provides more details on the AP124 keys outlined below:

arrow keys	Provide cursor movement around screen.
HOME	Moves cursor to first unprotected field.

TAB	Moves cursor to next unprotected field.
BACKTAB	Moves cursor to previous unprotected field.
CTRL-ENTER	Moves cursor to first unprotected field on next line.
INSERT	Toggles insert/typeover state
DELETE	Deletes character at cursor position.
BACKSPACE	Deletes one character left of cursor position.
CTRL-e	Erases to end of field.
CTRL-c	Sends attention to SHARP APL interpreter.
CTRL-l	Refreshes current display.

The following are data entry keys for submitting an amended screen to AP124:

ENTER	Enter
PF keys 1 to 12	PF keys 1 to 12
SHIFT PF1, PF2 , PF3	PA1, PA2, and PA3
SHIFT PF4	Clear

National Use Characters

Like most environments that use one-byte character codes, SHARP APL lacks official encodings for international characters. The only fully correct solution to this problem is to adopt complete support for the ISO/IEC standard IS 10646 (also known as ***Unicode 1.1***) when it becomes available. Currently SHARP APL addresses this issue by reserving some character positions, known as ***National Use*** characters, whose display is permitted to vary.

National Use characters are fully documented in "Appendix A. National Use Characters" . However, if you require customised display or entry of National Use characters, you should contact the Soliton Associates Technical Support group for assistance.

Keyword Substitutes for APL Primitives

If APL character support is unavailable on your terminal, the system command `)sesm trans` allows you to type keyword replacements for missing APL primitive characters; the process is called ***transliteration*** (available for the `apl-8` interpreter only). The *trans* argument turns transliteration on or off, sets the transliteration keyword escape character, and reports the display status regarding transliteration. More information about the `)sesm` command is provided in the *System Guide, Chapter 6*.

Keywords only exist for primitives that cannot be found in your terminal's character set. Table 4.2 lists the APL primitives that have keywords under SHARP APL for UNIX.

Table 4.2. Transliteration keywords.

APL	Keyword	APL	Keyword	APL	Keyword	APL	Keyword
⍤	*rank*	⍺	*alpha*	⊂	*lshoe*	∈	*member*
⍳	*iota*	∩	*cap*	∘	*jot*	⍴	*rho*
∪	*cup*	∨	*or*	⍵	*omega*	×	*x*
∆	*delta*	|	*abs*	⌊	*min*	○	*circle*
⊤	*represent*	⍤	"	⌷	*quad*	⌶	*ibeam*
⌿	*overf*(irst)	⍀	*scanf*(irst)	→	*goto*	⊣	*lev*
÷	*divide*	⍲	*nand*	←	*is*	◇	*diamond*
≤	*le*	≠	*ne*	≡	*match*	≥	*ge*
⍪	*catenate*	⊆	*find*	⍳	*first*	⍝	*comment*
⍤	*on*	⍱	*nor*	⍋	*upgrade*	⊢	*dex*
↑	*take*	↓	*drop*	⍉	*transpose*	⊖	*theta*
⍟	*log*	⍫	*phi*	⍤	*upon*	⍕	*format*
⊃	*link*	∇	*del*	⌈	*max*	⊥	*base*
⍞	*inline*	⍢	*locked*	⍙	*deltabar*	⍒	*downgrade*
⍎	*do*	⍇	*domino*				

Each keyword must be followed by a ***space*** unless it is the last keyword before a carriage return. Due to the frequency of the use of ⎕, the interpreter allows the keyword escape character (e.g., `#`) to be used as a short cut for the keyword #*quad*; i.e., `#ai` is equivalent to `#quad ai` or `⎕ai`.

X Window Systems

For a terminal using an X Window system, each invocation of SHARP APL starts a separate `xterm` window. The subsequent line-by-line session is conducted in this window. Any SHARP APL application that uses AP124 to drive the screen in a full-screen manner uses this same window.

Naturally, several separate SHARP APL sessions can be conducted simultaneously, each in a separate `xterm` window. In fact, SHARP APL for UNIX's Host AP (AP11) can be used by a SHARP APL application to start separate SHARP APL windows.

Such a SHARP APL session conducted in an `xterm` window supports the display and entry of APL characters. It also supports the display of color by AP124 applications, provided that the chosen `xterm` program properly honors the ANSI-standard escape sequences for color and other display attributes.

Note: Some `xterm` programs honor only attributes such as reverse video, bold, underscore, etc. Some `xterm` programs support color as well but do not properly support all attributes, giving rise to limitations on how well color is supported by AP124 applications.

X Window APL Font

For workstations using an X Window system, it is necessary to install a font on the X server machine and a keyboard mapping on the X client machine (i.e., runs the `xterm` program). If you use an X server application for MS Windows, see "X Window Emulator" on page 4-8 for further information.

The APL font and keyboard mapping are supplied with SHARP APL for UNIX and are normally installed during the ***initial SHARP APL installation***. However, if the APL font does not display, execute the X Window APL font script (via `root`), as described in the following procedure:

- ❑ Sign on to user `root`
- ❑ `cd $SAXDIR/etc`
- ❑ Execute the script `xfonts.install`

This procedure installs the fonts `saxlarge`, `saxmedium`, and `saxsmall` in the system library `$OPENWINHOME/lib/fonts`.

Window Resizing

Once an `xterm` window is created for a SHARP APL session it can be resized within your terminal. However, the maximum number of print positions on a line is still the visible value of the system function □*pw* (printing width). By default, □*pw* is set to `80`; you may set it as low as `30` or as high as `255`; or you can indicate an infinite width by setting □*pw* to `0`.

AP124, which allows SHARP APL applications to drive the screen as a 3270-style full-screen device, is not resizable. However, the session manager is resized upon returning from AP124. Resizing during □*arbin* signals `event 1004 , input interrupt`. The session log is sometimes cut to 4 screen sizes on resizing; only the last 4 pages are kept (the top portion of the buffer is cut off). Resizing is deferred until the full-screen editor buffer directory screen (TAB *r*) is exited.

X Window Emulator

If you use an X server application for MS Windows to access the X client (e.g., Hummingbird's *eXceed*), you can access SHARP APL in the same manner. However, you must install the SHARP APL font on the PC acting as the X server to get APL character support.

eXceed provides an application that enables you to convert SHARP APL font files (`*.bdf`) into a format used by eXceed. Three files are required for this support: `saxlarge.bdf`, `saxmedium.bdf`, and `saxsmall.bdf`. These font files can be found in the***Terminal Emulators*** directory on the SHARP APL for UNIX distribution CD. Follow the steps outlined below to add these files to the eXceed font database.

Retrieve the Font Files

- ❑ Create a SHARP APL font directory on your PC, for example, `c:\exceed\font\sax.`
- ❑ Copy the three SHARP APL (`sax`) font files to this directory from the distribution CD.

Compile the Font Files

- ❑ Run the *Xconfig* application in eXceed.

- ❑ Double-click the *Font* icon.
- ❑ Click on *Compile Fonts...* in the *Font Settings* dialog box.
- ❑ Select the SHARP APL font directory (e.g., `c:\exceed\font\sax`) as the directory for your input files and select the three SHARP APL (`sax`) font files, type *BDF Files (*.BDF)*.
- ❑ Select an Output Font Directory (e.g., `c:\exceed\font\sax`)
- ❑ Click on *Options...* and ensure the settings displayed are correct.
- ❑ Click on *Compile* to start compiling. After compiling, if the output directory you specified is not in the font database, you will be prompted to add it. Follow the instructions provided.

View the Font Database

If you were successful, you will be able to display the SHARP APL (`sax`) font in the eXceed font list.

- ❑ Run the *Xconfig* application in eXceed.
- ❑ Double-click the *Font* icon.
- ❑ Click on *Font Database...* in the *Font Settings* dialog box.
- ❑ Select the SHARP APL font directory (e.g., `c:\exceed\font\sax`) from the Font Database list provided.
- ❑ Click on *Font List for Directory...* to display the list of fonts and aliases in that directory. The three new SHARP APL (`sax`) fonts should be listed here.

Note on FONT names

The distributed SHARP APL font files were given the same names as their UNIX counterparts. If `saxmedium.bdf` is too long a file name for your MS Windows operating system, you may want to create an alias `saxmed` mapped to `saxmedium`. When you run the `$SAXDIR/bin/sax` script, it checks the value of the environment variable `APLFONT` to select the right font from the X server. The default is `saxmedium`.

Using Terminal Emulators

If you do not have an X server application for the PC, SHARP APL for UNIX supports two other methods for conducting a session from your MS-DOS/Windows system: PC108, which emulates an HDS108 terminal, or the VT100/VT320 terminal emulator within MS-DOS Kermit.

Both emulators are able to support the display and entry of APL characters, and the display of color by AP124 applications. In both cases, the PC may be connected to the UNIX host by an asynchronous serial line or, under certain conditions, by TCP/IP networks.

To run under ***Windows NT***, the emulator software must support flow control and should support `550 FIFO`; Kermit does but PC108 does not.

Note: You can access your SHARP APL system using other terminal emulation software (e.g., Hyperterminal or `telnet`) but you will be ***without*** APL character support.

Asynchronous Serial Line Connection

For asynchronous serial line connections, note the following:

- PC108 can only be used over an asynchronous serial line configured as a 7-bit line (i.e., it cannot be used over one configured as an 8-bit line).
- MS-DOS Kermit can be used over either a 7-bit or 8-bit line.

TCP/IP Connection

For PCs running a version of MS Windows that uses a 3rd party (i.e., non-Microsoft) TCP/IP stack, there are some circumstances where it is possble to use PC108 or Kermit over a TCP/IP connection instead of an asynchronous line. Contact Soliton if you wish to investigate such an arrangement.

For PCs running a version of MS Windows uses the Microsoft TCP/IP stack (Windows 95/98 or Windows NT), it is not generally possible to use PC108 or Kermit over a TCP/IP connection. For all cases of TCP/IP connection, the recommended arrangement is to use an X server product.

File Transfer

Both terminal emulators have file transfer capabilities that make it possible to transfer files between the MS-DOS and UNIX systems:

- With Kermit, file transfer between MS-DOS/Windows files and UNIX files is available via Kermit's file transfer protocol (provided Kermit for UNIX is installed).
- With PC108, file transfer between MS-DOS/Windows files and SHARP APL files is available (via PC108's upload/download facility), using a SHARP APL workspace (`1 pc108`).

Note: Since the SHARP APL for UNIX Host AP (AP11) provides for conversion between UNIX files and SHARP APL files, both of the above can indirectly achieve the other's capability.

`egaset` *APL Font*

You must have a monitor that has EGA capability (at least) in order to display APL characters while using a PC terminal emulator (in MS-DOS mode). This is achieved by executing the program `egaset` to load the APL screen font. Loading instructions for `egaset` are provided in the section "Terminal Emulators Installation and Configuration" on page 4-12. See "SHARP APL Character Support" on page 4-2 for more information about entering and displaying APL characters.

Problems Using `egaset`

If you have problems getting APL characters in an MS-DOS window using `egaset` under Windows 95, check your `autoexec.bat` file for a couple of lines such as:

```
mode con codepage prepare=((850) c:\windows\command\ega.cpi)
mode con codepage select=850
```

These should be commented out or deleted altogether. This could possibly cause a problem if you are using a code page other than US English.

Support for Color/Highlighting in AP124

Both the Kermit and PC108 terminal emulators support the use of color by AP124, but all combinations of color and extended highlighting are not properly supported: PC108 does not support high intensity, blinking or underscored fields but converts them to some use of color instead; Kermit does not support underscored fields so they appear as blinking fields.

Terminal Emulators Installation and Configuration

This section provides some basic instructions for installing and configuring the terminal emulation software for both types of connections: PC108 asynchronous or Kermit asynchronous.

Retrieve the Font Files

Before a terminal emulator can be installed, your PC must be equipped with SHARP APL character support. The necessary software is provided in the ***Terminal Emulators*** directory on the distribution CD.

- ❑ Copy the `SAXPC` subdirectory to your hard disk, retaining the directory structure:

```
SAXPC
    EGASET
    KERMIT
    PC108
```

- ❑ Choose a suitable font size for your monitor screen. There are three sizes distributed with SHARP APL for UNIX: 8×12, 8×14, and 8×16. You can find the appropriate font by experiment. For example, the following will load the 8×12 font:

```
C:\> cd \saxpc\egaset
C:\> aplchar 812
```

PC108 7-bit Asynchronous Connection

PC108 can only be used over an asynchronous serial line configured as a 7-bit line.

- ❑ The first time you use PC108, reconfigure the distributed copy with the correct

setting of the COM port number and speed of the line:

```
C:\> cd \saxpc\pc108
C:\> mode comn:bbbb,e,7,1
```

where *n* is the port number, and *bbbb* is the line speed; for example:

```
mode com1:9600,e,7,1
```

- ❑ Execute the following command:

  ```
  C:\> pc8sax fff x1 dnnn
  ```

  where *fff* is the font size, and *nnn* is 123 for COM1 or 213 for COM2; for example:

  ```
  pc8sax 816 x1 d123.
  ```

  This should start a PC108 session. If your line is a dialup line, you will need to type the appropriate dial command.

- ❑ Once logged in to your UNIX system, you should be able to start a SHARP APL session using the `saxpc` script, as follows:

  ```
  {unix:1} setenv TERM pc108
  {unix:2} saxpc
  ```

  (or `saxpc m` for a monochrome monitor)

- ❑ Once in the SHARP APL session type:

  ```
  )load 1 pc108
  savelx
  ```

  This will save the current settings in the PC108 configuration file `\saxpc\pc108\pc108.nvm`. Subsequent executions of the PC108 emulator can then proceed normally.

Kermit Asynchronous Connection (7-bit or 8-bit)

- ❑ Configure the correct setting of the COM port number and speed of the line by editing the file `\saxpc\kermit\serport.cfg`.

  ```
  C:>\cd \saxpc\kermit
  C:>\kersax fff asy7
  ```

where *fff* is the font size. (e.g., `kersax 816 asy7`). This should start a Kermit session. If your line is a dial-up line, you will need to type the appropriate dial command.

- Once logged into your UNIX system, you can start a SHARP APL session using the `saxpc` script as follows:

```
{unix:1} setenv TERM vt100
{unix:2} saxpc
```

The default values in most of the configuration files are set so that you can access SHARP APL for UNIX. The `kersax.bat` file has a file path which may need to be adjusted. If other changes are required, refer to the Kermit documentation.

The following is a sample session:

`kersax 812 asy7`	`kersax` is the `.bat` startup file. `812` refers to the APL font size. `asy7` is the type of connection, 7- bit async.
`atdt1234567`	dials the number on a touch tone phone.
`CONNECT xxxx`	Continue signing on using standard procedures.

5
Printing

APL PostScript Printing

SHARP APL for UNIX provides PostScript printer support for APL characters via a utility called `wys` which can either be used as a UNIX command or called from SHARP APL programs. In order to make `wys` available for use, your system administrator should take the normal steps to activate an `lp` scheduler and start the printer. Configuration and maintenance documentation for `wys`, `lp`, and APL character support is distributed with the SHARP APL for UNIX installation.

The `wys` Utility

The utility `wys` is provided for printing text files. Comprehensive print options are available, and multiple file names may be specified.

For example, enter the following at the UNIX prompt to perform a background print of file `abc` and all files with names beginning with `def`, printing two copies of each file, 66 lines per page:

```
wys -b1 -c2 -L66 abc def*
```

Installed Files

`wys` is installed automatically, as part of SHARP APL. Table 5.1 lists the `wys`-related files that must be in place for operation of this utility.

Table 5.1. `wys` *files.*

File	*Description*
`$SAXDIR/bin/wys`	`wys` command (linked to `/usr/bin/wys` if desired)
`$SAXLIB/local/etc/wysrc`	Site environment file
`$SAXDIR/etc/wysrc`	Distributed environment file
`$SAXDIR/lib/lp/ehandler.ps`	PostScript error handler
`$SAXDIR/lib/lp/lwreset`	Reset printer
`$SAXDIR/lib/lp/psload` †	Load PostScript font
`$SAXDIR/lib/lp/qav`	PostScript APL diagnostic
`$SAXDIR/lib/lp/wys`	`wys` command
`$SAXDIR/lib/lp/wyscov`	`wys` cover script
`$SAXDIR/lib/lp/wysopts`	`wys` option overview
`$SAXDIR/lib/lp/wystest`	`wys` test data

† `LW` *is the default printer unless the* `Prt` *environment variable is defined.*

Installing a Postscript Error Handler

Because the UNIX print scheduler is autonomous, diagnostic messages produced by the PostScript interpreter are not displayed unless a PostScript error handler is installed. An error handler is provided which causes the printer to print the error messages it would normally transmit to the controlling shell.

Normally, installation of the PostScript error handler is not necessary, but if it is required for diagnostic purposes, it can be installed as follows:

```
cd $SAXDIR/lp
cat ehandler.ps | lp
```

Accessing the wys Script

So that users only need to type wys at the UNIX prompt, standard UNIX procedures can be used to make the script, $SAXDIR/bin/wys available.

There are two suggested ways to do this. Your system administrator can choose the method most suited to your system:

- Link the script into a standard path (i.e., an entry in the PATH environmental variable)

  ```
  ln –s $SAXDIR/bin/wys /usr/bin/wys
  ```

- If there is a global user shell login file, you can define the following environmental values

  ```
  setenv SAXDIR /usr/sax/rel
  setenv PATH $SAXDIR/bin/:$PATH
  ```

  Substitute the path for SHARP APL for UNIX on your system in the appropriate places.

These actions may have already been done in order to set up the sax script. If neither of these methods is implemented on your system, users can define SAXDIR and PATH in their shell login script in their home directory.

Operation

Activating the Printer

In order to make wys available for use, the UNIX system administrator should take the normal steps to activate an "lp" scheduler and start the printer.

Reseting the Printer

The printer can be reset with the command:

```
$SAXDIR/lib/lp/lwreset
```

Environment Variables

All wys environment variables are explained in Table 5.2. System defaults for each environment variable are set in the file $SAXDIR/etc/wysrc, when wys is installed. A different set of defaults can be saved by the system's UNIX administrator in $SAXCNF/etc/wysrc.

wys users can establish their own default environment variables by saving a modified copy of the file $SAXDIR/etc/wysrc as the file .wysrc in their $HOME directory.

Options

wys options can either be specified as environment variables or passed as arguments to the wys command. There is a specific order of precedence used by wys for selecting option values:

1. The wys command
2. The wys user environment file $HOME/.wysrc
3. The wys site environment file $SAXCNF/etc/wysrc
4. The distributed wys environment file $SAXDIR/etc/wysrc

A short overview of wys options can be obtained online by printing file $SAXDIR/lib/lp/wysopts.

Table of Options

Table 5.2, below, describes all wys options. The first column shows the name of the option when passed as a parameter to wys. The second column shows the name of the environment variable. The value shown after the name of the environment variable is the default set in $SAXDIR/etc/wysrc. The third column explains the purpose of the option.

Table 5.2. wys options. *(continues for 3 pages)*

Script	*Environment Variable (+ default)*	*Description*
wys -		Print the contents of standard input.
wys -b	BGflag=1	Print in the background. 0 = no 1 = yes

Script	*Environment Variable (+ default)*	*Description*
`wys -c`	`Copy=1`	The number of copies to be printed.
`wys -C`	`Charff="@"`	Specify a form-feed (new page) character. See the `n` option.
`wys -d`	`Prt=LW`	The name of the destination printer, if a printer *other* than the default printer is required.
`wys -f`	`Font=/APLitalic`	The PostScript font name. The possible values for this option depend on type of PostScript printer installed. See the printer documentation for more information. Font APLitalic selects the `wys` APLfont (if installed); otherwise Courier.
`wys -F`	`Filecov=0`	Select leading and trailing covers for each file printed by a `wys` request. See also the `j` option. 0 = none 1 = leading file cover 2 = trailing file cover 3 = leading and trailing file covers
`wys -h`	`Headflag=7`	Specify the page header text, if any, to appear on each page printed by `wys`. The page header can consist of some or all of: 0 = no header text 1 = page-header text. The system default is `$LOGNAME@$Machine`. User supplied header text can be specified using the `H` option. 2 = the name of the file printed. 4 = the date and time of printing. These options can be combined by adding the option numbers. Thus `h5` selects `h` options 1+4 (user supplied text, and date and time stamp). See also the `N` and `P` options.
`wys -H`	`Head=$LOGNAME@Machine`	Specify page-header text for use with the `h` option.
`wys -i`		Obtain list of names of files to be printed from standard input.
`wys -j`	`Jobcov=0`	Job cover specification. See also the F option. 0 = none 1 = leading job cover 2 = trailing job cover 3 = both job covers
`wys -l`	`Lansdcape=0`	Orientation of page. 0 = portrait 1 = landscape
`wys -L`	`Lines=99999`	Maximum number of lines printed per page. This count includes header lines, if present. The standard header generated by the `h` and `P` options occupies 3 lines.

Script	*Environment Variable (+ default)*	*Description*
`wys -m`	`Message=0`	This option is provided for use when `wys` is called from other software. If it is selected, `wys` returns a confirmation message to the calling shell each time a `wys` print request is successfully submitted. 0 = no 1 = yes
`wys -n`	`Newpage=0`	Select form-feed check at print position 1. When this option is selected a form-feed (new page request) is issued if a user supplied form-feed character is found in print position 1. See the C option. 0 = no 1 = yes
`wys -N`	`Name=""`	Specify the name of a header file, for example Name=header. If option `-h` is non-zero and `-N` is selected, the content of the file specified by the `N` option option is used as a page header.
`wys -o`	`Outfile=""`	Name of output file. This option allows `wys` PostScript output to be saved in a file, instead of being printed immediately. The `wys` command overwrites `Outfile`, if it already exists; but if multiple file names are specified in a single `wys` command, the output for the second and subsequent files is catenated to `Outfile`.
`wys -p`	`Points=10`	Specify the character size, in points. (72 points = approximately 1 inch or 2.54 centimeters).
`wys -P`	`Pagenr=0`	Select page numbering. This option causes a page number to appear at the top right of each printed page. See also the `h` and `S` options. 0 = no 1 = yes
`wys -q`		Display a summary of `wys` options.
`wys -r`	`Req=/dev/null`	Specify the name of a spool request file. When the option b=1 is selected, `wys` invokes the UNIX print scheduler from a background task. This means that the print request number that the scheduler sends to standard output is not available to the calling shell. If the `r` option is selected, the print request number is placed in the specified file and can then be used by the calling application. When the option `b=0` is selected, `wys` invokes the print scheduler in the foreground, and the `r` option is not required.
`wys -R`	`Remove=0`	Remove print file after printing.
`wys -s`	`STSCflag=0`	Select STSC APL character ROM translation. 0 = no 1 = yes
`wys -S`	`Startpg=1`	Specify start page number. This option controls the start page number used when the P option is selected, for example `-S2` causes page numbering to start at 2.

Script	*Environment Variable (+ default)*	*Description*
`wys -t`	`Tabs=8`	Specify tab stops. This option causes `wys` to use `tab` characters in the output, where possible.
`wys -v`	`Verbose=0`	Select verbosity. This option causes `wys` to return a user-friendly message to the calling shell each time a print request is submitted. 0 = no 1 = yes
`wys -w`	`Wrap=0`	Select line wrap. This option causes `wys` to wrap lines according to the values of the `W` and `X` options. See also the `Y` option. 0 = no 1 = yes
`wys -W`	`Wrapl=0`	Specify wrap length (0 means use site default) for use by w option.
`wys -x`	`Wrapc=\\`	Specify wrap character, for use by w option.
`wys -y`	`Wrapi=6`	Specify indentation after wrap, for use by w option.

Print Formats

The number of character positions per page that `wys` can print depends on the character size chosen. Table 5.3 shows the depth and width of the print area for the most common character sizes.

Table 5.3. `wys` ***print formats.***

Orientation	*Points*	*Print Depth*	*Print Width*
Portrait	10	73	87
Portrait	11	66	79
Landscape	10	51	133
Landscape	11	46	121

`wys` file `$SAXDIR/lib/lp/wystest` is a text file with all character positions filled. Printing it is a convenient way to investigate the effect of `wys` parameter settings.

Example:

```
wys -l1 -L56 -p12 -P1 $SAXDIR/lib/lp/wystest
```

This command prints `wystest` using 12-point characters printed 56 lines per page in landscape format with page numbering.

wys *Examples*

Display (query) `wys` options.

```
wys -q
```

Print two copies each, of files `abc` and `def`.

```
wys -c2 abc def
```

Print, using a list of file names obtained from standard input (thus, using a list of file names generated by another command). For example, print all files called `.wysrc` in all directories within `/usr`.

```
find /usr -name .wysrc -print | wys -i
```

Print the contents of standard input. For example, use `wys` to print the summary of `wys` options produced by the command `wys -q`.

```
wys -q | wys -
```

Print using `wys` called from an APL application. For example, use `wys` within a simple user-defined function that prints the contents of a function.

```
    ∇ print fns;H
[1]   ⍝11 ⎕svo 'H'⍝ open communication with the Host AP
[2]   fns←,>⍕1⍤⎕fd⍤>⍥⍕1 fns⍝ generate the text file
[3]   H←'echo "',fns,'" ¦ wys -f/APLitalic -' ⍝ print the text file
    ∇
```

A
National Use Characters

Like most environments that use 1-byte character codes, SHARP APL for UNIX lacks official encodings for many characters that are useful (or even essential) for some customer requirements. The only fully correct solution to this problem is to adopt complete support for the ISO/IEC standard IS 10646 (also known as ***Unicode 1.1***) when it becomes available. However, we recognize that customers need a way to circumvent the limitations of the single-byte character set.

SHARP APL for OS/390 addressed this issue by reserving 14 character positions whose display was permitted to vary from site to site. These character positions are known as ***National Use*** characters. However, this mechanism is inherently deficient; for example, it is possible for two SHARP APL/MVS sites communicating via e-mail to display the same text in two very different ways.

Despite these shortcomings, the National Use approach is currently the only acceptable solution. SHARP APL for UNIX therefore has the same type of mechanism incorporated into it.

But the UNIX product has a character set based on the ASCII codeset, which, coupled with the need to provide connectivity with the OS/390 product, adds to the problems involving site-specific interpretations of character codes. One OS/390 APL character that is considered National Use, for example, is the "dollar sign" (in typical North American National Use mapping), which is a "shell meta character" in UNIX, and therefore must be treated specially in some SHARP APL interface code.

The compromise adopted for SHARP APL for UNIX is as follows:

1. The mapping used by the NSVP/SAMI interface to MVS will be configurable. $\square av m$ must always return the mapping presently in use, so this configuration is potentially hazardous to change. Any site that requires this configuration option should contact their SHARP APL support person for assistance.

2. Eight of the characters that are present in the SHARP APL/MVS character coding only as North American variants of the National Use characters are considered to be ordinary members of the SHARP APL for UNIX character set. These characters are detailed in Table A.1 (assuming the standard $\square av m$ mapping).
3. Twenty characters have been set aside in the SHARP APL for UNIX character set for site use as National Use characters. These characters are detailed in Table A.2 (assuming the standard $\square av m$ mapping).
4. The terminal translation tables used by SHARP APL for UNIX fall into several categories, but most of them may be customized by a site if this turns out to be necessary. Any site that finds a need to customize display or entry of National Use characters should contact their SHARP APL support person for assistance.

Table A.1. Ordinary characters.

MVS $\square av$ NU positions				***Corresponding UNIX $\square av m$ codes***		
NU code	***$\square av$ index (dec)***	***(hex)***	***North American meaning***	***UNIX $\square av$ mapping***	***UNIX $\square av$ index (dec)***	***(hex)***
QNU1	5	05	Cent Sign	ZCENT	250	FA
QNU4	4	04	Dollar Sign	ZDOLLAR	36	24
QNU5	196	c4	PL/1 Not	ZPL1NOT	182	B6
QNU6	197	c5	Split Bar	ZSTILE	124	7C
QNU7	198	c6	Accent (Grave)	ZBQUOTE	96	60
QNU8	199	c7	Number Sign	ZHASH	35	23
QNU9	200	c8	At Sign	ZAT	64	40
QNU10	201	c9	Double Quote	ZDQUOTE	34	22

Table A.2. National Use characters.

MVS □av positions				*Corresponding UNIX NU codes*		
Mnemonic Q-code	*□av index (dec)*	*(hex)*	*North American meaning if NU*	*UNIX □av Mnemonic*	*UNIX □av index (dec)*	*(hex)*
QNU2	193	c1	Vertical Bar	ZNU1	159	9F
QNU3	194	c2	Exclamation	ZNU2	156	9C
QNU11	202	ca	Not (High Tilde)	ZNU3	158	9E
QNU12	203	cb	Left Brace	ZNU4	155	9B
QNU13	204	cc	Right Brace	ZNU5	157	9D
QNU14	205	cd	Backslash	ZNU6	128	80
QCCONST	13	0d		ZNU7	160	A0
	23	17		ZNU8	163	A3
	195	C3		ZNU9	165	A5
	192	C0		ZNU10	167	A7
	24	18		ZNU11	170	AA
	227	E3		ZNU12	181	B5
	228	E4		ZNU13	183	B7
QTDELTA	84	54		ZNU14	184	B8
QSDELTA	85	55		ZNU15	186	BA
	221	DD		ZNU16	192	C0
	224	E0		ZNU17	198	C6
	225	E1		ZNU18	203	CB
	226	E2		ZNU19	208	D0
	229	E5		ZNU20	209	D1

SHARP APL *for UNIX*

Auxiliary Processors Manual

JUMP TO ...

CHAPTER 1. OVERVIEW
CONTENTS
PREFACE
MASTER INDEX
USING THIS DOCUMENTATION

HANDBOOK
AUXILIARY PROCESSORS MANUAL
FILE SYSTEM MANUAL
INTRINSIC FUNCTIONS MANUAL
LANGUAGE GUIDE
SYSTEM GUIDE
SVP MANUAL

© 2000 by Soliton Associates Limited

44 Victoria Street, Suite 2100, Toronto, Ontario, Canada M5C 1Y2

Pubcode: UW-033-0501 (0007)

SHARP APL *for UNIX*

DEVELOPER LIBRARY

Auxiliary Processors Manual

SOLITON ASSOCIATES

© 2000 by Soliton Associates Limited–Printed in Canada
44 Victoria Street, Suite 2100, Toronto, Ontario, Canada M5C 1Y2

All rights reserved. No part of this work may be reproduced in any form or by any means without prior written permission of the copyright owner.

Soliton Associates Limited shall not be liable for errors contained herein or for incidental consequential damages in connection with the use of this material. The information contained within this document is subject to change without notice.

Acknowledgments

SHARP APL is a registered trademark of Soliton Associates Limited.

SVP and NSVP are trademarks of Soliton Associates Limited.

IBM and AIX are registered trademarks of International Business Machines Corporation.

RISC System/6000 is a trademark of International Business Machines Corporation.

UNIX is a registered trademark of UNIX System Laboratories, Inc.

Sun Workstation is a registered trademark of Sun Microsystems, Inc.

Sun, Solaris, SunOS, Open Windows, and SunView are trademarks of Sun Microsystems, Inc.

All other products referred to in this document are trademarks or registered trademarks of their respective trademark holders.

Publication code: UW-033-0501 (0007)

Contents

Tables and Figures vii

Preface

Introduction ix
Chapter Outlines ix
Conventions x
Documentation Summary x
Contacting Soliton Associates xi

1. Overview

What is an Auxiliary Processor? 1-1
The AP Daemon 1-1
The SVP and NSVP 1-2
Shared Variables/Names 1-3
Access Control 1-3
Offers to Share 1-4
System Auxiliary Processors 1-4

2. AP1: S-task

Introduction 2-1
Initiating an S-task 2-1
Modes and Controls 2-2
Duration of an S-task 2-2
Prefix Codes 2-3
Sending Prefix Codes 2-3
Receiving Prefix Codes 2-4

Freestanding S-task 2-7
Default Value of Independence Bit 2-8
Reconnecting to a Freestanding S-task 2-8
Sample Uses of the S-task Facility 2-9
Utility Functions `talk` and `apl` 2-9

3. AP11: Host AP
Introduction 3-1
AP11 Commands 3-2
Public Workspace `1 hostap` 3-2
Return Codes from Host AP 3-3
The Utility `host` 3-3
Commands for Closed Files 3-4
Commands on Open Files 3-6
Other Commands 3-10

4. AP124: Full-Screen AP
Introduction 4-1
AP124 and the Session Manager 4-2
Full-Screen Dialogue 4-2
Passing Instructions to AP124. 4-3
Communicating with AP124 4-7
Establishing Contact 4-7
Using the Shared Variable to Control the Screen 4-8
AP124 Return Codes 4-10
Refreshing the Screen 4-12
The Geography of Fields 4-12
Adjacent Fields 4-12
Row-Structure of a Field 4-13
The Format Table 4-16
AP124 Commands 4-20
Set Field Formats 4-21
Reset Field Attributes 4-22
Transmit Data for Display 4-23

Set Cursor Position 4-23
Sound Alarm 4-24
Repaint the Screen 4-24
Set Character Attributes 4-24
Write-Wait-Read 4-25
Receive Data from Display 4-26
AP124 Differences Under UNIX and OS/390 4-28
Public Workspace `1 ap124` 4-29

Index IX-1

Tables and Figures

Figure 2.1. Shared variable interactions between tasks. 1-3
Table 3.1. AP11 basic commands. 3-2
Table 3.2. AP11 return codes. 3-3
Table 3.3. The open command control options. 3-7
Table 3.4. Mode values for the *lseek* command. 3-9
Table 3.5. The *fcntl* command. 3-9
Table 3.6. Values for *lock* ***type*** parameter. 3-10
Table 4.1. AP124 return codes. 4-11
Table 4.2. AP124 format table. 4-17
Table 4.3. AP124 commands. 4-20
Table 4.4. Reset field attributes. 4-22
Table 4.5. Skip and trailing characters. 4-22
Table 4.6. Writing to the screen. 4-23
Table 4.7. Writing character attributes. 4-25
Table 4.8. Releasing the screen. 4-25
Table 4.9. Reading from the screen. 4-26

Preface

Introduction

This document discusses SHARP APL *auxiliary processors* (APs). An AP is a task that accepts commands and returns results via the *Shared Variable Processor* (SVP). Most APs operate under a command language unrelated to APL, doing work for an APL task that it is not itself equipped to do. SHARP APL for UNIX is supplied with three auxiliary processors, AP1, AP11, and AP124, which are documented in separate chapters in the manual.

Rather than reproduce existing material, references to other SHARP APL for UNIX publications are supplied where applicable.

Chapter Outlines

This document is organized into the chapters described below.

Chapter 1, "Overview," introduces the concept of an Auxiliary Processor, discusses the shared name communication, and lists the three system APs supplied with SHARP APL.

Chapter 2, "AP1: S-task," discusses the APL auxiliary processor.

Chapter 3, "AP11: Host AP," discusses the Host auxiliary processor.

Chapter 4, "AP124: Full-Screen AP," discusses the Full-Screen auxiliary processor.

Conventions

The following conventions are used throughout this documentation:

`⎕io←0`	Although the default value for `⎕io` in a clear workspace is *one*, all examples in this manual assume that the index origin is *zero*.
α and ω	Two APL symbols are used in syntax descriptions of monadic and dyadic functions; *alpha* (α) identifies the left argument and *omega* (ω) identifies the right argument.
`constant width`	Examples of non-APL input and output are represented in a `constant width` typeface.
default (%) prompt	Examples in this book assume you are using the C shell. The default prompt for the C shell is the *percent sign* (%).
passno	Argument variables appear in *sans-serif italic* type.
[ ]	Optional arguments are shown between square brackets. For example, in the following APL statement, the argument *passno* is optional:

'*filename*' `⎕stie` *tieno [,passno]*

These default environment variables represent frequently used pathnames in SHARP APL for UNIX documentation and scripts:

`$SAXLIB`	→	`/usr/sax`
`$SAXDIR`	→	`/usr/sax/rel`
`$SAXCNF`	→	`/usr/sax/local`
`$HOME`	→	home directory of the current user.

Documentation Summary

The SHARP APL for UNIX reference documentation is published in several volumes. One or more of the following may be referenced in this manual.

SHARP APL for UNIX,

- *Handbook,* publication code UW-000-0401

- *Language Guide,* publication code UW-000-0802
- *System Guide,* publication code UW-000-0902
- *SVP Manual,* publication code UW-001-0501
- *File System Manual,* publication code UW-037-0501
- *Intrinsic Functions Manual,* publication code UW-007-0501

For a complete list of SHARP APL publications, please refer to the ***Support*** link on the Soliton Associates website: `http://www.soliton.com.`

Contacting Soliton Associates

Problems or questions regarding your SHARP APL system or this guide should be directed to the Soliton Associates Technical Support group:

`support@soliton.com`

Requests for additional technical documentation should be directed to the Soliton Associates Sales group:

`sales@soliton.com`

1
Overview

What is an Auxiliary Processor?

A SHARP APL ***auxiliary processor,*** AP, is a program that provides a specialized service to an APL task, doing work that the APL task is not itself equipped to do. The AP is not usually written in APL and often operates under a command language unrelated to APL.

APs communicate with APL tasks using *shared names* via multi-user processors called the *shared variable processor* (SVP) and the *network shared variable processor* (NSVP), installed as pseudo-device drivers in the UNIX kernel.

The AP Daemon

When you start an APL session, a special process (called the AP daemon) is intialized on your behalf. As soon as you offer to share a name with an AP, the special process starts the AP so that it will reciprocate your offer to share. If you do not require any of the services APs provide, no AP tasks are started for you.

The AP daemon makes initial startup of an APL session faster than it would be if all APs were started for every session. However, because the AP daemon must start the AP before the AP can reciprocate offers to share with your task, your first transaction with an AP may be slightly slower than any that follow.

The SVP and NSVP

A name that two programs share is like a window between them; indeed, you can say that a shared name is a variable that exists in two different environments at the same time. This interface is provided by a pseudo-device installed in the SHARP APL kernel called the ***Shared Variable Processor***, SVP: it manages the transfer of data and arbitrates access control between the cooperating programs.

The SVP permits one APL task to share a variable (and thereby communicate) with other APL tasks, or with an auxiliary processor. It provides the link through which an APL workspace is able to start and then communicate with subordinate APL tasks, or to communicate with the independent APL tasks of other users. Shared name communication can be between any tasks which follow SVP protocol, including those outside of APL.

The Network Shared Variable Processor, NSVP, enables program-to-program communication between connected systems using an extended version of the SVP protocol. Information is transferred along a communications line between NSVP processors residing on each of the systems. Communication between the NSVP and the local task is handled by the SVP.

Shared name communication involves:

- managing all name sharing for your task and for all other tasks running under the same UNIX system at the same time
- maintaining tables of shared names and the tasks (auxiliary processors or other APL tasks) that are sharing them
- providing intermediate storage for the values set by one partner but not yet used by the other
- enforcing rules that prevent conflict over whose turn it is to set or use a shared name.

The SVP, the NSVP, and the system variables and functions that facilitate the use of shared names are fully described in The SHARP APL for UNIX *SVP Manual*.

Shared Variables/Names

Shared variables provide a way for APL programs to communicate with other programs whether written in APL or not. When two programs have a variable in common, either program can *set* the value so the other can then *use* it. The term *shared name* originated from the realization that it is possible to share functions and other named objects between users.

For example, suppose Task 1 and Task 2 represent the active workspaces of two APL users who have agreed to share the name `x`. The left side of Figure 2.1 shows events in Task 1's workspace, and the right side events in Task 2's workspace. The fact that `x` is shared leads to the following simple example . Task 2 assigns to `x` the value `abc def`. Almost simultaneosly, Task 2 evaluates `2×x` and gets the answer `10`.

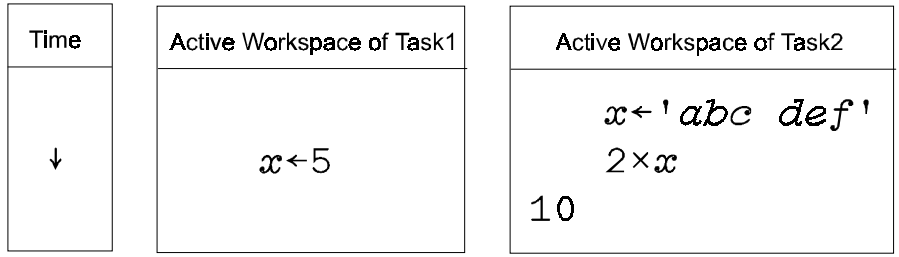

Figure 2.1. Shared variable interactions between tasks.

Of course the reason for this result is that after Task 2 gave `x` the value `abc def`, Task 1 gave `x` a new value. Since `x` is shared, Task 2 now sees the ***new*** value.

Access Control

To prevent the conversation from getting "out of sync", a mechanism called *access control* can impose delays to prevent (for example) two successive uses of the shared variable by a program without the partner having set it between them. When an attempt which would violate the specified access control is made, the inhibited program waits at the point of setting or using the shared variable until the partner has taken the appropriate complementary action.

The rules that govern access control are well-documented in the SHARP APL for UNIX *SVP Manual, Chapter 2*.

Offers to Share

A conversation using shared variables is initiated by one of the two partners by making an *offer to share* a name for a variable that will serve as the channel of communication between them. At this point the *degree of coupling* for this variable is one. When the other partner makes a matching offer, the connection is complete: the variable is shared, and the degree of coupling is two. Any user of shared variables can inquire about offers to share available to that user and about the degree of coupling of a particular variable. The system functions □*svo* and □*svq* make this possible.

System Auxiliary Processors

The auxiliary processors described below are supplied with SHARP APL for UNIX. Usually an AP is uniquely identified by its own number, which is generally but not always below 100

AP1 — **The APL, S-task, auxiliary processor**. By sharing a name with AP1, one APL task can start and control another APL task. Input and output is handled via a shared name. AP1 is documented in "Chapter 2. AP1: S-task".

AP11 — **The Host auxiliary processor (Host AP)**. By sharing a name with AP11, an APL task can pass commands to and receive responses from a UNIX shell. Commands may be *simple,* for transmitting a character vector, or *enclosed,* for transmitting an enclosed array. See "Chapter 3. AP11: Host AP" for further information.

AP124 — **The full-screen auxiliary processor**. By sharing a name with AP124, an APL task can pass commands and data to a full-screen manager and capture data from fields on the screen. These full-screen capabilities are similar to those provided by SHARP APL for MVS. "Chapter 4. AP124: Full-Screen AP", describes the underlying concepts of this auxiliary processor.

2
AP1: S-task

Introduction

An S-task is an APL task controlled by way of a variable shared with AP1, the APL auxiliary processor. An S-task has no screen or keyboard. It is initiated and controlled not by a person but by another APL task. Usually, it exists to provide some special service either to the task that created it or to a group of cooperating tasks.

A task controlled from a terminal or keyboard is called a T-task. The human user who controls a T-task types commands at the keyboard, and receives replies at a screen or printer. The task that controls an S-task also gives commands and receives replies. The commands may be quite similar or even identical to those a human user would type. What is different is the way the controlling task sends commands and receives replies.

The controlling task formulates commands as *character vectors* that contain characters a human user might type at a keyboard. The controlling task tells the S-task what to do by assigning the character string to a variable that the controlling task shares with AP1. It receives its replies in the same way: AP1 sets the shared variable with a character string that represents the resulting display.

Initiating an S-task

The controlling task initiates an S-task by offering to share a name with AP1, as follows:

```
1 ⎕svo 'sh'
```

When AP1 accepts the offer to share, the controlling task sees the result of `⎕svo 'sh'` become `2`. AP1 immediately starts an S-task under the same UNIX account as the controlling task. The S-task session starts in a clear workspace, which (unless indicated otherwise) has the same workspace size as the task that initiated it.

This procedure differs slightly from that used in SHARP APL for MVS. There, the controlling task must sign on, supplying a user ID and password. Under UNIX, the situation is different. Once the UNIX system has accepted a login, no further sign-on is required for any APL task, whether T-task or S-task.

Modes and Controls

During the life of a T-task, the SHARP APL system may at various times be in one of several modes. The human user can deduce which mode the APL session is in from its behavior and history. In the control of an S-task, these modes are made explicit. Each time the controlling task sets a new value for the variable it shares with AP1, it must include a `4`-byte *prefix*. The normal prefix consists of `(4⍴0)@ ⎕av`. Similarly, each time AP1 sets a new value for the variable that it shares with the controlling task, it includes a `4`-byte prefix to describe the rest of the transmission. The prefixes sent by AP1 contain a lot of information; see the section “Prefix Codes,” below.

Duration of an S-task

An S-task continues until terminated, much as a T-task does. If the controlling task fails to supply input to an S-task, it may wait indefinitely (much as a T-task waits if its human user goes away from the keyboard without logging off). Ordinarily, an S-task requires that the controlling task continue to share the controlling variable with AP1. However, an S-task can be made *freestanding*. The controlling task notifies AP1 that it will retract the name it shares with AP1, but that the S-task is not to be terminated. After receiving acknowledgment from AP1, the controlling task may then retract the name it shares with AP1. Usually this is done after starting execution of a server program in the S-task.[1]

Once started, an S-task continues until one of the following occurs:

1. A server program is basically a loop; it waits for requests to be sent to it, handles requests when it receives them, and otherwise waits for an external event to drive it. Server programs are briefly discussed later in this chapter.

- The controlling task sends a sign-off command, which (with the normal prefix) is:

```
sh←((4⍴0)@ ⎕av),')off'
```

- The controlling task ceases to share the controlling variable by explicitly retracting it, expunging it, clearing the workspace in which it exists, etc., ***without*** first notifying AP1 that the S-task is to remain active without a controlling variable (see "Freestanding S-task," in this chapter).
- An external command terminates it. This might be the UNIX command `kill` or the APL system function ⎕*bounce*.

Prefix Codes

Each character vector sent to or received from AP1 must begin with a prefix.

Sending Prefix Codes

Each time the controlling task sets the variable it shares with AP1, the value it sets must be a character vector representing one line of input. The line need not include a trailing newline or carriage return character.

There are five possible prefix codes. One of these, the code for APL input, precedes any line sent to the session manager or the APL interpreter, including system commands, lines entered for immediate execution, lines in response to ⎕- or ⍞-input, lines sent in response to ⎕*arbin*, and lines sent to the ∇-editor. The other four codes stand alone. If you send one of the other codes as part of a longer line, the balance of the line is ignored. The five prefix codes are as follows:

`(0 0 0 0 @⎕av)`**,text**	**APL input.**
`0 0 1 0 @⎕av`	**Not meaningful** in SHARP APL for UNIX. In SHARP APL for MVS, this is a request for accounting information, and the values returned are those in `3↑⎕ai`. In the UNIX version, the values returned are `3⍴¯1`.
`0 0 2 0 @⎕av`	**Interrupt signal.**

`0 0 3 0 @ ⎕av`	**Turn *independence bit* on.** It notifies AP1 that you permit the S-task to become freestanding; that is, to keep running after you retract the controlling variable that you share with AP1. (See "Freestanding S-task," below)
`0 0 4 0 @ ⎕av`	**Turn *independence bit* off.** It notifies AP1 that you assume responsibility for your S-task, so that it may survive only as long as you continue to share with AP1 the variable by which you control it. See "Freestanding S-task," below.

Receiving Prefix Codes

Input to an S-task differs from output from an S-task in two important respects:

- Whereas almost all forms of *input* are sent in the same way (with prefix `0 0 0 0@ ⎕av`), *output* from an S-task carries with it information about the mode in effect when it was sent, and information about the mode that will govern interpretation of the next input. This information is carried in the fourth byte of the prefix of each segment. It is up to the controlling task to interpret it.
- Whereas input to an S-task (like input to a T-task) is sent one line at a time, output from an S-task (like output for a T-task) may be spread over many lines. For this reason, output is *segmented.* Each value to which AP1 sets the shared name may contain not just one line of response, but many segments. Within the character string set by AP1, each segment starts with `2` bytes indicating the length of the segment. A single setting of the shared name may contain only a single segment, or may consist of many segments.

A complete reply from the S-task may span many sets of the shared variable. It is up to the controlling task to unpack the segments within a single value of the shared variable, and to read the shared variable again when the prefix of the last unpacked segment indicates that there is more to come.

Separating Segments

Each segment consists of 2 bytes of length information, 4 bytes of prefix code, and the text of the segment. The length is encoded by treating the first 2 bytes as a 16-bit binary number. The value written there is the gross length, including the 2 bytes of length information and 4 bytes of prefix.

The following excerpt shows how the controlling task might separate segments from a single setting of the shared variable, here given the name *sh*. Note that *sh* is interlocked to prevent rereading it, so the program contains only one use of *sh*, to capture its value and assign it to *s*.

```
      s←sh
      text←''
more: l←256⊥⎕av⍳2↑s
      prefix←4↑2↓s
      text←text,6↓l↑s
      s←l↓s
      prefix interpret text   ⍝ Act upon text and prefix
      →(0≠⍴s)/more
```

Prefix Codes from AP1

The following are 0-index *indices* into ⎕*av*:

`0 0 1 0`	Your preceding entry was invalid.
`0 0 2 0`	The APL interpreter lacked space to process your preceding entry.
`0 0 3 0`	Although AP1 is active and able to reply, no S-task has been started and so no reply from APL is possible. This option is not implemented, and is available in SHARP APL for MVS only.
`0 0 0` *mode*	Your preceding entry has been accepted. Byte 3 describes the text that accompanies this reply, and/or the S-task's present state.

Binary Representation of Prefix Codes

Let the variable *mode* be defined as follows:

```
mode←(8⍴2)⊤⎕av⍳3@ prefix⊣⎕io←0
```

A `1` in each of the eight positions within *mode* has the following significance:

`0@` *mode*	The S-task is now in immediate execution mode.
`1@` *mode*	The accompanying text is a blot. This value is not used in SHARP APL for UNIX, which does not use blots.
`2@` *mode*	The S-task has not been started.
`3@` *mode*	The S-task is in the state in which its keyboard appears locked to the controlling task; in order to request input, you must send an interrupt signal. This value is never used in the UNIX environment of SHARP APL, since UNIX generally assumes full-duplex communication, rather than the reversing half-duplex for which this control is intended.
`4@` *mode*	An output interrupt has been received. You should reset the S-task's virtual cursor to the beginning of the next line. This control is necessary because the buffering of transmissions between APL and the end user may leave the APL interpreter unable to predict the state of the device it is addressing.
`5@` *mode*	Arbitrary I/O. This bit is used both to describe the output that accompanies the prefix and (where input is expected) to indicate that your next transmission will be interpreted as arbitrary input.
`6@` *mode*	The accompanying text is a prompt for input. This is your cue that it is your turn to send text to AP1.
`7@` *mode*	The accompanying text is output, and there is more output to follow (either a continuation of this display or the prompt for input). This is your cue to read another segment, or (if necessary) again use the shared variable to receive another segment.

Common Values for Prefix Codes

The vector that follows treats the last byte of the prefix as an integer code. For completeness, it includes some codes associated with the timesharing system of SHARP APL for MVS. Because that system handles its own logons for remote

users, it includes provision for the generation of a blot, used as part of the sign-on procedure, and also elicitable during an APL session. The UNIX version does not handle sign-on, and does not support the blot.

00000001	1	Normal output, more to come.
00000010	2	Normal output, final (i.e. prompt for ⍞ or ∇-editor input).
00000100	4	□*arbout* or □*arbin*, more to come.
00000101	5	□*arbout*, final.

Codes 4 and 5 are indeed defined and used as shown despite the apparent reversal of the meaning of the last bit.

00000110	6	The accompanying text is a prompt for □*arbin* input.
00001001	9	APL requests the controlling task to reset the position of its cursor, following an interruption of output.
00010001	17	The S-task is ready for input, but its keyboard is in a locked state, awaiting an interrupt signal. Cannot arise.
00100000	32	S-task not signed on to APL.
01000001	65	Part of a blot; there is more output to come. Not used.
10000010	130	Prompt for immediate-execution input; this is the standard prompt for a new input from the controlling task when APL has completed execution of the line previously sent to it.
11000010	194	Blot prompt for immediate-execution input. Cannot arise.
11100010	226	Sign-on blot prompt. Not used.

Freestanding S-task

It is possible for the controlling task to retract the offer to share the controlling variable with AP1. When it does so, the S-task is left with no means of control, and is said to be *freestanding*.

A freestanding S-task may or may not continue to run. Its fate depends on the current value of the *independence bit* that AP1 uses for each S-task. The controlling task turns the S-task's independence bit on by the transmission `(0 0 3 0)@ ⎕av`, and turns it off by the transmission `(0 0 4 0)@ ⎕av`.

The independence bit affects the S-task as follows:

Independence bit on The S-task continues indefinitely, executing the last instruction given to it while it was in immediate execution mode. When it completes execution (and enters any input mode for a new instruction), it terminates. A freestanding S-task provides a service equivalent to that of an N-task. Note that the instruction the S-task is executing may be a program that loops indefinitely, perhaps causing the S-task to act as a resident server that runs until some prearranged criterion is met, or until you terminate it from outside.

Independence bit off The S-task continues indefinitely as long as the controlling task continues to share the controlling variable with AP1. The S-task may return to immediate execution mode (awaiting a new instruction). While the controlling task sends it no new instruction, it simply waits, inactive, until it receives one. When the controlling task retracts the name it shares with AP1 (for any reason, by any means), AP1 immediately terminates the S-task.

Default Value of Independence Bit

SHARP APL assumes that by default the independence bit is *off*. A freestanding S-task cannot survive unless you first explicitly set the independence bit *on*.

Reconnecting to a Freestanding S-task

You may need to reconnect to an S-task when the S-task is still running but you are no longer sharing with AP1 the name formerly used to control it.

To reconnect, you must have the same user number and clone ID as you had when you started the S-task; and, you must offer the same name. When AP1 accepts the share, you are again connected to the S-task.

If you left the task running an endless program that waits until it receives a request (i.e., when the S-task is a server), you may have to transmit an interrupt before you regain direct control of it.

Sample Uses of the S-task Facility

The distributed workspace `1 tools` includes the utility functions *apl* and *talk*. Each uses an S-task (starting it if it is not already running). The utility *apl* takes as its argument one *line* of APL and passes it to the S-task to execute. The niladic function *talk* conducts an entire interactive session with an S-task; the utility function *yes* prompts the user where needed.

Utility Functions *talk* and *apl*

The utility *talk* is intended as an exercise to illustrate several capabilities of the S-task interface. For production work, a utility such as *apl* is more likely to be useful than this model of *talk*, which does not attempt to handle all possible scenarios of input and testing of tasks.

In the definition of *apl*, when the argument starts with the character ←, the function returns control to the calling environment immediately, without waiting for a result from the S-task. That is the appropriate way to launch a freestanding server.

```
   ∇ Δ←yes;⎕ec
[1]    ⍝ require a reasonable attempt at yes or no
[2]    L1:→(1≠⍴Δ←((4 4 ⍴'yes y   no  n   ')∧ . =4↑⍞)/ 1 1 0 0)↓0
[3]   'Please answer Yes or No' ◇ →L1
    ∇

    ∇         talk;now;num;sav;var;⎕io;⎕trap;⎕PR
[1]         ⍝ Manage an S-task session
[2]           ⎕io←0 ◇ ⎕trap←'⋄1002 1003 1005 c →Break' ⎕PR←' '
[3]           var←'APL',⍕num←⍳,>(2=⎕nc'SOX')@ '⍳0'⊃'SOX'
[4]           →(0=⎕svo var)↓Initial
[5]           ('Cannot share ',var) ⎕signal (0≠(>,<num,1 0)⎕svo var)↓11
[6]           ¬0 1 1 1 ⎕svc var ◇ →Receive
[7]    Input: now←⍞ ◇ sav←(now ≠ ' ')/now
[8]           →(('')exit'⊃')off')⍳<sav) @  Exit, Off, Send
[9]     Send: ⍎var,'←(0 0 0 0@ ⎕av),now'
[10] Receive: now←⍎var
[11] Display: num←256⊥⎕av⍳2↑now ◇ ⍞←4↓sav←2↓num↑now
```

```
[12]          now←num↓now ◇ →(6>⍴now)↓Display
[13]          →(0@  (2⍴2) ⊤ ⎕av⍳3@ sav)↓Receive ◇ →Prompt
[14]   Break: 'Send an interrupt to the S-task?'
[15]          →yes↓Initial ◇ ⍎var,'←0 0 2 0@ ⎕av' ◇ ⊣⍎var ◇ →Receive
[16] Initial: ⎕←6⍴' '
[17]  Prompt: →Input
[18]     Off: ⍎var,'←0 0 4 0@ ⎕av' ◇ ⊣ ⎕ex var
[19]    Exit:
    ∇

    ∇         Δ← ⍺ apl ⍵;lin;now;num;off;var;⎕io
[1]         ⍝ Send one line to an S-task
[2]           ⎕io←0 ◇ ⍎(2=⎕nc'⍺')/'SOX←⍺'
[3]           var←'APL',⍕num←⍎,>(2=⎕nc'SOX')@ '⍳0'⊃'SOX'
[4]           Δ←'' ◇ →(off←0=⎕svo var)↓Command
[5]          ('Cannot share ',var) ⎕signal (0≠(>,<num,1 0)⎕svo var)↓11
[6]           ⊣0 1 1 1 ⎕svc var ◇ →Receive
[7]  Command: now←(⍵ ≠ ' ')/⍵
[8]           →((')break'⊃')off')⍳<now) @  Break, Off, Send
[9]     Send: now←'←'=1↑⍵ ◇ ⍵←now↓⍵
[10]          ⍎var,'←(0 0 0 0@ ⎕av),⍵' ◇ →now⍴0
[11] Receive: now←⍎var
[12] Display: lin←256⊥⎕av⍳2↑now
[13]          ⎕signal (0=⎕av⍳4@ now)↓11 ◇ num←(8⍴2)⊤⎕av⍳5@now
[14]          Δ←Δ,6↓lin↑now ◇ now←lin↓now
[15]          →(6>⍴now)↓Display ◇ →(3@ num)↑On
[16]          →(6@ num)↓Receive
[17]       On: →off↓0 ◇ off←0 ◇ Δ←'' ◇ →Command
[18]   Break: ⍎var,'←0 0 2 0@ ⎕av' ◇ now←⍎ var ◇ →0
[19]     Off: ⍎var,'←0 0 4 0@ ⎕av' ◇ ⊣⎕ex var ◇ →0
    ∇
```

3
AP11: Host AP

Introduction

AP11 is the *Host AP* that provides for communication between an APL task and the UNIX operating system. It has two modes of operation:

Simple You transmit a *character vector* to AP11, which treats the characters as a command in a Bourne shell, and passes it unedited to UNIX. When you next use the variable that is shared with AP11, it contains the response from the shell as a character vector with embedded newline characters as needed. For the vocabulary of commands you can transmit in this manner, and the interpretation of the responses, consult the User documentation for your UNIX system.

Enclosed You transmit an *enclosed array* to AP11 to facilitate processing of a small vocabulary of commands for access to UNIX files and other UNIX facilities.

You can change from one mode to the other without having to reshare the variable. Sending a simple command gets a simple reply from UNIX. Sending an enclosed command gets a enclosed reply containing the requested data in the second element; the first element of the reply is usually `0`, but can contain a return code to indicate the nature of any problem AP11 encounters carrying out your command.

AP11 Commands

When you send AP11 an enclosed command, the first element contains the command name, and subsequent elements contain additional data appropriate to each command.

The basic commands, listed in Table 3.1, fall into two groups. In the first are the commands that apply to closed files: in each case, the command opens a file, performs its action, and closes the file when the action is complete. In the second group are commands that give you access to useful system facilities.

Table 3.1. AP11 basic commands.

Command	*Purpose*
`get`	Read contents of UNIX file.
`put`	Write character list to UNIX file.
`size`	Inquire about size of UNIX file.
`chdir`	Change current directory.
`shell`	Execute command in a UNIX shell.

A third group of AP11 commands allows an APL program to control a UNIX file while placing all responsibility for the use of the file on the APL program. These facilities require a thorough knowledge of the UNIX system calls they represent.

Public Workspace `1 hostap`

You will find the public workspace `1 hostap` extremely useful when working with UNIX files. The utility function `host` is available here (as well as in the public workspace `1 tools`). Cover functions for each AP11 command provide a shorthand notation for construction of application models. For example, instead of the expressions

```
      11 ⎕svo 'HOST'
1
      1 ⎕svc 'HOST'
1
      HOST←'size'⊃'myfile'
      r←>'⍴1↓HOST
```

you can enter

```
r←host 'size'⊃'myfile'      ⍝ shares HOST if not shared already
```

or

```
r←size 'myfile'
```

Return Codes from Host AP

A return code is an integer scalar. A positive value is an error code as generated by the UNIX operating system.[1] The unique negative return codes are described in Table 3.2.

Table 3.2. AP11 return codes.

Value	*Meaning*
0	Normal return; operation successful.
¯1	Data too large for shared names interface.
¯2	Data type is inappropriate; *character vectors* only.
¯3	Unknown command.
¯4	Argument error in local command.
¯5	Unable to grab keyboard.
¯6	Error occurred during *tty*.

The Utility *host*

The workspace `1 tools` contains the definition of a function *host*. Its argument is either a character vector which is passed directly to UNIX for execution, or an enclosed command to the Host AP (as described in the foregoing sections). When the enclosed action is to retrieve data from a UNIX file, its result is the requested

1. See `/usr/include/errno.h`.

data. Otherwise, its result is an empty table. The function makes use of the global name *HOST*; when the name is not already shared, the function takes care of offering it and setting the appropriate interlocks.

The function *host* separates the elements of the result for you, and opens them to return the requested data, for example:

```
    ∇         ∆← ⍺ host ⍵ ;sav;var
[1]         ⍝ Execute one Host AP request
[2]           ⍙ (2=⎕nc'⍺')/'SOX←⍺'
[3]           var←'HOST',⍕sav←⍙,>(2=⎕nc'SOX')@'⍳0'⊃'SOX'
[4]           →(0 = ⎕svo var) ↓ Set
[5]           ('Cannot share ',var) ⎕signal (0≠(>,<sav,11 0)
               ⎕svo var)↓11
[6]           ⊣0 1 1 1 ⎕svc var
[7]      Set: ∆←'' ◇ ⍙var,'←⍵'
[8]      Ref: →(sav ≡ ⊃sav←⍙var) ↓ Keep
[9]           →((,<0)≡ 1↑sav)↑Just
[10]          ('Host error ',⍕>0@ sav) ⎕signal 541
[11]    Just: sav←>1@ sav
[12]    Keep: ∆←∆,sav
[13]          →('' ≡ sav) ↓ Ref
    ∇
```

An optional left argument allows for a *sign-on index* obtained from a remote host. The default sign-on index when you don't use a left argument is the one that applies to the machine on which your active workspace is running. See ⎕*svn* in *SVP Manual, Chapter 3*. A second global name *SOX* is associated with this usage.

Commands for Closed Files

The commands in this group are *get*, *put*, and *size*. Each of these commands opens a file, performs its action, and closes the file when the action is complete.

'*get*' ⊃ *file* [⊃ *length* [,*offset*]]

The *get* command may be a 2 or 3-element enclosed array.

- The first element contains the command '*get*' as a character vector.

- The second element contains the name of the *file* from which data is to be read, as a character vector in standard UNIX form. It may be relative to your current directory (from which you started your APL task) or absolute (a complete UNIX path from the root directory). If it refers to an APL file, it should include the suffix `.sf` or `.sw` (unlike APL system commands and file system functions).
- The optional third element indicates a specific range of bytes to be read, that is the *length*. It is a `1`- or `2`-element integer vector consisting of the number of bytes to be read and, if not from the beginning of the file (for which the default offset is `0`), the *offset* from the start of the file. When you omit the third element, the reply contains the entire file.

AP11 takes care of opening and closing the file; no separate command is required to do so. You receive the response when you next use the variable shared with AP11. The response is a `2`-element enclosed array. The first is a return code, and the second is the accompanying data.

The function `host` in workspace `1 tools` provides a tool that includes both *set* and *use* of the shared name and unpacks the result. With it, you can execute statements similar to the following:

```
x←host 'get'⊃'file'
```

and leave the sharing, setting, and using of the shared name to the utility.

`'put'` ⊃ *file* ⊃ *data* [⊃ *offset*]

The `put` command may be an enclosed arrary of either three or four elements.

- The first element contains the command `'put'` as a character vector.
- The second contains the name of a *file*, in the same form as for the command `get`. AP11 takes care of opening the file and closing it afterward. If the file does not exist, AP11 creates it with permission `600`, or as modified (for this entire session) by the value of the UNIX environment variable `umask`.
- The third is the *data* to be written, as a character vector. No other form of data is supported. The total size of the variable presented to the variable shared with AP11 must be less then `32768`. If the array you want appended to a UNIX file is longer than this limit, use the utility function `putu` in the public workspace `1 pc108`.

- The optional fourth element contains an integer scalar identifying the *offset* from the start of the file at which the data is to be written. Use this array only when you want to overwrite a portion of the file. Do not include the length, which the command *put* infers from the length of the list included in the third array. When you omit the fourth array, data is appended to the end of the file.

The result obtained from your next use of the variable shared with AP11 is a return code indicating the success or failure of your request to put data in a file.

The function *host* in workspace `1 tools` provides a tool that includes both *set* and *use* of the shared name and unpacks the result. With it, you can execute statements similar to the following:

```
      host 'put'⊃'file12'⊃'Testing the workings of AP11'
```

`'size' ⊃ file`

The *size* command is a 2-element enclosed array.

- The first element contains the characters `'size'` as a character vector.
- The second is the name of a *file*, in the same form as described for the command *get*.

The result obtained from your next use of the variable shared with AP11 is a 2-element enclosed array containing a return code and the size in bytes of the file about which you inquired; for example:

```
      host 'size'⊃'/usr/sax/rel/bin/apl'
490042
```

Commands on Open Files

This group of AP11 commands allows an APL program to control a UNIX file while placing all responsibility for the use of the file on the APL program. Do not attempt to use these facilities without a thorough knowledge of the UNIX system calls represented here.

You (or your program) can open a file that UNIX considers a *device* which is, in turn, attached to external hardware, such as a modem. In this fashion, an APL task may gain contact with another APL task active in a separate environment.

'open' ⊃ file ⊃ control [⊃ mode]

The AP11 command *open* requires the name of a UNIX file as a character list in the second cell of its argument, and in the third an integer scalar or 1-element array that controls which access levels and other status flags are applied when the file is opened. Sum the appropriate UNIX system values for all flags you need. Consult your UNIX user's manual, or use the values preset for you in the public workspace `1 hostap`. Their names are familiar to those who know the UNIX environment.

Table 3.3 provides a list of APL variables in workspace `1 hostap` that correspond to UNIX options available.

The optional fourth element contains a *mode* when the third element calls for creation of a file through the control flag *O_CREAT*.

The return for this command (or explicit result of the utility *open*) contains a UNIX file descriptor that is a positive integer intended for use with other commands (*write*, *read*, etc.).

Table 3.3. The open command control options.

Value	*Meaning*
O_RDONLY	Open for reading only.
O_WRONLY	Open for writing only.
O_RDWR	Open for reading and writing.
O_NDELAY	Open returns without waiting for a carrier.
O_APPEND	Sets pointer to end of file prior to each write.
O_TRUNC	File truncated to length 0.
O_CREAT	Create file.

'close' ⊃ descriptor

The command *close* requires the file descriptor of an open UNIX file. Its side effect is to close the file that the descriptor points to. The return from this command (or the explicit result of the utility *close*) contains an empty list.

'ioctl' ⊃ *descriptor* ⊃ *cmd* ⊃ *arg*

The AP11 command *ioctl* requires an open file descriptor in the second cell of its argument, an integer *ioctl* request in the third cell, and a character scalar or list, ***arg***, in the fourth cell. ***arg*** is taken as a string of bytes and coerced to the format required by the `ioctl` request ***cmd***. Consult your UNIX documentation for appropriate `ioctl` requests.

'read' ⊃ *descriptor* ⊃ *bytes*

The command *read* requires the file descriptor of an open UNIX file and the number of bytes to be read from the current pointer. The location of the file pointer may be set through the AP11 command *lseek*. Immediately after the use of the *open* command, the pointer is at the start of the file unless you use *O_APPEND*. After the *read* command is complete, the pointer is advanced by the number of characters read. An empty list is returned if the end of the file has been reached. The reply (or the explicit result of the utility *read*) contains a character list of bytes read from the file.

'write' ⊃ *descriptor* ⊃ *data*

The command *write* requires a file descriptor of an open UNIX file and a character list to be written to that file. The position of the file pointer (which indicates the offset in bytes from the start of the file where the current *write* operation will occur) can be set through the AP11 command *lseek*. The pointer is automatically advanced by the number of bytes written to the file. The return of the *write* command (or the explicit result of the utility *write*) contains an empty list.

'lseek' ⊃ *descriptor* ⊃ *offset* ⊃ *mode*

The command *lseek* sets a pointer in an open UNIX file. The placement of the pointer is controlled by the value of the flag ***mode***, which may have one of the three values described in Table 3.4. The values above have names in BSD UNIX all of which begin with *L_*.

Table 3.4. Mode values for the `lseek` command.

Value	*Meaning*
SEEK_SET	Pointer set to *offset*.
SEEK_CUR	Pointer set to current location plus *offset*.
SEEK_END	Pointer set to size of the file plus *offset*.

'getenv' ⊃ var

This command retrieves the value of the environmental variable *var*. If the variable does not exist, a null string is returned; hence, there is no way of distinguishing between an unset variable and an empty variable.

'putenv' ⊃ var ⊃ value

This command assigns an environmental variable *var* the value of *value*.

'fcntl' ⊃ descriptor ⊃ cmd ⊃ arg

This command performs a variety of functions on open file descriptors. Table 3.5 describes the possible actions.

Table 3.5. The `fcntl` command.

cmd	*Meaning*	*Result*
F_DUPFD	New file descriptor.	New file descriptor.
F_GETFD F_SETFD	Get close-on-exec flag. Set close-on-exec flag.	Value of flag (low-order bit). Value other than ¯1.
F_GETFL F_SETFL	Get file status flags. Set file status flags.	Value of file flags. Value other than ¯1.

'lock' ⊃ descriptor ⊃ type [⊃ length [, offset]]

This command will apply a file lock on the specified descriptor, determined by the *type* parameter. Should the lock be successfully applied, *lock* will return 0. A failure is indicated by ¯1. If offset is not specified, it defaults to 0, or to lock from the beginning of the file. *length* specifies the number of bytes to be locked at position *offset*. If *length*, *offset* are not specified, then the entire file is locked.

The values *type* can have are listed in Table 3.6. Note that should shared locks not be supported on a system, they are automatically upgraded to exclusive locks.

Table 3.6. Values for lock type parameter.

Value	*Meaning*
LOCK_UN	Remove the specified lock.
LOCK_RB	Blocking shared lock.
LOCK_RN	Non-blocking shared lock.
LOCK_WB	Blocking exclusive lock.
LOCK_WN	Non-blocking exclusive lock.

Other Commands

The commands in this group, *break*, *chdir* and *shell* give you access to useful system facilities.

'break'

This command sends an interrupt to AP11 and applies to a halted *fcntl*, *ioctl*, *open*, *read*, or *write* command. The result from the break command is unpredictable if work is in progress.

'chdir' ⊃ path

The *chdir* command is a 2-element enclosed array.

- The first element contains the command 'chdir' as a character vector.

- The second contains a UNIX path as a character vector.

The effect is to change the apparent *path* for references to file names. However, the new path affects references to files made by way of AP11 *only*, and then only during the time that the name you have shared with AP11 remains shared. It has no effect on the directories used by the APL system or those that appear in APL system commands or in the arguments to APL file primitives.

The result obtained from your next use of the variable you shared with AP11 is a 2-element enclosed array containing a return code and an empty character vector.

`'shell' ⊃ command ⊃ arg1 ⊃ arg2 ⊃ . . .`

The *shell* command is an enclosed arrary of at least two elements.

- The first element contains the command `'shell'` as a character vector.
- The second element contains the name of the UNIX command as a character vector.
- The third and following elements contain character vectors which embody the arguments to the UNIX command you wish to execute. There is no limit to the number of arguments as far as AP11 is concerned. You can split up the arguments in whatever way provides the clearest treatment of internal blanks, and of other characters special to the shell.

This command lets you invoke a UNIX shell to execute a command. For example, to invoke the UNIX editor `vi` on a set of files, use *host* this way:

```
host 'shell'⊃'vi'⊃'file1'⊃'file2'
```

You can use CTRL-L to refresh the screen if your environment does not provide a window manager.

Note: Remember to exit from UNIX back to SHARP APL when you have completed your use of the *shell* command.

4
AP124: Full-Screen AP

Introduction

In SHARP APL you can write a program to generate a display that occupies the entire screen and permits users to write instructions or responses anywhere on the screen. For example, you might want to maintain a journal of transactions using a program that has this type of interface:

- permanent headings to indicate the various categories of information
- blank areas to receive new entries (or areas that display existing information that can be edited)
- programmable function keys to trigger instant actions or menus that provide more options
- automatic cursor placement that singles out regions on the screen to indicate menu choices or erroneous entries.

You can install these features and many more using the *full-screen auxiliary processor,* AP124, provided with SHARP APL for UNIX.

This AP is functionally equivalent to AP124 as distributed with the OS/390 version of SHARP APL, although certain minor exceptions exist based on differences between the UNIX and the OS/390 operating systems.

AP124 under OS/390 is used to manage the screens of IBM 3270-style display stations. A program that runs properly in SHARP APL for OS/390 should run as well in the UNIX version, unless it exploits the light pen (available for 3270 terminals only). However, the converse is not true because AP124 under UNIX is more tolerant than the OS/390 version. For example, SHARP APL for UNIX permits one field to be right next to another, whereas the 3270 requires a horizontal separation of at least one position between fields.

The UNIX version supports both a `1`-variable interface to AP124 (with *ctl*) and a `2`-variable interface to AP124 (with *ctl* and *dat*).

Both SHARP APL versions of AP124 (UNIX and OS/390) are upwardly compatible with AP124 as provided by IBM. A program written for the IBM version will run under both SHARP APL versions of the AP. However, the converse is not necessarily true, since both these versions have enhancements not included in the IBM product.

AP124 and the Session Manager

To introduce the concepts that underlie AP124, look at how things work when you rely on the session manager to prepare an entry for the SHARP APL interpreter. You *release* an entry to the session manager by pressing RETURN. Each time you release an entry, the session manager first updates the session log by appending your entry and APL's response, and then moves the log up so that your entry and its response are the last items showing.

When you use AP124 to produce a display, however, it seizes control of the entire screen and replaces whatever the session manager had displayed with a display of its own. AP124 works with a succession of screens each of which occupies the entire area available. AP124 is able to display characters anywhere on the screen. It's also able to make parts of the screen brighter, blinking, underlined, reversed, or to appear in color on a terminal equipped for it.

Full-Screen Dialogue

Like the session manager, AP124 conducts a dialogue—alternating between displays it generates and responses made from the keyboard. However, AP124 provides much more.

The whole screen is your stage. Instead of restricting input to a single logical line, programs that use AP124 allow you write anywhere on the screen.

AP124 provides more than one way to release an entry for processing. You can release new and edited information to AP124 by pressing RETURN and by pressing any of the programmable function keys labelled F1 to F12. Each application assigns its own meanings to function keys: the function key may be designed to write information to the screen, to release information to AP124, or to do both automatically.

When you release a screen to AP124, it keeps a complete record of what you did. It can report which key you used to release the screen, where on the screen the cursor was located when you did so, and which parts of the screen you altered. You can obtain what was written in any field at the time you release the screen.

Once a program has caused AP124 to write to the screen, AP124's display remains visible until the session manager requests input or output. This request may not occur until the expression or program that produced the AP124 display reaches a normal end, encounters an untrapped error, or comes to a request for □- or □-input or output.

When any of those events occur, the normal session manager takes over and the session log reappears. However, AP124 retains the internal description of its displays, and can still respond to inquiries about them, or, when commanded to do so, can again seize control of the screen and regenerate its display.

Passing Instructions to AP124.

AP124 is an auxiliary processor. Your APL workspace communicates by sharing a variable with it. This section describes the individual steps that transpire as you communicate with AP124. Probably you'll incorporate these steps into user-defined functions, or take advantage of the utility functions supplied in the public workspace `1 ap124`. That will automate and conceal the mechanism, and spare you from having to write it all out in detail each time.

The following sequence gives you a general explanation of what's going on.

1. You (or your program on your behalf) offer to share a variable with AP124.
2. Your program sets the shared variable (assigns a value to it). AP124 then uses that value. (See the *SVP Manual, Chapter 2* for a discussion of communication by shared variables.) The values that you set tell AP124 how you want the screen to appear, what modifications from the keyboard to permit, or what information you want back.
3. AP124 sets the value of the shared variable to report that it has carried out your command, or, if it was unable to carry it out, what sort of trouble prevented it. It also sets the shared variable to report where on the screen you wrote, what you wrote there, and what key you pressed to release the screen.
4. Your program examines the reports it receives from AP124 and decides what to do next. It may decide to request a new display, to modify the one just displayed, or even to continue the display unchanged; some applications

display screen after screen, not returning to immediate execution (and hence not returning to the session manager) until much later in the work session. Or your program may decide to do something else entirely.

Throughout its use of AP124, your program alternates between setting the shared variable with commands, and then using the shared variable to find out what happened.

The Internal Record and the Screen You See

AP124 maintains tables in its own area of storage. These tables describe the appearance of the screen for each variable you share with AP124. Programmers refer to these as the screen buffers. They determine what will be displayed when a display is requested. Most of AP124's commands affect the stored tables rather than the display itself. A program that uses AP124 has to start by setting up those tables.

Certain commands cause AP124 to place the current contents of its tables on the screen. If you use the keyboard to alter the screen by inserting or deleting characters, AP124 immediately posts those changes into its internal tables. When you ask to read what is on the screen, you are really getting a report of what's now in the tables that AP124 maintains. Indeed, it is possible to set up the tables and then read what is in them without ever actually displaying their contents to the screen. This approach takes advantage of the formatting capabilities of AP124 without recourse to display; the contents of the internal screen buffers can be the basis for a report that is printed without ever appearing on a screen.

Some commands simply change the data in the buffer tables. Others cause AP124 to generate a new display. A display always copies to the screen the data from *all* the screen buffers. Some of that data may have been placed in the buffers much earlier, and remained there unchanged; some of it may have been placed there only a moment before, as part of the series of commands that elicits the display itself.

Even when AP124 relinquishes control of the screen, its buffers remain intact. You can still read what is in them, or use them to regenerate the screen. Only when the shared variable link is broken does AP124 clear its screen buffers.

Reading and Writing Fields

AP124 divides the screen into *fields.* A field is a rectangular area of the screen defined by the coordinates of its *top left corner* and its *shape* (number of rows and columns it occupies). To have your program display something on the screen, you must tell AP124 three things:

1. Where each field is located (and certain other attributes that control whether the field is keyable, its color, etc.). AP124 keeps its own record of the specifications you provide.
2. What data you want displayed in each field. AP124 also keeps its own record of the data for each field.
3. That you want the screen displayed. AP124 displays the entire screen, including the data for all defined fields. If some fields have not changed since your preceding display, you don't have to specify them again; at each display, all fields appear, showing the new data where you have supplied it, or repeating the old where you have not.

When you use the keyboard to alter or write on a screen display, AP124 permits you to do so only within an already specified field, and then only when the description of that field permits you to write there. Although you can move the cursor to a position that is not within any field, AP124 won't let you write there.

The Keyboard with AP124

While AP124 has control, most keys have the same meanings as under the session manager, but a few are adapted to the new environment. The keys whose meanings are unchanged are as follows:

Graphics	Letters, numbers, blank, etc.
BACKSPACE	The destructive backspace, including ALT-BACKSPACE.
CTRL-C	To signal attention.
CTRL-P	To form overstrikes.
CAPS LOCK	To shift meanings of keys.
INS	For toggling *insert* versus *replace* mode (choice between them is same as with the session manager).

Keys with Different Full-Screen Meanings

← → ↑ ↓	The cursor keys on the numeric key pad move the cursor on the screen and have no effect on the log. Other keys on the numeric key pad have no function in full-screen mode.
HOME	The HOME key moves the cursor to the beginning of the keyable field nearest the top left corner of the screen (the first field you encounter scanning along successive rows from left to right).
TAB	The TAB key moves forward along the row you are now on to the leftmost position in the next keyable field; if it doesn't find another keyable field on this row, it continues on down through successive rows until it finds one.
DEL	The DELETE or DEL key deletes the character at the current cursor position.
Function keys	While AP124 is in control, the first twelve programmable function keys do not retain the meanings you may have assigned to them under the session manager. Instead, they pass a *release* signal to AP124. Since AP124 can report which key was pressed, each program that uses AP124 can adopt its own scheme, interpreting each key as a distinct signal.

Correspondence to 3270 keys

IBM 3270 display stations come with a wide variety of keyboards. They may have keys labelled PA1, PA2, and PA3 (forms of program attention), a key labelled CLEAR, and some number of keys labelled PF1 through PF24. AP124 encodes the function keys in UNIX so that their identifying codes correspond to keys identified in the keyboard map shown in the SHARP APL for UNIX *Handbook, Chapter 4*.

The keys PA1, PA2, PA3, and CLEAR are intended as interrupt signals. They release the screen; but you should not depend on being able to read back what you may have written on the screen before you pressed one of them.

DUP and FIELD MARK produce displayable characters to which a program may attach special significance; keying one of them does not release the screen.

Status Line during Full-Screen Mode

AP124 has its own status line at the bottom of the screen. It identifies AP124 as the process in control of the display, and indicates two modes: the mode of execution, which may be either RUN or INPUT; and the mode of data entry, which is either replacement (no symbol) or insertion (a caret mark appears). As long as a variable is shared with AP124, the status line remains displayed. This feature is terminal dependent.

No Type-Ahead during Full-Screen Mode

It is not possible to type your next entry before the computer has finished executing the preceding one.

Communicating with AP124

Your link to AP124 is a variable shared between your active workspace and AP124. By sharing several variables with AP124, you can establish several different screens, and display one or another depending on which of the shared variables you use to control the display.

You provide AP124 two kinds of information: the *command* indicating the action you want it to take, and any *data* that command requires. You link a command and its data in a `2`-element enclosed array. Only this approach, which allows one shared variable per screen, is available. You can link multiple commands in a single array through which you can pass it a sequence of several different commands and any data each one may require.

The value that the SVP sets in the shared variable is also a boxed array containing a return code and any data produced by the commands you issued.

Establishing Contact

Let the shared variable, which you offer to share with processor `124`, be called *`ctlS`*:

```
      124 ⎕svo 'ctlS'
1
```

The result is 1, which indicates that you made an offer that has not yet been accepted. You then set access control on the shared variable:

```
      1 ⎕svc 'ctlS'
1 1 1 1
```

The result shows full interlock in effect. It is important that *ctlS* be interlocked before you refer to it. Although AP124 will set the interlock itself, remember that any AP is scheduled independently of your workspace, and it may not have had time to set access control before your program tries to set the shared variable. Therefore, your program should always set the interlock as soon as you have made your offer.

Using ⎕*svo*, you can find out whether AP124 has accepted the share:

```
      ⎕svo 'ctlS'
2
```

The result indicates that AP124 has accepted the offer to share *ctlS*. It's prudent to verify that sharing has been established before you proceed with sending commands to AP124. Within a function, you can verify that the variable is shared by using a test such as the following:

```
      →(2=⎕svo 'ctlS')/ok
```

Using the Shared Variable to Control the Screen

Once you have shared a variable with AP124, the usual sequence of steps is as follows:

1. You set *ctlS* to be an array each of whose elements is formed by boxing a command code and the data that the command needs. When you transmit only one command and it doesn't require data, *ctlS* may be either an item or a 1-element array containing the boxed command code.

   Suppose *c*1 contains your first command, and *d*1 is the data it needs; suppose *c*2 is your second command and doesn't require data; suppose *c*3 is your third command, and *d*3 is its data. You transmit all three commands to AP124 by an expression such as:

   ```
   ctlS←c1 ⊃ d1 ⊃ c2 ⊃ c3 ⊃ d3
   ```

2. AP124 opens the first item of *ctlS* and interprets the command it finds there. If that command calls for data, AP124 opens the next item of *ctlS* to get that data.

3. If your variable *ctlS* has more items, AP124 opens the next one to find the next command, and, if need be, the one following to find its accompanying data. AP124 continues in this fashion until it has evaluated all the commands in *ctlS*.

4. When AP124 has completed work on all the commands contained in *ctlS* or has given up because it was unable to execute a command, it sets *ctlS* to be a list of boxed items. The first item contains a 2-elemtent integer array. When all the commands were executed successfully, that first item consists of

   ```
   <0 ,n
   ```

   where *n* is the number of commands that AP124 executed.

5. When AP124 was unable to complete work on all the commands in *ctlS*, it sets the first item of *ctlS* to consist of:

   ```
   <tr ,n
   ```

   where *tr* is a code describing the trouble AP124 experienced, and *n* is the (0-origin) number of the command on which AP124 was working when it was unable to continue. AP124 sets the second item of *ctlS* to contain an error message that corresponds to the trouble code *tr*. This feature of the SHARP APL for UNIX version of AP124 removes the need for users to look up error text. For example:

   ```
         r←ctlS⊣ctlS←1 1⊃1 1 3 90
         r⊣⎕ps←2/¯1 3
   |____| |______________________________________|
   |32 0| |defined field extends beyond the screen|
   |____| |______________________________________|
   ```

6. Your program uses the value of *ctlS*. Start off by storing it under some other variable. This is important because the shared variable interlock on *ctlS* does not allow direct reference to *ctlS* again until you have set it (presumably to convey your next command). Since you will need to refer separately to the various items that AP124 returned in *ctlS*, you need to have *ctlS*'s value available to you without having to use the variable *ctlS* again. In the sample function *cmd* you'll see that the value of the shared variable is assigned to *r* so that subsequent references are to *r* rather than to the shared variable.

7. Your program checks the first item of the value returned in *ctlS* to find out whether your command sequence has been carried out.
8. The succeeding items in the value returned in *ctlS* contain the various result arrays, one for each command that returns a result. It's up to you to keep track of which commands called for results, and hence which component contains which result.

   Whenever you have set *ctlS*, you should use it in order to see how things turned out. In the definition of *cmd* that follows, to make sure that set and use don't get out of step, the two are combined in a single statement:

```
    ∇ r←cmd ⍵
[1] ⍝ Execute multiple AP124 commands
[2]   r←ctlS ⊣ ctlS←⊃⍵
[3]   →(×''⍴>'⍴r)↓End                ⍝ Unless error report
[4]   (errormsg>''⍴r) ⎕signal 124    ⍝ Interpret error code
[5]   End:  r←1↓r                    ⍝ What's left is data
```

AP124 Return Codes

As explained earlier, a non-zero return code means that an error has occurred during the operation. Table 4.1 provides a list of these return codes and their meanings.

Table 4.1. AP124 return codes.

Value	*Meaning*	*Value*	*Meaning*
0	Operation successful.	52	Device is not a 3270.
11	Control variable rank error.	53	Required shared storage unavailable.
12	Control variable length error.	54	Printer not available.
13	Control variable domain error.	89	Unknown shared variable return code.
14	Invalid command.	91	Physical field table overflow.
15	Position cursor in undefined field.	92	Physical field table error; interrupt.
21	Data variable rank error.	94	Device not available.
22	Data variable length error.	95	Unexpected IO error.
23	Data variable domain error.	96	Chained ccw string not complete.
24	Data variable not shared.	97	Bad 3270 orders in output data.
30	Invalid field number.	98	Full-screen support not available.
32	Defined field extends beyond the screen.	99	Unknown 3270 device error.
33	Reference outside field definition.	201	Invalid color.
35	Light pen field starts in column 1.	202	Invalid highlight.
36	Light pen field not contained in 1 line.	204	Invalid program symbol set.
37	Invalid field type.	205	Internal buffer overflow.
38	Invalid field intensity.	206	Invalid skip column in format table.
41	Data variable not given in correct seq.	207	Invalid nulls column in format table.
42	Data variable referenced by incorrect seq.	301	Cannot open font file.
43	Invalid translate table code.	302	Cannot create inverse font.

Refreshing the Screen

In certain situations, the smooth contact between the session manager and the session may be disrupted by the intricacies of screen control required to support AP124. The session manager lets you use CTRL-L to refresh your screen, which is most often needed when you return to immediate execution following a session controlled by AP124.

The Geography of Fields

Think of your screen as a table. The most common screen sizes are `24`-by-`80` and `32`-by-`80`. AP124 automatically selects the screen size appropriate to your terminal. Within the screen area, each field is defined as a rectangle described by four numbers: the coordinates of its *top left corner* (row and column, in `1`-origin), and its *shape* (rows and columns).

If you want a field to lie across the top of your screen, define it as a `1`-row field that occupies the entire width of the screen. It may be useful for the entry of commands or to display a caption. Its top corner is at row `1`, column `1`, and its shape is `1`-by-`80`, so its shape and position are described by the four numbers `1 1 1 80`.

The definition of the format of each field is a vector of up to `11` integers. Vectors that define several fields are organized into a format matrix, which has as many rows as fields, and as many as `11` columns. The first four elements in a field definition are always the position and extent of the field in `1`-origin. The *row number* (also in `1`-origin) of a field definition in a format table is regarded as the *number of the field* in all subsequent references to the display controlled by that format table.

Since the one row caption field happens to be specified by the first row of the format table that describes the screen, the program refers to it as “field `1`” (you can't tell that from the appearance or behavior of the display).

Adjacent Fields

Terminals of the IBM 3270 family require that some of the space otherwise devoted to storing characters for the screen be allocated instead to attributes for each row of each field. In practice, this means that one position on each row of each field is unavailable for entry or display: you can't place a field so that a row

is adjacent (horizontally) to a row of another field. Leave at least one unused position between them. The same constraint also applies to positions that straddle an edge of the screen; when one field uses the right-most position on a row, the left-most position of the next row can't be in a different field. Moreover, the screen positions wrap around so that the bottom right corner and the top left corner cannot belong to different fields.

The adjacency restriction derives from the IBM 3270 hardware and not from AP124 itself. There is no such restriction when you use AP124 in UNIX. However, if your program might also be used on a 3270 terminal, you must bear the restriction in mind and refrain from laying out fields so that they touch horizontally. The examples shown in this guide are consistent with this restriction.

Row-Structure of a Field

When a field has multiple rows, in some ways it is treated as a collection of separate rows, and in others as a single vector, reshaped to fit into its total rectangular space.

When you write data to a field, you must supply a *list of characters.* Even if you have a table of characters whose shape exactly matches the shape of the field in which you want to display it, you still send AP124 the *ravel* of that table. For example, if a field is defined to have `3` rows and `20` columns, to fill it you send AP124 a `60`-character vector. AP124 looks up its record of the field's shape, and fills the positions within the field from the data you have supplied. Suppose you call the shape of a field *`fs`*, and the data to be displayed there *`d`*. Then AP124 displays the data so that it appears as *`fs`*`⍴(×/`*`fs`*`)↑,`*`d`*.

When you supply more data than will fit in the field, AP124 discards the excess. When you supply fewer characters than are needed to fill the field, only the characters you supply are displayed. The remaining positions are filled with blanks. Similarly, when you read a field from the screen, you receive it as a single-row table representing the entire field. If you want to treat the data from a many-rowed field as a table having the same shape that you see on the screen, it's up to you to keep track of the field's shape, and reshape the data you receive accordingly.

Although a field is written and read as if it consisted of a single row, there are nevertheless ways in which a field acts as a collection of separate rows. The term *field row* is used to refer to the portion of a field lying on a single row of the screen.

Within a Field Row

For AP124, the keyboard is in replacement mode by default. When you move the cursor to a place where a character already appears and then key another character, the new character replaces the first. But when you put the keyboard into insert mode by pressing the INSERT key, newly keyed characters are inserted ahead of the characters already on the screen; already displayed characters are pushed to the right.

AP124 lets you do that only until you have filled the field row. Blanks can be pushed off the end of a field row (and are discarded), but other characters can be deleted only by the DELETE key. AP124 in the UNIX version of SHARP APL permits blanks to be pushed off the end of a field row, but 3270 display stations do not; see "Blanks and Nulls," below. Characters that are displaced cannot be pushed into the next field, and cannot be pushed into a lower row of the same field.

The DELETE key can be used to delete characters in a field row. Pressing DELETE deletes a single character. When you delete, the remaining characters further to the right in the same field row move leftward to fill the vacated positions. But characters are not pulled backward from beyond the boundaries of the individual field row. AP124 fills positions vacated at the end of a field row with blanks. This is not necessarily true for AP124 supported by the OS/390 version of SHARP APL. See "Blanks and Nulls," below.

Blanks and Nulls

When you delete characters in a field row, the characters further to the right move leftward to close up the gap. What takes their place at the end of the field row? AP124 always fills the space vacated at the end of a field row with the *space* character. This is not the situation with AP124 in SHARP APL for OS/390, which offers the choice of a space or a *full-screen null* character, which displays as a blank, but can be distinguished from the space you get by pressing the space bar. The UNIX version makes no such distinction.

In SHARP APL for OS/390, when AP124 transfers data from your workspace to its storage buffers for display, you may specify that it automatically replace the trailing blanks at the end of each field row with nulls.

Whether AP124 does so is controlled by the blank processing attribute which is part of each field's description. And when you read from the screen and return the field data to the workspace, if AP124 has been instructed to retain trailing nulls for this field, that portion of the result vector appears exactly as it did on the screen. You can also elect to have nulls on the screen converted to blanks.

Effect of Trailing Nulls on Insertion

Trailing nulls are a vital concern for programs written for IBM 3270 display stations. That's because 3270s do not permit trailing blanks at the end of a field row to be pushed out by insertions (whereas SHARP APL for UNIX does). If your program also runs under the OS/390 version, and you wish to permit insertions into a field row that has space at the right ends of its rows, it's essential to specify that the trailing blanks on each row be displayed as nulls. When you read back the contents of a field, AP124 gives you a similar option to have nulls at the end of each field row converted to blanks.

As displayed, a field may have some trailing nulls, either because AP124 wrote them there when it first transferred the data from the workspace, or because you deleted characters and didn't replace them. Those trailing nulls may or may not be important to your application. If you need to make an explicit check for net deletions, you should elect to have AP124 return the nulls unchanged. But if nulls are of no interest, you can direct AP124 to replace them by blanks in the data that it delivers in response to any of the commands that read the contents of fields back to the workspace.

Blank-Processing Attribute

One of the attributes of a field (controlled by the value in column `11` of the format table) instructs AP124 how you want trailing blanks handled. The four options, encoded by the numbers `0`, `1`, `2`, and `3`, are shown in Table 4.2.

The full-screen null character is not supported in AP124 (it ignores column `11` in the format table) or in $\square av$ (there is no character in the character set of SHARP APL for UNIX that corresponds to `254@`$\square av$ in the OS/390 version).

Cursor Skip at End of Field Row

One of a field's attributes established in the format table (column `10`) is control over where the cursor is allowed to go when it reaches the end of a field row. Setting and changing this attribute is described in the section on field formats.

When the skip attribute is *autoskip,* AP124 disregards positions that are in display-only fields, or in no field at all, and moves immediately to the next *keyable* field row.

When the cursor reaches the end of a field row, AP124 moves the cursor to the keyable position that is next on the screen. If there's an available position in another field and it is reachable by going rightward on the same row, that's where the cursor goes. It doesn't return to the field it was in until it comes down to the beginning of the next row and, moving from left to right in order, again encounters that field.

AP124 calculates which position comes next by considering them in row-major order. That's the order in which you read a book: left to right along each row, then to the beginning of the row below, and so on. If AP124 scans to the end of the screen without finding a keyable field row, it wraps around and continues its search from the top left corner.

When the format table specifies *no skip,* AP124 advances the cursor to the next position on the screen regardless of what field it's in or whether you are permitted to key there.

The Format Table

AP124 describes the format of its fields by an `11`-column table of nonnegative integers having one row for each field. When you set or reset the format table, you can supply `4`, `5`, `6`, or all `11` columns. If you specify a table with fewer than `11` columns, you get the default values for the columns you omit. A `1`-row table may be presented as a vector.

See Table 4.2 and the explanation that follows.

Table 4.2. AP124 format table.

1–2	*3–4*	*5*	*6*	*7*	*8*	*9*	*10*	*11*
Upper Left	*Shape*	*Type*	*Intensity*	*Color*	*PS*	*Highlight*	*Field End*	*Trailing Write/Read*
		0–I/O	0–Invisible	0-Green	0	0–None	0–Skip	0–B,B
		1–Num I/O	1–Normal	1-Blue		1–Blink	1–Don't	1–N,B
		2–Output	2–High	2-Red		2–Reverse		2–B,N
				3-Pink		4–Underline		3–N,N
				4-Green				
				5-Turquoise				B=Blank
				6-Yellow				N=Null
				7-White				

Columns 1 – 4. ***Position and shape.*** A field is defined by the 1-origin coordinates of its first (upper left) position, and its shape (number of rows and columns). You must explicitly specify these four columns for any field you define.

Column 5. ***Type.*** Controls whether a field is for display only, or will also permit entries from the keyboard.

There are three options:

0 – ***Keyable*** (unprotected).

1 – ***Keyable*** (numeric only). Aside from blanks, the only characters you can type in a numeric field are ¯−+.0123456789. Setting this attribute doesn't cause AP124 to return numeric data; it merely restricts input to numeric characters. Your program must determine what numbers those characters represent.

2 – ***Display only*** (protected). This is the default

Column 6. ***Intensity.*** Instead of the normal intensity, you may elect to have the characters in a field displayed with high intensity, or not displayed at all. An undisplayed field is useful for passwords or similar confidential information. An AP124 command will still report what you key there, even though nothing is visible on the screen.

There are four intensity display values:

`0` – ***Invisible***.

`1` – ***Normal*** (default).

`2` – ***High*** intensity.

Column 7. ***Color.*** This attribute governs all characters displayed in a field (except as overridden for individual characters by the *Write Character Attributes* command). For a terminal without color, the information in this column is ignored.

There are seven colors available:

`0` – ***Green*** (default).

`1` – ***Blue***.

`2` – ***Red***.

`3` – ***Pink*** (`1+2D`).

`4` – ***Green***.

`5` – ***Turquoise*** (`1+4D`).

`6` – ***Yellow*** (`2+4D`).

`7` – ***White*** (`1+2+4D`).

Column 8. ***Program Symbol Set.*** Some 3270 display stations permit a program to supply an alternate set of glyphs called a *program symbol set* in place of the regular character set. However, this feature is not supported for SHARP APL. For compatibility with IBM's AP124, this field exists but its value must always be `0`.

Column 9. ***Highlighting.*** This attribute governs all characters displayed in a field (except as overridden for individual characters by the *Reset Character Attributes* command).

There are four highlighting values:

`0` – ***None*** (default).

`1` – ***Blink***.

`2` – ***Reverse video.***

`4` – ***Underline***.

Column 10. ***Skip.*** This attribute determines where the cursor moves next after reaching the last character of a field row.

There are two options:

`0` – ***Skip to beginning of next keyable field row*** (default).

`1` – ***Advance to next physical position*** (regardless of field).

Column 11. ***Trailing Characters.*** This column controls how AP124 displays unused positions at the end of a field row, and how it reads back nulls. It is ignored in SHARP APL for UNIX; however, the available options are as follows:

`0` – ***Read blank, write blank.***

`1` – ***Read blank, write null.***

`2` – ***Read null, write blank.***

`3` – ***Read null, write null.***

AP124 Commands

The commands recognized by AP124 are described in this section. A summary of these commands is provided in Table 4.3. The variable of the corresponding utility function from the public workspace `1 ap124` is shown with each command.

Table 4.3. AP124 commands.

Command	*Data*	*Description*	*Utility Function*
`0`		No operation.	
`1`	Format table.	Format screen into fields.	*format*
`1,fld`	Format table.	Reformat selected fields.	*reformat*
`2,fld`	Data.	Display data in field immediately.	*immwr*
`3`		Display,wait for input.	*readscreen*
`4,fld`	Data.	Supply data for later display.	*write*
`5,fld`		Read selected fields.	*getfields*
`6,fld`	Type code.	Change field type.	*fieldtype*
`7,fld`	Intensity code.	Change field intensity.	*intensity*
`9`		Read format table.	*readformat*
`11`		Set alarm to sound at next display.	*alert*
`12`	Field,row,col.	Set cursor position	*setcursor*
`16`	Attribute.	Change field attributes.	*fieldattr*
`20`		Repaint screen.	*erasescreen*
`23`		Read attributes for selected fields.	*getattrs*
`24,fld`	Color code.	Change field color.	*color*
`26,fld`	Highlight code.	Change field highlighting.	*highlight*
`30`		Get data for modified fields.	*readdata*
`31`		Get boxed data for modified fields.	*readdatax*
`35,fld`	Attributes.	Change attributes at next display.	*writeattr*
`37,fld`	Attributes	Change attributes immediately.	*immwrattr*
`38`		Get attributes for modified fields.	*readattr*
`39`		Get boxed attributes for modified fields.	*readattrx*
`40`		Read entire screen.	*getscreen*
`41`		Read attributes for entire screen.	*getscreenx*

Set Field Formats

This command can be used either to set all field formats, or to reset the format of selected previously defined fields. To provide a complete new set of formats, the command is `1`. To provide new formats for selected fields, the command is a vector consisting of `1` followed by the numbers of the fields to be revised.

Set New Field Formats

```
ctlS←1 ⊃ fmt
r←ctlS          ⍝ boxed return code
```

The variable ***fmt*** is the format table, having `1` row for each field, and `4`, `5`, `6`, or `11` columns. If you supply a table with fewer than `11` columns, default values are assumed for the columns you omit. The utility function is *`format`*.

Reformat Selected Fields

```
ctlS←(1,n) ⊃ fmt
r←ctlS          ⍝ boxed return code
```

The variable ***n*** is a list of the numbers of the previously defined fields whose formats are to be changed. The variable ***fmt*** is a format table in the same form as for a new format table. The utility function is *`reformat`*.

Changing the format of a field has no effect on the data available for display in it. If you change a field's format by allowing it less space, AP124 displays as much data as will fit, but the remainder is not discarded from AP124's buffers and will again become visible if you subsequently change the format to provide sufficient space to contain it.

Read Current Format Table

```
ctlS←<9
r←ctlS          ⍝ return code linked to format table
```

Returns the current format table, with one row for each field that has been defined. The utility function for this command is *`readformat`*.

Reset Field Attributes

Each of the commands 6, 7, 16, 24, and 26 (defined in Table 4.4) resets the values of a particular attribute for selected previously defined fields.

```
ctlS←(c,n) ⊃ a
r←ctlS          ⍝ boxed return code
```

The variable *c* is a command number, the variable *n* is a list of field numbers, and the variable *a* is a list of the same length as n containing new values for the attribute.

Table 4.4. Reset field attributes.

Value	*Meaning*	*Utility Function*
6	Reset field type.	*fieldtype*
7	Reset field intensity.	*intensity*
24	Reset field color.	*color*
26	Reset field highlighting.	*highlight*
16	Reset field skip and trailing blank treatment.	*fieldattr*

The possible values for controlling skip ~2|*a* and trailing blank and null treatment ⌊*a*÷2 are combined into a single variable, as described in Table 4.5.

Table 4.5. Skip and trailing characters.

Value	*Skip*	*Write*	*Read*
0	No-auto	Blanks	Blanks
1	Auto	Blanks	Blanks
2	No-auto	Nulls	Blanks
3	Auto	Nulls	Blanks
4	No-auto	Blanks	Nulls
5	Auto	Blanks	Nulls
6	No-auto	Nulls	Nulls
7	Auto	Nulls	Nulls

Transmit Data for Display

Two commands permit you to specify the characters to be displayed in selected fields:

```
ctlS←(c,n) ⊃ dmat
r←ctlS          ⍝ boxed return code
```

The variable *c* is the command number (2 or 4); *n* is a vector of field numbers; and *dmat* is a table of character data to be displayed, having one row for each of the fields identified in *n*, and sufficient columns to contain the data of the largest field. Note that a row of *dmat* contains the *ravel* of the data to be displayed; the division of a field into rows and columns is not apparent in the data provided in *dmat*. A row containing data for a small field must nevertheless have the same length as all the other rows transmitted together in *dmat*. However, AP124 ignores the positions unneeded for the particular field, so it doesn't matter how the row is padded. The data display commands are defined in Table 4.6.

Table 4.6. Writing to the screen.

Value	*Meaning*	*Utility Function*
4	Respecify characters for fields indicated; do not yet generate a new screen display.	*write*
2	Respecify characters for fields indicated generate a new display of all fields.	*immwr*

Set Cursor Position

```
ctlS←12 ⊃ fld,row,col
r←ctlS          ⍝ boxed return code
```

Positions the cursor at the indicated position at the next *Write-Wait-Read* command. A cursor position is described by three numbers: field number, and row and column within that field (all in 1-origin). You can move the cursor to a position that is not within any field by making the first of those numbers 0, in which case the row and column coordinates are the 1-origin coordinates of the location on the screen as a whole. The utility function is *setcursor*.

Sound Alarm

```
ctlS←<11
r←ctlS          ⍝ boxed return code
```

At the next *Write-Wait-Read* command, sound the beeper once after displaying the screen but before accepting input from the keyboard. The utility function is *alarm*.

Repaint the Screen

```
ctlS←<20
r←ctlS          ⍝ boxed return code
```

Rewrites the screen from the current contents of AP124 internal buffers, ignoring anything that may have been typed on the screen. Leaves the cursor where it was. The utility function that embodies this command is *erasescreen*.

Set Character Attributes

Two commands permit you to specify how the *color* and *highlighting* of individual character positions depart from the color and highlighting specified for the field as a whole:

```
ctlS←(c,n) ⊃ table
r←ctlS          ⍝ boxed return code
```

The variable ***n*** is a vector containing the field numbers of the fields whose attributes are to be reset character by character, and the variable ***table*** is a table in the same form as for command 4 (write): one row for each field in ***n***, and sufficient columns to provide data for the largest field mentioned in ***n***. The data consists of characters, each of which encodes the information regarding color and highlighting. The color and highlighting of a character are encoded by:

```
(8 4⊥color, highlighting)@ ⎕av
```

The possible values for color are 0 through 7, and the possible values of highlight are 0 through 3. Color values have the same significance as described in the section on the format table and field attributes. Highlight values have the same significance as described in the section on the format table and field attributes; however, 3 *is the value for underline* (*not* 4). The commands for writing character attributes are defined in Table 4.7.

Table 4.7. Writing character attributes.

Value	*Meaning*	*Utility Function*
35	Reset attributes individually for all positions in the indicated fields but do not generate a new display.	*writeattr*
37	Reset attributes individually for all positions in the indicated fields and generate a new display for all fields.	*immwrattr*

Write-Wait-Read

```
ctlS←<3
r←ctlS          ⍝ return code linked to result vector
```

This command does three things:

1. It generates a new display for all fields (based on the data already in AP124's buffers).
2. It waits for input from the keyboard, permitting editing of any fields except those restricted to display only. Input continues until you release the screen by pressing RETURN or any of the function keys.
3. It returns a result indicating how you released the screen, where on the screen the cursor was when you did so, and a list of the fields in which you made any entry or deletion.

The result of write-wait-read command is an integer vector of more than five items, as defined in Table 4.8.

Table 4.8. Releasing the screen.

1	*2*	*3*	*4*	*5*	*[. . .]*
		Position of cursor			
Release Code	Modifier	Field	Row	Column	Modified Fields
0–ENTER	0	**n**	**n**	**n**	**n . . .**
1–PF key	1 to 24	**n**	**n**	**n**	**n . . .**
4–PA key	1 to 3				
5–CLEAR					

The terms PA, PF and CLEAR are borrowed from 3270 display stations; in SHARP APL for UNIX, they refer to keys that change from terminal to terminal.

This command does not report directly on the data on the screen, but indicates that the screen has been released, and which fields you changed before you released it. Use the commands described in the next section to obtain the data from the fields displayed on the screen.

Receive Data from Display

Four commands, defined in Table 4.9, provide alternative ways of reading the contents of AP124's buffers (hence, of reading what's on the screen).

Table 4.9. Reading from the screen.

Value	*Meaning*	*Utility Function*
5	Read specified fields as rows of a table.	*getfields*
23	Read attributes for selected fields.	*getattrs*
30	Read fields modified during last *Write-Wait-Read* command as rows of a table.	*readdata*
31	Read fields modified during last *Write-Wait-Read* command as enclosed arrays.	*readdatax*
38	Get attributes for modified fields.	*readattr*
39	Get boxed attributes for modified fields.	*readattrx*
40	Read entire screen as single table.	*getscreen*
41	Read attributes for entire screen as single table.	*getscreenx*

Read Specific Fields

```
ctlS←<5,n
r←ctlS          ⍝ return code linked to result table
```

n is a vector of field numbers. The result is a table with one row for each element of *n*, and as many columns as necessary to contain the largest field. Rows representing shorter fields are padded with blanks.

Read Attributes for Specific Fields

```
ctlS←<23,n
r←ctlS          ⍝ return code linked to result table
```

n is a vector of field numbers. The result is a table with one row of the attribute characters for each element of *n*, and as many columns as necessary to contain the largest field. Rows representing shorter fields are padded with zeros.

Read Modified Fields

```
ctlS←<30   or   ctlS←<31
r←ctlS          ⍝ return code linked to result table
```

For either of these commands, the result contains data from those fields in which you did any keying, insertion, or deletion before you released the last screen, and for which you have not yet read the data by an earlier use of these commands. The result does not indicate from which fields the data came; you have to keep track yourself of which fields were modified from the result of the *Write-Wait-Read* command.

Command `30` returns a character table result, with one row for each modified field, and as many columns as needed to contain data from the largest modified field. Rows representing smaller fields are padded with either blanks or nulls, as indicated by the value of column `11` of the format table for those fields. The data contains no indication of the shape of the field from which it comes; you have to reshape either by noting the shape when you established the field, or by reading the current shape from the current format table. (See "Read Current Format Table" earlier in this section.)

Command `31` returns a vector of enclosed arrays, each of which contains the ravelled data from a field that was modified.

Read Attributes for Modified Fields

```
ctlS←<38    or    ctlS←<39
r←ctlS          ⍝ return code linked to result table
```

Command `38` returns an attribute table result, with one row for each modified field, and as many columns as needed to contain attribute data from the largest modified field. Rows representing smaller fields are padded with nulls. The data

contains no indication of the shape of the field from which it comes; you have to reshape either by noting the shape when you established the field, or by reading the current shape from the current format table.

Command `39` returns an array of enclosed attribute data, each element of which contains the ravelled attribute data from a field that was modified.

Read Entire Screen

```
ctlS←<40
r←ctlS          ⍝ return code linked to result table
```

The result is a table having the same shape as the entire screen, containing the characters at each position on the screen. The utility function is *getscreen*.

Read Attributes for Entire Screen

```
ctlS←<41
r←ctlS          ⍝ return code linked to result table
```

The result is a table having the same shape as the entire screen, containing the attribute characters for each position on the screen. To decode the characters, see the section "Set Character Attributes" earlier in this section. The utility function is *getscreenx*.

AP124 Differences Under UNIX and OS/390

These differences arise principally from differences in the two operating systems, and in the types of terminals available in each environment.

- OS/390 does not permit different fields to occupy horizontally consecutive positions. UNIX has no such restriction.
- UNIX ignores both blanks and nulls when characters are to be pushed out of a field row by insertions further to the left. OS/390 permits only nulls to be pushed out in that fashion.

Public Workspace `1 ap124`

Workspace `1 ap124` contains utilities to facilitate work with AP124. Those that return an explicit result are shown with `r←`. They all use a single shared variable `ctlS` and a buffer `bufS` to store commands until one occurs that requires display, at which point the entire buffer is assigned to the shared variable. If an error occurs, the utility signals an error to the user-defined function that called it using event number `124`.

`format` ***fmt***	Command `1`. The variable ***fmt*** is the format table containing `4`, `5`, `6`, or `11` columns.
fld `reformat` ***fmt***	Command `1`. The format table ***fmt*** is applied to field numbers ***fld***.
`r←readformat`	Command `9`. The result is the current format table.
fld `fieldtype` ***typ***	Command `6`. The *type* values for field numbers in ***fld*** are reset to the values in ***typ***.
fld `intensity` ***int***	Command `7`. The *intensity* values for field numbers in ***fld*** are reset to the values in ***int***.
fld `color` ***clr***	Command `24`. The *color* values for field numbers in ***fld*** are reset to the values in ***clr***.
fld `highlight` ***hilt***	Command `26`. The *highlight* values for fields in ***fld*** are reset to the values in ***hilt***.
fld `fieldattr` ***att***	Command `16`. The *attribute* values in ***att*** are applied to field numbers in ***fld***.
fld `write` ***data***	Command `4`. Characters in a row of the table ***data*** are written to the buffer of the field whose number is in the corresponding position of ***fld***.
fld `writeattr` ***att***	Command `35`. The encoded character attributes in a row of ***att*** are written to the buffer of the field whose number is in the corresponding position in ***fld***.
fld `immwr` ***data***	Command `2`. Write ***data*** to the buffers of the fields indicated in ***fld*** in the same manner as ***write***, and generate a new display of all fields.

fld `immwrattr` *att*	Command `37`. Write the attributes in *att* to the buffers for the fields indicated *fld* in the same way as for `writeattr` and generate a new display of all fields.
`r←readscreen`	Command `3`. Generate a new display for all fields, wait for keyboard input, and report in the result how the screen was released, the position of the cursor at release, and a vector of fields that were modified.
`r←getfields` *fld*	Command `5`. Return data as a character table for field numbers in *fld*.
`r←getattrs` *fld*	Command `23`. Return attribute data as a character table for field numbers in *fld*.
`r←readdata`	Command `30`. Return data as character table for fields which were modified during the last `readscreen` operation and have not yet been read by `readdata` or `readdatax`.
`r←readdatax`	Command `31`. Same as `readdata` except that instead of data for fields being returned as rows of a table, they are returned as a vector of enclosed arrays.
`r←readattr`	Command `38`. Return attribute data as character table for fields which were modified during the last `readscreen` operation and have not yet been read by `readattr` or `readattrx`.
`r←readattrx`	Command `39`. Same as `readattr` except that instead of attribute data for fields being returned as rows of a table, they are returned as a list of boxed lists.
`r←getscreen`	Command `40`. The result is a character table representing the current characters on the entire screen.
`r←getscreenx`	Command `41`. The result is a character table representing the attributes of current characters on the entire screen. See the utility `writeattr`.
fld `setcursor` *row*`,`*col*	Command `12`. Set the position (*row*`,`*col*) at which the cursor will appear in field *fld* at the next use of `readscreen`.

alert	Command `11`. Set a flag so that at the next use of *readscreen* the alarm will sound.
erasescreen	Command `20`. Repaint the screen with the current contents of internal buffers, ignoring any modifications that were made since the last write operation. Leave the cursor where it is.

SHARP APL *for UNIX*

File System Manual

JUMP TO ...

CHAPTER 1. OVERVIEW
CONTENTS
PREFACE
MASTER INDEX
USING THIS DOCUMENTATION

HANDBOOK
AUXILIARY PROCESSORS MANUAL
FILE SYSTEM MANUAL
INTRINSIC FUNCTIONS MANUAL
LANGUAGE GUIDE
SVP MANUAL
SYSTEM GUIDE

© 2000 by Soliton Associates Limited
44 Victoria Street, Suite 2100, Toronto, Ontario, Canada M5C 1Y2

Pubcode: UW-037-0501 (0007)

SHARP APL *for UNIX*

DEVELOPER LIBRARY

File System Manual

SOLITON ASSOCIATES

© 2000 by Soliton Associates Limited–Printed in Canada
44 Victoria Street, Suite 2100, Toronto, Ontario, Canada M5C 1Y2

All rights reserved. No part of this work may be reproduced in any form or by any means without prior written permission of the copyright owner.

Soliton Associates Limited shall not be liable for errors contained herein or for incidental consequential damages in connection with the use of this material. The information contained within this document is subject to change without notice.

Acknowledgments

SHARP APL is a registered trademark of Soliton Associates Limited.

SVP and NSVP are trademarks of Soliton Associates Limited.

IBM and AIX are registered trademarks of International Business Machines Corporation.

RISC System/6000 is a trademark of International Business Machines Corporation.

UNIX is a registered trademark of UNIX System Laboratories, Inc.

Sun Workstation is a registered trademark of Sun Microsystems, Inc.

Sun, Solaris, SunOS, Open Windows, and SunView are trademarks of Sun Microsystems, Inc.

All other products referred to in this document are trademarks or registered trademarks of their respective trademark holders.

Publication code: UW-037-0501 (0009)

Contents

Tables and Figures v

Preface

Introduction vii
Chapter Outlines vii
Conventions viii
Documentation Summary viii
Contacting Soliton Associates ix

1. Overview

Introduction 1-1
The File System 1-1
Workspaces 1-2
Files 1-2
Non-APL Files 1-3
File Names Under UNIX 1-3
Concealed Suffix 1-4
File Utilities 1-4
The File System Server 1-5

2. APL Component Files

APL Files 2-1
APL Component File Concepts 2-2
File Names 2-3
Tieing a File 2-4
Control of Access to a File 2-4

File System Functions 2-6
File System Functions for Inquiry 2-8
File System Functions for Existence 2-10
File System Functions for Linkage 2-13
File System Functions for Reading and Writing 2-15
File System Functions for Description and Limits 2-18
File System Functions for Access Control 2-23
File System Functions for Concurrency 2-24
File Space Management 2-28

3. File Utilities
Introduction 3-1
The `tiers` Utility 3-2
The `holders` Utility 3-5
The `flockers` Utility 3-7
The `sift` Utility 3-9
When to use `sift` 3-10
The `sortio` Utility 3-10
The `cfcheck` Utility 3-13
`cfcheck` Messages 3-15

4. The File System Server
Overview 4-1
Fileserver's Features 4-2
Fileserver Operations 4-3
The `fsutil` Command 4-3
Connecting the Interpreter 4-7
Fileserver Restart 4-7
Installation and Configuration 4-8
Parameter List 4-10
Technical Notes 4-13

Index IX-1

Tables and Figures

Table 2.1. File system functions with permissions.. 2-7
Figure 4.1. SHARP APL File Management via Fileserver. 4-1
Table 4.2. Fileserver Restart Actions. 4-8

Preface

Introduction

This manual discusses the SHARP APL file system, APL component files, related component file utilities, and the SHARP APL for UNIX Fileserver, a high-performance alternative to the conventional file system that runs independent of the interpreter.

Rather than reproduce existing material, references to other SHARP APL for UNIX publications are supplied where applicable.

Chapter Outlines

This document is organized into the chapters described below.

Chapter 1, "Overview," explains, in general terms, the SHARP APL file system, APL component files, related component file utilities, and the SHARP APL for UNIX Fileserver.

Chapter 2, "APL Component Files," discusses component file concepts and reviews the system functions used to access APL component files.

Chapter 3, "Component File Utilities," documents utilities used to report the state of, and optimize the storage/retrieval of, APL component files.

Chapter 4, "The File System Server," discusses the Fileserver, a file management system that brokers transactions between the UNIX file system and the conventional SHARP APL file system.

Conventions

The following conventions are used throughout this documentation:

`⎕io←0`	Although the default value for `⎕io` in a clear workspace is *one*, all examples in this manual assume that the index origin is *zero*.
α and ω	Two APL symbols are used in syntax descriptions of monadic and dyadic functions; *alpha* (α) identifies the left argument and *omega* (ω) identifies the right argument.
`constant width`	Examples of non-APL input and output are represented in a `constant width` typeface.
default (%) prompt	Examples in this book assume you are using the C shell. The default prompt for the C shell is the *percent sign* (%).
passno	Argument variables appear in *sans-serif italic* type.
[ ]	Optional arguments are shown between square brackets. For example, in the following APL statement, the argument *passno* is optional:

'filename' ⎕stie tieno [,passno]

These default environment variables represent frequently used pathnames in SHARP APL for UNIX documentation and scripts:

`$SAXLIB`	→	`/usr/sax`
`$SAXDIR`	→	`/usr/sax/rel`
`$SAXCNF`	→	`/usr/sax/local`
`$HOME`	→	home directory of the current user.

Documentation Summary

The SHARP APL for UNIX reference documentation is published in several volumes. One or more of the following may be referenced in this manual:

SHARP APL for UNIX,

- *Handbook,* publication code UW-000-0401

- *Language Guide,* publication code UW-000-0802
- *System Guide,* publication code UW-000-0902
- *SVP Manual,* publication code UW-001-0501
- *Auxiliary Processors Manual,* publication code UW-033-0501
- *Intrinsic Functions Manual,* publication code UW-007-0501

For a complete list of SHARP APL publications, please refer to the ***Support*** link on the Soliton Associates website: `http://www.soliton.com.`

Contacting Soliton Associates

Problems or questions regarding your SHARP APL system or this guide should be directed to the Soliton Associates Technical Support group:

`support@soliton.com`

Requests for additional technical documentation should be directed to the Soliton Associates Sales group:

`sales@soliton.com`

1 Overview

Introduction

The SHARP APL file system provides a means to store data in a form directly usable by APL, but outside the workspace. It supplies a set of system functions that allow you to create and manage auxiliary storage as well as transfer data to and from the active workspace. SHARP APL also permits you to access data that has been stored by other users, as well as share a common collection of data with other users.

The File System

In SHARP APL, various elements are required for the storage and retrieval of data—these elements comprise the ***file system***:

Account number	UNIX user identification number that also serves as your APL library number.
Active workspace	Environment where variables, functions, and packages are understood, and the working storage area where calculations take place.
Component file	Collection of data stored independently but accessed by an active workspace.
Fileserver	A file management system that runs independent of the APL interpreter.
Saved workspace	Copy of the active workspace saved in storage.

Library Group of saved workspaces or component files belonging to one account number.

The SHARP APL file system writes to or reads from UNIX files saved with an APL-specific internal format.

Workspaces

APL systems have naming conventions that are much different from those of the underlying operating system's: the primary unit of storage is called the ***saved workspace,*** identified solely by the owner's account number and a single-level name. It contains all the variables and functions present in the ***active workspace*** at the time it was saved. Sets of workspaces having common account numbers are called workspace libraries.

For more information on the basic organization of the SHARP APL system see the *System Guide, Chapter 2*.

Files

In the APL environment, a file is defined as a collection of data (variables) stored external to the APL workspace but directly accessible to expressions executed from within the workspace. Each file is organized as a list of *components.* The system functions that deal with component files transfer data between the file and the active workspace one component at a time, identifying it by its position in the list. These functions are described in "Chapter 2. APL Component Files".

The same naming convention for APL workspaces is extended to APL component files: each has a two-part name consisting of a numeric account number and a single-level name. A set of files having a common account number is called a file library.

The APL file system establishes a common administration for all use of APL component files by any APL task. It not only manages a task's separate use of private files, but permits files to be shared and manages concurrent read/write sharing of the same file by several concurrent tasks.

Non-APL Files

SHARP APL can also access non-APL files, such as standard UNIX files. A non-APL file appears to APL as a continuous character vector. Users access these files by sharing a variable with the Host AP, AP11.

You may read or write the entire file as a single object. Alternatively, you may read or write a block of consecutive bytes. A block is described by its length and its offset from the beginning of the file. Characters are passed back and forth between the UNIX file and the active workspace without translation or modification.

More information on this subject is available in the *Auxiliary Processors Manual, Chapter 3*.

File Names Under UNIX

SHARP APL for UNIX offers two alternative styles of naming files and workspaces. One is consistent with the naming conventions of earlier APL systems, and the other is consistent with the current practice of the underlying operating system, UNIX:

APL library style	The name is written in the form `123 data`
UNIX path style	The name is written in the form `/usr/sue/data`

These are different ways of referring to the same file or workspace. The APL system maps its numeric account and library information to the UNIX system's directory of names.

Examples:

To load the workspace called *utils* from `fred`'s home directory, you write a command of the form

```
)load /usr/fred/utils
```

If `fred` is user `101`, then the expression

```
)load 101 utils
```

loads the same workspace.

If the directory `/usr/fred` is the current directory when you start your session, then you can type

```
)load ./utils
```

from SHARP APL, which has the same effect as the above two commands.

Concealed Suffix

The SHARP APL file system identifies the UNIX files that contain APL data by their suffixes: the suffix `.sf` is attached to the name of a UNIX file which is structured as an APL component file; and, the suffix `.sw` is attached to the name of a UNIX file that contains a saved APL workspace. These qualifying suffixes are visible from the UNIX shell, or from commands executed for you through AP11. However, when you refer to an APL file with the APL file functions or to a workspace using APL system functions or system commands, you neither see the suffix nor write it.

File Utilities

There are 5 utilities available for monitoring and optimising SHARP APL files:

- `tiers` reports all ties of a list of APL files.
- `holders` reports all holds (via □*hold* or □*fhold*) on a list of APL files.
- `flockers` reports the owners of APL's low-level UNIX locks (as obtained by `fcntl()`) on a list of APL files, and is mainly of interest to Soliton support personnel.
- `sift` compacts APL files to remove unused space.
- `sortio` reduces processing time by sorting all APL files of a specialized format.
- `cfcheck` reports on the structural integrity of an APL component file.

These utilities are fully documented in "Chapter 3. File Utilities".

The File System Server

The File System Server, Fileserver, which can be installed and run independent of the interpreter, interposes itself between UNIX and the conventional SHARP APL file system, which is "hardwired" into the interpreter. It operates more like a database management system where there are client processes (SHARP APL interpreter) and server processes (`fsios`) communicating via a shared memory segment.

Although it is delivered as an optional installation, and its operation is transparent to the SHARP APL interpreter, the Fileserver offers certain clear advantages over the unassisted SHARP APL file system:

Efficiency Additional I/O buffering in shared memory is available. Tie and hold tables are not stored with the file - they are maintained in shared memory. Reads and writes can operate at the same time - □*fhold* operations no longer interfere with them. Expedient disk space management is maintained.

Convenience Fileserver can be customized via configuration file parameters and then restarted without affecting the interpreter processes.

Security Files can be protected from non-SHARP APL access by augmented UNIX storage permissions of `0400` or even `0000` in directories with permissions of `0111`.

Reliability Fileserver's procedure for writing to the disk maintains data consistency and reduces the possibility of a corrupted file due to system failure.

The installation, configuration, and operation of the Fileserver is covered in "Chapter 4. The File System Server".

2
APL Component Files

In APL terminology, a file is a collection of data outside the workspace but directly accessible to statements executed from within the workspace. Access to files is provided in two ways:

- You can access standard UNIX files by way of the Host auxiliary processor (AP11), as described in the *Auxiliary Processors Manual, Chapter 3*.
- You can access APL component files by way of APL system functions provided for that purpose.

This chapter explains component file concepts and includes descriptions of all system functions used to access APL component files.

APL Files

An APL file contains within it a directory and descriptive information so that, to the APL user, it appears to be organized into *components.* The system functions that deal with APL files transfer data between the file and the active workspace one component at a time. Each component contains a value that can be assigned to any variable: an array of any type, rank, or shape; or a package. A component is identified by its position within the file (for example, component `1`, component `200`, and so on).

The APL file system provides

- indexed access to components
- an access matrix that governs access for other APL users
- for the entire file, a record of the user who created, last set access, or last modified its contents, and when this occurred

- for each component in the file, a record of its size in bytes, who last wrote it, and when it was last written.

APL Component File Concepts

An APL component file differs both from standard UNIX files and from an APL workspace in a number of ways.

Components as variables. Each component of a file is a variable. Like a variable in a workspace, a component may be an array of any type, rank, or shape; or a package. A component may contain the definition of a function either in its character representation (see `⎕fd` and `⎕cr` in the *System Guide, Chapter 5*) or as the referent of a name in a package.

Indexed access. A component in a file is referred to by its *component number.* Component numbers are consecutive positive integers ranging from the number of the first component up to the number of the last component; components are numbered in `1`-origin so that the lowest possible component number is component `1`. A component does not have a name, but the value it contains may be assigned to a name when transferred to an active workspace.

Open-ended size. A file may be of any size. Data is transferred between file and workspace one component at a time (so a component may not be larger than a variable in the workspace that wrote that component). There is no limit to the total size of a file other than that imposed by the physical limits of the UNIX file system that contains it. A common use of files is to partition data into components which may be processed serially to reduce the amount of data stored in the workspace at a given moment. In similar fashion, a large application may partition the set of function definitions it uses into those required at different phases of the work. A master program reads in the definitions needed for a particular phase and then overlays them with those needed for the next phase.

File functions. An APL component file is manipulated by system functions that you can invoke from the keyboard or execute within APL programs. In contrast, certain workspace actions require use of *system commands,* which are not permitted within a user-defined function, and so may be entered only from the keyboard or from the controlling shared name of an S-task.

Concurrent use. A task may use many files at the same time, whereas it has only one active workspace (apart from any S-tasks it controls).

Shared use. Several tasks (including tasks running for different users and on different hosts) may have access to the same file at the same time. The file system includes provision for interlocks to permit serialization.

Access control by user and function. The owner of a file may specify access controls for it. Every file has an access matrix that identifies which user accounts may tie it and which file functions each user account may employ with respect to it. At creation, a file has an empty access matrix that permits its owner to do anything to the file and denies any form of access to any other user. The access matrix may require the task that ties it to supply a *passnumber* (a numeric password) in order to tie the file.

Reservation. You may assign a file a space reservation. The default file size is 100,000 bytes. The reservation is a limit beyond which the file system will accept no further actions that would increase the file's size. The owner (or others with the necessary permission) may change the reservation using `⎕resize`(described on page 2-21) so the limit functions primarily as a safeguard against unintended growth. Regardless of the limit set, no space is allocated until actually used; the existence of a high limit is no assurance that space is actually available.

Component information. The file system automatically records for every component of every file the size of the component (in bytes), the account number of the task that wrote it, and the time when that write occurred. See the description of `⎕rdci` on page 2-23.

File information. The file system automatically records the creator, the last modifier of the access matrix, and the last modifier of data in the file, and when these events last occurred. See the description of `⎕rdfi` on page 2-20.

File compaction. Any user can maintain APL files from UNIX through the use of the command `sift`, which compresses unused space from a component file. See "Chapter 3. File Utilities".

File Names

The name of a component file is formed from the same characters as those permitted for names of APL objects: any of the `52` characters `a` to `z` and `A` to `Z`, the characters `Δ`, `∆`, or `_`, or the numerals `0` to `9`. The name may not, however, start with an underbar or a numeral; the length of the name may not exceed `11` characters; and, the name must not currently be in use by another file in the directory specified.

Tieing a File

All use of a component file is by way of a *tie number.* To make use of a file, an APL task must first tie it. The act of tieing both opens the file and assigns it a number. The tie functions ⎕*tie* and ⎕*stie* associate the file's name with an arbitrary positive integer. All other functions refer to the file by its tie number rather than by its name (except for ⎕*erase*, which requires both). The use of tie numbers serves primarily to help the system manage resources efficiently.

Duration of a File Tie

Once you have tied it, a file remains tied until untied with the system function ⎕*untie* or until the end of the task, whichever occurs sooner. A file tie is not affected by changes to the active workspace; in particular, it is unaffected by such actions as clearing the active workspace, loading another saved workspace into the active area, or using ⎕*run* to start another task. In that regard, a file tie is like a session variable and unlike a shared variable.

A file tie does not continue into another session. When an APL task ends, SHARP APL automatically unties any files that were tied, and does not keep any record of what files were tied or what tie numbers were used to tie them.

Control of Access to a File

Each file has an *access matrix*, which has three columns. Each `3`-column row of the access matrix reports information concerning the use of a file:

Account number The first column of the access matrix holds the account number (the UNIX user ID) of the user that seeks access to the file. The number `0` is used to mean any account.

Permission code The second column of the access matrix holds a number indicating which file functions may be used by the account identified in the first column. Each individual permission code is a power of `2`. Any combination of permissions may be expressed by the sum of the individual codes. The binary representation of the permission code shows a `1` for each permitted function and a `0` for each that is not permitted. The number `¯1` (binary representation all `1`s) is used to mean any file function.

Passnumber The third column of the access matrix holds an arbitrary integer that a file user must supply when tieing the file and as part of the argument of each function used with the file (with a few exceptions). Passnumber `0` is equivalent to no passnumber.

If you are not the owner of the file, you may tie it or use file functions with it only as provided in the access matrix.

A single account number may be mentioned several times in the access matrix, perhaps with different passnumbers. The permission in effect for you is the first occurrence (scanning from the top) of your account number paired with the passnumber you used to tie the file. If your account number does not appear, paired with the appropriate passnumber, then you get the permission associated with user number 0 in the access matrix.

If you are the owner of the file, as long as your account number is not mentioned in the first column, you may do anything with the file. It is as though every access matrix has a hidden last row containing your account number followed by `¯1 0` (meaning all functions, no passnumber).

If your account number does appear in the access matrix, even though you are the owner, the same access rules apply to you as apply to any user. You are restricted to the actions described by the permission code on the row that contains your account number and the passnumber with which you tied the file.

Warning: In order to permit access to APL files from within APL, it is necessary to grant read/write access to the corresponding UNIX files. Consequently, when `⎕stac` is used to set a non-empty access matrix, and if `-Yfcreatesw=0` is specified as an APL startup option, the SHARP APL file system grants to all users read/write access to the corresponding UNIX file.

Tieing a File with a Passnumber

You are always free to pick the tie number yourself. But your tie number must be accompanied by a passnumber that matches the one that the access matrix specifies for your account. For example, the access matrix of an APL file called `myfile` might contain the following row:

```
114  4609  38787
```

If your account number is `114`, you may tie the file using any tie number and the passnumber `38787`. For example, to give the file the tie number `44`, you execute the following:

```
'myfile' ⎕stie 44 38787
```

Doing so gives you permission to use with this file all the functions encoded by the number `4609`. Having tied the file with a passnumber, you must include the passnumber with the tie number in every use of the file.

Supplying a (nonzero) passnumber when none is required is just as much an error as failing to supply one when it is required. Either action, or using a function excluded by the permission matrix, is rejected with the message *`file access error`*.

Warning: It is possible to alter a file's access matrix in such a way that it becomes impossible to do anything useful with it (at least from APL). Of course, all these security measures refer only to access by way of the APL interpreter. The files still exist in the UNIX environment, whose requirements are not necessarily those of APL. A utility function called *`filehelper`* (available in the public workspace distributed as `1 tools`) can extricate you from this predicament. Only the owner of a file can execute it, and then only once. The side effect of executing this function is to remove the access matrix from the file, thus restoring full access to its owner.

File System Functions

The file system functions listed in this chapter are summarized in Table 2.1. The convention for describing rank is to present the monadic, dyadic left, and dyadic right ranks in that order. For functions that are strictly monadic or dyadic, only the first item or the last two items in the rank list are shown. Two additional symbols are used: ∞ to indicate infinite rank, and ⋆ to indicate undefined rank.

Table 2.1. File system functions with permissions.

Category	*Name*	*Permissions*
Inquiry	`⎕avail`	
	`⎕lib`	
	`⎕names`	
	`⎕nums`	
	`⎕paths`	
Existence	`⎕create`	
	`⎕rename`	`128`
	`⎕erase`	`4`
Linkage	`⎕tie`	`2`
	`⎕stie`	any permission
	`⎕untie`	any tied file
Reading/Writing	`⎕read`	`1`
	`⎕append`	`8`
	`⎕appendr`	`16384`
	`⎕replace`	`16`
Description and Limits	`⎕rdci`	`512`
	`⎕rdfi`	`65536`
	`⎕size`	`1` or `32768`
	`⎕resize`	`1024`
	`⎕drop`	`32`
Access Control	`⎕stac`	`256` or `8192`
	`⎕rdac`	`256` or `4096`
Concurrency	`⎕hold`	`64`
	`⎕fhold`	`2048`
Maintenance	†	

† *See* `sift` *in "Chapter 3. File Utilities"*
`⎕read` *and* `⎕size` *share permission* `1`
`⎕rdac` *and* `⎕stac` *share permission* `256`

File System Functions for Inquiry

The following system functions deal with inquiries about APL component files or the file system as a whole.

⎕*avail* File Availability

Rank: `*`

In SHARP APL for OS/390 , this monadic system function reports whether the APL component file system is active. The function ⎕*avail* was retained in the UNIX version for the convenience of those who transfer applications from other systems. However, the file system is an integral part of SHARP APL for UNIX, so the only possible result is a numeric `1` (indicating that the file system is available). The only acceptable argument is an empty vector.

```
      ⎕avail ⍳0
1
```

⎕*lib* File Library

Rank: `1`

The monadic system function ⎕*lib* returns a character matrix containing the names of files in the single directory indicated by the argument.

Right argument: A directory path or library number. To identify a directory in path format, supply a character constant containing the UNIX path to the directory, either from the root directory or relative to your current directory. To identify a directory in library format, supply a library number as an integer item (or a `1`-element integer vector). You may specify your own library number (assigned by the system administrator and returned as the result of `0@ ⎕ai`). All private libraries are automatically available for a session. Public libraries are recognized when the UNIX script from which you started APL includes the parameter `-L` followed by a list of public library numbers and the directories to which they refer. (See the *System Guide, Chapter 3* for more information on libraries and directories)

Result: A character matrix with one row per file. It includes all APL component files in the indicated directory.

When your argument uses path format, each row contains only the name of the file without the path to the directory that contains it. The entries are left justified, and rows are right-padded with blanks to give a width as wide as the longest file name. The suffix `.sf` is not shown.

```
      ⎕lib '/usr2/joe'
tst
bench
```

When your argument is a library number or an empty vector (which signifies your account number), the result is a `22`-column matrix. The first `10` columns contain the library number, right-justified with leading blanks. The next column is blank. The last `11` columns contain the file name, left-justified with trailing blanks.

```
      ⎕lib ''
 103 utl
 103 tst
```

⎕names Library Names of Tied Files

The niladic system function ⎕*names* provides a matrix containing the file names of files currently tied by this task, one file per row. The result is in library format. Provided you do not create, tie, or untie a file between invoking ⎕*nums* and ⎕*names*, these two functions have their values in the same order. For a row in ⎕*names*, the corresponding item in ⎕*nums* shows the number to which that file is tied.

```
      ⎕names
 106 utils
 103 tools
```

⎕nums Tie Numbers of Tied Files

The niladic system function ⎕*nums* provides a list containing the tie numbers of files tied by the current task, one file per row. Provided you do not create, tie, or untie a file between invoking ⎕*names* and ⎕*nums*, these two functions refer to files in the same order. For a row in ⎕*names*, the corresponding item in ⎕*nums* shows the number to which that file is tied:

```
      ⎕nums
1 541
```

⎕*paths* Pathnames of Tied Files

The niladic system function ⎕*paths* provides a character vector of enclosed arrays containing the full UNIX pathnames of all currently tied files, one file per item. The result is in path format. Provided you do not create, tie, or untie a file between invoking ⎕*nums* and ⎕*paths*, these two functions refer to files in the same order. For an item in ⎕*paths*, the corresponding item in ⎕*nums* shows the number to which that file is tied.

```
      ⎕ps←2/¯1 3
      ⎕paths
|-----------------| |-----------------|
|/usr2/fred/utils| |/usr2/bill/tools|
|_________________| |_________________|
```

File System Functions for Existence

The following functions deal with creating, renaming, and deleting APL component files.

⎕*create* Create a File

Rank: `1 1`

The dyadic system function ⎕*create* creates a new APL component file.

Right argument: An integer scalar or vector of one or two integers. The first number must be a positive integer less than 2^{31} that is not currently in use as a tie number. Once the file you name in the left argument is created, this integer becomes the file's tie number. The second number, if supplied, must be a positive integer less than 2^{31}. This integer will be used as the starting component number for the new file. You cannot specify a passnumber when you create a file.

Left argument: A character vector (or scalar) containing the name proposed for the new file. Where appropriate, the name may be preceded by a directory path (either stated as an absolute path from the root directory, or relative to your home directory). A file name in library form may include the character representation of a library number, separated from the file name by a blank. If no library number appears, the named file is presumed to be created in your current directory. The destination library must already exist, and you must have access suitable to create a file in it.

Names for APL Files

The name of a component file is formed from the same characters as those permitted for names of APL objects: any of the 52 characters *a* to *z* and *A* to Z, the characters Δ, ⍙, or _, or the numerals 0 to 9. The name may not, however, start with an underbar or a numeral; the length of the name may not exceed 11 characters; and the name must not currently be in use by another file in the directory specified.

Size Limit

You may optionally include a file size limit when you create a file. This number, in character form, follows the file name and must be separated from the name by one or more blanks. The file system honors any request to append one new component or replace an existing one with a larger array than it now contains, even a request that extends the file beyond an existing size limit. When you include no limit, the value 100000 is assumed by default.

Effect: A new file is created, provided the arguments are valid. The new file has no components. It has a 0-by-3 access matrix granting unlimited access to the file's owner, and no access to anyone else. (See the section "Control of Access to a File," earlier in this chapter.) The new file is share-tied to the tie number indicated in the right argument.

Result: An empty matrix.

For example, the statement

```
'/usr2/bill/newfile 200000' ⎕create 12
```

creates the file *newfile* in the directory /usr2/bill. Once the file reaches 200,000 bytes, it may not grow further until the limit is increased with ⎕*resize*. The file is share-tied to the tie number 12.

⎕*erase* Erase a Tied File

Rank: 1 1

The dyadic system function ⎕*erase* lets you remove a tied file from its file library (if you have permission to erase it).

Right argument: The tie number of a tied file. If you specified a nonzero passnumber when you tied the file, it must appear as the second item of the right argument.

Left argument: A character scalar or vector containing the name of the file to be erased. The left argument is intentionally redundant. Information on the name of the file is available from the tie number. However, the interpreter requires you to provide both forms of identification to make it less likely that you'll inadvertently erase the wrong file.

Permission: `4`

Effect: The file ceases to exist. The number formerly tieing it is again free. The file's name disappears from the result of `⎕lib` and from `⎕nums` and `⎕names`. The UNIX system marks the space the file occupied as garbage, and in the course of its other work will sooner or later overwrite it.

Result: An empty matrix.

```
'file' ⎕tie 1
'file' ⎕erase 1
```

⎕rename Rename a Tied File

Rank: `1 1`

The dyadic system function `⎕rename` lets you change the name of a tied file.

Right argument: The tie number of a tied file. If you included a nonzero passnumber when you tied the file, that passnumber must appear here as the second item of the right argument.

Left argument: A character vector containing the new name proposed for the file; it must be a well-formed name. It may be preceded by a library number or directory path. You can change the name so that the file is moved from the directory in which it first resided to your own directory, or to a directory to which you have write access, provided that the file's access matrix includes permission for you to rename it, and the UNIX system permits you to create a file in the target directory and to delete a file in the source directory.

It is not usually possible to transfer a file to someone else's private library with the system function `⎕rename`. If you want to give a file to someone else, first give that person `⎕rename` permission for that file, then ask that person to rename it (so that `⎕rename` is used by the recipient rather than the donor).

As with the system function ⎕*create*, you may also include a size limit in the left argument.

Permission: `128`

Effect: The file is renamed as indicated. It now belongs either to your file library, or to the file library you indicated in the left argument. Either way, you are now the owner of the file.

Result: An empty matrix.

```
'106 yourfile' ⎕stie 12
'myfile' ⎕rename 12
```

Note: The behavior of ⎕*rename* under UNIX differs with the behavior under OS/390. If another task (either an S-task running on your account or any task running on another account) has permission to rename a file you own, and your current task has the filed tied at the time the rename from the other task occurs, your active workspace will not see the new name until you untie the file. One effect of this behavior is that you cannot erase the renamed file using the (now obsolete) file name as it is known to your current task. In general, a file which you know is to be shared by other tasks should be erased only under ***carefully controlled conditions***.

File System Functions for Linkage

For most file functions, a file is not named directly, but instead referred to indirectly by a tie number. To initiate use of a file, it must first be tied to establish the linkage between its name and the tie number.

Note: There is a limit imposed by UNIX on the total number of files tied system-wide (summed over all users). When you attempt to tie a file that would exceed that limit, the system rejects the statement with the message *file system ties used up*.

⎕*stie* Share-tie and ⎕*tie* Exclusive-tie

Rank: `1 1`

The dyadic system functions ⎕*stie* and ⎕*tie* let you establish the linkage between the name of an APL component file and the tie number used to refer it.

Right argument: The proposed tie number followed, if required, by the passnumber called for in the file's access matrix. The proposed tie number must be a positive integer less than 2^{31} and not be in use by this task as a tie number for another file. Note that it does not matter what the tie number is, provided that you choose a number that is not in use to tie some other file, and (as a practical matter) that applications referring to the file have a way to know what the tie number is.

Note: It is poor practice to have the tie number appear as a constant in your programs. To do that would run the risk that a tie number appearing in one program conflicts with a tie number of some other program that you want to have running at the same time. It is preferable to generate a tie number arbitrarily and then store it where it can be found by the functions requiring it; for example, as a global variable in the workspace. You can also invoke a function that in effect refers to files by name, matching the name you propose with an entry in ⎕*names* or ⎕*paths*, and then selecting an appropriate item from ⎕*nums*.

You can generate an arbitrary nonconflicting tie number by an expression such as the following:

```
tieno←(~⎕io)+((⍳1+⍴⎕nums)∈0,⎕nums)⍳0
```

When the right argument of a tie function has two items, the second is the passnumber. When the argument is a scalar or `1`-element vector, a passnumber of `0` is understood.

Left argument: A character scalar or an vector containing the name of an existing file. The first part of the left argument may contain a path to the file (if it is in a directory other than your home directory) or a library number. If you omit indication of a directory or a library, your own current directory is assumed.

Permission: If you have permission to invoke any other function with a file, you can also share-tie it. However, unless you are the owner of the file, you need explicit permission to use the exclusive form, ⎕*tie* (permission code `2`).

Effect: When you make a request to tie a file with ⎕*tie*, the file system ties the file for you only if no other task has tied it. You then have exclusive access to the file, and no other task may tie it (either by ⎕*tie* or by ⎕*stie*) until you untie it.

When you use ⎕*stie*, the file system ties the file for you provided no user (including you) has tied it exclusively (that is, with ⎕*tie*). While you have the file share-tied, another task may share-tie it also. That's the meaning of share-tieing: your tie is not exclusive and does not prevent others from share-tieing the same file at the same time. When you share-tie a file, the system does not inform you whether other tasks have share-tied the file.

Result: An empty matrix, for either tie function. In the event that conflicting use by another task prevents you from tieing the file, your statement containing ⎕*tie* or ⎕*stie* is rejected with the message *file tied*.

```
      ⎕nums

      '103 file' ⎕stie 1 12
      ⎕nums
1
```

⎕*untie* Untie Tied Files

Rank: 0

The monadic system function ⎕*untie* unties the files whose tie numbers appear in its right argument.

Right argument: An integer scalar or vector of one or more tie numbers. Passnumbers do not apply to ⎕*untie*, and you do not include a passnumber even if you specified one to tie the file. Any number appearing in the right argument is assumed to represent a tie number.

The behavior of ⎕*untie* is permissive: integers that are not tie numbers are ignored.

Effect: Files whose tie numbers appear in the right argument are untied. There is no way to refer to the files that were untied until you subsequently tie them again.

Result: An empty matrix.

```
      ⎕nums
1 3 999
      ⎕untie 3
      ⎕nums
1 999
```

File System Functions for Reading and Writing

The following system functions let you have access to the data in a file after the linkage between file name and tie number is established.

⎕*append* and ⎕*appendr* Append Variable to Tied File

Rank: ∞ 1

The dyadic system functions ⎕*appendr* and ⎕*append* are used to place an array in a new component, which is appended to the end of a tied file. ⎕*appendr* differs from ⎕*append* in that it also returns the new component number as a result.

Right argument: The file's tie number followed, when required, by a passnumber.

Left argument: A variable to be appended as a new component at the end of the file.

Permission: 16384 for ⎕*appendr* and 8 for ⎕*append*.

Effect: Provided that the number of bytes already in the file (see ⎕*size*) does not exceed the file's size limit, a new component is appended. That component contains the variable in the left argument together with its description, which indicates type, rank, and shape. The component also contains information indicating the number of bytes of storage that the variable requires, the account number of the task that appended the new component, and a time stamp (see the description of the file system function ⎕*rdci*).

Result: An integer item (for ⎕*appendr*) whose value is the component number of the newly appended component. The result of ⎕*append* is an empty matrix.

```
     'new' ⎕create 1
     ⎕size 1
1 1 1056 100000
     (3 3⍴⍳9) ⎕appendr 1
1
     ⎕size 1
1 2 1116 100000
```

⎕*read* Read a Component of a Tied File

Rank: 1

The monadic system function ⎕*read* lets you read the array stored in a component of a tied file.

Right argument: Identifies a single component of a tied file as a vector of two or three integers. The first item contains the tie number of a tied file. The second item contains the number of an existing component within that file. If a

component exists, its component number is less than the second item of the result of ⎕*size*, but not less than the first. The system rejects an attempt to refer to a nonexistent file component with the message *file index error*. The third item, where required, is the passnumber.

Permission: 1.

Result: The value of the variable stored in the indicated file component.

```
'tools' ⎕stie 12
x←⎕read 12 1
```

⎕*replace* Replace Variable Stored in a Component

Rank: ∞ 1

The dyadic system function ⎕*replace* lets you replace the contents of a component of a tied file.

Right argument: A vector of two or three integers identifying the tie number, component number, and passnumber, if required. The first element in the vector contains the tie number of a tied file. The second item contains the number of an existing component within that file (see the discussion of the second item of the argument of ⎕*read*). Where required, the third item contains a passnumber.

Left argument: The variable whose value is to replace the value previously stored in that component. It may have any size, shape, or type. There is no requirement that it have the same shape or type as the variable it replaces.

Permission: 16.

Effect: The variable formerly stored as the specified component of the file is replaced by the new variable. For the system to make the replacement, one of the following must be true:

- The storage already in use for the file does not exceed the file's size limit.
- The new component requires no more space than the old one.

Following a successful replacement, the component contains the data itself, together with a description indicating the type, rank, and shape of the data. It also contains information indicating the number of bytes required to store the

object, the account number of the task that did the replacement, and a time stamp indicating the date and time at which the replacement took place (see the description of the file system function □*rdci*).

Result: An empty matrix.

```
      x←3 3⍴⍳9
      'file' □tie 1
      □size 1
1 5 6420 100000
      x □replace 1 2
      x≡□read 1 2
1
```

File System Functions for Description and Limits

The following file functions report information regarding a file or components within it, and establish limits on its growth.

□*drop* Drop Components from a Tied File

Rank: 1

The monadic system function □*drop* lets you drop components from the beginning or end of a tied file.

Right argument: A 2- or 3-element integer vector. The first element contains the tie number of a tied file. The absolute value of the second item indicates the number of components to drop. When this number is positive, components are dropped from the beginning of the file; when negative, they are dropped from the end. If you attempt to drop more components than exist, the system rejects the expression with the message *file index error*. When required, the third item of the argument is a passnumber.

Permission: 32.

Effect: The indicated number of components is dropped from the file. You can drop a block of consecutive components either from the beginning or from the end of a file. That reduces the range of valid component numbers. Suppose a file has 100 components, numbered from 1 to 100:

```
      ⎕size tn
1 101 156842 200000
```

If you drop the last `10` components,

```
      ⎕drop tn,¯10
      ⎕size tn
1 91 146248 200000
```

the file will then have `90` components, numbered from `1` to `90`. But if instead you drop the first `10`,

```
      ⎕drop tn,10
      ⎕size tn
11 101 136482 200000
```

the file will have `90` components numbered from `11` to `100`. Dropping components from a file does not alter the numbers by which you refer to the components that remain.

If you drop all components from a file, whether from the beginning or the end, the number of bytes used by the file will not drop to `0`: some overhead is required for directory and access information. For information on file maintenance and reclamation of space, see "Chapter 3. File Utilities".

Result: An empty matrix.

⎕*rdci* Read Component Information

Rank: `1`

The monadic system function ⎕*rdci* returns information about a particular component of a tied file.

Right argument: Refers to particular component in a file and is a vector of either two or three integers. The first contains the tie number of a tied file. The second is the number of a particular component in that file. If you specified a nonzero passnumber to tie the file, there must also be a third item containing that passnumber.

Permission: `512`.

Result: A `3`-element numeric vector containing the following information about the indicated file component:

```
r←⎕rdci ⍵,1
```

`0@ r` Size of the file component in bytes, including both the data and the data description (which specifies its type, rank, and shape).

`1@ r` Account number of the task that wrote the current value of the component.

`2@ r` Time stamp showing when the current value of the component was written; time is recorded in 60^{ths} of a second since 1960/3/1 at 00:00:00.

The representation in 60^{ths} may be converted to days, hours, minutes, and seconds by the expression:

```
0 24 60 60 60⊤2@ r
```

The function *timen* in the public workspace `1 tools` converts the time stamp into ⎕*ts* form:

```
     timen 2@ r
2000 4 11 15 23 4 17
```

⎕*rdfi* Read File Information.

Rank: `1`

The monadic system function ⎕*rdfi* returns information about a tied file.

Right argument: Refers to a particular file and is a vector of either one or two integers. The first contains the tie number of a tied file. If you specified a nonzero passnumber when you tied the file, there must also be a second item containing that passnumber.

Permission: `65536`.

Result: A `4`-by`2` integer matrix containing information for the file you specify. Each row contains a user number and a time stamp (presented in the same format as for the last item in the result of ⎕*rdci*).

```
r←⎕rdfi ⍵
```

`0@ r` The first row of information relates to the creation of the file using the ⎕*create* system function. It contains the user number that created the file, followed by the time stamp of when the file was created.

`1@ r` The second row relates to the latest setting of the file access matrix using the ⎕*stac* system function. It contains the user number that last set the access matrix, followed by the time stamp of when the access matrix was set.

`2@ r` The third row is reserved for future use, and is currently set to `¯1`.

`3@ r` The fourth row relates to the latest alteration of the file, using the system functions ⎕*append*, ⎕*appendr*, ⎕*drop*, ⎕*rename*, ⎕*replace*, or ⎕*resize*. It contains the user number that last altered the file, followed by the time stamp associated with that alteration. This information might be used to implement a file caching scheme for shared files.

⎕*resize* Set File Size Limit

Rank: `1 1`

The dyadic system function ⎕*resize* lets you change the file size limit for a tied file.

Right argument: The file's tie number followed, when required, by a passnumber.

Left argument: The desired new file size limit. Its value must be greater than or equal to zero.

Permission: `1024`.

Effect: The file's reservation is revised as indicated.

Result: An empty matrix.

```
     ⎕size n←12
1 5 2096 100000
     50000 ⎕resize n
     ⎕size n
1 5 2096 50000
```

□*size* Size of Tied File

Rank: `1`

The monadic system function □*size* returns information about the size of a tied file.

Right argument: The file's tie number followed, when required, by the passnumber.

Permission: `1`, which also gives permission to use □*read*; or `32768`, which is specific to □*size*.

Result: A `4`-element integer vector describing the file's size as follows:

```
      ⊢s←□size ⍵
1 5 2096 100000
```

`0@ s` Lowest component number for this file (`1` for a file from which no leading components have been dropped).

`1@ s` Number of next component to be appended (`1` more than the number of the last existing component).

`2@ s` Total space now occupied by the file, in bytes. This number includes the file's internal directories and may include dead space when a file has grown by replacing some components with arrays larger than they previously contained.

`3@ s` File size limit, in bytes.

If ⍵ is the tie number of a newly created (and therefore empty) file having a size limit of `100000` bytes, then its size would be reported as follows:

```
      □size ⍵
1 1 1056 100000
```

The number of components in the file tied to ⍵ is given by:

```
      -⊂/2↑□size ⍵
```

File System Functions for Access Control

The file system provides detailed control of access to a file both by the account number of the task doing the action and the specific action that is permitted. The following functions set or report on access controls.

⎕*rdac* Read Access Matrix

Rank: `1`

The monadic system function ⎕*rdac* returns the access matrix of a tied file.

Right argument: The file's tie number followed, when required, by a passnumber.

Permission: `256` or `4096`.

Result: The `3`-column integer library access matrix of the file. The first integer in each row is an account number, the second is a permission code, and the last is a passnumber. The significance of the columns in the access matrix is discussed in the section “Control of Access to a File,” earlier in this chapter.

⎕*stac* Set Access Matrix

Rank: `2 1`

The dyadic system function ⎕*stac* lets you set the access matrix of a tied file.

Right argument: The file's tie number followed, when required, by a passnumber.

Left argument: The proposed access matrix. It is a `3`-column matrix of integers or, if you are forming a `1`-row matrix, a `3`-element vector of integers. The interpreter requires that the argument have this form and contain integers, but does not otherwise verify that the values in it are valid.

Permission: `256` or `8192`.

Effect: The left argument becomes the new access matrix for the file. Note that, unlike the OS/390 version, SHARP APL for UNIX does not reevaluate permissions to currently tied files. The new access matrix only affects new ties to the file.

Result: An empty matrix. The following gives to account `154` unlimited access with no passnumber, and to all accounts permission `3` (that is, actions `1` and `2`) provided they use passnumber `99`:

```
     ⊢mat←2 3⍴154 ¯1 0 0 3 99
154 ¯1  0
  0  3 99
     mat ⎕stac 4
```

Note: If you invoke ⎕*stac* in such a way that you exclude yourself from ⎕*stac* access to the file, you cannot execute ⎕*stac* again to restore it. See the utility function *filehelper* in the public workspace `1 tools`.

File System Functions for Concurrency

The following functions deal with coordinating use of a set of files some of which are in use by different tasks at the same time.

⎕*fhold* Hold Tied Files

Rank: `1 2 1`

⎕*hold* Hold Tied Files

Rank: `2`

The system functions ⎕*fhold* and ⎕*hold* coordinate concurrent access to shared files. They permit one APL task to request temporary exclusive use of all or part of a file that is at the same time share-tied by other tasks.

⎕*fhold* is the newer and more general form. ⎕*hold* is unable to accept a left argument, and is unable to include passnumbers or component ranges in its right argument. However, for uses that do not require these, ⎕*hold* and ⎕*fhold* have the same effect. In particular, they operate on a common queue of held files, and in that respect are interchangeable.

Right argument: An array of tie numbers. For ⎕*fhold*, this may also be a matrix of one to four rows. When it is a `4`-row matrix, the rows are used as follows:

`0@⍵`	Tie numbers. A `1`-row matrix has the same effect as a vector.

`1@ω`	Passnumbers. Passnumber `0` is equivalent to "no passnumber." Because `⎕hold` is unaffected by passnumbers, a file used for an application requiring passnumbers is usually given an access matrix prohibiting use of `⎕hold`.
`2@ω`	Lower bound of a hold range.
`3@ω`	Upper bound of a hold range. The effective hold range includes the lower and upper bounds. Thus, if the lower and upper bounds are the same, the range is one component. If the lower bound is greater than the upper bound, the range is empty.

When the right argument has fewer than four rows, the following defaults apply:

`3` *rows:*	The hold range is one number, identified in row `3`.
`2` *rows:*	The hold range is not specified, and applies to all components.
`1` *row:*	No passnumbers (equivalent to passnumber `0`). Hold applies to all components.

The *hold* mechanism permits cooperating applications to pass information to each other by way of a file that both have tied.

Where they wish to hold specific ranges within a file, the application seeking to hold a range executes `⎕fhold` with a `3`-row or `4`-row argument.

When no component range is stated, the hold applies to the entire range of integers -2^{31} through $2^{31}-1$ (and thus to all possible components when the range is interpreted as a block of consecutive components). The components described by the range need not exist. This allows `⎕fhold` to be used as a general enqueue mechanism.

Monadic ⎕fhold and ⎕hold

Monadic use of either file holding function has the following effect:

- All preceding holds (if any) are released. It does not matter whether the new request to hold applies to the same files or to different ones. Each execution of either of the hold functions cancels any preceding use.

- Execution of your task goes into an indefinite delay until all the files you have asked to hold are free of holds set by other tasks.

  A keyboard interrupt cancels a pending hold even when the interrupt is trapped. If the interrupt is not trapped, control returns control to the keyboard.

- Once all the files you have requested are free, your hold takes effect. Execution of your task proceeds. See "Duration of a File Hold," below.

When the interpreter receives from another task a request to hold any of the files you are holding, or for a range that overlaps a range you are holding, that task goes into an indefinite delay until either you release your hold on the files the other task has asked for, or you untie them. SHARP APL does not inform you of a pending request by another task.

Permission: `2048` for `⎕fhold`; `64` for `⎕hold`.

Result: An empty matrix.

Dyadic `⎕fhold`

Left argument: An item, either `1` or `¯1`, with effects as follows:

`1` ***Selective hold***. The files identified in *ω* are added to the hold queue without releasing any files previously held.

`¯1` ***Selective release***. The files identified in *ω* are released without releasing any other files.

Permission: `2048`.

Result: A *list* of the tie numbers of the files held following its execution, or a `3`-row matrix showing the tie numbers of the held files together with the first and last members of the held range of each file.

When *ω* has two rows or fewer (that is, refers to files but not to held ranges), the result is a vector, with one element for each held file, containing their tie numbers.

When *ω* has three or four rows, the result is a `3`-row matrix, in which the first row contains the tie numbers of held files, and rows `2` and `3` contain the first and last members of their held ranges.

When ⍵ is empty, there is no change to the way files are held, so any of

```
1 ⎕fhold ''  or  1 ⎕fhold 1 0⍴''  or  1 ⎕fhold 2 0⍴''
```

returns a list of the tie numbers of held files, and either of

```
1 ⎕fhold 4 0⍴''  or  1 ⎕fhold 3 0⍴''
```

returns a matrix showing the tie numbers of all holds.

Duration of a File Hold

A file hold, once in effect, is ended by any of the following:

- `¯1 ⎕fhold` for that file.
- Any use of monadic `⎕fhold` or `⎕hold`.
- The end of your task (sign-off, bounce, crash, etc.).
- Untieing the affected file or files.
- Any return to immediate execution (including one trapped as event `2001`, since the only actions possible with event `2001` lead ultimately to immediate execution). However, returning to the keyboard for ⎕- or ⍞-input does not terminate a hold.

Note: You cannot experiment with file holds by entering statements one after another in immediate-execution mode. Any hold you enter that way terminates immediately when the system returns for the next entry from the keyboard. In that situation, the most you could do would be to put, into a single line of entry (separated by diamonds) both the file hold and the statement you want executed during the hold.

Note: Inside the block spanned by a file hold and its subsequent release, make sure than none of the functions invoked there issues a monadic hold, since such a monadic hold would thereby undo your earlier file hold when it issued its own.

File Space Management

A UNIX command called `sift` is supplied to let you compact a component file to recover space no longer used. Such wasted space may come about when you replace components with arrays larger than their previous contents, or when you drop components with the system function □*drop*.

You can invoke the command `sift` on any component file to which UNIX grants you write access. The command preserves the entire file as it compacts it — the access matrix, ownership, and contents of existing components are preserved unchanged.

You can use the command `sift` either directly from the UNIX shell, or through AP11 using a variable shared between SHARP APL and UNIX. For more information, see "Chapter 3. File Utilities"

3
File Utilities

Introduction

This chapter documents *utilities* used to monitor or optimise APL files. These utilities, located in `usr/sax/rel/bin`, are invoked from the UNIX command line:

`tiers`	reports all ties of a list of APL files.
`holders`	reports all holds (via □*hold* or □*fhold*) on a list of APL files.
`flockers`	reports the owners of APL's low-level UNIX locks (as obtained by `fcntl()`) on a list of APL files.
`sift`	compacts APL files to remove unused space.
`sortio`	reduces processing time by sorting all APL files of a specialized format.
`cfcheck`	reports on the structural integrity of an APL component file.

Note: If the Fileserver is installed, some of these utilities may not work or may not be required. See "Chapter 4. The File System Server" for more information.

The `tiers` *Utility*

`tiers [-u]` *filelist*

The `tiers` utility reports the status of all ties of a list of APL files. Although it does not actually write to the files, the `tiers` utility requires read and write access to all the files in its ***filelist*.**

The argument ***filelist*** is a list of UNIX pathnames separated by spaces, where each pathname is the name of an APL component file with or without the `.sf` suffix.

Note: `tiers` does not workwhen the Fileserver is installed. It simply reports errors. As explained in "Chapter 4. The File System Server", `fsutil -T` should be used instead.

"Unlocked" Option (`-u`*)*

The option flag `-u` indicates that the `tiers` utility runs unlocked (i.e., without obtaining the APL file operation lock `FOPLOCK`). Use of the `-u` flag minimizes the detrimental impact of the `tiers` utility upon the performance of the APL processes that have the files tied. If this flag is specified, there is a small possibility that the result of the `tiers` utility could be abnormal due to race conditions between `tiers` and APL.

Furthermore, if an APL process is in a hung state with the `FOPLOCK` for a file held, use of the `-u` flag allows `tiers` to produce its report, which can aid in identifying the hung process. However, the `flockers` utility is more useful for identifying such processes. Use of the `-u` flag does not cause damage to the APL file.

`tiers` *Report Summary*

The result of invoking the `tiers` utility is a multiple line report for each file in the ***filelist***. The first line is a title listing the name of the file, the last line reports the status of the file locking semaphore set for that file, and the intervening lines list all of the current ties of the file. The reports for each individual file are separated from each other by blank lines.

If there are no ties of the file, the report for that file is the title line followed by the message `no ties found`.

File Tie Information

The information reported for each file tie is as follows:

`Host id` Identifier of the host machine on which the APL process that tied the file is running. This is the result of the `gethostid()` system call formatted in hexadecimal.

`Process id` Process ID of the APL process that tied the file.

`User id` User number of the APL process that tied the file.

`Type` Type of file tie. There are three types of file tie, denoted numerically as `0`, `1`, or `2`, corresponding to the values of the `-Yusefilelocks` parameters of the APL sessions that have tied the file.

Type `0` indicates that no file locking is employed by that tier of the file. Type `1` indicates that the standard host file locking primitive (typically `fcntl()`) is in use by that tier of the file. This is the default type of tie. Type `2` indicates that the semaphore file locking mechanism is being employed by that tier of the file. The existence of a type `2` tie on a file implies the existence of a file locking semaphore set for that file.

It is not possible for both type `1` and type `2` ties to simultaneously exist for the same file. Although type `0` ties can co-exist with either type `1` or type `2` ties, this can affect the integrity of the file.

File Locking Semaphore Information

If there are more than one type `2` ties of a file then there is a file locking semaphore set associated with that file. If there are no type `2` ties of a file, the file usually does not own a file locking semaphore set, though it may do so as a result of abnormal termination of an APL process that previously had a type `2` tie on the file (an unusual situation that is not dangerous).

The next type 2 tie issued for that file uses the existing file locking semaphore set rather than allocate a new one. The non-existence of a file locking semaphore set can be assumed when zero information is reported for the file locking semaphore set. The information reported for the file locking semaphore set is as follows:

`Semaphore host id` — Identifier of the host machine where the file locking semaphore set resides. All APL processes which have type 2 ties on the file must be on this host machine. The host ID is the result of the `gethostid()` system call, formatted in hexadecimal.

`Semaphore owner` — User number of the process that currently owns the semaphore set. The semaphore set is initially owned by the user represented by the first process to issue a type 2 tie of the file. When the current owner of the semaphore unties the file, the semaphore is "passed on" to another user that has a type 2 tie on the file. This mechanism is necessary so that the last process that has a type 2 tie on a file can deallocate the semaphore set when it unties the file.

Only the owner of a semaphore set can deallocate the semaphore set. It is therefore necessary to ensure that the owner of the semaphore set is always a user that currently has a type 2 tie on the file.

This mechanism can fail when the owner of the semaphore set terminates abnormally, resulting in the semaphore set being left in existence and associated with the file. This situation is corrected when that user issues another type 2 tie for the file or when the machine is rebooted. The semaphore set can also be manually deallocated via the `ipcrm` system command, but this command should not be used unless it can be guaranteed that no APL process attempts to tie the file while the semaphore set is being deleted.

Meanwhile, for as long as the semaphore set remains in an "orphaned" state, all file ties and unties continue to proceed normally, except that the semaphore set will not be deallocated when no tiers of the file remain. There is one exception to this. Any attempt to tie a file from a host machine other than the one where the semaphore set resides will fail if the file possesses an orphan semaphore set.

`Semaphore key` — Key of the semaphore set, used to gain access to the semaphore via the `semget()` system call.

Example:

```
tiers masdir msgfile

TIES FOR FILE masdir.sf:
Host id: c72b0263, Process id: 24590, User id: 116, Type: 1
Host id: 7230b679, Process id: 13090, User id: 104, Type: 1
Semaphore hostid: 0, Semaphore owner: 0, Semaphore key: 0

TIES FOR FILE msgfile.sf:
Host id: c72b0263, Process id: 22307, User id: 116, Type: 2
Semaphore hostid: c72b0263, Semaphore owner: 116, Semaphore key:
 2103720d
```

The file `masdir.sf` is tied by two APL processes: process `24590`, owned by user number `116` on the host machine identified as `c72b0263`, and process `13090`, owned by user number `104` on the host machine identified as `7230b679`. Both file ties are type `1`, which specifies that the standard host file locking mechanism is in use. As expected, this file has no file locking semaphore set, indicated by the zero values for all the semaphore fields.

The file `msgfile.sf` is tied by one APL process, process `22307`, owned by user `116`, on the host machine identified as `c72b0263`. The tie is type `2`, indicating use of the semaphore file locking mechanism. As expected, the file locking semaphore set exists on the same host machine as the only tier of the file, and the owner of the semaphore set is the only tier of the file. The file locking semaphore set key is `2103720d`.

The `holders` *Utility*

`holders [-u]` *filelist*

The `holders` utility reports the status of all current holds obtained via `⎕hold` or `⎕fhold` on a list of APL files. Although it does not actually write to the files, the `holders` utility requires read and write access to all files in its ***filelist***.

The argument ***filelist*** is a list of UNIX pathnames separated by spaces, where each pathname is the name of an APL component file with or without the `.sf` suffix.

Note: `holders` does not workwhen the Fileserver is installed. It simply reports errors. As explained in "Chapter 4. The File System Server", `fsutil -H` should be used instead.

"Unlocked" Option (`-u`)

The option flag `-u` indicates that the `holders` utility runs unlocked (i.e., without obtaining the APL file operation lock `FOPLOCK`). If this flag is specified, there is a small possibility that the result of the `holders` utility could be inaccurate due to race conditions between `holders` and APL. Use of the `-u` flag minimizes the detrimental impact of the `holders` utility upon the performance of the APL processes that have the files tied.

Furthermore, if an APL process is in a hung state with the `FOPLOCK` for a file held, use of the `-u` flag allows `holders` to produce its report, which can aid in identifying the hung process. However, the `flockers` utility is more useful for identifying such hung processes. Use of the `-u` flag does not cause damage to the APL file.

`holders` *Report Summary*

The result of invoking the `holders` utility is a multiple line report for each file in the *filelist*. The first line is a title listing the name of the file and the remaining lines list all of the current holds on the file. The reports for each individual file are separated from each other by blank lines.

If there are no holds on the file, the report for that file is the title line followed by the message `no ties found`.

File hold Information

The information reported for each file hold is as follows:

`Host id` Identifier of the host machine on which the APL process that tied the file is running. This is the result of the `getgid()` system call formatted in hexadecimal.

`Process id` Process ID of the APL process that held the file.

`User id` User number of the APL process that held the file.

`Range` The lower and upper bounds of the range held. The lower and upper bounds are always `INT_MIN` and `INT_MAX` for `⎕hold`.

Example:

```
holders msgfile masdir

HOLDS FOR FILE msgfile.sf:
Host id: c72b0263, Process id: 24590, User id: 116, Range: 5 6
Host id: 7230b679, Process id: 13090, User id: 104, Range: 7 9

HOLDS FOR FILE masdir.sf:
Host id: c72b0263, Process id: 22307, User id: 116, Range: -2147483648 2147483647
```

The file `msgfile.sf` is held by two APL processes. Process `24590`, owned by user number `116` on the host machine identified as `c72b0263` has the range `[5,6]` held, and process `13090`, owned by user number `104` on the host machine identified as `7230b679` has the range `[7,9]` held.

The file `masdir.sf` is held by one APL process. Process `22307`, owned by user `116` on the host machine identified as `c72b0263` has the range `[-2147483648,2147483647]`. This hold is likely the result of executing `⎕hold`.

The `flockers` *Utility*

`flockers` *filelist*

The `flockers` utility reports the holders of the APL low-level file operation locks on a list of APL files. These locks are not available to APL programmers, and are only of interest when attempting to diagnose file system problems. Although it does not actually write to the files, the `flockers` utility requires read and write access to all the files in its ***filelist*.**

The argument ***filelist*** is a list of UNIX pathnames separated by spaces, where each pathname is the name of an APL component file with or without the `.sf` suffix.

Note: `flockers` reports `FSERVLOCK` if tied via the Fileserver. See "Chapter 4. The File System Server" for more information.

`flockers` *Report Summary*

The result of invoking the `flockers` utility is a two line report for each file in *filelist*. The first line is a title listing the name of the file and the second line lists the holder of each of the low-level APL file locks. The reports for each individual file are separated from each other by blank lines.

Lock Holder Information

The information reported for each file is as follows:

`FOPLOCK` Process ID of the process that currently holds the file operation lock (`FOPLOCK`). The `FOPLOCK` is normally obtained for every file operation. However, if the file is tied by type 2 ties, indicating use of the semaphore file locking mechanism, then the `FOPLOCK` is only obtained for □*create*, □*tie*, □*stie*, and □*untie* operations. If a process has the `FOPLOCK` held for an extended period, then that process is almost certainly hung and is blocking access to the file by other processes.

`FHOLDLOCK` Process ID of the process that currently holds the file hold lock (`FHOLDLOCK`). The `FHOLDLOCK` is normally obtained for □*hold* and □*fhold* operations. However, if the file is tied by type 2 ties, indicating use of the semaphore file locking mechanism, then the `FHOLDLOCK` will not be used at all. If a process has the `FHOLDLOCK` held for an extended period, then that process is almost certainly hung and is blocking other processes from performing □*hold* and □*fhold* on the file.

`FTTABLOCK` Process ID of the process that currently holds the file tie table lock (`FTTABLOCK`). The `FTTABLOCK` is obtained for a brief interval during every □*create*, □*tie*, □*stie*, and □*untie* operation. Furthermore, if the `FTTABLOCK` is held, then the `FOPLOCK` must be held as well. If a process has the `FTTABLOCK` held for an extended period, then that process is almost certainly hung and is blocking access to the file by other processes.

In all cases, a reported process ID of `-1` indicates that no process holds the lock in question.

Example:

```
flockers msgfile masdir

FCNTL() LOCKS FOR FILE msgfile.sf:
FOPLOCK: -1, FHOLDLOCK: -1, FTTABLOCK: -1

FCNTL() LOCKS FOR FILE masdir.sf:
FOPLOCK: 22307, FHOLDLOCK: -1, FTTABLOCK: 22307
```

No processes hold any of the low-level APL file locks for the file `msgfile.sf`. Process `22307` holds both the `FOPLOCK` and `FTTABLOCK` for the file `masdir.sf`. This indicates that process `22307` is in the process of either creating, sharing, tying, or untying the file.

The `sift` *Utility*

`sift` *filelist*

Through use of □*replace* and □*drop*, an APL file may develop pockets of unused space. You can compact your files, or arrange for the system administrator to run a periodic job using the command `sift`. The argument to `sift`, *filelist*, is the name of one or more APL component files, including the suffix `.sf`:

```
% sift *.sf
```

You can activate `sift` directly from the UNIX shell, or by using AP11 from an active workspace. `sift` supplies a return code of `0` for success or `1` for failure.

The file that `sift` produces has exactly the same content as the file you supply as an argument: the total number of components is the same; the contents of those components is unchanged; their component numbers are unaffected; the file's access matrix is reproduced intact; and internal timestamp information is preserved. Only the amount of space required to store the data is modified.

Note: When the Fileserver is installed, `sift` will only work if the file is not already tied.

When to use `sift`

One way to determine if a file may benefit from the `sift` utility is to run `cfcheck` first. When you use `cfcheck`, it can identify problems with a file's structural integrity, including size, and it will issue a warning message. Some of the `cfcheck` messages described in the section "The `cfcheck` Utility" on page 3-13, indicate when a file might require sifting.

The `sortio` Utility

`sortio` [`-N`*rows*][`-S`*bytes*][`-aAibf`] *file1 file2* [*sort keys*]

The command `sortio` is available to provide UNIX sorting of APL data stored in a specialized format in an APL component file. If your application produces appropriate component files, this command can reduce your processing time significantly.

File Format

Each component of the APL file you wish to sort must contain an array of the same structure: a list of boxed simple tables. Corresponding tables in successive components must be of the same APL data type. All tables within a component must have the same number of rows. All tables in all components (except the last) must have the same number of rows as those in the first component; tables in the last component may all have fewer rows than those in preceding components. A table containing numeric data must have exactly one column; a table containing character data may have as many columns as required to store the data.

If these conditions are met, then the command `sortio` can be applied to the file to produce a sorted APL file of the same structure in which the arrays are sorted based on arguments to the command.

`sortio` *Arguments*

The only arguments required for `sortio` are the names of the source APL file and the sorted APL file, both expressed without the `.sf` suffix. If the sorted file already exists, it is overwritten. If you invoke `sortio` from within SHARP APL and the sorted file exists and is tied, it is overwritten and remains tied to the same tie number.

The `sortio` optional arguments control the file structure in three ways:

- *size* of the sorted components (controls either the number of rows per table or the size of the array in each component)
- *data type* of the tables within each component
- *sorting keys*, which UNIX applies using the `sort` command.

The options, if used, must appear ordered with size and data type options before the names of the files, and with sorting keys after the names of the files; as shown in the following syntax example,

`% sortio [-N`*rows*`] [-S`*bytes*`] [-aAibf]` *file1 file2* [*sort keys*]

If you do not use any options, the structure (size, type, and order of tables) of the first component of the source APL file (*file1*) `file1.sf` becomes the basis for the structure of all components (except possibly the number of rows in the last component) of the sorted APL file (*file2*) `file2.sf`.

Size

The two options that control the size of components in the sorted APL file are mutually exclusive. Use only one at a time.

`-N` specifies the number of rows in each table of each component (except possibly the last) of the sorted file. For example,

```
% sortio -N100 file1 file2
```

sorts the file `file1.sf` and places the sorted version in `file2.sf`. The tables within each component (except possibly the last component) of `file2` have 100 rows.

`-S` specifies the size (in bytes) of each component of the sorted file. For example,

```
% sortio -S20000 file1 file2
```

sorts the file `file1.sf` and writes the sorted version in `file2.sf`. The size of the entire array in each component will not exceed 20000 bytes.

Data Types

`sortio` recognizes five data types:

`a` alphabetic (character)

`A` dictionary sorted (character); only letters, digits, and blanks (spaces and tabs) are significant in comparisons

`b` Boolean

`i` integer

`f` floating point

If you use this option, you must specify a list of types that has exactly the same number of items as the number of tables within each component of the source APL file. You can coerce the data types of the tables in the sorted file to be those you name, provided any changes you introduce do not cross the numeric/character boundary. For example, you can force rounding of a table by specifying *integer* (or even *Boolean*) where there was *floating point* before.

If the components of `file1.sf` contain tables with types character, Boolean, integer, and floating point (in that order), then the following expression has the same effect as using the default types for the file:

```
% sortio -abif file1 file2
```

You could round the floating point table and impose dictionary sort order on the character table with the following expression:

```
% sortio -Abii file1 file2
```

Sorting Keys

The command `sortio` goes through three steps as it sorts a component file. First it converts the boxed tables to an internal character form (fields within records). Second, it passes the internal character form to the UNIX `sort` command. Finally, it converts the character output from that command back to the APL tables you expect in the components of your sorted file.

Your knowledge of the UNIX `sort` command is at your disposal; any options you specify are passed through to `sort` without inspection by `sortio`.

The command `sortio` uses the TAB character as its field delimiter for the internal character form. You can specify another delimiter character for the `sort` command as you specify other `sort` options and your delimiter setting will override TAB, but for purposes of sorting only. Your delimiter will not cause removal of TAB characters which are used in the final step of `sortio` to restore the APL tables from the character records.

Suppose we wish the floating point table in `file1.sf` to be the primary sort field, the character table to be the secondary sort field, and the output to be in descending order. The UNIX `sort` command treats fields as numbered in zero-origin. Keys are expressed by the use of `+` to begin and `-` to end a field.

The floating table is the last of four, while the character table is the first. The option `-r` reverses the sense of the sorting, thereby enabling descending order. The entire argument to `sortio` is then the following expression:

```
% sortio -Abii file3 file4 -r +3 -4 +0 -1
```

The `cfcheck` *Utility*

`cfcheck` [`-s`|`-v`|`-V`|`-c`|`-a`|`-f`] *filelist*

The SHARP APL component file check utility, `cfcheck`, determines if a file is damaged as far as its directory/component structure is concerned. This allows the UNIX system administrator to detect problems in time to prevent further damage to files and workspaces. It is a good idea to run `cfcheck` on all files from time to time, especially after a power failure or operating system crash.

If all specified files in *filelist* exist, can be opened for reading and are undamaged, `cfcheck` supplies a return code of `0`; otherwise the return code is `1`.

`cfcheck` does not tie files. It will yield random results if used on tied files that are being updated. Generally, `cfcheck` should be run on files that are not tied.

Note: Unless the `-c` option is specified, `cfcheck` only checks the structural integrity of the file. It reads the file header, component directories and free space bitmaps but not the components themselves. Therefore it is usually quite fast even on a very large file. However, if `-c` is specified, `cfcheck` reads all data in the file, which may take a considerable amount of time .

`cfcheck` *Arguments*

filelist, is the name of one or more APL component files, including the suffix `.sf`. If no options are specified, `cfcheck` displays warning/error messages for files that may be damaged or cannot be accessed.

Example:

```
% cfcheck /dir/file1.sf /dir/file2.sf
File /dir/file1.sf OK
ERR0110  /dir/file2.sf: file not found
```

List of Options

The following options control how the check utility is run and determines the format of the messages displayed:

`-V` displays SHARP APL version information and exits without checking any file. This does not require the *filelist* argument, and ignores it if supplied.

Example:

```
cfcheck -V
```

```
cfcheck version 4.7.00000000  1995/07/18  IBM/6000
SHARP APL Component File Versions 2 to 3
```

The second line of the output indicates that the current version of `cfcheck` understands the format of component files version 2 and 3 but not version 1.

`-s` runs `cfcheck` in ***silent*** mode. This does not return anything except an exit code.

`-v` runs `cfcheck` in ***verbose*** mode. This outputs more warning messages, internal file header information, and a file map. The information returned is useful for Soliton technical support.

`-c` checks component internals. If this option is specified `cfcheck` reads component data as well as checks its internal structure. Because of the extra data, `cfcheck` may take longer to produce a result for large files.

`-a, -f` apply only to the version 1 component file structure and are irrelevant for the current structure (version 3).

`cfcheck` *Messages*

When `cfcheck` returns messages for files that may be damaged or cannot be opened for reading, they appear in the following format:

XXXnnnn filename:*text*

where,

XXX either `MSG` (for warnings) or `ERR` (for errors)

nnnn is the message number

Warning Messages - `MSG` *Prefix*

`0704` `Long component: n1 bytes required, n2 used, component n3 offset n4 size n5`

This appears only if `-v` is specified and usually accompanies `0705.`

`0705` `rank 0 and 0 elements, component n1 offset n2 size n3`

This usually indicates that the component is all binary zeroes including descriptor, rank, shape and value.

`0721` `component` *ccc* `(offset` *n2* `size` *n3*`) data alignment error`

A variable inside the component is not aligned at the doubleword boundary.

`0800` `block` *n1* `not used and not marked as free`

A block is 256 characters long. Such a block can only be recovered by `sift`. Unless the `-v` option is specified this message is only displayed once for the first block found.

Error messages - ERR *Prefix*

`0110` `file not found`

`0120` `file open error` *errno*

Indicates the UNIX error number value after a UNIX system call, in this case `open()`.

`0130` `file fstat() error` *errno*

`0140` `unable to allocate file maps for` *n* `blocks, error` *errno*

`malloc()` was not able to allocate memory. *n* is the number of blocks in the file. File maps require approximately 6 bytes of memory per block.

Note: Errors `ERR0110`, `ERR0120`, `ERR0130`, and `ERR0140` indicate that `cfcheck` hasn't been able to check the current file.

`0210` `lseek() error` *errno*`, offset` *n1*`,` *descr*

Provides the offset at which `lseek()` failed and a description of the of the part of the file `lseek()` tried to locate. Example descriptions are "`file header`" or "`bitmap directory`".

`0220` `read error` *errno*`, offset` *n1*`, size` *n2*`,` *descr*

An error occured while reading `size` *n2* bytes at `offset` *n1*.

0221 `unexpected EOF, offset n1, size n2, read n3, descr`

`read n3` is the actual number of bytes read.

0310 `magic number error: expected n1, got n2`

Most likely not a SHARP APL component file at all.

0311 `component number offset error: is n1, must be >= 0`

Invalid value for file header field '`cmpoff`'.

0312 `file stage error: is n1, must be between n2 and n3`

Invalid value for file header field '`stage`'.

0313 `file version error: is n1, must be between n2 and n3`

Invalid value for file header field 'version'

0314 `component range error: n1 to n2, must be <=`

In the file header, the last component number is less than the first component number.

0315 `first cmp. number error: is n1, must be > cmpoff (n2)`

In the file header, the first component number must be greater than '`cmpoff`'.

0327 `file size error: is n1, must be >= n1`

0328 `file size error: header size n1 is significantly greater than real size n2`

The file size as specified in the file header and the real file size don't agree. If this is the only problem `sift` will probably fix the file.

0329 `access table size 0. Offset is n1, pointer offset n2`

If this is the only problem the file can be fixed with *`filehelper`* from workspace `1 tools` which will erase the damaged access table.

0410 <*descr*> error: offset *n1* size *n2* pointer offset *n3*: negative offset

0411 <*descr*> error: offset *n1* size *n2* pointer offset *n3*: misaligned offset

0412 <*descr*> error: offset *n1* size *n2* pointer offset *n3*: offset must be < both real file size (*n4*) and header file size (*n5*)

0420 <*descr*> error: offset *n1* size *n2* pointer offset *n3*: size <= 0

0421 <*descr*> error: offset *n1* size *n2* pointer offset *n3*: offset + size must be <= both adjusted real file size (*n4*) and header file size (*n5*)

Note: Errors ERR0410 to ERR0421 indicate that there is a problem with file part *descr* which is supposed to reside at offset *n1* and have size *n2*. The pointer offset *n3* is the file offset at which the offending size and/or offset is specified.

0500 <*descr1*> error: offset *n1* size *n2* pointer offset *n3*: block *n4* overlaps with <*descr2*>

0501 <*descr1*> offset *n1* to *n2* overlaps with <*descr2*> offset *n3* to *n4*

This is an equivalent of message ERR0500 for version 1 files.

Note: Error ERR0500 means that file block *n4* appears to be used by both *descr1* and *descr2*.

0610 component number range invalid: zcn *n1*, cmpoff *n2*, stage *n3*

File appears to have more components than is possible for its 'stage'.

0700 unknown descriptor *n1*, component *n2* offset *n3* size *n4*

Component's entry descriptor byte value is unknown to SHARP APL.

0701 negative rank *n1*, component *n2* offset *n3* size *n4*

0702 different leading length *n1* and trailing length *n2*, component *n3* offset *n4* size *n5*

0703 short component: *n1* bytes required, *n2* on disk, component *n3* offset *n4* size *n5*

Less data on disk than is necessary to store a variable of this type, rank and shape

0720 component *n1* (offset *n2* size *n3*) data error: *txt*

txt is a message from `spring()`; e.g., '`object not within linear array`'

4
The File System Server

Overview

The SHARP APL file system server, Fileserver, is a file management system that runs independently of the interpreter, acting as a transparent broker between the UNIX file system and the conventional SHARP APL file system.

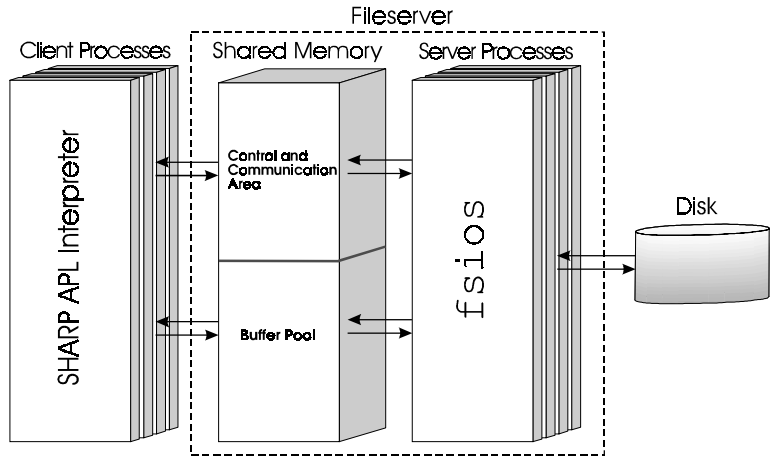

Figure 4.1. SHARP APL File Management via Fileserver.

The Fileserver concept is similar to a database management system where there are client (SHARP APL interpreter) processes and server (`fsios`) processes communicating via a shared memory segment—SHARP APL performs all file operations by posting requests to the server process. Without the Fileserver, SHARP APL reads/writes directly to the disk.

Fileserver's Features

The Fileserver is designed to improve the performance and reliability of the conventional SHARP APL file system without changing any of the functionality:

Efficiency

- Additional I/O buffering in shared memory is available.
- Tie and hold tables are no longer stored with the file - they are maintained in shared memory.
- Reads and writes can operate at the same time - `⎕fhold` operations no longer interfere with them. Write operations are asynchronous.
- Expedient disk space management is provided.

Convenience

- Fileserver can be customized via configuration file parameters then restarted without affecting the interpreter processes.

Security

- Files can be protected from non-SHARP APL access by augmented UNIX storage permissions of 0400 or even 0000 in directories with permissions of 0111.

Reliability

- Fileserver's procedure for writing to the disk maintains data consistency and reduces the possibility of a corrupted file due to system failure.

Fileserver Operations

Fileserver binaries (`fsutil` and `fsios`) and configuration file (`saxfs.cnf`) must be properly installed during SHARP APL installation (refer to the section "Installation and Configuration" for more information). The Fileserver then requires two startup operations before service requests can be directed to it:

1. The Fileserver itself must be started independently by a system administrator or at UNIX boot time using the command `fsutil -i`.
2. Once the Fileserver is running, the SHARP APL interpreter may then be connected by entering the parameter `usefilelocks=3` at the command line or in the environment variable `APLPARMS`; for example:

```
APLPARMS="usefilelocks=3";EXPORT APLPARMS
```

If `usefilelocks=3` is specified, the SHARP APL interpreter attempts to access all files via the Fileserver. It does not tie a file if it is already tied by another SHARP APL interpreter without the Fileserver. Therefore, for S-tasks to work with the Fileserver you must pass the `usefilelocks` value to the parent T-task through `APLPARMS` rather than via the command line.

The intrinsic function, `unix.usefilelocks`, may be used to determine the value of `usefilelocks` at run time. It is bound to the name *`usefilelocks`* in `1 unix.` This IF requires a right argument that must be an empty string.

Details on `fsutil` command options, connecting the SHARP APL interpreter, and Fileserver shutdown and restart operations are provided in the sections that follow.

The `fsutil` *Command*

This command accepts several options that direct it to perform different Fileserver-related operations including start, stop, and abort. These options are outlined in Table 4.1. The primary operations (options `-i`, `-s`, `-a`, `-q`, and `-c`) are explained in more detail later in this section.

Note: `fsutil -i`, `-s`, `-a`, and `-q` operations are non-interruptible. `fsutil -i`, `-s`, `-a`, and `-c` operations must be run by `root` (superuser) only. Only one option can be used per `fsutil` operation ("`y`" is the only exception).

It is possible to automatically start up and shut down the Fileserver from UNIX scripts run during UNIX system startup and shutdown. In this case, the non-interactive (`-y`) form of `fsutil` should be used; e.g.,

```
fsutil -iy, fsutil -sy
```

Table 4.1. `fsutil` ***Options.***

Option	*Description*
`-i`	Starts Fileserver (see below). Additional "y" option turns off "are you sure" question: e.g., `fsutil -iy`. (must be executed by `root`)
`-s`	Stops Fileserver (see below). Additional "y" option turns off "are you sure" question: e.g., `fsutil -sy`. (must be executed by `root`)
`-a`	Aborts Fileserver (see below). Additional "y" option turns off "are you sure" question: e.g., `fsutil -ay`. (must be executed by `root`)
`-y`	Turns off confirmation on start/stop/abort operations (i.e., `-iy`, `-sy` or `-ay`).
`-q`	Queries Fileserver state: `Online/Offline/Abort/Inconsistent` (see below).
`-c`	Clears transition semaphore (see below). (must be executed by `root`)
`-r`	Displays UNIX shared memory, semaphores and message queue requirements.
`-H`	Displays holds for the file, should be used with the Fileserver instead of `holders` utility.
`-T`	Displays ties for file, should be used with the Fileserver instead of `tiers` utility.
`-v`	Reports the current version of the Fileserver.
`-D`	Dumps (prints out to screen) current Fileserver shared memory.
`-d`	Displays various performance-related statistics.
`-m`	Allows modify certain Fileserver parameters at run time.
`-tB`	Displays buffer usage statistics.

Some `fsutil` options, such as `-T` and `-H`, require additional parameters. A full list of `fsutil` options and their parameters may be obtained by running `fsutil` without any parameters. Currently, `fsutil` options `-D`, `-d`, `-m` and `-tB` are meant to be used primarily by Soliton support staff.

Activating the Fileserver (`fsutil -i`)

The Fileserver is activated by the system administrator, or at UNIX boot time, via the command `fsutil -i`. This option reads the configuration file, creates semaphore sets and shared memory, initializes shared memory and starts the I/O server executable, `fsios`. The `fsios` processes run with an effective user ID of `root`.

`fsutil` must know the location of `fsios`. The default UNIX path and file name is `$SAXDIR/bin/fsios`; however, this can be changed by setting environment variable `SAXFSIOS`.

Shutting Down the Fileserver (`fsutil -s`)

When the Fileserver is shut down via this option (`fsutil -s`), it ensures that the following processes are performed:

- Fileserver is marked `offline` in shared memory, thus notifying the interpreter and `fsios` processes that it is being shut down.
- All current SHARP APL clients are permitted to finish file access operations, but denied further access.
- `fsios` processes are allowed to process all outstanding write requests then exit.
- IPC resources are removed.

Fileserver Abort (`fsutil -a`)

This option simply removes any Fileserver-related IPCs that still exist. It is intended to be used *only* if the Fileserver is reported to be in `Abort` or `Inconsistent` mode (see `fsutil -q`). After `fsutil -a` is applied, `fsutil -q` reports Fileserver as `Offline` and any SHARP APL processes that hang as a result of the abort will wake up signalling the event `FS not available`.

`fsutil -a` does not check if `fsios` processes are still running. There shouldn't be any; however, as a precaution, issue the command `ps -ef |grep fsios`, then use `-9` to kill any `fsios` processes that remain.

Fileserver Modes (`fsutil -q`)

When querying the state of the Fileserver using `fsutil -q`, there are four possible modes that could be reported:

`Online`	Fileserver is up and running.
`Offline`	Fileserver is not running and no Fileserver IPCs exist.
`Abort`	Fileserver is not running, Fileserver's shared memory segment exists, and the mode field in shared memory reads `Abort.`
`Inconsistent`	Fileserver is in an inconsistent state.

`Online` and `Offline` modes are *normal* conditions. `Inconsistent` and `Abort` modes are reported only under certain *abnormal* conditions, as follows:

- `Inconsistent` mode occurs when `fsutil -s` or `fsutil -a` is killed with kill `-9` or the Fileserver's IPCs have been tampered with.
- `Abort` mode occurs if the Fileserver's I/O server (`fsios`) is killed with `-9`, or a SHARP APL client process is killed with `-9` or `SIGQUIT` or otherwise dies during a critical section of Fileserver code. In such a case, an `fsios` process will detect this in a matter of seconds, switch the shared memory mode field to `abort`, then exit. Other `fsios` processes will exit once they notice the Fileserver is down; Fileserver clients will signal the event `file system not available` or hang until `fsutil -s` is issued.

Clearing the Transition Semaphore (`fsutil -c`)

Before changing or querying the Fileserver, `fsutil -i`, `-s`, `-a` and `-q` operations create a semaphore called the "transition semaphore." This semaphore is automatically removed once `fsutil` finishes its task.

However, if this semaphore is not removed for some reason (i.e., `fsutil` is killed with `-9`), the next operation will fail when it sees that a transition semaphore already exists. When this happens, the transition semaphore must be deliberately cleared, using `fsutil -c`, to allow subsequent use of `fsutil`.

Note: If two `fsutil -q` operations are run at the same time, one of them would return `Inconsistent` with an explanatory message "`FS in transition or being queried`."

Connecting the Interpreter

As mentioned on page 4-3, the SHARP APL interpreter is instructed to use the Fileserver by specifying the parameter `usefilelocks=3` at the command line or in the environment variable `APLPARMS`. Once this parameter is set, all files are accessed via the Fileserver.

An interpreter process then does the following:

- Attaches itself to shared memory (once, on the first file operation).
- Establishes a communication area (CA) in shared memory, on the first tie. Whenever the last tie is dropped by the process, it releases its CA, making it available to other processes. On the next tie it gets another CA (not necessarily the same one). SHARP APL remains connected to the Fileserver provided it holds at least one tie.
- On read, tries to find a component in the shared memory buffers. If it's not there, it posts a requests reading from the disk into shared memory by one of the I/O servers (`fsios` processes). Semaphores are used to synchronize this activity.
- On write, updates a file by copying data to shared memory and requesting an I/O server to copy it onto disk.

Note: Simultaneous access to a file from a Fileserver and non-Fileserver task is not possible and the error returned is `file tied`.

Fileserver Restart

Currently, a Fileserver restart is not transparent to SHARP APL processes. In short, the simplest way to handle this situation is to restart the Fileserver, then restart all SHARP APL applications and reload workspaces. SHARP APL processes not restarted report `file system not available` and `file tie error` events. Restart actions are described in Table 4.2.

Table 4.2. Fileserver Restart Actions.

Shutdown Event:	*Restart Actions:*

`)off`	`fsutil -s`
	`fsutil -i`
	
	Sign on to SHARP APL again.

`⎕untie ⎕nums`	`fsutil -s`
	`fsutil -i`
	
	Tie files again.

`fsutil -s`	`⎕avail ''`
	`0`
	`fsutil -i`
	`⎕avail ''`
	`1`
	Tie files again.

Note: It is a good idea to save `⎕names` and `⎕nums` values and use them in retying the files.

Installation and Configuration

The script `saxfs.install` must be run during SHARP APL installation to set proper owner, mode, and permission for the two binaries that comprise the Fileserver, `fsutil` and `fsios`.

The configuration file, `saxfs.cnf`, is read by `fsutil` and SHARP APL clients when the Fileserver is initialized. It provides the following important numeric parameters:

- The starting key for IPC resources; this is the only way SHARP APL determines which semaphore(s) and shared memory segment to use.
- Values used by `fsutil` for shared memory initialization (maximum number of connected processes, number of tied files system wide, and the number of shared memory buffers used to hold component data).

`saxfs.cnf` is read only once and changes to it do not affect Fileserver operations until it is shut down and restarted. The default UNIX path and file name is `$SAXDIR/etc/saxfs.cnf.` This can be changed by setting environment variable `SAXFSCONF`.

`fsutil` checks obvious things on initialization and reports errors. For parameters that affect shared memory size there is a shared memory structure allocated for each unit. The current estimates for these structures is as follows:

`MAXPROC`:	400 bytes
`COMPONENTS`:	40 bytes each
`BUFFERS`:	8 bytes + 256 bytes of the buffer itself
`MAXFILE`:	1400 bytes each.

The following UNIX system parameters are most likely to be affected by the Fileserver's configuration:

- maximum number of open files per process (should be at least `MAXFILE`)
- shared memory size per segment
- shared memory size total
- maximum number of semaphore sets (Fileserver currently requires 4)
- maximum number of semaphores in set
- maximum number of semaphores (system-dependent).
- maximum number of message queues (should be at least `IOSERVERS`).

`fsutil -r` reads the appropriate configuration file and prints out default or computed values for the parameters not specified in the file as well as estimated UNIX shared memory, semaphore and message queue resource requirements for this configuration.

Note: It doesn't matter whether Fileserver is running or with what configuration. The only source of information for `fsutil -r` is the configuration file, `$SAXDIR/etc/saxfs.cnf` or as defined by `$SAXFSCONF`.

Parameter List

Although all parameters can be changed by the SHARP APL site, most are assigned defaults or are computed based on other parameters' values. The default configuration distributed with SHARP APL for UNIX, should be adequate for most Fileserver requirements provided the *mandatory* parameters (`MAXPROC`, `MAXFILE`, `BUFFERS`) are correctly adjusted for your site.

Mandatory Changes

The following Fileserver configuration file parameters must be changed to values that are appropriate for your site.

`MAXPROC` Maximum number of SHARP APL processes connected to the Fileserver at any given moment. ***Default value*:** `32`.

"Connected" effectively means "holding a tie." `fsutil -i` allocates the `MAXPROC` communication areas in shared memory and if none are available at connection time, the event *`file system ties used up`* is signalled.

`MAXFILE` Maximum number of files tied by at least one SHARP APL process at any time, system wide. ***Default value:*** `64`.

Since every I/O server may have to open all of them, it is necessary to set the same value in UNIX for the maximum number of files opened by a process, at least.

`BUFFERS` Number of component buffers (256 bytes each). Component data is buffered here. ***Default value:*** `40000`.

The more buffers available, the more efficient file access will be. Therefore, this value should be as large as it can be without obstructing memory usage by other processes or leading to excessive swapping. Setting this parameter to a small value is not recommended.

Optional Changes

The following parameter defaults may be substituted, if your site prefers an alternate configuration.

`SAXUSER` Owner of SHARP APL's file directories. ***Default value:*** `-4`.

Using Fileserver, any user can create a file in any directory owned by `SAXUSER`. The default is -4 which should be a non-existent user.

`SAXGROUP` ID number of the group that all users must belong to in order to access Fileserver. ***Default value:*** `-1`.

The default, `-1`, means users do not have to belong to any particular group since Fileserver IPCs are created with UNIX permissions of `0666`. Otherwise, for security, the permissions are `0660`.

`HOLDS` Maximum number of component ranges held (`⎕hold` / `⎕fhold`) for all files by all processes at any time. ***Default value:*** *Computed*.

The computed default value is slightly greater than `MAXPROC` × *avg(holds-per-process)*. This allows for an unlikely case where each process holds a component range for every file tied by it. Site administrators may want to decrease this value according to the average number of ranges held by a process in the site's application environment.

`COMPONENTS` Maximum number of components buffered system wide at any time. ***Default value:*** *Computed*.

The computed default value is `MAXPROC` × `MAXFILE` × `4`, but no less than `1024` and no greater than `BUFFERS ÷ 4`. Remember that file control structures such as the component directories and access table are stored internally as components.

`IOSERVERS` Number of I/O servers (`fsios` processes) that `fsutil -i` starts to process disk read/write requests. ***Default value:*** `6`.

If this value is too small, SHARP APL waits too long for the completion of requests from other SHARP APL processes assigned to the same I/O server.

Each component file is assigned an I/O server to process all write requests for this file. For this reason an ideal value for IOSERVERS is the number of files used write-intensively. It is also important in an application to tie such files first, followed by the files used only for occasional writes or read-only.

Defaults Only

The following Fileserver configuration file parameters do not require site changes (unless under certain unusual circumstances).

KEYBASE — IPCs key base. ***Default value:*** 100.

Keys of shared memory segments and semaphore sets used by Fileserver are KEYBASE, KEYBASE+1, KEYBASE+2. The default should be sufficient; however, there may be two reasons to adjust the value of this parameter on site: if this value interferes with IPC keys used by other applications; or, if you wish to run several instances of the Fileserver on one machine.

SHMADDR — Shared memory base address. ***Default value:*** O*perating system dependent*.

Never change this parameter. This defines the address at which SHARP APL and Fileserver processes attach Fileserver shared memory.

Note: SHMADDR and KEYBASE are the only parameters used by apl. Other parameters are read from the configuration file by fsutil on initialization and put into a structure in shared memory. apl processes then use only shared memory values.

FILENAMES — Space for tied file names, in 256 byte blocks. ***Default value:*** *Computed.*

There must be enough space to accommodate full pathnames of all tied files plus 6 bytes of internal information per file. The computed default value is derived from MAXFILE using a formula that allows for a 64-character average file name length.

`RESIDENT` Makes shared memory segment resident in physical memory. ***Valid values:*** `1` or `0`. ***Default value:*** `1`.

`HQUEUES` Number of queues to keep unused `⎕hold`/`⎕fhold` structures. ***Default value:*** *Computed.*

Unused hold entries are stored in these queues. The computed default value is `MAXPROC` adjusted to be within the range `16` to `256`.

Technical Notes

This section explains, in more detail, Fileserver's UNIX characteristics and discusses the differences between SHARP APL's behavior *with/without* Fileserver connected.

`⎕avail`

Without Fileserver, `⎕avail` always returns `1`. *With* Fileserver, if the Fileserver goes down or there are too many processes connected, `⎕avail` returns `0` and unties all files tied by the process; once the Fileserver is reinitialized, the next call to `⎕avail` returns `1`.

APL event 28 – `file system not available`

If the Fileserver goes down, all file operations except `⎕avail` signal APL event 28, `file system not available` and untie all files.

APL event 30 – `file system ties used up`

If the Fileserver runs out of a shared memory resource or there are too many processes connected, the SHARP APL file operation signals APL event 30, `file system ties used up`.

User Privileges

The Fileserver provides for UNIX user and group level security. All SHARP APL Fileserver users must belong to the group with the numeric ID specified by SAXGROUP in the configuration file. If SAXGROUP is -1, the default, users don't have to belong to a particular group: group security is disabled.

Files can be created, erased, or renamed by a user in:

- directories to which the user has write access;
- directories owned by the user with the numeric ID specified by SAXUSER in the configuration file. The user only needs search permissions for these directories.

If SAXUSER is set to an invalid ID (such as -4, the default value), files can be created, erased, and renamed only in the directories that are writable by the SHARP APL user who initiated the session.

Asynchronously Reported Write Errors

With the Fileserver, file write operations are asynchronous: a SHARP APL process checks for file full, index errors and access violations, updates data in shared memory, queues a request to an I/O server and returns success. If the I/O server can't perform the operation due to a hardware error or a no disk space condition, it will notify the SHARP APL process which will raise an appropriate event on the next file system operation.

Examples:

```
1 □append 1
□read 1,25
34 file system no space
```

This would be impossible without the Fileserver. With the Fileserver it means □*append* failed for lack of disk space.

```
1 □append 1
2 □replace 1,25
34 file system no space
```

Without the Fileserver this means not enough space for □*replace*. With the Fileserver it means there was not enough space for □*append*.

fcreatesw

Files are created with UNIX permissions determined by the fcreatesw value; however they are not changed on □stac if the Fileserver is used.

UNIX File Access

fsios processes are run by root and all read/write operations are done by them, not by SHARP APL. *With* the Fileserver, SHARP APL can read/write only under the following circumstances:

- if the file access table allows the specified operation
- if SHARP APL has UNIX search permissions on the file directory
- if the file is readable (writable) by root.

So files can have permissions of 0000 and be stored in directories with permissions of 0111.

□create, □erase, and □rename directory update operations are allowed by SHARP APL using the Fileserver under the following circumstances:

- the directory is writable by the user who runs SHARP APL; or
- the directory's owner is SAXUSER and the user who runs SHARP APL has search access to it.

The idea here is to allow users to create files in some but not all directories where they do not have UNIX write access. For □rename, source and target directories are either writable by the user, or they are owned by SAXUSER.

File Utilities

Of the utility programs documented in "Chapter 3. File Utilities", flockers, tiers, and holders are inappropriate for use with the Fileserver.

The Fileserver administrates locks independent of the UNIX operating system file locks, therefore flockers would not provide useful information. In the traditional component file system, information about file ties (tiers) and holds (holders) is recorded in the physical file itself. Under the Fileserver, such information is not stored in the physical file. However, to maintain compatibility, the Fileserver still reads the file tie and hold information from the physical file when opening a component file.

If the Fileserver is installed use the `fsutil` command:

- `fsutil -T` *filename* determines who has a file tied.
- `fsutil -H` *filename* determines who has issued holds on a file.

For further information on `fsutil` options, refer to "The fsutil Command" on page 4-3.

SHARP APL *for UNIX*

Intrinsic Functions Manual

JUMP TO ...

CHAPTER 1. OVERVIEW
CONTENTS
PREFACE
MASTER INDEX
USING THIS DOCUMENTATION

HANDBOOK
AUXILIARY PROCESSORS MANUAL
FILE SYSTEM MANUAL
INTRINSIC FUNCTIONS MANUAL
LANGUAGE GUIDE
SVP MANUAL
SYSTEM GUIDE

© 2000 by Soliton Associates Limited
44 Victoria Street, Suite 2100, Toronto, Ontario, Canada M5C 1Y2

Pubcode: UW-007-0501 (0007)

SHARP APL *for UNIX*

DEVELOPER LIBRARY

Intrinsic Functions Manual

SOLITON ASSOCIATES

© 2000 by Soliton Associates Limited–Printed in Canada
44 Victoria Street, Suite 2100, Toronto, Ontario, Canada M5C 1Y2

All rights reserved. No part of this work may be reproduced in any form or by any means without prior written permission of the copyright owner.

Soliton Associates Limited shall not be liable for errors contained herein or for incidental consequential damages in connection with the use of this material. The information contained within this document is subject to change without notice.

Acknowledgments

SHARP APL is a registered trademark of Soliton Associates Limited.

SVP and NSVP are trademarks of Soliton Associates Limited.

IBM and AIX are registered trademarks of International Business Machines Corporation.

RISC System/6000 is a trademark of International Business Machines Corporation.

UNIX is a registered trademark of UNIX System Laboratories, Inc.

Sun Workstation is a registered trademark of Sun Microsystems, Inc.

Sun, Solaris, SunOS, Open Windows, and SunView are trademarks of Sun Microsystems, Inc.

All other products referred to in this document are trademarks or registered trademarks of their respective trademark holders.

Publication code: UW-007-0501 (0009)

Contents

Preface v
Introduction v
Chapter Outlines v
Conventions vi
Documentation Summary vi
Contacting Soliton Associates vii

1. Overview
What is an Intrinsic Function? 1-1
Supplied System Intrinsic Functions 1-2
Supplied User-Defined Intrinsic Functions 1-3
IF System Commands 1-3

2. User-Defined Intrinsic Functions
Overview 2-1
Specifications 2-3
The `.saxif` file 2-3
Argument Mapping 2-5
Result 2-7
Building Intrinsic Functions 2-8
Using `ifgen` and `ifcc` 2-8
Restrictions 2-10
Limitations 2-10
Programming Notes 2-11
Side Effects 2-11

Debugging 2-12
association error 2-13
Intrinsic Functions Within Packages 2-13

3. System Intrinsic Functions
sapl.retcode 3-1
sapl.singleref 3-2
sapl.translate 3-3
sapl.type 3-4
sapl.version 3-5
system.bind 3-6
system.bindrep 3-8
system.def 3-9
system.erase 3-11
system.list 3-12
unix.errno 3-12
unix.filesync 3-13

4. IF System Commands
Introduction 4-1
)binds Display Bound Intrinsic Functions 4-1
)ifdef Display IF Definition/Redefine IF 4-2
)iferase Remove IF Definitions 4-3
)iflist List IF Groups and Their Contents 4-4

5. Supplied IF Facilities
SHARP APL's UNIX System Function Facility 5-1
Miscellaneous Intrinsics 5-1
Message Queue Intrinsics 5-2
Semaphore Intrinsics 5-4
Shared Memory Intrinsics 5-7
Related APL Functions. 5-10
SHARP APL's Socket Facility 5-11

Index ix-1

Preface

Introduction

An intrinsic function (IF) is a function defined and stored outside of SHARP APL that can be associated with (or *bound* to) a name in the APL workspace. This volume explains how to build and run user-defined intrinsic functions, documents the set of supplied intrinsic functions, and explains the system commands used to administer IFs during an APL session.

Rather than reproduce existing material, references to other SHARP APL for UNIX publications are supplied where applicable.

Chapter Outlines

This document is organized into the chapters described below.

Chapter 1, "Overview," introduces the concept of an intrinsic function, discusses the two classes of intrinsic functions, and lists the types of intrinsic function available with SHARP APL for UNIX.

Chapter 2, "User-Defined Intrinsic Functions," explains in detail how to build and run user-defined intrinsic functions.

Chapter 3, "System Intrinsic Functions," describes the system intrinsic functions that are supplied with SHARP APL.

Chapter 4, "IF System Commands," describes the system commands that allow users to manipulate available IFs after the APL session has started.

Chapter 5, "Supplied IF Facilities," lists the sets of user-defined intrinsic functions that are distributed with SHARP APL for UNIX.

Conventions

The following conventions are used throughout this documentation:

`⎕io←0`	Although the default value for `⎕io` in a clear workspace is *one*, all examples in this manual assume that the index origin is *zero*.
α and ω	Two APL symbols are used in syntax descriptions of monadic and dyadic functions; *alpha* (α) identifies the left argument and *omega* (ω) identifies the right argument.
`constant width`	Examples of non-APL input and output are represented in a `constant width` typeface.
default (%) prompt	Examples in this book assume you are using the C shell. The default prompt for the C shell is the *percent sign* (%).
passno	Argument variables appear in *sans-serif italic* type.
[ ]	Optional arguments are shown between square brackets. For example, in the following APL statement, the argument *passno* is optional:

'filename' ⎕stie tieno [,passno]

These default environment variables represent frequently used pathnames in SHARP APL for UNIX documentation and scripts:

`$SAXLIB`	→	`/usr/sax`
`$SAXDIR`	→	`/usr/sax/rel`
`$SAXCNF`	→	`/usr/sax/local`
`$HOME`	→	home directory of the current user.

Documentation Summary

The SHARP APL for UNIX reference documentation is published in several volumes. One or more of the following may be referenced in this manual.

SHARP APL for UNIX,

- *Handbook,* publication code UW-000-0401

- *Language Guide,* publication code UW-000-0802
- *System Guide,* publication code UW-000-0902
- *SVP Manual,* publication code UW-001-0501
- *File System Manual,* publication code UW-037-0501
- *Auxiliary Processors Manual,* publication code UW-033-0501

For a complete list of SHARP APL publications, please refer to the ***Support*** link on the Soliton Associates website: `http://www.soliton.com.`

Contacting Soliton Associates

Problems or questions regarding your SHARP APL system or this guide should be directed to the Soliton Associates Technical Support group:

`support@soliton.com`

Requests for additional technical documentation should be directed to the Soliton Associates Sales group:

`sales@soliton.com`

1 Overview

What is an Intrinsic Function?

An intrinsic function (IF) is a function that is defined and stored outside the APL workspace. It can be associated with (or *bound* to) a name in the APL workspace using the system function ⎕*bind*, and that name can be called like any other APL function.

There are two classes of intrinsic functions: those that are delivered as part of the APL interpreter and, like APL primitives, cannot be modified by the user; and, those that are ***user-defined***, created for specific user tasks, by Soliton, or by the users themselves.

For example, *sapl.version*, a ***system*** intrinsic functions that is defined in the APL interpreter, can be bound to a name such as *foo* in the APL workspace. The newly defined object *foo* is called like any APL primitive function.

User-defined intrinsic functions, however, reside outside both the workspace and the interpreter. These intrinsic functions give the APL programmer access to system calls, C library routines, and user-written C functions. The executable that contains the functions (or references to the functions) is built as a separate file from the interpreter, and is dynamically loaded when the function is first referenced. Using the intrinsic function facility, an APL application can have direct access to services offered in the operating system and to custom-written routines coded in C. See the file `$SAXDIR/xmpls/if` for an example of a user-defined intrinsic function group.

Workspace `1 if`

The workspace `1 if` contains the latest online documentation of the IF facility, and of the intrinsic functions included as part of the base SHARP APL for UNIX

system. It also contains names that are bound to the functions *system.bind* (*IFbind*), *system.bindrep* (*IFrep*), *system.list* (*IFlist*), *system.def* (*IFdef*), and *system.erase* (*IFerase*).

Supplied System Intrinsic Functions

The following system intrinsic functions are supplied with SHARP APL for UNIX. These are fully described in "Chapter 3. System Intrinsic Functions".

sapl.retcode	Allows the APL interpreter's return (or exit) code to be set or queried.
sapl.singleref	Isolates APL data so that it can be safely modified by an intrinsic function.
sapl.translate	Translates data using a specified translate table.
sapl.type	Returns the data type of the specified variable.
sapl.version	Returns information about the current version of the APL interpreter.
system.list	Returns a list of the IF groups that are defined for the current APL session, together with their members.
system.def	Returns definitions of IF groups defined for the current APL session, together with the definitions of their members; or to establish the definitions of IF groups and their members.
system.erase	Removes IF groups defined for the current APL session, or members of defined IF groups.
system.bind	Creates an IF bind or queries association level (identical to the functionality provided by the system function □*bind*.)
system.bindrep	Creates a display form of an IF bind.
unix.filesync	Allows user to control the APL interpreter's use of `fsync()` and `sync()`.
unix.errno	Returns the current value of *errno*.

Supplied User-Defined Intrinsic Functions

As well, some user-defined intrinsic functions are distributed with SHARP APL:

- The SHARP APL *UNIX* System Function Facility is a set of intrinsic functions that correspond to some of the standard UNIX system functions.
- The SHARP APL Socket Facility is a set of intrinsic functions that correspond to the Berkeley socket facility and are called by cover functions defined in an APL workspace.

The intrinsic functions used in these facilities are listed in "Chapter 5. Supplied IF Facilities".

IF System Commands

The following system commands allow users to manipulate the sets of intrinsic functions available after their APL session has started, rather than being required to start the APL session in a suitably configured way. These are described in "Chapter 4. IF System Commands".

`)binds`	Displays a list of bound intrinsic functions in the workspace.
`)iflist`	Lists the IF groups that are defined for the current APL session, together with their members.
`)ifdef`	Displays the definitions of IF groups defined for the current APL session, together with the definitions of their members, or establishes definitions of IF groups and their members.
`)iferase`	Removes IF groups defined for the current APL session, or members of defined IF groups.

Refer to "Chapter 3. System Intrinsic Functions" for a list of the system IFs associated with these commands.

2
User-Defined Intrinsic Functions

Overview

When started, the interpreter reads configuration files that describe all the user-defined intrinsic functions it will be able to access. The configuration describes where the files that contain the intrinsic functions reside, as well as how the interpreter calls each function.

The file that describes intrinsic functions is called `.saxif`. This file may reside in a number of different directories, depending on whether the directories describe functions distributed as part of the base APL system, whether they provide local services, or whether they are specific to a single user. For example, both the `$SAXCNF` directory and the user's `$HOME` directory are examined for `.saxif` files.

The definition of a particular function in the `.saxif` file resembles an ANSI C prototype. For example, the `strcmp()` function could be described with the following line in `.saxif`:

```
strcmp  Cfunction       int strcmp(char *, char *)
```

Specifically, that the `strcmp()` function takes two arguments (both pointers to characters) and returns an integer as a result.

When associated with a name in a workspace (for example, *cmp*), the C function can be called by calling *cmp* with an enclosed array right argument, with each element corresponding to one of the arguments to the C function; for example:

```
      cmp 'compare this string' ⊃ 'with this one'
¯1
```

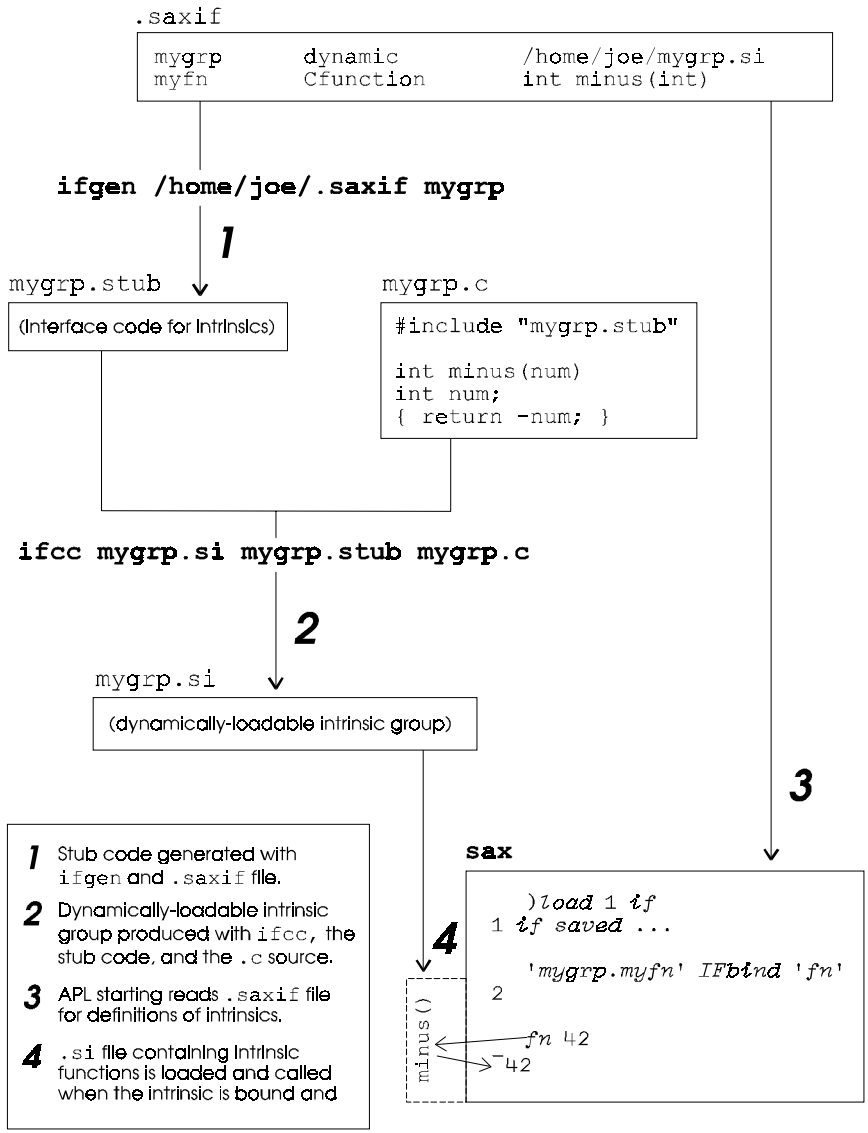

Figure 2.1. Generating and running a user-defined intrinsic function.

The result of the C function is returned as the result of the APL intrinsic. In addition, C functions that modify objects to which they are given pointers cause those modifications to be reflected in the workspace as side-effects of calling the APL intrinsic.

A set of related intrinsic functions are combined in a group, and are built and loaded together. They are executable files built with the linkage editor and having a nonstandard entry point. By convention, these files have a `.si` suffix.

A group of intrinsics is built using two utilities provided with SHARP APL: `ifgen`, which reads a `.saxif` file for function definitions and produces a `.stub` file containing interface code; and `ifcc`, which compiles and links a group of intrinsic functions from their `.c` source.

Figure 2.1 illustrates the steps involved in producing an intrinsic function which simply returns the negative of its integer argument.

Specifications

The `.saxif` file

The `.saxif` file contains the descriptions for intrinsic functions, including the location and contents of the `.si` files that contain groups of intrinsic functions. Two kinds of records exist in the `.saxif` file, `dynamic` and `Cfunction`. A `dynamic` record names an intrinsic group and the `.si` file. The `dynamic` record is then followed by any number of `Cfunction` records, each of which describes the name of the intrinsic as known by APL, and the prototype of the C function that will be called when the intrinsic is invoked. Finally, lines whose first non-blank character is a hash mark (#) are treated as comment lines. Comment lines and blank lines are ignored.

A `dynamic` record has the following form:

groupname `dynamic` *path*

groupname is the name of the group as it will be known to APL. This corresponds to the first segment in an intrinsic function name. *path* is the full pathname of the intrinsic group file name.

A `Cfunction` record has the following form:

name `Cfunction` *prototype*

name is the second segment of the intrinsic function name as it is known to APL. *prototype* contains the C function name and a complete description of the arguments the function takes and the result it returns.

This prototype is based on ANSI C prototypes. For example,

```
adder   Cfunction       int adder(int, int *)
```

describes a function which takes two arguments, one an integer, the other a pointer to an integer. The result of the function is also an integer.

The prototype part of a `Cfunction` record differs from ANSI C in the following respects:

- Data types are single words. No storage class or other modifiers are permitted. The following data types are recognized:

```
char   short   int     long
uchar  ushort  uint    ulong   (u...=unsigned ...)
void   float   double  ldouble   (ldouble=long double)
```

- Structures are represented by their members separated by commas and enclosed in parentheses. For example, the C function

```
int foo(arg)
struct test { char *n;   int i; } *arg;
{. . .}
```

  is represented in the `.saxif` prototype as

```
foo     Cfunction    int foo( {char *, int}* )
```

Consider the following example `.saxif` file:

```
#
#  Definition for example group of intrinsic functions
#

example    dynamic    $SAXDIR/lib/example.si
version    Cfunction  char *version(int)
noop       Cfunction  int noop()
```

```
adder      Cfunction  int adder(int, int *)
uid2name   Cfunction  char *uid2name(int)
name2uid   Cfunction  int name2uid(char *)
gettimeofday Cfunction int gettimeofday({int,int}*, {int,int}*)
strcmp      Cfunction   int strcmp(char *, char *)
```

The first non-comment line defines the group `example`, and aims the interpreter at the file `$SAXDIR/lib/example.si` for the dynamically loadable code. The subsequent lines describe the functions that can be found in this file, and how they are to be called. This example file allows an APL task to create binds to intrinsic functions named *`example.version`*, *`example.noop`*, *`example.adder`*, etc.

The definitions of the functions continue until another group is found (by another `dynamic` statement), or until end-of-file.

When started, the APL interpreter will search three locations for `.saxif` files:

```
$HOME/.saxif
$SAXCNF/.saxif              (default SAXCNF=/usr/sax/local)
$SAXDIR/etc/.saxif          (default SAXDIR=/usr/sax/rel)
```

The first is user-specific intrinsics, the second is for site-local intrinsics, and the third is intended to describe loadable intrinsic groups that have been supplied by Soliton.

Argument Mapping

There are a few rules that govern how an APL array right argument is mapped onto the arguments of the C function being called. The interpreter is responsible for presenting proper arguments to the C function from the argument the APL intrinsic has been given. If APL is unable to map the argument onto the function, a *`domain error`* is reported. The rules governing how the argument is mapped follow:

- APL intrinsic functions that are bound to C functions are always called monadically, with an array right argument containing enclosed elements. Each element of the right argument is mapped onto a separate argument of the C function. APL integers, floating point numbers, and characters map directly onto their C counterparts.

- When an argument to a C function contains a pointer to an object, it is represented in APL as either an array of the appropriate data type, or as a single element of that type enclosed.
- In addition to an array or an enclosed element mapping onto a pointer, a simple scalar zero (0) may be used to pass a `NULL` pointer to the function.
- C structures are represented by a vector of enclosed arguments, each one mapping onto its corresponding field in the `struct`.
- Finally, if the APL intrinsic function is called with a simple, nonempty right argument *ω*, it is treated as *,<ω*. This permits C functions with a single argument to be called without having to explicitly build an enclosed right argument to the intrinsic.

Examples:

`foo(int, char *)`

Can be called by:
foo `(<42), <'Array of characters'`
or by: *foo* `(<42), <<'A'`

In both cases the function foo receives the integer `42` and a pointer to the '*A*' in the right argument.

`goo({char,double}*)` Can be called by: *goo* `,<< 'x' ⊃ 3.14`

Note that `'x' ⊃ 3.14` forms the array to be mapped onto the structure. It is enclosed once to make it a pointer (because the function takes a pointer to such a structure). It is enclosed a second time and the entire argument ravelled to make the right argument an enclosed array with one element for the one argument of `goo`.

`hoo(int *)`

Can be called by:
hoo `,<3 4 5` (or even *hoo* `3 4 5`)
hoo `,<<3`
hoo `,<0` (or even *hoo* `0`)

In the first two cases, the C function is given a pointer to the `3` in the right argument. In the last case a `NULL` pointer is passed.

`name2uid(char *)`	Can be called by: *`name2uid ,<'joe'`* or simply by: *`name2uid 'joe'`*

When considering how arguments will be mapped onto C functions, it is important to understand how the APL interpreter stores various data types internally. Knowing if and when data can be coerced is necessary when calling intrinsic functions, especially if the C function is going to alter data to which it has been given a pointer.

Internally, the APL interpreter stores characters as 1-byte sequences, with no terminating null byte. Booleans are stored as individual bits, APL integers are stored as `long`s, and APL floating point numbers are stored as `double`s.

If you call a function with an APL argument whose representation differs from that expected by the C function, the interpreter will attempt to coerce the argument into the expected form. If it cannot, a *`domain error`* will result. If it can, the coerced value is presented to the C function being called. Once execution of the C function completes, the (possibly altered) value is again coerced back into the original argument.

A specific and common case of the above occurs when a C function must alter a `short`, `int`, or `long` data type. The argument passed to the intrinsic function should ensure that the APL data types of those pieces of the argument are APL integers, and not APL Booleans. Thus, when assigning a variable that will be passed to a C function as a `short`, `int`, or `long` data type and altered as a side-effect, you should create the variable using a constant such as `999`. This will ensure that the variable is stored internally as an APL integer. In contrast, the values `0` and `1`, for example, may be stored as Boolean and therefore should not be used.

Result

The result of the C function being called is returned as a result of the APL intrinsic function. This must be a simple data type, such as `int`, which is returned as integer; or `double`, which is returned as floating point; or a pointer to a `char`, which is returned as a character vector. Other types of results, such as pointers to structures, cause a *`result error`* to be signalled.

Building Intrinsic Functions

Files that contain intrinsic functions are built using two utilities, `ifgen` and `ifcc`. The first, `ifgen`, reads a `.saxif` file and builds stub code for a specific group. The second, `ifcc`, is a cover for the native `cc` and `ld` commands, and is used to compile and link the functions into a dynamically loadable intrinsic group.

The intrinsic function itself consists of two pieces: the stub code; and the "target," which is the code you wish to execute.

The stub code built with `ifgen` is C source that provides the interface between the interpreter and the target. For a given target, the code generated in the `.stub` file contains an entry point with the same name as the target but with a `SAXIF_` prefix. The `SAXIF_xxx` function, which is called by the interpreter, references the target. This design allows the compiler, linker, and loader to resolve these references appropriately. The actual code generated may be very different depending on whether the target is a system call, a reference to a function in a shared library, or to static code. In any case, the interpreter simply has to call the `SAXIF_xxx` entry point and let the compiled code there call the target.

Once the stub code is built, C source must be written for the intrinsics. This source must `#include` the stub code, and must further ensure that all required prototypes for the functions referenced in the stub code are included. Finally, the C source must include the actual code for any intrinsic the user will write himself.

The `ifcc` script is then run to produce a `.si` file, which can be loaded by the interpreter. The `.stub` file and all `.c` files that contain the source must be included as arguments. This script compiles all the source and links the final `.si` file.

Using `ifgen` *and* `ifcc`

The use of the two utilities `ifgen` and `ifcc` is described below.

`ifgen` *.saxif_file grpname*

- *saxif_file* names the file containing the group and function definitions for which stub code is to be generated.
- *grpname* is the specific name of the group within the `.saxif` file.

The utility will generate a *grpname*.stub file, which must be included when linking the .si file with ifcc.

ifcc [-g] *grpname*.si *grpname*.stub *source1*.c [*source2*.c ...]

- *grpname*.si is the name of the output file.
- *grpname*.stub is the name of the stub code.
- *source1*.c is the name (or names) of the .c source for the intrinsic.
- -g indicates that symbolic debugging information should be included in the resulting file. (See the description of the -g argument to cc.)

Example:

An example of building an intrinsic that gives access to three functions follows. The functions are strcmp (standard C library routine), kill (system call), and minus (user-written). This illustration will be shown as if the intrinsics were being written by and for a user named joe.

First, write the definitions of the group and functions in the .saxif file.

.saxif

```
joestuff  dynamic     /home/joe/joestuff.si
strcmp    Cfunction   int strcmp(char *, char *)
kill      Cfunction   int kill(int, int)
minus     Cfunction   int minus(int)
```

Generate the stub code:

```
ifgen /home/joe/.saxif joestuff
```

(produces joestuff.stub file)

Next, write the C source:

joestuff.c

```
/* Source for joestuff group of intrinsics */
#include <string.h>                    for strcmp()

int minus(num)                         user-written
int num;
{ return -num; }
```

```
#include "joestuff.stub"                    stub code
```

Now compile and link the dynamically loadable intrinsic group:

```
ifcc joestuff.si joestuff.stub joestuff.c
```

Finally, run APL and run an intrinsic:

```
'joestuff.strcmp' ⎕bind 'compare'
2
     compare 'this string' ⊃ 'with this one'
¯1
     'joestuff.minus' ⎕bind 'minus'
2
     minus 6
¯6
```

Restrictions

The intrinsic function facility allows an APL program to pass native C data types, structures containing native C data types, and pointers to these. There is no facility to permit it to recognize bit fields, unions, storage classes in declarations, or pointers to functions.

When using this facility it is important to realize that the target code will be running in the same process as the rest of the interpreter, and that it is entirely possible for the actions of an intrinsic function to affect the operation of the interpreter. Specifically, care must be taken to avoid closing files that the interpreter is using, altering signal handlers (including `SIGUSR1`, which is used by the shared variable processor), forking, freeing memory, or altering any of the region "owned" by the interpreter. Also, despite running in the interpreter environment, user-defined intrinsic functions are not interruptable as are other parts of the interpreter.

Limitations

Some areas of the intrinsic function facility are defined but are not yet available. These areas are listed below.

- APL Boolean arrays are not mapped onto `int`, `short`, or `long` pointers. A Boolean scalar can, however, be mapped onto an `int`, `short`, or `long`.
- An APL Boolean that is mapped onto a `short`, `int`, or `long` is not updated if the argument is altered by the C function.
- The result type of the C function must be a simple native C type, or a character pointer. Pointers to structures are not yet permitted.
- The following data types are not yet supported:

```
void
float
ldouble
```

(Note that an APL floating point is stored internally as a C `double`.)

Programming Notes

Side Effects

The APL interpreter maintains one copy of data with multiple references to it where possible. For example,

```
a ← 'string'
b ← ,a
c ← 42 ⊃ b
```

would be stored as

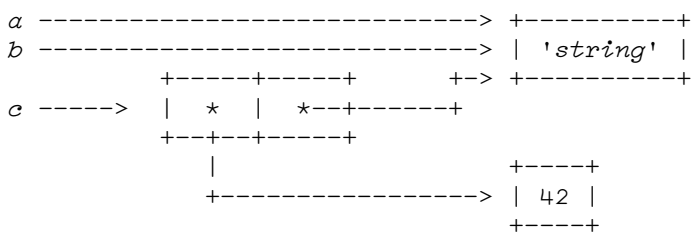

In the above example, the data for both *a* and *b* and for part of *c* would be shared. If any of these variables were used in an intrinsic that altered its argument, the changes would be reflected in all three variables.

It is the responsibility of the user to ensure that, if one of the arguments is to be altered as a side-effect of a function call, and it is not clear whether there are multiple references to that data, that a copy of the data with a single reference is passed to the intrinsic.

A built-in intrinsic function, *sapl.singleref*, is supplied to provide this ability.

If you were to provide the variable *c* above as an argument to an intrinsic and wanted to ensure that the changes made to the data were only visible through the name *c*, the following would accomplish that:

```
     'sapl.singleref' ⎕bind 'sref'
     d ← sref c
     ifun d
     d
99 new
     c
42 string
```

The APL interpreter will make multiple references to a single piece of data in several cases. Functions that can wind up trivially returning one of their arguments-for example, ravel, lev, and dex-often simply make multiple references rather than copying data.

Another common place where the interpreter will make multiple references to the same data is when constant scalar values between `¯1` and `99` are used. Internally the interpreter generates multiple references to fixed areas of data in the workspace to avoid wasting a lot of space on very common constants. You should use *singleref* if there is a chance that some or all of the data being provided in the right argument of an intrinsic is one of these constants. Any attempt to pass an argument with a reference to one of these interpreter constants to a function which could alter them will result in a *domain error*.

Debugging

The APL interpreter itself is stripped of symbol table information, but the `-g` parameter can be used on the `ifcc` utility to generate this for the dynamically loadable intrinsic group. It is recommended that a trivial function be added to each user-written group which simply returns a zero. This no-op function can be used to ensure that the group is loaded before running any of the code in the

intrinsic group that should be run under a debugger. Once this function has run successfully, a `dbx` process can be attached to the APL process to permit symbolic debugging of the loaded intrinsics.

association error

There are many conditions which can give rise to an *association error* when attempting to call an intrinsic. The following should be checked if you are getting this error:

- Full name of intrinisic: The full name of the intrinsic (which can be determined by using *IFrep* on the APL intrinsic function) must either match a built-in intrinsic (which are listed in the workspace `1 if`), or be an intrinsic named in any of the `.saxif` files that are read when APL is started. The segments of the full name (**group**.**function**) must match the group name in a `dynamic` statement and the function name in a `Cfunction` statement which follows. If a name match is found, the result of an IF bind of that name will be `2`.
- Dynamic group is not loadable: This can be caused if the pathname on the `dynamic` statement for this group does not name a dynamically loadable file. The `ifcc` script should be used to build this file.
- Function not in group: When running `ifcc` to compile and link a dynamic group, the output lists each function name (prefixed with `SAXIF_`) that is in the group. This should include all functions named in the prototype part of the `Cfunction` statements.

Intrinsic Functions Within Packages

Bound intrinsic functions can be included in an APL package. For example:

```
      ⎕nc 'func1'          ⍝ 'func1' is a bound IF
4
      p←⎕pack 'func1 a b c' ⍝ include 'func1' in a package
      ⍎⎕ex 'func1'          ⍝ erase 'func1'
      ⎕nc 'func1'          ⍝ it has gone
0
      ⎕pdef p
      ⎕nc 'func1'          ⍝ it has returned
4
```

Because of this feature, it is possible to store IFs in APL files, or transfer them via shared variables to other processes (including processes on other machines). However, whenever the IF is brought into an active workspace, it must have been defined (using *IFdef*, *)ifdef*, or by means of the SHARP APL configuration documents) in order to be executable.

3
System Intrinsic Functions

The intrinsic functions documented in this chapter are supplied as part of the base SHARP APL system. As is the case with all intrinsic functions, these must first be bound to a name in the active workspace using the system function ⎕*bind* . For example, the following expression binds the *unix.filesync* intrinsic function to the name *fs*. The result of 2 indicates that the bind was successful.

```
      'unix.filesync' ⎕bind 'fs'
2
```

sapl.retcode

The *sapl.retcode* intrinsic function enables the APL interpreter's return (or exit) code to be set and/or queried. When used dyadically, *sapl.retcode* also optionally causes APL to terminate immediately. (See "Side Effects," below.) As is the case with all intrinsic functions, *sapl.retcode* must first be bound to a name in the active workspace before it can be used.

Right Argument: The right argument to *sapl.retcode* must be either an integer-valued singleton vector or an empty vector. If the right argument is nonempty, the single element must be integer-valued, nonnegative, and less than or equal to 255.

Left Argument: The left argument of *sapl.retcode*, if present, must be a Boolean-valued singleton vector. If the left argument is elided, it is treated as zero.

Result: The result of *sapl.retcode* is an integer scalar whose value is the APL interpreter's return (or exit) code as it was before *sapl.retcode* was invoked.

Side Effects: If a nonempty right argument is supplied, the APL interpreter's return (or exit) code is set to the value of the right argument.

If the value of the left argument is `1`, the APL interpreter is terminated immediately, using the return (or exit) code supplied in the right argument, or the most recently supplied return (or exit) code if the right argument is empty. If the left argument is elided or has a value of `0`, execution continues normally. Note that when a left argument of `1` is supplied, APL terminates shortly thereafter, so that the line of code that invoked *`sapl.retcode`* may or may not complete execution before APL terminates. The timing of APL termination is somewhat unpredictable, though it will occur before the next line begins execution.

The APL interpreter's return (or exit) code can be tested if APL is invoked from a shell script, using the $? keyword within the script.

Example:

```
      'sapl.retcode' ⎕bind 'rc'
2
      rc ⍱0                 ⍝ Query return code.
0
      rc 4                  ⍝ Set return code to 4.
0
      0 rc ⍱0               ⍝ Query return code dyadically.
4
      1 rc ⍱0               ⍝ Terminate with current code.
interrupt
      1 rc 99               ⍝ Terminate with code 99.
interrupt
```

sapl.singleref

The *`sapl.singleref`* intrinsic function is a monadic function that returns its argument unchanged as far as one can tell with APL primitives; that is,

```
      ⍵ ≡ singleref ⍵
```

The result of *singleref*, however, is a separate copy of its right argument. Other than the result of this function, there is no other place in the workspace which references this data. This is useful when writing applications which call user-defined intrinsic functions, which may alter their arguments as a side-effect to their execution.

The result of *singleref* may be safely passed to an intrinsic function which will alter its arguments without any exposure that the data itself may be referenced from multiple places in the workspace.

sapl.translate

The *sapl.translate* intrinsic function is used to translate a character vector using a specified translate table. *sapl.translate* can be used either monadically or dyadically.

Right Argument: The right argument is an arbitrary-rank simple character vector, or a simple numeric vector all of whose elements are in the inclusive range [0–255].

Left Argument: The left argument, if provided, is a 256-byte character vector specifying the translation to be performed. If this argument is elided, ⎕*av* will be assumed.

Result: The result of *sapl.translate* always has the rank and shape of the right argument, and is always character. If a translation table is provided, one level of translation is performed.

Example:

```
      'sapl.translate' ⎕bind 'xlt'
2
      xlt 86 87 88          ⍝ Simple index into ⎕av
abc
      (5⍬⎕av) xlt 'abc'     ⍝ Translate them this time
fgh
```

sapl.type

The *sapl.type* intrinsic function returns information about the objects in the right argument. *sapl.type* can be used either monadically or dyadically.

Right Argument: The right argument is any APL value (e.g., simple, enclosed, package). Note that it is *not* a name.

Left Argument: The left argument, if provided, is the length of the result to be returned. The default is 3, which is the maximum length. (There is no execution penalty for longer results.) Note, however, that the maximum length may change; so if you plan to use an expression that relies upon the length of the result, specify an explicit length.

Result and Side Effects: The result of *sapl.type* is always a (possibly empty) numeric vector, the elements of which are described in Table 1.1. The element number ⎕*io*+2, "storage class," may be useful in some applications, but should be used as sparingly as possible. Storage class is assigned at the whim of the system, and the storage class of an object may change in unexpected situations. As well, any code dependent upon storage class is highly non-portable.

Table 1.1. Result of *sapl.type*.

Element 0 – Structure		*Element 1 – Type*		*Element 2 – Storage Class*	
¯1	package	¯1	package	¯1	package
0	?	0	?	0	?
1	simple	1	numeric	1	Boolean – 1 bit/element
2	enclosed	2	character	2	integer – 4 bytes/element
¯1	package	3	enclosed	3	floating – 8 bytes/element
				4	character – 1 byte/element
				5	(unused)
				6	enclosed
				7	complex – 16 bytes/element

Example:

```
'sapl.type' IFbind 'type'
2
      type 'abc'           ⍝ Check character
1 2 4                      ⍝ simple, character, character
      type ⎕pack ''        ⍝ Check package
¯1 ¯1 ¯1                   ⍝ package, package, package
      2 type <\3           ⍝ Check characters
2 3                        ⍝ enclosed, enclosed
```

sapl.version

The *sapl.version* intrinsic function returns information regarding the current version of the APL interpreter. As is the case with all intrinsic functions, *sapl.version* must first be bound to a name in the active workspace before it can be used.

Right Argument: The right argument of *sapl.version* must be an empty vector.

Left Argument: No left argument is permitted.

Result: The result of *sapl.version* is a 4-element enclosed array whose elements are as follows:

- 0@result: A character vector whose value is the name of the executable file containing the currently running APL interpreter.
- 1@result: A floating-point scalar whose value is the version number of the currently running APL interpreter.
- 2@result: An integer vector whose value is the ⎕*ts*-style timestamp of the release date of the currently running APL interpreter.
- 3@result: A character vector whose value is the name of the hardware platform upon which the APL interpreter is currently running.

Example:

```
      'sapl.version' ⎕bind 'ver'
2
      ⎕ps←¯1 1 3 3            ⍝ For pretty display.

      ver ''                  ⍝ Display version information.

------------------------ --------- ---------------------- ---------
| /usr/sax/rel/bin/apl | | 5.0.0 | | 1999 3 14 0 0 0 0 | | SPARC |
------------------------ --------- ---------------------- ---------
```

system.bind

The *system.bind* intrinsic function is used to create an IF bind or to query association level. This system intrinsic function can be used either monadically or dyadically. *system.bind* is bound to *IFbind* in the workspace *1 if*. The same functionality is provided by the system function ⎕*bind* .

Right Argument: The right argument is a character vector or scalar containing the name of a single APL user identifier (reserved names are not permitted). The dyadic use of *system.bindrep* requires that the name be locally undefined, and creates an IF bind from the bind display. The monadic use is a query only.

Left Argument: The left argument, if provided, is a character vector or martix containing a valid IF bind display (see below). Note that only line 0 (*group.function*) of the display form is required by *system.bindrep*; that is, the vector '*system.bindrep*' is a valid left argument to *system.bind*.

Result and Side Effects: The result of both monadic and dyadic *system.bind* is a single integer reporting the association state of the named symbol. The association state describes the usability of an IF bind, as described below:

- `0` Monadic use only. The symbol is not an IF bind.
- `1` The symbol is a valid IF bind, but it refers to an IF that is not present in this system.
- `2` The symbol is a valid IF bind to a known and available IF.

Association state `1` is provided so that a system can construct and process IF binds for or from other systems with a different set of available IF routines. An attempt to use an IF bind whose association state is not `2` will result in an *`association error`* being signalled (event number `10`).

Examples:

```
      'system.bindrep' IFbind 'brep'    ⍝ Form IF bind
2                       ⍝ System reports association level 2

      ⎕bind 'brep'      ⍝ Check association level
2                       ⍝ Still 2
      x←brep 'IFbind'   ⍝ Get display form of IFbind

      x ⎕bind 'bind'    ⍝ Try another bind to same IF
2                       ⍝ It worked
```

IF Bind Display Representation

An IF bind may be a CR-delimited character vector, or a character matrix. In either case it is regarded as a series of lines with the following interpretation:

Line `0` — Line `0` is always and only the **group.function** name of the intrinsic function to which the bind refers. This routine may or may not exist on the system doing the bind. (See the discussion of association state above.)

Line `1` — This line identifies the bind type (only IF bind is currently possible) and reports the time that the bind was created. The double `⍝` symbols in this line distinguish this line from user commentary. (See example below.)

Lines `2` **–** ***n*** — Any bind may include supplied user commentary text. This text is maintained as part of the bind data structure and will always be returned as part of the display form by *`IFrep`* (*`system.bindrep`*). Note that user commentary begins with a single `⍝` symbol. This symbol is required; lines beginning with two such symbols will not be included as part of the user commentary.

Note that the definition of lines 1 and 2–***n*** result in the preservation of user commentary intact across display/rebind operations.

Example:

```
     IFrep 'example'
joestuff.eg
⍝ if created 1994-03-26 04:08:22
⍝ This is user commentary associated
⍝ with this bind
```

system.bindrep

The *system.bindrep* intrinsic function is used to create a display form of an IF bind. *system.bindrep* can be used either monadically or dyadically. This function is bound to *IFrep* in the workspace 1 *if*.

Right Argument: The right argument is a character vector or scalar containing the name of an IF bind.

Left Argument: The left argument, if provided, is a numeric singleton with a value of 1 or 2. If 1 is supplied, the result will be a CR-delimited character vector. If 2 is supplied, the result will be a character matrix. If the left argument is omitted, 1 is assumed.

Result: The result of *system.bindrep* is a character display representation of the indicated IF bind. (See *system.bind*, above, for a description of the display form). This is guaranteed to be a valid input to *system.bind*.

Example:

```
     IFrep 'IFrep'          ⍝ Display IFrep bind.
system.bindrep
⍝⍝ ifn bind, created 1991-03-11 22:08:24

     ⍴IFrep 'IFrep'        ⍝ Check default shape
56
     ⍴2 IFrep 'IFrep'      ⍝ Check matrix shape
2 40
```

system.def

The *system.def* intrinsic function is used to manipulate the definitions of both a group of functions and the individual functions in the group. In its monadic form it is used to obtain the current definition of a group or function within a group; in its dyadic form it is used to establish the definition of a group or function. This function can be used to establish the definitions of intrinsic functions from within an APL application, instead of relying on the startup processing associated with an APL session. Thus, an application that needs certain intrinsic functions can define them when needed. This system intrinsic function is bound to *IFdef* in the workspace `1 if`.

Right Argument: The right argument is a character vector containing the name of an intrinsic function group, or the name of an intrinsic function within a group.

Left Argument: The left argument, if provided, is a character vector containing the definition of the intrinsic function group or function within the group. These definitions must be valid definitions according to the rules given for `.saxif` definition files, explained earlier in this section.

Result: The result of *system.def* is an empty vector in the dyadic case; in the monadic case, it is a character vector containing the current group or function definition.

Example:

```
    'dynamic /home/joe/ifs/exgrp.si' ifdef 'exgrp'  ⍝ define a group
    ifdef 'exgrp'⍝ verify that it is defined
dynamic /home/joe/ifs/exgrp.si
    'cfunction int exfn1(int,char*)' ifdef 'exgrp.exfn1' ⍝ define 1st function
    'cfunction int exfn2(int*)' ifdef 'exgrp.exfn2' ⍝ define 2nd function
    ifdef 'exgrp.exfn1'  ⍝ verify that they are both defined
cfunction exfn1(int,char*)
    ifdef 'exgrp.exfn2'
cfunction exfn2(int*)
```

This example is equivalent to processing the following *.saxif* file at APL session startup:

```
exgrp    dynamic /home/joe/ifs/exgrp.si
exfn1    Cfunction int exfn1(int,char*)
exfn2    Cfunction int exfn2(int*)
```

Note: The corresponding system command, *)ifdef*, is described in "Chapter 4. IF System Commands"

system.erase

The *system.erase* intrinsic function is used to remove the definitions of both a group of functions and the individual functions in the group. This function can be used to remove the definitions of intrinsic functions from within an APL application when they are no longer required, or possibly in order to redefine them. This system intrinsic function is bound to *IFerase* in the workspace *1 if*.

Right Argument: The right argument is a character vector containing the name of an intrinsic function group, or the name of an intrinsic function within a group.

Left Argument: None.

Result: The result of *system.erase* is an empty vector.

Example:

```
     IFerase 'exgrp.exfn1'    ⍝ remove both functions
     IFerase 'exgrp.exfn2'
     ⍴ IFdef 'exgrp.exfn1'    ⍝ verify that they have been removed
0
     ⍴ IFdef 'exgrp.exfn2'
0
     IFerase 'exgrp'          ⍝ remove the group definition
     ⍴ IFdef 'exgrp'          ⍝ verify that it has been removed
0
```

Note: The corresponding system command, *)iferase*, is described in "Chapter 4. IF System Commands"

system.list

The *system.list* intrinsic function is used to list the names of defined intrinsic function groups and the names of functions within an individual defined group This function is bound to *IFlist* in the workspace *1 if*.

Right Argument: The right argument is either an empty vector, or a character vector containing the name of an intrinsic function group.

Left Argument: None.

Result: If the right argument is an empty vector, the result is a character matrix containing the names of all defined intrinsic groups; if the right argument is the name of a defined intrinsic group, the result is a character matrix containing the names of the defined functions within the group.

Example:

```
     IFlist ''       ⍝ list all defined groups
Sunix
sapl
socket
system
unix
     IFlist 'unix'   ⍝ list all functions in the 'unix' group
errno
filesync
```

Note: The corresponding system command, *)iflist*, is described in "Chapter 4. IF System Commands"

unix.errno

The *unix.errno* intrinsic function is used to report the current value of the C language variable `errno` within the APL interpreter program. This is useful in combination with other intrinsic functions that rely on the `errno` variable to

report reasons for failure. It is also sometimes useful in being able to diagnose failures of an APL program for which the normal APL error messages do not completely specify the problem.

Right Argument: The right argument is an empty vector.

Left Argument: None.

Result: The result is an integer scalar containing the current value of the `errno` C-variable.

Example:

```
      'unix.errno' ⎕bind 'errno'
2
      errno ''
22
```

(The standard meaning of this value can be looked up in the UNIX file `/usr/include/sys/errno.h`)

unix.filesync

The *unix.filesync* intrinsic function permits the APL interpreter's use of `fsync()` and/or `sync()` to be controlled by the user. APL uses `fsync()` or `sync()` to ensure that data written to APL files is written immediately to disk rather than being buffered, since the data can potentially lie in the buffers indefinitely before being written to disk. APL's choice of `fsync()`, `sync()`, or no synchronization at all is normally controlled by the start parameter `-Yfsync`. (See the *Handbook, Chapter 3*.) The *unix.filesync* intrinsic function allows this parameter to be altered during execution, and also allows the parameter to be set differently for individual APL files. As with all intrinsic functions, *unix.filesync* must first be bound to a name in the active workspace before being used.

Right Argument: The right argument to *unix.filesync* must be either an integer-valued vector or an empty vector. If the right argument is nonempty, each element is interpreted as an APL file tie number.

Left Argument:

The left argument of *unix.filesync*, if present, may take several forms, depending on the nature of the right argument. If the right argument is empty, the left argument may only be an integer-valued singleton vector whose value is `0`, `1`, or `2`. If the right argument is a nonempty vector, the left argument may be an empty vector. It may also be an integer-valued vector of the same length as the right argument, where each element has the value `¯1`, `0`, `1`, or `2`. In this case the left argument may instead be an integer-valued singleton vector with a value of `¯1`, `0`, `1`, or `2`, in which case it is treated as an vector of the same length as the right argument with the same element repeated as many times as necessary.

Results and Side Effects:

If the left argument is elided and the right argument is an empty vector, there is no side effect and the result is an integer scalar containing the APL interpreter's global `filesync` setting, as follows:

- `0`: Perform no file synchronization.
- `1`: Use `sync()` to perform file synchronization.
- `2`: Use `fsync()` to perform file synchronization.

If the left argument is an integer-valued singleton vector and the right argument is an empty vector, the result is the APL interpreter's global `filesync` setting (as detailed above) before the execution of *unix.filesync*. As a side effect, APL's new global `filesync` setting becomes the value of the left argument.

If the left argument is an empty vector, and the right argument is an integer-valued vector of tie numbers, there is no side effect and the result is an integer vector of the same length as the right argument. Each element of the result is the file-specific `filesync` setting for the corresponding tie number in the right argument. In addition to the three global `filesync` settings, there is an additional setting that can be applied to specific files:

```
¯1:  Use global filesync setting for this file.
 0:  Perform no synchronization on this file.
 1:  Use sync() to synchronize this file.
 2:  Use fsync() to synchronize this file.
```

If both the left and right argument are integer-valued vectors of the same length (though the left argument may be an integer-valued singleton vector), the result is an integer vector of the same length as the right argument. Each element of the result is the file-specific `filesync` setting, as detailed above, for the corresponding tie number in the right argument, before execution of *`unix.filesync`*. As a side effect, the file-specific `filesync` settings for the tie numbers in the right argument are set to the corresponding values in the left argument.

If the left argument is elided, and the right argument is an integer-valued vector, the result is a Boolean vector of the same length as the right argument. Each element of the result is `1` if the corresponding element of the right argument is a valid APL file tie number, and `0` otherwise. As a side effect, an `fsync()` is performed upon all files tied to valid tie numbers in the right argument.

When APL is started, its initial global `filesync` parameter is specified by the `-Yfsync` parameter, which defaults to `2`, if not specified. Each APL file, when tied or created, starts off with a file-specific `filesync` value of `¯1` (use the global `filesync` parameter for that file). These can subsequently be changed by use of the *`unix.filesync`* intrinsic function.

It should be noted that a file is exposed to a significant risk of data loss in the event of a UNIX crash if it is not being synchronized. However, synchronization does make writing to a file much slower, so an application doing many consecutive writes may often wish to disable synchronization for a specific file, and then perform explicit `fsync()`'s upon that file after a certain number of writes to the file. This combines most of the data security of automatic synchronization with greatly improved performance.

Example:

```
     'unix.filesync' □bind 'fs'
2
     fs ¹0                  ª Query global filesync value.
2
     0 fs ¹0                ª Turn off global syncing.
2
     2 2 fs 1 2             ª Use fsync() for files 1 and 2.
¯1 ¯1
     (¹0) fs 1 2 3          ª Query file-specific settings.
2 2 ¯1
     update 3               ª Many writes to file 3.
update complete
     fs 3                   ª Explicitly synchronize file 3.
1
```

4
IF System Commands

Introduction

The system commands, `)binds`, `)iflist`, `)ifdef`, and `)iferase`, allow users to manipulate available IFs after the APL session has started (without being required to start the APL session in a suitably configured way). The IFs corresponding to these commands enable manipulation of IF groups and their members to be done under program control. In particular, this allows an APL application to establish the IF groups that it needs as part of its own functioning, rather than relying on the user to start the APL session with suitable configuration parameters.

More information on SHARP APL system commands is available in the *System Guide, Chapter 5*.

`)binds` Display Bound Intrinsic Functions

`)binds` *[nm1]*

This command causes SHARP APL to display a list of names whose visible referent is a bound intrinsic function, in alphabetical order, horizontally across the terminal display. For alphabetization, the collating sequence is sorted based on `⎕av`. When the command is followed by a name, the list shows only those names that occur in the alphabetized list at or after the position of the sample name. The names listed are the same as those returned in the result of `1 ⎕ws ¯1` (See the *System Guide, Chapter 5*, for more information on `⎕ws`).

)ifdef Display IF Definition/Redefine IF

)ifdef nm1[.nm2] [defn]

This command manipulates the definitions of either a group of functions or the individual functions in the group.

)ifdef *nm1*	obtains the current definition of the group of intrinsic functions whose name is *nm1*.
)ifdef *nm1.nm2*	obtains the current definition of the intrinsic function *nm2* within the group *nm1*.
)ifdef *nm1 defn*	establishes the definition of the intrinsic function group *nm1*.
)ifdef *nm1.nm2 defn*	establishes the definition of the intrinsic function *nm2* within the group *nm1*.

In these last 2 forms, *defn* must be a definition of the intrinsic function group or function within the group that is valid according to the rules given for `.saxif` definition files in "Chapter 2. User-Defined Intrinsic Functions".

This command can be used to establish the definitions of intrinsic functions after starting an APL session, instead of relying on the startup processing associated with an APL session.

Examples:

To define a group,

```
      )ifdef exgrp dynamic /home/joe/ifs/exgrp.si
```

To verify that it is defined,

```
      )ifdef exgrp
dynamic /home/joe/ifs/exgrp.si
```

To define two functions within the group,

```
      )ifdef exgrp.exfn1 Cfunction exfn1(int,char*)
      )ifdef exgrp.exfn2 Cfunction exfn2(int*)
```

To verify that they are both defined,

```
     )ifdef exgrp.exfn1
Cfunction exfn1(int,char*)
     )ifdef exgrp.exfn2
Cfunction exfn2(int*)
```

Note: The corresponding IF, *system.def*, is descriped in "Chapter 3. System Intrinsic Functions".

)iferase Remove IF Definitions

)iferase *nm1[.nm2]*

This command removes the definition of either a group of functions or an individual function in the group.

)ierase *nm1*	removes the current definition of the group of intrinsic functions whose name is *nm1*.
)iferase *nm1.nm2*	removes the current definition of the function *nm2* within the group of intrinsic functions whose name is *nm1*.

This command can be used to remove the definitions of intrinsic functions after an APL session has started, when they are no longer required, or possibly in order to redefine them.

Examples:

To remove both functions from a group,

```
     )iferase exgrp.exfn1
     )iferase exgrp.exfn2
```

To verify that they have been removed,

```
     )ifdef exgrp.exfn1
     )ifdef exgrp.exfn2
```

To remove the group definition,

```
     )iferase exgrp
```

To verify that it has been removed,

```
      )ifdef exgrp
```

Note: The corresponding IF, *system.erase*, is descriped in "Chapter 3. System Intrinsic Functions".

)iflist List IF Groups and Their Contents

)iflist [nm1]

This command displays a list of names of the currently defined intrinsic function groups or the list of names of currently defined functions within an individual defined group, in alphabetical order, horizontally across the terminal display. For alphabetization, the collating sequence is sorted based on □*av*.

)iflist	lists the names of all currently defined intrinsic function groups.
)iflist nm1	lists the names of all currently defined intrinsic functions within the group *nm1*.

Examples:

To list all defined groups,

```
      )iflist
Sunix  sapl   socket system unix
```

To list all functions in the `unix` group,

```
      )iflist unix
errno    filesync
```

Note: The corresponding IF, *system.list*, is descriped in "Chapter 3. System Intrinsic Functions".

5
Supplied IF Facilities

Two sets of user-defined intrinsic functions are included with the distribution of SHARP APL for UNIX:

- The SHARP APL *UNIX* System Function Facility
- The SHARP APL Socket Facility.

These facilities are described in this chapter. As is the case with all intrinsic functions, these must first be bound to a name in the active workspace using the system function □*bind* .

SHARP APL's UNIX System Function Facility

SHARP APL for UNIX is distributed with a set of intrinsic functions (and a few user-defined functions) that provide access to numerous C library and system calls available in UNIX.

Miscellaneous Intrinsics

The following intrinsic functions are mapped to UNIX system functions of the same names and are fully documented by `man` pages:

access	*chdir*	*chmod*	*chown*	*clearerr*
close	*creat*	*ctermid*	*cuserid*	*dup*
dup2	*fchmod*	*fchown*	*fclose*	*fcntl*
fdopen	*feof*	*ferror*	*fflush*	*fgetc*

fgets	*fileno*	*fopen*	*fpathconf*	*fputc*
fputs	*fread*	*freopen*	*fseek*	*fstat*
fsync	*ftell*	*ftok*	*fwrite*	*getcwd*
getegid	*getenv*	*geteuid*	*getgid*	*getgroups*
getlogin	*getpgrp*	*getpid*	*getppid*	*getuid*
getw	*initgroups*	*isatty*	*kill*	*link*
lseek	*lstat*	*mkfifo*	*mknod*	*open*
path	*pathconf*	*pclose*	*pipe*	*poll*
popen	*putw*	*read*	*rewind*	*setegid*
seteuid	*setgid*	*setpgid*	*setrgid*	*setruid*
setsid	*setuid*	*sleep*	*stat*	*sterror*
sysconf	*system*	*time*	*times*	*ttyname*
umask	*ungetc*	*unlink*	*utime*	*write*

These function names are also supplied in the form of IF binds in the workspace `1 unix`.

Message Queue Intrinsics

The IFs described in this section represent the group of UNIX message queue functions as they are implemented in SHARP APL. For further information, consult your UNIX `man` pages on `msgget`, `msgctl`, `msgrcv`, and `msgsnd`.

Note: Arguments to the following IFs are expressed in a monospaced font to indicate ***UNIX-style syntax***. This ensures that the descriptions of these functions are consistent with the `man` page descriptions of their UNIX counterparts.

msqcreate (int key, int msq_flag)

Creates a new message queue data structure for the given key and sets its permissions according to msq_flag; if a message queue for Key = key already exists, it acts according to IPC_ bits of the msq_flag bitmask. This function returns the message queue ID associated with Key = key. On error it returns −1 and errno is set to indicate the error.

UNIX equivalent: msgget (key, msq_flag)

msqsnd (int msq_id, long type, char* pdata, uint size, int flag)

Writes to the queue specified by msq_id a message 'pdata' of the type 'type', a process must have write permission to perform this operation. This function returns 0 on success, −1 on error and errno is set to indicate the error.

UNIX equivalent: msgsnd (msq_id, (const void*) ptr, size, flag).

msqrcv (int msq_id, long* ptype, char* pdata, uint size, long type, int flag)

Reads from the queue a message specified by arguments msq_id and type (of interest) and stores it into ptype and pdata; a process must have read permission to perform this operation. This function returns the number of bytes actually stored in the data buffer. On error it returns −1 and errno is set to indicate the error.

UNIX equivalent: msgrcv(msq_id, (void*)ptr, size, (long int) type, flag)

msqget_stat (int msq_id, int* perm, ulong* stat1, long* stat2)

Fetches the current values of the fields of msqid_ds data structure associated with msq_id and stores them in the arrays perm, stat1, stat2 mapped to the fields of msqid_ds structure:

perm - to msg_perm
stat1- to msg_cbytes, msg_qnum, msg_qbytes
stat2- to msg_lspid, msg_lrpid, msg_stime, msg_rtime, msg_ctime

This function returns 0 on success, −1 on error and errno is set to indicate the error.

UNIX equivalent: msgctl(msq_id, IPC_STAT, struct msqid_ds *buf).

msqset_perm (int msq_id, int uid, int gid, int mode)

Sets the values of {uid, gid, mode} fields of msqid_ds data structure associated with msq_id to the passed values. This function returns 0 on success, −1 on error and errno is set to indicate the error.

UNIX equivalent: msgctl(msq_id, IPC_SET, struct msqid_ds *buf).

msqremove (int msq_id)

Removes the message queue from the system; can be executed only by a process whose effective uid equals sem_perm.cuid or sem_perm.uid, or by a process with superuser privileges

Note: This removal is immediate; any process still using that message queue will get an error of EIDRM on its next attempted operation on it. This function returns 0 on success, −1 on error and errno is set to indicate the error.

UNIX equivalent: msgctl(msq_id, IPC_RMID, struct msqid_ds *buf).

Semaphore Intrinsics

The IFs described in this section represent the group of UNIX semaphore functions as they are implemented in SHARP APL. For further information, consult the UNIX man pages on semget, semctl, and semop.

Note: Arguments to the following IFs are expressed in a monospaced font to indicate ***UNIX-style syntax***. This ensures that the descriptions of these functions are consistent with the man page descriptions of their UNIX counterparts.

semcreate (int key, int number_of_sem, int sem_flag)

Creates a new semaphore data structure for the given key with number_of_sem semaphores in the set, and sets its permissions according to sem_flag. If a semaphore for Key = key already exists, it acts according to IPC bits of sem_flag bitmask. This function returns the semaphore ID associated with Key = key. On error it returns −1 and errno is set to indicate the error.

UNIX equivalent: semget(key, number_of_sem, sem_flag).

Note: The following six *semget* functions return a specific field of a `sem` structure if the current process has read permission.

semget_val `(int sem_id, int sem_number)`

Returns the semaphore value `>= 0`. On error it returns `-1` and `errno` is set to indicate the error.

UNIX equivalent: `semctl(sem_id,semnum,GETVAL)`.

semget_pid `(int sem_id, int sem_number)`

Returns `pid` for the last operation. On error it returns `-1` and `errno` is set to indicate the error.

UNIX equivalent: `semctl(sem_id, semnum, GETPID)`.

semget_ncount `(int sem_id, int sem_number)`

Returns the number of processes awaiting `semval > currval`. On error it returns `-1` and `errno` is set to indicate the error.

UNIX equivalent: `semctl(sem_id,semnum,GETNCNT)`.

semget_zcount `(int sem_id, int sem_number)`

Returns the number of processes awaiting `semval = 0`. On error it returns `-1` and `errno` is set to indicate the error.

UNIX equivalent: `semctl(sem_id,semnum,GETZCNT)`.

semget_all `(int sem_id, ushort* array)`

Fetches all semaphore values in the set into the array. This function returns `0` on success, `-1` on error and `errno` is set to indicate the error.

UNIX equivalent: `semctl(sem_id,0,GETALL, union semun arg)`.

semget_stat `(int sem_id, int* sem_perm, int* sem_stat)`

Obtains status information about the semaphore identified by `sem_id` by fetching the `semid_ds` structure for this set.

Note: 'sem_base' field of semid_ds structure points to the first semaphore in the set, it is valid in the kernel space only; for this reason the semget_stat call does not return that field; sem_perm vector is mapped to the members of ipc_perm structure, sem_stat vector is mapped to three remaining fields of semid_ds data structure.

This function returns 0 on success, −1 on error and errno is set to indicate the error.

UNIX equivalent: semctl(sem_id,0, IPC_STAT, union semun arg).

semset_val (int sem_id, int sem_number, int value)

Sets the value of semval for the set member sem_number if the current process has write permission. This function returns 0 on success, −1 on error and errno is set to indicate the error.

UNIX equivalent: semctl(sem_id,semnum,SETVAL,union semun arg).

semset_all (int sem_id, ushort* array)

Sets all the semaphore values in the set to the values pointed to by 'array' if the current process has write permission. This function returns 0 on success, −1 on error and errno is set to indicate the error.

UNIX equivalent: semctl(sem_id,0,SETALL,union semun arg).

semset_perm (int sem_id, int uid, int gid, int mode)

Sets the owner user and group IDs, and the access permissions for the set of semaphores associated with sem_id; can be executed only by a process whose effective uid equals sem_perm.cuid or sem_perm.uid, or by a process with superuser privileges. This function returns 0 on success, −1 on error and errno is set to indicate the error.

UNIX equivalent: semctl(sem_id,0,IPC_SET,union semun arg).

semoper (int sem_id, int* sem_num, int* sem_op, int* sem_flg, uint num_of_semop)

Performs operations on the set of semaphores associated with sem_id; each semaphore operation is specified by sem_op with sem_flg for sem_num. Acts on the basis ALL or NONE, thus preserving consistency of semaphore operations on the whole set. Vectors sem_num, sem_op, sem_flg must have length equal to num_of_semop.

Operations are based on the value of `sem_op`:

`sem_op`	description
`<0`	grab the resource (lock)
`>0`	release the resource (unlock)
`=0`	check the resource availability.

This function returns `0` on success, `−1` on error and `errno` is set to indicate the error.

UNIX equivalent:
`semop (sem_id, struct sembuf *sembufp, size_t num_of_semop)`.

`semremove` `(int sem_id)`

Removes the semaphore set from the system; can be executed only by a process whose effective `uid` equals `sem_perm.cuid` or `sem_perm.uid`, or by a process with superuser privileges.

Note: This removal is immediate; any process still using this semaphore will get an error of `EIDRM` on its next attempted operation on one of the semaphores of this set.

This function returns `0` on success, `−1` on error and `errno` is set to indicate the error.

UNIX equivalent: `semctl (sem_id, 0, IPC_RMID)`.

Shared Memory Intrinsics

The IFs described in this section represent the group of UNIX shared memory functions as they are implemented in SHARP APL. For further information, consult the UNIX `man` pages on `ftok`, `shmget`, `shmop`, and `shmctl`.

Note: Arguments to the following IFs are expressed in a monospaced font to indicate ***UNIX-style syntax***. This ensures that the descriptions of these functions are consistent with the `man` page descriptions of their UNIX counterparts.

`ftok` `(char *filename, char flag)`

Returns the key value based on filename and flag. On error it returns `−1` and `errno` is set to indicate the error.

UNIX equivalent: `ftok (filename, flag)`.

shmgetexist (int keyval)

Returns the ID of a shared memory segment with key = keyval. The segment must already exist. On error it returns −1.

UNIX equivalent: shmget(key, 0, 0)

shmcreate (int keyval, int segsize, int perm)

Returns the ID of a shared memory segment with key = keyval. If the segment doesn't exist it creates a new segment of size segsize and sets its permissions to perm. On error it returns −1.

Note: Shared memory permissions are similar to those of a file but only read and write bits are relevant. Octal 0624 is read-write for owner, write-only for group, and read-only for others.

UNIX equivalent: shmget(key,segsize, IPC_CREAT|perm).

shmcreatex (int keyval, int segsize, int perm)

Creates a new shared memory segment with key = keyval of size segsize and sets its permissions to perm, return its ID. If a segment with key = keyval already exists, it returns -1. It returns −1 on error.

UNIX equivalent: shmget(key,segsize, IPC_CREAT | IPC_EXCL | perm).

shmlock (int id)

Locks shared memory segment with specified ID in memory (Solaris only). This function returns 0 on success (or on AIX systems). On error it returns −1

Note: Only root can do it.

UNIX equivalent: shmctl(id, SHM_LOCK, 0).

shmset (int id, int uid, int gid, int perm)

Sets owner (uid), group (gid) and permissions (perm) for shared memory segment with specified ID. If you don't want to change a particular value pass it as -1:

shmset (*id*, 182, −1, −1) will change owner only.

Only the owner, creator or `root` can do it. The user and group are specified as numeric IDs, not names. This function returns `0` on success, `-1` on error.

UNIX equivalent: `shmctl(id, IPC_SET, ...)`.

shmat `(int id, uint addr, int flag)`.

Use only as *shmat*`(`*id*`,0,0)`. It attaches a shared memory segment with the specified ID to a process's memory. The segment must already exist. This function returns the address at which the segment has been attached expressed as an integer. It returns `-1` on error.

Note: In APL this address will likely appear as a negative integer.

UNIX equivalent: `shmat(id,addr,flag)`.

shmdt `(uint addr)`

Detaches the shared memory segment attached at `addr`. This does not destroy the segment. All shared memory segments are detached automatically if the process terminates. This function returns `0` on success, `-1` on error.

UNIX equivalent: `shmdt (addr)`

shmremove `(int id)`

Removes the shared memory segment with the specified ID. This does not detach the segment. In fact the segment isn't removed until after all attached processes detach it. This function returns `0` on success, `-1` on error.

UNIX equivalent: `shmctl(id, IPC_RMID, 0)`.

To move data between shared memory and workspace:

```
shm2ws(char *aplvar, uint shmaddr, int len)
ws2shm(char *aplvar, uint shmaddr, int len)
```

`aplvar`	workspace variable.
`shmaddr`	address - presumably an address in shared memory, based on the value previously returned by *shmat*.
`len`	length of data to be copied.

This ***always*** returns 0. It ***does not check*** the address and length in any way. An improper address or length will at best cause SIGSEGV.

To move data between file and shared memory:

```
shm2file(int1, uint2, int3)
file2shm(int1, uint2, int3)
```

int1	file descriptor.
uint2	address - presumably an address in shared memory, based on the value previously returned by shmat.
int3	length of data to be copied.

This returns whatever read or write return. It ***does not check*** the address and length in any way. An improper address or length will at best cause SIGSEGV.

Related APL Functions.

The workspace 1 *unix* includes a set of user-defined APL functions that help SHARP APL interact with UNIX. Althouth these are ***not IFs,*** and do not require ⎕*bind*, they are grouped in this section because they serve a related purpose.

constant	provides a platform-independent way to retrieve the value of a UNIX flag such as O_RDONLY: *fd←open 'filename'⍪(constant 'O_RDONLY')⍪0*
constants	returns a list of available constants.
cputime	returns the total user and system time accumulated by this process, and by all children processes. It returns the cpu times in seconds. It is a cover to the time function.
errno	returns the most recent value of errno.
groups	returns an integer vector containing all the group IDs that this user belongs to. It is a cover to the getgroups system call.

Iand *Ior*	perform the ***integer and*** and ***integer or*** on their arguments. These corresponds to the & and | operators in C, and allow you to easily set, clear, and test bits in an integer.
path	returns the UNIX path name of a component file (as in the notation of the left argument to □*stie*, □*create* or □*rename*. This path name is the same as would be present in □*paths* if the file were tied.
pwd	Returns a character vector containing the current working directory for this process. It is a cover to the `getcwd` library function.
*select*1	selects using lists of file descriptors rather than masks.
*unix*Δ*noop*	returns 0.
*unix*Δ*version*	returns version information for this facility.

SHARP APL's Socket Facility

SHARP APL for UNIX provides a number of intrinsic functions that correspond to the socket primitives available through the standard Berkeley UNIX system calls. These functions are documented by `man` pages. For more information on the Berkeley socket facility, see *UNIX Network Programming,* by W. Richard Stevens.

These socket intrinsics can be called directly by existing IF binds which are available in the workspace `1 socket`.

accept	*bind*	*close*	*connect*	*constants*
errno	*fcntl*	*genaddr*	*getpeername*	*getsockname*
getsockopt	*ioctl*	*listen*	*recv*	*recvfrom*
select	*send*	*sendto*	*setsigpipe*	*setsockopt*
socket	*version*			

Some functions have not been implemented: *readv*, *writev*, *sendmsg*, *recvmsg*, *get*, *peername*, *getsockname*, *shutdown*.

SHARP APL *for UNIX*

Language Guide

JUMP TO ...

CHAPTER 1. LANGUAGE OVERVIEW
CONTENTS
PREFACE
MASTER INDEX
USING THIS DOCUMENTATION

HANDBOOK
AUXILIARY PROCESSORS MANUAL
FILE SYSTEM MANUAL
INTRINSIC FUNCTIONS MANUAL
LANGUAGE GUIDE
SYSTEM GUIDE
SVP MANUAL

© 2000 by Soliton Associates Limited

44 Victoria Street, Suite 2100, Toronto, Ontario, Canada M5C 1Y2

Pubcode: UW-000-0802 (0007)

SHARP APL *for UNIX*

DEVELOPER LIBRARY

Language Guide

SOLITON ASSOCIATES

© 2000 by Soliton Associates Limited–Printed in Canada
44 Victoria Street, Suite 2100, Toronto, Ontario, Canada M5C 1Y2

All rights reserved. No part of this work may be reproduced in any form or by any means without prior written permission of the copyright owner.

Soliton Associates Limited shall not be liable for errors contained herein or for incidental consequential damages in connection with the use of this material. The information contained within this document is subject to change without notice.

Acknowledgments

SHARP APL is a registered trademark of Soliton Associates Limited.

SVP and NSVP are trademarks of Soliton Associates Limited.

IBM and AIX are registered trademarks of International Business Machines Corporation.

RISC System/6000 is a trademark of International Business Machines Corporation.

UNIX is a registered trademark of UNIX System Laboratories, Inc.

Sun Workstation is a registered trademark of Sun Microsystems, Inc.

Sun, Solaris, Open Windows, and SunView are trademarks of Sun Microsystems, Inc.

All other products referred to in this document are trademarks or registered trademarks of their respective trademark holders.

Publication code: UW-000-0802 (0009)

Contents

Tables and Figures xi

Preface

Introduction xiii
Chapter Outlines xiii
Conventions xiv
Documentation Summary xiv
Contacting Soliton Associates xv

1. Language Overview

The Interpreter and Immediate Execution 1-1
Language Features 1-2
Like Math, Unlike Other Languages 1-5
Syntactic Classes in APL 1-6
The Workspace 1-7
Arrays 1-8
Functions and Operators 1-9
Treatment of Operators 1-11
User-Defined Functions 1-12
Rank, Frame, and Cell 1-14
Deprecated Primitives: Braces 1-17
SAXBRACELOG 1-18

2. Naming

Names for User Objects 2-1
Assigning Names to Variables 2-1

The Result of Assignment 2-2
Indirect Assignment 2-3
Multiple Assignment 2-3
Reassignment and Name Conflicts 2-4
Indexed Assignment 2-4
Localization of Names 2-5

3. Data

Rank and Shape 3-1
Reporting Shape and Rank 3-2
Arrays, Cells, and Frames 3-2
Complementary Partition 3-4
Types of Array Elements 3-5
Numeric Elements 3-5
Character Data 3-7
APL Character Set 3-8
Enclosed Elements 3-9
Heterogeneous Arrays 3-10
Packages 3-11
Input and Output 3-12
Representation of Constants 3-12
Higher-Rank Arrays 3-14
Heterogeneous or Enclosed Arrays 3-14
Display of Arrays 3-15
Empty Arrays 3-21

4. Functions

Categories of Functions 4-1
Arithmetic Functions, Scalar 4-5
+ω Conjugate/Identity 4-6
α+ω Plus/Addition 4-6
−ω Minus/Negate 4-6
α−ω Minus/Subtraction 4-6
×ω Signum 4-6

α×ω Times/Multiplication 4-7
÷ω Reciprocal 4-7
α÷ω Divide/Division 4-7
⋆ω Exponential 4-7
α⋆ω Power 4-8
⍟ω Natural Logarithm 4-8
α⍟ω Base-α Logarithm 4-8
|ω Magnitude/Absolute Value 4-9
Residue/Remainder 4-9
!ω Factorial 4-10
α!ω Out-Of/Combinations 4-11
⌊ω Floor 4-12
α⌊ω Minimum 4-12
⌈ω Ceiling 4-12
Fuzzy Floor and Ceiling 4-13
α⌈ω Maximum 4-13
○ω Pi Times 4-13
α○ω Circle Functions 4-14
α∧ω And/Least Common Multiple (LCM) 4-17
α∨ω Or/Greatest Common Divisor (GCD) 4-17
Fuzzy GCD and LCM 4-18
α⍲ω Nand 4-18
α⍱ω Nor 4-18
~ω Not 4-19
?ω Roll 4-19
Arithmetic Functions, Non-scalar 4-20
⌹ω Matrix Inverse 4-20
α⌹ω Matrix Divide 4-20
α?ω Deal 4-22
α⊤ω Encode/Representation 4-23
α⊥ω Decode/Base-10 Value 4-24
Relational Functions 4-25
α=ω Equals 4-25

⍺≠⍵ Not Equals 4-27
⍺<⍵ Less Than 4-27
⍺≤⍵ Less Than or Equal 4-27
⍺≥⍵ Greater Than or Equal 4-27
⍺>⍵ Greater Than 4-27
⍺≡⍵ Match 4-28
⍺∈⍵ Member 4-28
⍺⍷⍵ In 4-30
Indexing Functions 4-31
⍺@⍵ From 4-31
⍳⍵ Count 4-35
⍺⍳⍵ Index Of 4-36
⍺⌷⍵ Index 4-37
⍋⍵ Numeric Grade Up
⍒⍵ Numeric Grade Down 4-38
⍺⍋⍵ Character Grade Up
⍺⍒⍵ Character Grade Down 4-39
⍺[⍵] Bracket Indexing 4-42
Structural Functions 4-45
⍴⍵ Shape 4-46
⍺⍴⍵ Reshape 4-46
⍺↑⍵ Take 4-49
⍺↓⍵ Drop 4-52
↓⍵ Raze 4-53
<⍵ Enclose/Box 4-56
⊃⍵ Conditional Enclose/Box 4-58
⍺⊃⍵ Link 4-58
>⍵ Disclose/Open 4-59
,⍵ Ravel 4-61
⍪⍵ Table 4-61
⍺,⍵ Catenate
⍺⍪⍵ Catenate-Down 4-62
⌽⍵ Reverse 4-66

⊖ω Reverse-down 4-66
α⌽ω Rotate 4-67
α⊖ω Rotate-down 4-68
⍉ω Monadic Transpose 4-69
α⍉ω Dyadic Transpose 4-71
Miscellaneous Functions 4-73
⊢ω Pass 4-73
⊣ω Stop 4-74
α⊣ω Left 4-74
α⊢ω Right 4-74
≠ω Nubsieve 4-75
↑ω Nub 4-76
=ω Nubin 4-76
α~ω Less 4-77
@ω All 4-78
⍕ω Monadic Format 4-78
α⍕ω Dyadic Format 4-84
⍎ω Execute 4-92

5. Operators

Precedence of Operators 5-3
Valence of Operators 5-3
Syntactic Class of Arguments to Operators 5-3
Consecutive Operators 5-4
Look-ahead 5-5
Separating Side-by-Side Constants 5-5
Monadic Operators 5-6
f⌿ ω f-Reduce
f≠ ω f-Reduce-down
f\ ω f-Scan
f⍀ ω f-Scan-down 5-8

⍵ Compress/Copy
⍵ Compress-down/Copy-down
⍵ Expand
⍵ Expand-down 5-11
⍵ Swap 5-13
⍺ ⍵ Swap 5-13
⍵ Select 5-14
⍺ ⍵ Merge 5-14
Dyadic Operators 5-15
Cases of Dyadic Operators 5-15
⍺ ⍵ Tie 5-15
⍵ Ply 5-18
⍺ ⍵ Inner Product 5-18
⍵ Alternant 5-20
⍵ Upon
⍺ ⍵ Upon
⍵ On
⍺ ⍵ On
⍵ Under
⍺ ⍵ Under 5-21
⍵ Rank
⍺ ⍵ Rank 5-25
⍺ ⍵ Cut 5-28
⍵ Cut 5-32
⍵ With
⍵ With 5-32

6. Syntax

Line: The Unit of Execution 6-1
Lines, Statements, and the Diamond Separator 6-2
Grouping and Separators Within a Statement 6-3
Quotes Delimit a Character Constant 6-3
Parentheses 6-4
Brackets Delimit an Index 6-6
Semicolon Separates Index Expressions 6-7

The Branch Arrow as a Grouping Device 6-8
Colon and Label 6-9
Lamp and Comment 6-9

Index ix-1

Tables and Figures

Table 1.1. APL symbols and character names. 1-3
Table 1.2. APL symbols listed by description. 1-4
Table 1.3. APL symbols by syntactic class 1-6
Figure 3.1. Partitioning an array into cells. 3-3
Table 3.1. The APL character set. 3-8
Figure 3.2. An array that contains two enclosed elements. 3-10
Table 4.1. Ranks of primitive functions. 4-3
Figure 4.1. Schematic representation of residue. 4-10
Table 4.2. Families of circular functions. 4-14
Table 4.3. Individual circle functions. 4-14
Figure 4.2. Pythagorean and circular functions in terms of the unit circle. 4-15
Figure 4.3. Hyperbolic functions in terms of the unit hyperbola. 4-16
Table 4.4. Comparing characters with identical weights. 4-41
Figure 4.4. Effective weight ⌹ of ⍵ occurring at two positions in ⍺. 4-42
Figure 4.5. Take and drop. 4-50
Table 5.1. SHARP APL monadic operators. 5-2
Table 5.2. SHARP APL dyadic operators. 5-2
Figure 5.1. Three forms of composition. 5-22
Table 5.7. Composition templates. 5-23
Figure 5.2. Close or cellwise composition. 5-24
Table 5.9. Specifying the rank of a function. 5-26

Preface

Introduction

This document is intended for SHARP APL for UNIX application developers. It explains all features of the language interpreter with chapters on APL object names, data, functions, operators, and syntax. Although this volume describes the language with precise detail, it does try to explain how to design and implement the applications that are to be written in APL.

Rather than reproduce existing material, references to other SHARP APL for UNIX publications are supplied where applicable.

Chapter Outlines

The SHARP APL for UNIX System Guide is organized into the chapters described below.

Chapter 1, “Language Overview”, provides an overview of the elementary but fundamental points of the SHARP APL language.

Chapter 2, “Naming”, discusses how SHARP APL objects are named and how these names are chosen.

Chapter 3, “Data”, provides complete details on the structure of SHARP APL data, types of array elements, and the entry and display of data arrays.

Chapter 4, “Functions”, classifies and discusses all the primitive functions in SHARP APL for UNIX together with examples to illustrate their use.

Chapter 5, “Operators”, is the general reference for all primitive operators supported by SHARP APL for UNIX.

Chapter 6, “Syntax”, explains how APL symbols are arranged to form expressions, statements, and lines according to the rules of APL syntax.

Conventions

The following conventions are used throughout this manual:

`⎕io←0`	Although the default value for ⎕io in a clear workspace is one, all examples in this manual assume that the index origin is ***zero***.
⍺ and ⍵	Two APL symbols are used in syntax descriptions of monadic and dyadic functions; *alpha* (⍺) identifies the left argument and *omega* (⍵) identifies the right argument.
m and *n*	These letters are sometimes used as placeholders for the left and right arguments to an operator when those arguments are ***arrays***.
f and *g*	These letters are sometimes used as placeholders for the left and right arguments to an operator when those arguments are ***functions***.
@	The APL primitive function ***from*** (⍺@⍵) is used as shorthand to identify elements in a vector or matrix. For example, the first element in a vector *v* would be identified as `0@v`; the first row of a matrix *m* would be identified as `0@m`.

Documentation Summary

The SHARP APL for UNIX reference documentation is published in several volumes. One or more of the following may be referenced in this guide.

SHARP APL for UNIX,

- *Handbook,* publication code UW-000-0401
- *System Guide,* publication code UW-000-0902
- *SVP Manual,* publication code UW-001-0501
- *File System Manual,* publication code UW-037-0501
- *Auxiliary Processors Manual,* publication code UW-033-0501
- *Intrinsic Functions Manual,* publication code UW-007-0501

For a complete list of SHARP APL publications, please refer to the ***Support*** link on the Soliton Associates website: *http://www.soliton.com.*

Contacting Soliton Associates

Problems or questions regarding your SHARP APL system or this guide should be directed to the Soliton Associates Technical Support group:

support@soliton.com

Requests for additional technical documentation should be directed to the Soliton Associates Sales group:

sales@soliton.com

1
Language Overview

SHARP APL is designed to be independent of the environment in which it is interpreted and executed. However, because SHARP APL is embedded in the context of the UNIX environment, you rarely use the APL language to compute without also taking advantage of the supporting environment.

This chapter reviews elementary but fundamental points of the SHARP APL language.

The Interpreter and Immediate Execution

SHARP APL provides an interpreter for ***statements*** written in APL. APL statements, which resemble mathematical notation, may contain numbers, symbols, and alphanumeric names. (See the section "Language Features", below.)

The interpreter proceeds directly from APL statements to execution; there are no intermediate stages of compilation, assembly, or linkage. The default mode in which the system operates is called ***immediate execution.*** As soon as you supply a statement (for example, by typing it at the keyboard), the system interprets and executes it. There are other modes, such as those for editing rather than executing definitions, or supplying input as programs may request it; however, immediate execution is the default mode and the mode to which the system returns when other work is completed or interrupted.

Because the system interprets each statement as you present it, there is a continual alternation between two phases: you submit a statement, and the system executes it. When execution is complete, the system displays the result (if any) and awaits your next statement. The record of the transactions between user and system looks like the transcript of a dialogue. (The session manager is discussed in *System Guide, Chapter 7*.)

To distinguish your statement from the system's response, the system issues a prompt to solicit each statement. The prompt normally consists of a new-line character and six blanks, so that each statement you enter from the keyboard appears (on the screen or in the session log) indented from the left margin by six blanks. The system's response lacks that indentation; for example:

```
      3 2÷6
0.5 0.3333333
```

Language Features

In an APL statement, the special symbols represent ***primitives***. They are considered primitives because their meanings are permanently fixed and understood by the APL interpreter without further definition. Table 1.1 lists the APL symbols and their names. Each primitive ***function*** or ***operator*** is denoted by a single symbol. For example, to divide a by b, you would write:

```
a÷b
```

To calculate the number of digits required to represent a positive integer n in base b, you would write:

```
1+⌊b⍟n
```

To mark those elements in an array t that differ from the first element in t, you would write:

```
t≠1↑t
```

And to perform matrix division of a by the transpose of b, you would write:

```
a⌹⍉b
```

For more information on functions and operators, refer to the section "Syntactic Classes in APL", below.

Table 1.1. APL symbols and character names.

Symbol	Character Names	Symbol	Character Names	Symbol	Character Names
!	Exclamation point	&	Ampersand	(	Left Parenthesis
'	Quote mark	)	Right parenthesis	*	Star; Asterisk
+	Plus	-	Minus; Dash	.	Period
/	Slash	:	Colon	;	Semicolon
<	Less-than	=	Equal	>	Greater-than
?	Question mark	@	At	[	Left bracket
\	Back-slash	]	Right bracket	∧	And; Caret
_	Underbar	{	Left brace	}	Right brace
~	Tilde	¨	Dieresis	¯	High minus; Macron
≤	Less-than-or-equal	≥	Greater-than-or-equal	≠	Not equal
∨	Or; Inverted caret	×	Times	⍪	Comma-bar
⌹	Domino	⌿	Slash-bar	⍲	Nand
⍱	Nor	≡	Equal-underbar	⌷	Squish quad
⍺	Alpha	⊥	Up tack	∩	Cap
⌊	Down stile	∈	Epsilon	∇	Del
∆	Delta	⍳	Iota	∘	Jot
⍞	Quad	|	Stile	⊤	Down tack
○	Circle	⍴	Rho	⌈	Up stile
↓	Down arrow	∪	Cup	⍵	Omega
⊃	Right shoe; Link	↑	Up arrow	⊂	Left shoe
⊢	Right tack; Dex	⍀	Backslash-bar	⊣	Left tack; Lev
÷	Divide	⌶	I-Beam	⊖	Circle-bar
⍟	Pawn	⍝	Lamp	⍷	Epsilon-underbar
⍒	Del-tilde	⍙	Delta-stile	⍸	Iota-underbar
⍤	Paw	⍞	Quote-quad	⍨	Thorn
⍥	Hoof	⍟	Log; Circle-star	⍉	Circle-back-slash
←	Left arrow; Is assigned	⍙	Delta-underbar	→	Right arrow; Branch
◇	Diamond	⍒	Del-stile	,	Comma
⍎	Circle-stile				

Note: This table lists each APL symbol followed by its established character names. Refer to Table 1.2 for symbols listed by description. Table 1.3 lists syntactic classes. The complete APL character set (⎕*av*) is listed in Table 3.1, page 3-8.

Table 1.2. APL symbols listed by description.

Description	Symbol	Type	Description	Symbol	Type
absolute value	|	function	addition	+	function
all	@	function	alternat	.	operator
and	∧	function	assignment	←	punctuation
base-10 value	⊥	function	brackets	[]	punctuation
branch	→	punctuation	catenate	,	function
catenate-down	⍪	function	ceiling	⌈	function
character grade down	⍒	function	character grade up	⍋	function
character input/output	⍞	data	colon	:	punctuation
comment	⍝	punctuation	conditional enclose	⊃	function
conjugate	+	function	copy	/	operator
copy-down	⌿	operator	count	⍳	function
cut	⍤	operator	deal	?	function
decode	⊥	function	del edit	∇	function
diamond separator	⋄	punctuation	disclose	>	function
divide	÷	function	drop	↓	function
enclose	<	function	encode	⊤	function
equals	=	function	execute	⍎	function
expand	\	operator	expand-down	⍀	operator
exponential	*	function	evaluated input/output	⍎	data
factorial	!	function	floor	⌊	function
format	⍕	function	f-reduce	/	operator
f-reduce-down	⌿	operator	from	@	function
f-scan	\	operator	f-scan-down	⍀	operator
grade down	⍒	function	grade up	⍋	function
greater than	>	function	greater than or equal	≥	function
greatest common divisor	∨	function	identity	+	function
in	∈	function	index of	⍳	function
inner-product	.	operator	label	:	punctuation
least common multiple	∧	function	left	⊣	function
left argument	α	variable	less	~	function
less than	<	function	less than or equal	≤	function
link	⊃	function	logarithm	⍟	function
magnitude	|	function	match	≡	function
matrix divide	⌹	function	matrix inverse	⌹	function
maximum	⌈	function	member	∈	function

Description	*Symbol*	*Type*	*Description*	*Symbol*	*Type*
merge	⍀	operator	minimum	⌊	function
minus	−	function	multiplication	×	function
nand	⍲	function	negate	−	function
nor	⍱	function	not	~	function
not equals	≠	function	nub	↑	function
nubin	=	function	nubsieve	≠	function
numeric grade down	⍒	function	numeric grade up	⍋	function
on	⍤	operator	open	>	function
or	∨	function	out	!	function
pass	⊢	function	Pi times	○	function
plus	+	function	power	*	function
quote	'	punctuation	rank	⍤	operator
ravel	,	function	raze	↓	function
reciprocal	÷	function	remainder	|	function
representation	⊤	function	reshape	⍴	function
residue	|	function	reverse	⌽	function
reverse-down	⊖	function	right	⊢	function
right argument	⍵	variable	roll	?	function
rotate	⌽	function	rotate-down	⊖	function
select	⍀	operator	semicolon	;	punctuation
shape	⍴	function	signum	×	function
stop	⊣	function	subtract	−	function
swap	⍨	operator	table	⍪	function
take	↑	function	tie	.	operator
times	×	function	transpose	⍉	function
under	⍢	operator	upon	⍤	operator
with	⍢	operator			

Like Math, Unlike Other Languages

In its use of symbols, APL resembles math and differs from almost all other programming languages, which have generally made do with the restricted set of characters found on conventional keyboards. In fact, an operation written in APL is sometimes called an ***expression*** because it looks like, and behaves like, a mathematical ***expression***.

Many of the symbols used in APL are actually borrowed directly from mathematics: "×" for ***multiply;*** "÷" for ***divide;*** and so on. Others have been adapted or coined. Where possible, visual symbolism is used. For example, "⌊" represents ***floor*** (the integer part) and "⌈" stands for ***ceiling;*** "○" denotes the circular (trigonometric) functions; "⌽" stands for ***rotate;*** and "⍉" stands for ***transpose.***

Syntactic Classes in APL

To understand a statement written in APL, you must first recognize how the APL interpreter groups individual characters into ***tokens***, and how tokens can be arranged to form statements.

A token is one or more characters that, in the interpretation of a statement, are treated as a unit. For example, any of the APL symbols for a primitive function or operator is a single token, since each stands alone. However, a ***name*** is just one token regardless of the number of characters required to spell it, just as a ***number*** is a single token regardless of the number of digits required to represent it.

Table 1.3. APL symbols by syntactic class

Punctuation	*Assignment*	*Function*	*Monadic Operator*	*Dyadic Operator*	*Data Variable*
( )	←	+ × − ÷ ⋆ ⍟	⌿ ⌿ &	⍤ ⍥ ⍣ .	⍺ ⍵
[ ]		⌈ ⌊ ⊥ ⊤ ⍋ ⍒	⍀ ⍀ ⊂		⍬ ⍞
;		< ≤ = ≥ > ≠			
'		∧ ∨ ⍲ ⍱ ∈ ⍴			
⍝		↑ ↓ ∼ ⍳ ○ ?			
◇		≡ ⍉ ⊖ ⌽ ⍋ ⍒			
:		∣ , ⍪ ⊆ ! ⌹			
→		⊃ ⊢ ⊣ ⌶ @			

Once you recognize tokens within a statement, you have to understand how they fit together according to the rules of syntax. In APL, a token or symbol may be from one of six syntactic classes:

Data Data is made up of numbers, characters, boxes, or some combination of these, arranged in an array (see "Arrays" below). Data may also be a ***package***–a collection of objects that can be manipulated using specialized system functions. However, a package is ***not*** an array.

Function	A function specifies an action to be taken to process data. The data on which a function is applied are called the ***arguments*** to that function. A function may take one argument (to its right) or two arguments (one on each side). A user-defined function may also be called a ***program.***
Monadic Operator	A monadic operator applies to a single argument, which may be an array or a function, to produce a new ***derived function***.
Dyadic Operator	A dyadic operator applies to two arguments, either of which may be an array or a function, to produce a new ***derived function***.
Assignment	The ***left arrow*** symbol (←) links a name to data being named.
Punctuation	Characters written in APL are grouped into tokens, statements, and lines. Parentheses are the principal punctuation in APL; as in conventional mathematics, they are used to alter the normal meaning of an expression by changing the order in which it is evaluated.

The Workspace

The APL system assigns each active user an amount of storage called a ***workspace.*** The workspace serves both as the place in which calculation occurs and the environment in which names are understood. Any function defined in a workspace can refer to any other function, or to any available variable. However, during execution, functions may establish their own ***local names,*** which are hidden from general view.

The concept of the workspace permits functions, each of which is defined separately, to be combined in a single statement or to invoke each other. The fact that they all reside in the same workspace makes the step of linking (required in some other languages to permit separate programs to work together) unnecessary in APL.

The APL interpreter permits you to save copies of everything in the ***active workspace***. Such a snapshot of the active workspace is called a ***saved workspace.*** Developing an APL application usually involves bringing together the various definitions it will use and then saving them together as a saved workspace.

Ways for saving and loading workspaces are described in the *System Guide, Chapter 2*.

Arrays

APL is an array-oriented language. In general, data in APL takes the form of an array[1] composed of any number of ***elements***, arranged along any number of ***axes*** or ***dimensions***. An array that contains only one element and has no axes at all is just a degenerate, albeit common, case of an array.

When a function is applied to an array, it is applied to it ***as a whole***. Depending on what the function is and how it is used (and how it may be modified by an operator), the function either treats the entire array collectively or considers the array to be a set of cells (see "Chapter 3. Data") to which the function applies in parallel fashion.

APL includes provisions for indexing (selecting individual elements from within an array), but APL programs make less use of indexing than would be required in languages whose functions are defined only on single elements. Moreover, the APL interpreter is designed for efficient processing of functions applied to an entire array, so that a program written to exploit the extensive array-handling powers of APL usually runs faster than those defined iteratively for individual elements.

Because of its array orientation, APL requires much less iteration, looping, testing, and branching than other programming languages. APL lacks constructs such as the `FOR` or `DO WHILE` instructions of some other languages. APL is rich in functions for the manipulation of arrays: ***shape***, ***reshape***, ***rotate***, ***transpose***, ***take***, ***drop***, ***compress***, ***expand***, ***replicate***, and so on. Because many of these array operations are not found in conventional mathematics, there are no established symbols for them; functions dealing with array manipulation account for many of APL's coined symbols.

Types of Arrays

An element in an APL array may be a character, a number, or ***enclosed***. An enclosed element is a single element that contains an array enclosed within it. Provided it remains closed, an enclosed element occupies a single position in an array, just as a single number or a single character does.

Enclosed elements permit irregularly shaped data to be stored in a regular rectangular framework. An array that contains enclosed elements is called a ***nested array***.

1. The only exception is a package, a kind of data whose elements are identified by name rather than by position in an array structure.

The function ***disclose,*** >, opens an enclosed array. When an enclosed element is opened, the array enclosed within it emerges and takes on whatever shape it had before it was closed. Disclose, applied to a nested array, returns an array of the opened elements to a conforming type and shape. Various dyadic operators may be used to combine disclose with other functions to operate independently on the contents of the elements of the nested array.

Packages

A package is a kind of data that in some respects resembles a miniature workspace. A package may contain an assortment of data, user-defined functions, and bound intrinsic functions. These are selected from the package by ***name*** rather than by indexing. Because packages are not arrays, they cannot be manipulated using functions that are defined on arrays. Rather, a set of specialized system functions are available for manipulating the contents of packages. Packages are described in detail in "Chapter 3. Data".

No Declarations

There is no provision in APL to declare in advance the type, shape, or size of data. The interpreter handles this for you. It determines the type and shape of data as it is being manipulated, and allocates space as required. The interpreter chooses various internal forms for the storage of numbers, but converts among them automatically. A single variable may contain elements of any type. However, a homogeneous array (one whose elements are all of the same internal type) may facilitate faster execution or take up less storage space than a heterogeneous array.

Functions and Operators

APL is an imperative, not a declarative language. A ***statement*** consists of an instruction to do something. A ***function*** specifies the action to be taken. A ***monadic operator*** modifies a function, specifying some aspect of how the function is to act; a monadic operator comes ***after*** the function it modifies. A ***dyadic operator*** joins the actions of two functions and comes ***between*** them. For further explanation, see the section "Treatment of Operators", below.

Precedence Rules

APL's rules of syntax are extremely simple: all functions have equal precedence, and all operators have equal precedence. However, operators have higher precedence than functions. This is because you have to know how a function has been modified or conjoined before the newly derived function can perform the appropriate action.

Order of Execution

A monadic operator takes as its argument the variable or function that it follows; a dyadic operator takes as its arguments the variable or function on either side of it. An operator produces a ***derived function*** as its result. Because operators have higher precedence than functions, as soon as you come to something that is modified by an operator, you have to evaluate the operator ***immediately.***

Example:

```
(a+b)×⍤(⍳,r)-(c+d)
```

Reading left to right, the quantity `(a+b)` is to be multiplied. But × (times) is modified by the dyadic operator ⍤ (rank) so you have to evaluate the rank operator before you can do anything about the multiplication. The operator's arguments are the function × on the left and the variable `(⍳,r)` on the right.

In order to evaluate ⍤, you have to evaluate its argument `(⍳,r)`. That done, you establish that the primary expression is *times-rank-l,r.* Its right argument is `-(c+d)`. Notice that although the right end of the statement contains the characters

```
(⍳,r)-(c+d)
```

`(c+d)` is *not* subtracted from `(⍳,r)`. The precedence rule means that `(⍳,r)` is the argument of the operator and gets evaluated before the subtraction is performed.

The fact that APL is read from left to right requires the interpreter, when it executes a statement containing several expressions, to evaluate the subordinate expressions to the right before it can interpret the primary expression (at the left). This is sometimes called the right-to-left rule, or leaf-to-root execution.

Infix Notation

In APL, as in arithmetic, a function may appear between its arguments. That is, the function is an ***infix.*** (In languages such as C or Lisp, almost all functions are prefixes, while in PostScript, all functions are postfixes.) When an APL function has only one argument, the function is a ***prefix*** (i.e. the argument comes after the function.

Ambivalence

Valence is the number of other things with which something combines. In SHARP APL, the same function can be used with two arguments, or with just one argument on the right. Since the same function can be used either with one argument or with two, functions are said to be ambi-valent or ***ambivalent.***

For example, the symbol "−"appears twice in the following statement:

```
−a−b
```

The first "−" denotes ***negation***: a quantity is to have its sign reversed. The quantity to be negated is the difference between `a` and `b`. The second "−" is used with two arguments and means ***subtraction***. As this example shows, the practice of using the same symbol with just one argument, or with arguments on both sides, is already familiar from arithmetic.

Monadic and Dyadic

When a function is used with arguments on both sides, its use is ***dyadic.*** Dyadic means "dealing with two arguments." When a function is used with only one argument, its use is ***monadic.***

Treatment of Operators

In his original paper, entitled ***Functions and Operators*** (1978), Iverson proposed the formalization in APL of a syntactic class called ***operators.***

An APL monadic operator can take a function as its argument, modify the function, and produce a new function (called a derived function). The newly modified function can then be applied to one or two data arguments to produce a data result. In similar fashion, a dyadic operator joins (and thus modifies) the action of two functions to produce a new derived function.

The concepts of operators had existed unrecognized in the earliest versions of APL. All implementations of APL support at least some of the operators, such as +/ for summation, +.× for inner product, and so on. At first, they were treated as ad hoc extensions—very useful, but lacking any general theory of structure or application. Formalizing them has permitted them to be used much more widely, and has led to major advances in the way APL handles the array structure of data. Use of the new operators depends heavily on the concepts of ***cell*** and ***frame*** in the representation of data. See the section "Rank, Frame, and Cell," below.

SHARP APL for UNIX extends the usefulness of operators by permitting an operator to modify any class of function: primitive, user, system, or derived. It also introduces operators that were not present in previous implementations of SHARP APL.

User-Defined Functions

A large part of programming in APL involves the creation of new functions, called user-defined functions. A ***program*** in APL is just another term for a user-defined function. Application design consists of selecting the data of interest, adopting an appropriate structure for it, and defining the functions that transform it. A user-defined function is denoted by a ***function name,*** beginning with at least one alphabetic character — these are distinct from primitives because each primitive has a symbol, whereas each user-defined function has a name.

Libraries of User-Defined Functions

User-defined functions are typically stored in a workspace for later use. A set of workspaces can also be collected and saved in a library.

An ***application*** is a set of related functions ready for use. Usually the members of such a set are saved together in a workspace ready for execution. It is possible to copy individual functions from a saved workspace without having to use the entire set.

A set of saved workspaces is called a ***library*** of workspaces. Ordinarily, you run an application by selecting from a library the workspace you wish to use. To build a new application, you may start by copying from one or more workspaces those functions you will use, then modifying them or adding new functions to obtain the behavior you want. For more information on libraries and workspaces see the *System Guide, Chapter 2*.

A Function's Definition

The definition of a function consists of

- a ***header*** stating its ***name*** and a ***description*** of its syntax, followed by
- a ***body*** consisting of zero or more ***lines*** each containing one or more statements.

For example, the definition of a function called *analyze* might start with a header such as *z←analyze ⍵*. This shows that, in use, *analyze* is to be followed by the data to be analyzed. Further, it shows that, inside the definition (but not necessarily when you actually use *analyze*), whatever value the argument then has will be referred to by the name *⍵*. Finally, it shows that the result that the function returns will, inside the definition, be called *z*.

Editing Functions

The SHARP APL system provides two editors for creating user-defined functions: a line-mode editor and a full-screen editor.

The line-mode editor, or ***del editor***, is invoked by typing the ***del*** character (∇). Because that symbol takes you into the del editor (and out again when you have finished editing), it was customary to show a del at the beginning and end of the display of a definition. The editor numbers the lines in the definition; it is also customary to show the numbers in the display.

The full-screen editor is available only from the interpreter called *apl*. This editor is invoked by the system command *)edit*. Instructions for using both these editors are available in the *System Guide, Chapter 8*.

One Definition May Use Another

The instructions that appear in a function's definition may themselves refer to other defined functions denoted by arbitrary names. For example, the definition of *analyze* might be:

```
    ∇ z←analyze ⍵
[1]   z←correlate annual adjusted average ⍵
    ∇
```

Presumably *correlate*, *annual*, *adjusted*, and *average* are defined names; some may be the names of functions, while others may be variables. Each time the interpreter executes the function *analyze*, it invokes whatever definitions it finds for the names in the statements contained in *analyze*.

At the time you write the definition for *analyze*, the interpreter does not need to know the meaning of names that appear in the definition. It is characteristic of an interpreted language that the interpreter looks up the meaning of each name not when you write the definition, but when it executes the definition. Indeed, at subsequent invocations of *analyze*, the definition of *correlate*, and that of other defined names, may have changed.

A function may invoke another, which in turn may invoke another, and so on. Functions may even invoke themselves (in which case, they are said to be recursive). Function calls can be nested to many levels. Since function can call another, it is easy for a program to build on functions already defined. Of course, everything must be defined in terms the interpreter can understand; therefore, the function definitions must ultimately refer to APL primitive functions and operators.

Modified Functions

Operators modify the way a function is applied. An operator can modify a primitive function, a user-defined function, a system function, or even another modified function. For example, to apply the user-defined function *analyze* to each of the two-dimensional matrices it finds within a data array called *x*, use the ***rank*** operator (*f*⍤*n*):

```
analyze⍤2 x
```

More generally, to apply *analyze* to each of the rank-*n* cells within *x*, write:

```
analyze⍤n x
```

Rank, Frame, and Cell

A function treats the array it works on as a set of one or more ***cells*** of data. The size of these cells depends on two things:

- the rank of the array (the total number of axes or dimensions the array actually has), and

- the rank of the function (the function's rules about the number of axes it expects for each cell of an argument array).

Rank of an Array

APL arrays often have several axes or dimensions. The ***rank*** of a variable is the number of axes it has. For example, a single element (say, the number `123.45` or the letter `'K'`) has rank ***zero*** because it has no axes. A ***vector*** (for example, the numbers `12 ¯24.3 4.75` or the letters `'Mary Jones'`) has rank ***one*** because its elements are spread along a single axis. A ***matrix,*** made up of rows and columns, has rank ***two***, and so on. In principle, there is no limit to the number of axes, but in practice the APL interpreter cannot handle an array with more than `127` axes.

The Shape and Rank of an Array

The ***shape*** function (⍴⍵) reports the ***length*** of each axis of its argument. Its result is a vector containing one number for each axis of the argument array. The value of each number indicates the length of the corresponding axis; that is, how many elements lie along that axis.

For example, for a matrix *m* with three rows and four columns, the shape function will return the shape of the matrix as follows:

```
      ⍴m
3 4
```

Applying the ***shape*** function to the result of ⍴*m* returns a single number indicating how many elements the shape of *m* contains:

```
      ⍴⍴m
2
```

In this sense, one use of the ***shape*** function returns the ***shape*** of an array, and two consecutive uses of the ***shape*** function returns the ***rank*** of an array.

Repetition of Cells in a Frame

In many situations, you may want to perform a calculation on a set of related data (for example, the various elements required to compute the value of a paycheck for an individual employee), then repeat the same calculation for many such transactions, each having essentially the same structure (such as calculating paychecks for all employees on the payroll).

When you are calculating the amount of each employee's check, you can view the data as a set of ***cells*** (one for each employee) embedded in a ***frame*** (that contains as many such cells as there are employees). A paycheck program can apply the same algorithm independently to each cell throughout the frame.

In other situations, you can apply a calculation not to the various cells separately, but to the totality of ***all cells.*** Consider tasks such as these: counting the number of employees, adding up the company's total payroll, or sorting the list of employees by the amount of their pay. They require you to consider the entire set of cells in an array. You cannot complete such tasks by independent analyses of the cells taken separately.

Implicit Iteration

Much of the repetitive nature of data processing is handled automatically in APL. There is often no need to use counters, tests, or loops because the interpreter takes care of those implicitly. You write as though the interpreter will carry out your calculation in a ***parallel fashion*** for all the cells throughout the frame of your data.

When you have a machine in which the hardware supports parallel processing, the processors work concurrently to handle many cells at one time. Such machines are still rare. In serial machines, the SHARP APL interpreter nevertheless makes work ***seem*** parallel: calculations seem to be executed concurrently for all the cells.

Quite apart from the speedups achieved by automatic iteration, it is also an advantage to write as though work were done in parallel. This is true even on a machine entirely lacking efficient array processing. Both as user and as programmer, you benefit from a notation that frees you from having to describe loops. It does not matter if, internally, the interpreter has to set up loops in order to treat each of the cells within a frame.

All–or–None Processing of Arrays

In APL, an array is treated as a whole. When you apply a function to an array, you apply it to the whole array, and you get a result for the whole array. Either the interpreter completes execution for the entire array, or for none of it. If the interpreter encounters an error during the processing of any cell within an array, it signals the error and halts execution of the function that encountered the error.

The interpreter follows the principle of all-or-none execution as far as it can, but there is no way it can do this rigorously when a function has ***side-effects*** (i.e. effects other than the explicit result returned). Most of the primitive functions in APL are free of side effects, and for them the all-or-none principle holds true.

User-defined functions are another matter: they may readily be written so that side effects occur. Similarly, side effects may accrue from using system functions that interact with structures outside the APL workspace, such as those involving files or shared variables. When the interpreter encounters an error while executing such a function, it still reports the error and halts. However, when you check the values of data in files, shared variables, or global variables, you may discover that the function may have executed for some cells but not for others. Responsible programmers take care to limit or control any side effects of their defined functions.

Deprecated Primitives: Braces

The left brace ({) and the right brace (}) are no longer used as primitives in SHARP APL for UNIX. Although braces may still be recognized in existing code, use of these characters as functions and operators is ***deprecated*** as of Version 6.0. The environment variable SAXBRACELOG can be used to locate and log the use of brace characters in APL expressions. Refer to the section "SAXBRACELOG" on page 1-18.

Functions that were represented by the left brace character ({), are now represented by the ***at*** sign (@). For more information see ***from***, documented on page 4-31, and ***all***, documented on page 4-78. Operators that were represented by the right brace character (}), are now represented by the ***ampersand*** (&). For more information see ***select***, documented on page 5-14, and ***merge***, documented on page 5-14.

SAXBRACELOG

If the environment variable SAXBRACELOG is set and is non-zero, SHARP APL for UNIX will write an entry in $SAXLOG each time one of the brace characters is found while parsing an APL expression. See the *Handbook, Chapter 3* for more information on environment variables.

An expression may not necessarily be reparsed after its first invocation in SHARP APL for UNIX Version 6.0. The log entry consists of a description of the event (braces found), the workspace and the top two elements of the stack. Stack elements are described by the function name or description of stack element type (immediate execution, $\square$*trap*, etc.) and, in the case of a function, the statement number in the function (the statement number may be different than the line number in a function; lines may contain multiple statements).

2
Naming

This chapter discusses how SHARP APL objects are named and how these names are chosen.

Names for User Objects

There are two main rules governing the choice of names for user variables and functions:

- A name must start with a letter of the alphabet or one of the characters ***delta*** (∆) or ***delta-underbar*** (⍙) .
- Subsequent characters in a name may be letters, ∆ or ⍙, numerals, or the underscore character, ***underbar*** (_). Since a name may not contain a blank, an underbar is sometimes used to form compound names such as `Net_Price`.

The symbols ***alpha*** (⍺) and ***omega*** (⍵) may also be used as names. These symbols are used *alone,* and not as part of a longer name. There are no rules regarding how you use the names ⍺ and ⍵, but by convention ***alpha*** is often used as a generic name for a function's left argument and ***omega*** for a function's right argument.

The mechanism for naming variables is discussed below. The mechanism for naming functions is discussed in the *System Guide, Chapter 8*.

Assigning Names to Variables

The ***left arrow*** (←) is used to assign names to variables. For example,

```
      x←analyze data
```

gives the name *x* to the result returned as the result of evaluating the expression *analyze data*. The left arrow plays the same role as the English copula "is"; an expression such as

```
      area←4×8
```

is often read as "Area is four times eight." Because the left arrow *assigns* a name to data, it is often called the ***assignment arrow***.

Once data has been named, further reference to the name is a reference to the variable's value; for example:

```
      area←4×8
      area
32
      area÷2
16
```

The Result of Assignment

If the name and left arrow occur leftmost in a statement, the value to the right is assigned to the name; but there is no other result and no display. However, when another function occurs to the left of the arrow in an expression such as

```
      3×x←analyze data
```

then the value returned by evaluating *analyze data* is passed through and becomes the right argument of ***times*** (×). Such a statement is read "Three times *x*, which is *analyze data*", and the resulting value is displayed.

The same rule applies to indexed assignment as well: if a function takes as its argument a name to which a new value is assigned, the argument is the value of the variable to the right of the arrow. For example, in the following statement, the resulting product is `3×a`:

```
      3×x[i]←a
```

Indirect Assignment

If the expression to the left of the assignment arrow is not a name, it is evaluated and the assignment is called ***indirect***. For the expression:

```
      (a) ← b
```

the result of *a* must be a whitespace delimited list of valid names or an enclosed array containing valid names. In both cases, if the number of names agrees with the frame of *b*, the opened elements of *b* are assigned to the corresponding names in *a*. For example:

```
      'abc def' ← 1 2
      abc
1
      def
2

      'ghi jkl' ← 1 ⊃ 2
      abc ≡ ghi ◇ def ≡ jkl
1
1
      'a b c'  ← 3
      a
3
      b
3
      c
3

      ∇z←name
[1]   z←'task',⍕''⍴⎕runs
[2]   ∇

      (name) ← ⎕ai
```

Multiple Assignment

A statement may contain several assignments in a row. For example,

```
x←y←z[i]←a
```

assigns the value of *a* to both the name *x* and the name *y*. It also replaces the value of the subarray `z[i]`. However, this does not make *x* and *y* synonymous: if you subsequently change the value of *x*, the value of *y* does not change.

Multiple Names for the Same Value

As a result of multiple assignment of the type illustrated above, or by a simple assignment of one name to another (such as `a←b`), two or more names may have the same value.

Similarly, when a user function is used so that its argument is a variable identified by its name (as in `analyze data`), there exist two names for the same array: the name that occurs in the calling statement (in the example, `data`) and the name internal to the definition of `analyze` by which `analyze` refers to its argument.

In such cases, the interpreter permits multiple names to point to the same object without making duplicate copies in memory. A new copy of the data is created only when one of the names is given a new value.

Reassignment and Name Conflicts

The name to the left of the left arrow must be *free for use*; that is, it must have no visible use as a function or as a label (see the *System Guide, Chapter 7*).

- When there has been no prior use, the left arrow creates a new variable having the name to the left of the arrow.
- When the name already has visible use as the name of a variable, the value to the right of the left arrow *replaces* the variable's former value.
- When the name has an existing visible use as a label or as the name of a function, the interpreter rejects the attempt to assign a new value to the name with the message `syntax error`.

Indexed Assignment

The name to the left of the left arrow may be modified by index brackets, in an expression of the form:

```
x[i]←a
```

The elements in array `a` are inserted at positions `[i]` within array `x`. In such a case, the following must be true:

- The array `x` must already exist.

- The expression within brackets must be appropriate to the rank of x (with the appropriate number of semicolons; see "Chapter 6. Syntax").
- The values of the indices must be valid for the shape of x.
- The shape of α must match the shape implied by the index expressions within the brackets (or be a single element, which will be inserted at all the locations indicated).

Obsolescence of Indexed Assignment

Indexed assignment does not appear in the *A Dictionary of APL.* It is considered obsolescent for reasons mentioned in the description of brackets and semicolons in "Chapter 6. Syntax". The effect of indexed assignment is to produce a new value for a variable by merging parts of an existing array with another array. This can be done in other ways more consistent with the rest of APL syntax. See the discussion of ***from*** ({) in "Chapter 4. Functions", and of ***merge*** (}) in "Chapter 5. Operators".

Localization of Names

A user-defined function may declare certain names to be *local* to itself. This means that the function (or any function it invokes) may make its own use of the names local to it, independent of any use those names may have outside the function. Localization has these effects:

- It assures the function's author that internal names needed by the function will be available to it (i.e. free for use).
- It protects names in the workspace from possible conflicting uses within the function's definition.
- It hides more global use of those names, if any, since from within a function (or any function it invokes), only the most local use of a name is visible.
- It cleans up after execution of a function. When execution of a function is complete, that ends its local use of names; the named objects it created vanish; and the interpreter frees any storage it had devoted to them.

A name is made local to a function by appearing in the function's *header.* The names used for the *arguments* and *result* of a user-defined function are automatically local to it. So are any *line labels* used in the function's definition.

When none of the currently active functions has localized a particular name, that name is *global* to the workspace. But when a name has been localized in a defined function currently being executed, that local use is understood. When a name is not local to the current function, but is local to a pendent function (a function whose execution has been started but not completed), the use local to the pendent function is understood.

A declaration of local use prevents seeing, changing, or creating a use of that name at a more global level. (It is the declaration that counts; it does not matter whether the localized name has actually been used.) Local use blocks the view of more global use. The most local use of a name is sometimes called the *visible* use of the name; a more global use is called a *shadowed* occurrence of the name.

3
Data

The fundamental data structure in SHARP APL is the rectangular array—an ordered collection of elements arranged along zero or more axes.

Rank and Shape

The elements within an array are arranged in a system of axes or Cartesian coordinates. The *rank* of an array is the number of axes it has. The simplest array has zero axes and is called a *scalar.* An array may have any number of axes; however, as a practical matter, the SHARP APL system permits no more than `127`.

A single scalar (say, the number `123.45` or the letter `'K'`) has rank *zero*; it has no axes. A *vector* (say, the numbers `12 ¯24.3 4.75` or the letters `'Mary Jones'`) has rank *one* because its elements are spread along a single axis. A *matrix,* made up of rows and columns, has rank *two*. And so on, for any number of axes.

Arrays are rectangular. The *shape* of an array is the number of positions along, or the length of, each axis. In a `3`-by-`4` matrix (having three rows and four columns) every row has the same number of elements (four), and every column has the same number of elements (three).

The axes of an array are identified by number. In the default APL system (that is, where `⎕io←1`), the first axis is by convention axis one, the next is axis two, and so on. In an environment where zero origin counting has been selected (that is, where `⎕io←0`), the first axis is axis zero, the next is axis one, and so on. See the discussion of index origin in *System Guide, Chapter 5*.

Reporting Shape and Rank

The ***shape*** function (⍴⍵) reports the ***length*** of each axis of its argument. Its result is a vector containing one number for each axis of the argument array. The value of each number indicates the length of the corresponding axis; that is, how many elements lie along that axis.

For example, for a matrix *m* with three rows and four columns, the ***shape*** function will return the *shape* of the matrix as follows:

```
      ⍴m
3 4
```

Applying the ***shape*** function to the result of ⍴*m* returns a single number indicating how many elements the shape of *m* contains:

Arrays, Cells, and Frames

An *array* may be considered to be a collection of *cells.* Each cell occupies some number of the array's trailing axes. The number of axes devoted to a cell is arbitrary, but they have to be at the *end* of the vector of axes.

Suppose you want each cell of an array to have *k* axes. Then the last *k* axes are said to be rank-*k* cells, or simply *k*-cells. For example, suppose the variable *x* is created as follows:

```
      x←2 3 4⍴'abcdefghijklmnopqrstuvwx'
      x
abcd
efgh
ijkl

mnop
qrst
uvwx
```

The vector `'abcd'` is a rank-`1` or `1`-cell of *x*; the two separate `3`-by-`4` matrices are `2`-cells of *x*; and each of the individual letters is a `0`-cell of *x*.

Whether you treat x as a single 3-cell, or a pair of 2-cells, or a 2-by-3 matrix of 1-cells, etc., does not depend on anything in x itself. How you look at x and partition its axes into cells depends on the function you apply to x. It also depends on the way the function is modified by the ***rank*** operator ($f\ddot{\circ}n$). (See "Chapter 5. Operators")

When you treat the last n axes as cells, the variable may have other axes left over. The axes that come *before* the n axes are said to be *frame* axes. The frame axes are the complement of the cell axes. For example, if the shape of b is `2 3 4`, then b has several possible frames:

Frame `2`	Relative to cells of rank `2`, each of which is a *matrix* of shape `3 4`.
Frame `2 3`	Relative to cells of rank `1`, each of which is a *vector* of shape `4`.
Frame `2 3 4`	Relative to cells of rank `0`, each of which is a *scalar*.
An empty frame	Relative to cells of rank `3`, the entire 2-by-3-by-4 array. Notice that the empty frame means that the entire array is treated whole, as a single cell.

Figure 3.1 illustrates the alternative ways of partitioning a 2-by-3-by-4 array into cells:

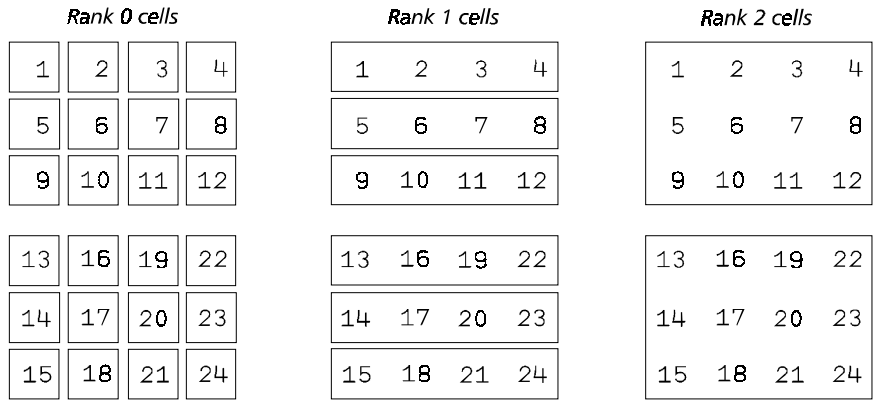

Figure 3.1. Partitioning an array into cells.

The number of cells that a frame contains is the product over the frame's shape. When one of the frame axes has length zero, the number of cells it contains is zero. Such a frame is said to be a *zero frame.* This is not the same as an empty frame:

- A zero frame is a frame that contains no cells.
- An empty frame contains one cell that is the entire array.

Complementary Partition

An array's cell rank is the complement of its frame rank. You can describe the partition by one of the following:

- the rank of each cell (in which case any leftover leading axes are part of the frame)
- the rank of the frame (in which case any leftover trailing axes are part of each cell)

To describe a partition by specifying the number of frame axes, use a *negative* number. The number of frame axes is the magnitude of that negative number; for example:

Rank `¯1`	The first axis is the frame. Each cell has whatever rank remains when the array's first axis is devoted to the frame and all the other axes are cell axes.
Rank `¯2`	The first two axes are the frame. Each cell has whatever rank remains.

For example, within the rank three array *x* mentioned earlier, the vector `'abcd'` may be called a `1`-cell of *x* or it may be called a `¯2`-cell of *x*. Similarly, the two `3`-by-`4` matrices within *x* may be called either `2`-cells or `¯1`-cells of *x*.

Major Cells

For any array, its `¯1` cells are its *major* cells. The number of major cells is the length of the array's first axis. A rank-`0` array (that is, a scalar, having no axes) has a single major cell: itself.

Infinite Rank

When a function has unbounded, or *infinite* rank, it treats its entire argument as a single cell, regardless of the number of axes the array may have. Several functions have infinite argument rank. All user-defined functions have infinite rank.

SHARP APL does not have a symbol for "infinity." When you need to specify that a function applies with infinite rank, any large number will do for the rank, provided it is not less than the array's actual rank.

Types of Array Elements

Elements within an array may be numbers, characters, or enclosed. These types are described in the sections that follow.

Numeric Elements

During input and output, a number is represented by digits, combined as needed with a decimal point (.), a high minus sign (¯), or the letter *e*.

A negative number has a leading high minus sign, as in `¯3.2` or `6.2e¯23`.

A number may be written in any of the following forms:

- ***integer***, for example, `2 ¯256`
- ***real***, for example, `2.1 0.009 ¯2.56`
- ***exponential***, for example, `1.6e9 6.2e¯23`.

In writing the various elements in a vector of numbers, it is permissible to intermix forms. For example, the following is perfectly valid:

```
x←1 ¯16 2.56 2.08e17
```

In similar fashion, when the APL interpreter displays a vector of numeric elements, it chooses the form independently for each element and thus may also intermix formats.

Although the interpreter permits mixed formats for entry or display, it stores all the members of an array of numbers in the same way. The effect of a statement such as the one just shown is to create an array all of whose elements are stored in the most demanding of the forms (see below).

Internal Representations of Numbers

SHARP APL uses three different internal forms for the storage of numbers. The forms differ in the space per element they require, so the choice of internal type may have important consequences for storage. The three types of representation are as follows:

Boolean Zero or one. Boolean representation is used for a single `0` or `1` typed from the keyboard, or for the result of a proposition, with `1` for true and `0` for false.

Storage: Bitwise, 0.125 bytes per element, plus overhead.

Integer A number with no fractional part. An integer is represented in `32` bits, permitting the representation of numbers from `¯2*31` to `¯1+2*31` (that is, `¯2147483648` to `2147483647`).

Storage: 4 Bytes per element, plus overhead.

Floating A number with a fractional part, or an integer outside the range representable in integer form. Also known as *real numbers.* Floating point numbers may range from `¯1.79e308` to `1.79e308`.

Storage: 8 bytes per element in IEEE (Institute of Electrical and Electronics Engineers) floating-point doubleword format, plus overhead.

In addition to the space required to store its elements, there is an additional overhead for system information, including the array's type, rank, and shape.

Within a numeric array, all elements have the same type of internal representation. For example, the vector `1 2` is stored entirely in integer representation, even though one of its elements is Boolean; similarly, the vector `2 2.3` is stored entirely as floating, even though one of its elements is an integer.

The interpreter initially assigns the most compact representation that will embrace all the elements in an array. Thereafter, upward conversion between types of numeric representation is automatic. For example, if you write

```
x←0 1 1 0
```

the interpreter stores *x* in Boolean representation. If you then catenate an integer to *x*, for example,

```
x←x,2
```

the interpreter converts the former value to integer representation before appending the integer `2`. Similarly, if you next catenate a floating-point number, for example,

```
x←x,1.9
```

the interpreter represents the entire array as floating point.

These conversions are automatic. No declaration of type is needed or allowed. There are no errors resulting from "mismatched" numeric types. The only practical consequences involve allocation of memory and perhaps execution speed.

However, downward conversion of type is not performed automatically. For example,

```
x←1 2 3-0 1 3
```

could be represented as Boolean; in fact, however, it will be represented as integer.

Character Data

To enter characters from the keyboard during immediate execution, you enclose them in quotes; for example, as `'ABC'` or `'Far away'` or `'3×$6'`. To represent the quote character itself, you type two consecutive quotes. (See the discussion of the quote delimiter in "Chapter 6. Syntax") However, in response to a request for input generated by ***quote-quad*** (⍞), whatever you type is accepted as characters. (See "Input and Output," later in this chapter.)

Character data has no significance as symbols, numbers, or names. However, the function ***execute*** (⍎⍵) instructs the interpreter to execute a vector of characters as an APL statement.

There is no implicit collating sequence for characters, so when you sort character data, you must provide a reference alphabet to control the order. See the discussion of ***character grade up*** (⍋) and ***character grade down*** (⍒) and in "Chapter 4. Functions"

APL Character Set

The SHARP APL system represents each character internally in one eight-bit byte. There are thus 256 possible characters. Not all of them are displayable or enterable from the keyboard.

Table 3.1. The APL character set.

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0		¯														
16	⌈	⊤	⌉	⊢	+	⊣	⌊	⊥	⌋	|	−					
32		!	"	#	$	%	&	'	(	)	×	+	,	−	.	/
48	0	1	2	3	4	5	6	7	8	9	:	;	<	=	>	?
64	@	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O
80	P	Q	R	S	T	U	V	W	X	Y	Z	[	\	]	^	_
96	`	a	b	c	d	e	f	g	h	i	j	k	l	m	n	o
112	p	q	r	s	t	u	v	w	x	y	z	{	¦	}	~	
128																
144																
160		¨	¯		≤		≥		≠	∨		×	÷	⍠	⍨	⍲
176	⍱	¡	€	£	¥		¬			⍣		⍥	⍂	≡	⍯	¿
192		⍺	⊥	∩	⌊	∈		∇	∆	⍳	∘		⍁	|	⊤	○
208			⍴	⌈		↓	∪	⍵	⊃	↑	⊂	⊢	⍀	⊣		÷
224	⌶	⊖	⍦	⍧		⍸	⍈	⍱	⍋	⍤	⍤		⍞		⍈	⍤
240	⍟				⍬		⍈	⍟			¢	←	⍙	→	◇	

† ***Control characters:*** 0=null; 8=backspace; 9=HT, horizontal tab; 10=LF, linefeed; 11=VT, vertical tab; 12=FF, formfeed, clear screen; 13=CR, carriage return. HT and FF are passed to the session manager, but not interpreted by APL. When passed to UNIX, either CR or LF produces a newline character.

The APL character set includes the standard printable characters of the ASCII (American National Standard Code for Information Interchange) character set, plus symbols specific to APL. It also includes characters used for the currency symbols ***dollar*** ($), ***pound*** (£), ***euro*** (€), and ***yen*** (¥), as well as graphic characters

used to construct draw lines and grids. Depending on the characteristics of your terminal, you may not be able to display all the characters in the system variable ⎕*av* (atomic vector), which is the list of all possible characters. The characters, and their positions within ⎕*av*, are shown in "Table 3.1. The APL character set." on page 3-8.

Enclosed Elements

An element within an array may be *enclosed.* An enclosed element is an element in the same sense as a single number or character: it has no axes and occupies a single position in the array's Cartesian framework.

However, an enclosed element may contain within it an array of any rank, shape, or type. It is thus a convenient way to place objects of arbitrary shape within the simple Cartesian framework of an array. To distinguish it from an array that contains enclosed elements, an array whose elements are not enclosed is said to be *simple, open* or *unenclosed.*

An enclosed array is created using the function ***enclose*** (<⍵). The result of enclose is always a scalar, and therefore has rank zero. For example:

```
      ⍴⍴x←<'Now is the time'
0
```

There are several ways to form an enclosed array (see "Chapter 4. Functions"). For example, the function ***link*** (⍺⊃⍵) joins a succession of open arrays to form a vector of enclosed elements. The link function joins its left and right arguments, first enclosing the left argument and enclosing the right only if it is not enclosed already.

```
      ⊢y←'Now'⊃'is'⊃'the'⊃'time'
┌───┐┌──┐┌───┐┌────┐
│Now│└is┘│the│└time┘
└───┘└──┘└───┘└────┘
      ⍴y
4
```

Figure 3.2 shows a 2-by-3-by-4 array of enclosed elements with two of them opened to display their contents.

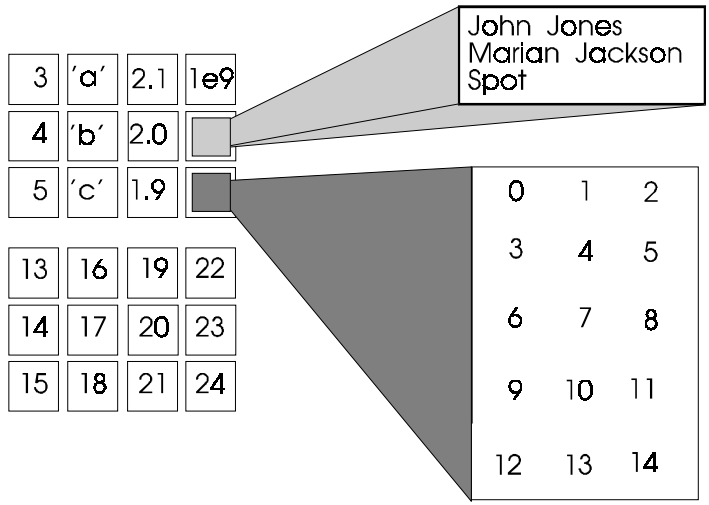

Figure 3.2. An array that contains two enclosed elements.

Heterogeneous Arrays

An array whose elements are all of the same type (all numeric, all character, or all enclosed) is said to be *homogeneous.* Homogeneous arrays offer some economies of storage. Moreover, since it may be complicated to write expressions which apply to an entire array when the elements within it are diverse, arranging data in a ***heterogeneous*** array may thereby sacrifice some of the economies of array processing that APL otherwise makes possible. However, heterogeneous arrays are particularly useful for displaying data, and for simplification of processing without regard for data type.

```
      ⊢x←1,'abc'
1 a b c
      ⍴x
4
      ⊢y←(<2 3 4),'abc' ⊣ □ps←2/-0 2
┌─────┐
│2 3 4│ a b c
└─────┘
      ⍴y
4
```

Packages

A *package* is a data type that in some respects resembles a miniature *workspace:*

- The elements within a package are selected by *name* rather than by position in a framework of axes.
- The names within a package refer to variables, user-defined functions, and bound intrinsic functions. A name in a package may even have *no* referent.
- A package is *not* an array. Hence, functions that are defined on arrays are *not* defined on packages.

A package is useful to provide *overlays*. These permit an application designer to materialize in a workspace—perhaps dynamically or within a local environment—whichever variables or functions the application requires, without altering the package from which they are taken. Often the source package resides not in the workspace at all but in a file outside the workspace.

A package is a data type independent of APL arrays, manipulable only by a set of specialized system functions (⎕*pack*, ⎕*pdef*, etc.). The contents of a package cannot be displayed directly. A statement whose result is a package produces as display only the message:

```
**package**
```

To see the definitions of the objects stored in a package, you have to extract the packed elements and then display them. The system functions for manipulating packages are listed below, with square brackets indicating optional arguments. For detailed descriptions of these functions, see the *System Guide, Chapter 5*.

`[⍺] ⎕pack ⍵`	Form a package.
`[⍺] ⎕pdef ⍵`	Define in the workspace objects from a package.
`⍺ ⎕pex ⍵`	Form a package by excluding some objects in a package.
`⍺ ⎕pins ⍵`	Form a package by merging two packages.
`[⍺] ⎕plock ⍵`	Lock the definitions of user-defined functions in a package.
`⎕pnames ⍵`	Return the names of objects in a package.
`[⍺] ⎕pnc ⍵`	Return the name-class of objects in a package.

`[α] ⎕ppdef ω`	Like `⎕pdef`, but protect existing objects by defining only those names that do not conflict with names outside the package.
`α ⎕psel ω`	Form a package by selecting a subset of the objects in a package.
`α ⎕pval ω`	Return the value of a variable in a package.

Input and Output

This section deals with the conventions for displaying data at the screen or printing terminal and with entering data from the keyboard. See the description of the function ***format*** (`⍕ω`) in "Chapter 4. Functions" and of the system function `⎕fmt` in the *System Guide, Chapter 5*. Also see the discussion of the "Session Manager" in the *System Guide, Chapter 7*.

Representation of Constants

A statement to be executed may contain characters that represent a number or a set of numbers. Similarly, a statement may contain characters that denote the characters themselves, as text. Such an entry is a *constant*: it states directly the value of a variable.

Constants may represent either numeric or character data. During immediate execution (or during an entry in response to ***quad*** (`⎕`)) you may type a numeric or character constant that is a scalar or a vector. Lines in the definition of a user function are constructed following the same rules.

An entry that consists of a single number or character is taken to be a *scalar.* An entry that consists of more than one number or character is a *vector.* For example, in the following, `s` is a scalar and `v` a vector:

```
s←126.2
v←126.3 14 127
```

Within a vector of numbers, the successive numbers are separated by blanks; extra blanks are not significant. You may write numbers within a vector in different formats: some using integer format without a decimal point, others using decimal format, yet others using exponential format as convenient or appropriate for the values you are writing. For example:

```
x←1 .3 0.0125 2.62e23 262e21
```

A negative number is preceded by the ***high minus*** symbol (¯). This symbol is ***not*** a function, but one of the characters used in writing certain numbers, just as the decimal point or the letter *e* occur in certain numbers. When several numbers are negative, each must have its own high minus. By contrast, the function ***minus*** (−⍵), like other functions, applies to an entire array. The expression

```
x←¯12 1.05 4000
```

gives the name *x* to a three element vector, only the first of which is negative, whereas

```
x←-12 1.05 4000
```

refers to a three element vector *all* of whose elements are made negative by the minus function.

To differentiate character from numeric data, you enclose a character scalar or vector in quotes. The beginning of a character constant is marked by a single quote mark. Everything you type after that (for the balance of the line) is treated as a character until you type a matching single quote to end the character string. Once you have started a quote, two consecutive quote marks are required to indicate the quote character itself. (See "Quotes Delimit a Character Vector" in "Chapter 6. Syntax")

In general, numbers are separated by blanks but characters are not. Each of the following creates and names a seven element vector:

```
x←25 ¯2.4 1e9 2.7e¯24 70.2 ¯2 0
 y←'Thanks!'
 z←'I can''t'
```

Within the zone enclosed by quotes, the blank is a character like any other. When the interpreter displays character data, it does not mark character data with quotes nor does it double the quote mark:

```
      z
I can't
```

A line containing an odd number of quotes (excluding quotes that occur within a comment) is invalid. The interpreter cannot tell which characters are inside and which are outside the quotation, and rejects such a line with a *syntax error* without executing any of it. If you write such a line with the function editor, the interpreter proceeds to establish the definition that contains it; you encounter the *syntax error* from the unmatched quote only when the interpreter attempts to execute the faulty line.

Higher-Rank Arrays

A constant entered from the keyboard is limited to a single scalar or a vector. To restructure the value into a matrix (or any array with more than one axis), you enter it as a vector and then apply a function such as ***reshape*** (⍺⍴⍵) to organize the elements into a higher-rank array; for example:

```
      t←3 4⍴'The fat cat.'
      t
The
fat
cat.
```

Heterogeneous or Enclosed Arrays

You can enter heterogeneous or enclosed data by including in the expression functions that will join homogeneous arrays to produce a heterogeneous one, or enclose open arrays to produce enclosed arrays. In the following example, the comma serves to join the four characters and three numbers into a seven-element, heterogeneous vector:

```
      m←'Paul',1933 8 24
```

In the following, the ***link*** function (⍺⊃⍵) encloses and joins the given vectors to form a three-element vector of boxes:

```
      n←'Paul'⊃'Hempstead'⊃1933 8 24
```

Display of Arrays

Unless the result produced by evaluating a statement is assigned a name, the interpreter displays the result. This is called *default display*. It occurs with no need for a "display" function and requires no explicit instructions regarding format. In the absence of other instructions about formatting, the interpreter formats each variable in the same way as for monadic use of the ***format*** function (⍕⍵); that is:

- display of numbers is controlled by the value of the system variable ⎕*pp* (printing precision)
- the position and spacing of heterogeneous or enclosed arrays is controlled by the value of the system variable ⎕*ps* (position and spacing).

The system displays a numeric vector horizontally, with a blank inserted to separate one element from the next. However, it does not put blanks between the elements of a character array. That way, characters can be placed side by side to form words, but distinct numbers (which often require digits) do not run together. For example:

```
      ⍳5
0 1 2 3 4
      h←'Hello!'
      h
Hello!
```

Position and Spacing in an Array

During display of heterogeneous and enclosed arrays, the system variable ⎕*ps* affects the position and spacing of the arrays. In a clear workspace, the value of ⎕*ps* is `¯1 ¯1 0 1`. (See the description of ⎕*ps* in the *System Guide, Chapter 5*.)

The first two elements of ⎕*ps* specify how an undersized element may be assigned to a display area with more rows or columns than the element would normally need. The default values of `¯1 ¯1` indicate the top left corner, with any extra space below and to the right.

The last two elements of ⎕*ps* control the amount of additional vertical and horizontal space that should separate enclosed or character elements in a heterogeneous array: the third element controls vertical space and the fourth controls horizontal space. The default values are zero lines vertically and one column horizontally.

Consider the horizontal spacing of character elements in the following arrays:

```
      h,⍳5
H e l l o ! 0 1 2 3 4
      h⊃h⊃⍳5
Hello! Hello! 0 1 2 3 4
```

The details of numeric formats and the arrangement of heterogeneous or enclosed elements in a grid for displaying arrays of rank two or higher are described in the section on ***format*** (⍕⍵) in "Chapter 4. Functions".

Separating Successive Subarrays

When it displays arrays with rank greater than two, the session manager (see the *System Guide, Chapter 7*) moves to a new position as follows:

On last axis *(each column)*	Move to the right (discussed below under "Preserving Columns").
On 2nd last axis (each row)	Start at the left margin one line down (single spaced between vectors).
On 3rd last axis (each matrix)	Start at the left margin two lines down (double space between matrices).
On *n*th last axis last	Start at the left margin ***n-1*** lines down.

Such additional blank lines are inserted by the session manager only when it generates a display; they do not appear as explicit newline characters in the result returned by the ***format*** function (⍕⍵) (although such a result, when displayed by the session manager, ***will*** include the additional blank lines).

Scissoring Wide Displays

An array returned explicitly as the result of the functions ***format*** (⍕⍵) or ⎕*fmt* is arranged without regard for the width available to display it. However, when you have output displayed on the screen or terminal printer, the session manager takes into account the length of the line available for display. The maximum number of print positions on a line is the value of the system variable ⎕*pw*

(printing width). By default, □*pw* is `80`. You may set it as low as `30` or as high as `255`; or you can indicate an infinite width by setting □*pw* to `0`. (For details, see the *System Guide, Chapter 5*).

The effect of □*pw* depends on the rank and type of the array being displayed.

Numeric vectors

In displaying a numeric vector, the system prints on each line as many numbers as will fit without splitting apart the digits that represent a single number. After it breaks a line that will not fit within □*pw*, the session manager continues with the next number on the line(s) below, indented by six positions.

Character vectors

In displaying a character vector, the system obeys newline characters (`10@`□*av* or `13@`□*av*) as they occur in the vector being displayed. Following each such explicit newline, it starts the next line at the left margin.

When the system finds it has already put □*pw* characters on a line and the next character is not a newline character, it forces a new line and indents the line that follows by six spaces. The system decides where to break a character vector solely on the basis of position on the screen, with no attempt to break words at blanks, etc.

Scissoring a Higher-Rank Array

When a matrix or higher-rank array does not fit within □*pw*, the session manager divides the image vertically. It first prints the amount that fits within □*pw* for all rows. When it has completed those, it moves to the n^{th} blank line below the preceding portion of display (where n is the rank of the array being printed), and then prints the next group of columns for all rows, with every line indented by six blanks, and so on until the display is complete.

Suppose that without □*pw*, a three row matrix could be printed as follows:

```
                    ▼
a1a2a3a4a5a6a7a8a9aAaBaCaDaEaF
b1b2b3b4b5b6b7b8b9bBbBbCbDbEbF
c1c2c3c4c5c6c7c8c9cCcBcCcDcEcF
```

When □*pw* is set to a value that forces a break at the position marked by▼, the array is broken into segments of consecutive columns, *as if cut by scissors*; the pieces are indented and displayed as a whole:

```
a1a2a3a4a5a6a7
b1b2b3b4b5b6b7
c1c2c3c4c5c6c7

      a8a9aAaB
      b8b9bBbB
      c8c9cCcB

      aCaDaEaF
      bCbDbEbF
      cCcDcEcF
```

Character matrices

When the array being displayed is a homogeneous array of characters, the session manager prints □*pw* characters for all rows in the first stage of display. Then it moves down to the n^{th} line and prints six blanks and the next □*pw*−6 characters for all rows. It repeats this process until all columns have been displayed.

Numeric matrices

The procedure is the same as that for character matrices, except that for each vertical segment, the session manager chooses a width that does not exceed □*pw* and does not split the display of a number.

Heterogeneous arrays

The procedure is generally the same as for numeric arrays, except that any row or column containing a character is separated from an adjacent row or column by a number of blanks specified in the last two elements of □*ps*.

Arrays Containing Enclosed Elements

The procedure is generally the same as for numeric and heterogeneous arrays, except that when the last two elements of □*ps* are negative, the enclosed elements may also be surrounded by characters to represent the enclosure. The session manager may find it impossible to apply to boxes the "no-split" rule that it uses for numbers.

Folding Wide Displays

If scissoring has been turned off, via the start-up option `-Yscissor`, the session manager displays each complete row of the array in sequence.

Folding Character Arrays, Heterogeneous Arrays, and Arrays of Enclosed Elements

In displaying a homogeneous character array, a heterogeneous array, or an array of enclosed elements, the session manager prints no more than □*pw* characters on a line. Once it has put □*pw* characters on the current line and the next character is something other that newline (`10@`□*av* or `13@`□*av*), it forces a new line. It indents six spaces and continues printing the current row of the array. The session manager performs this procedure for each row of the array, forcing a new line at the end of each row.

The session manager obeys newline characters as they appear in the array. That is, the next character is displayed at the left margin of the line below, and the length of the current line is counted from there.

Folding a Numeric Vector

The session manager displays a numeric vector such that each line contains as many numbers as will fit without separating the individual digits of a number.

When a line that does not fit within □*pw* is broken, the next line, and any subsequent lines that are continuations of the original line, are indented by six positions.

Folding a Numeric Table or Higher-rank Array

A numeric matrix, or a numeric array of any rank higher than one, is arranged so that all columns have the same width. The session manager figures out the number of print positions that would be required if all the elements were in a single column. It uses the field width thus calculated for all columns of the array (regardless of the actual widths of numbers in a particular column).

Using the field width just calculated, it starts at the beginning of each row and writes as many fields as will fit within ⎕*pw*. Then it moves to the next line, indents six positions, and prints as many of the remaining fields as will fit within ⎕*pw*−6. It continues in that fashion until the entire row has been printed, then prints the next row.

Folding an Array of Enclosed Elements

The display of an array of enclosed elements is treated in the same way as character data. Each line is broken after ⎕*pw* characters or at a newline character (whichever comes first), and the continuation is indented by six positions. The session manager makes no attempt to preserve the boundaries of enclosed elements when ⎕*pw* requires folding.

Preserving Columns and Integrity of Elements

In arranging its displays, the session manager applies two rules:

Preserve the alignment of columns

Alignment in columns refers to positions along the last axis of an array. When the element in a column is a number or enclosed, its display may require several characters. To distinguish columns of the data from position on the display, we call the latter *print positions.* The session manager assigns to each column of data sufficient print positions to accommodate the column's widest display.

Preserve the integrity of elements

In deciding where to break a line, the session manager attempts never to split the display of an element. If possible, it breaks the display only between columns. For numeric data, the minimum value of ⎕*pw* is sufficiently high, and the maximum value of ⎕*pp* sufficiently low such that a break between columns is always possible. Enclosed elements are more difficult. Where possible, the

session manager keeps them intact. But since an enclosed element may be of any size, it becomes difficult or impossible to find break points that keep all of them intact. When the width of an enclosed element is too large, the session manager abandons the effort to keep it together and cuts it arbitrarily, as it would a character matrix.

Empty Arrays

An empty vector appears as a blank line—a line on which nothing is printed. An empty matrix or an empty array of any rank greater than one, produces no display at all (not even a blank line). This is true even for an array which has a positive number of rows but a zero as a length of any of its other axes.

Some people write expressions such as `0 0⍴foo x` as a way of discarding the result returned by foo (presumably because they nevertheless want some side effect that foo produces). The function stop (⍝) does the same thing—it suppresses the display of a result from the expression to its right (see "Pass and Stop" in "Chapter 4. Functions").

4
Functions

This chapter presents all the primitive functions in SHARP APL for UNIX together with a few examples to illustrate their use. The descriptions often distinguish between the name for a function's symbol and the name for the monadic and dyadic use of the function. Such distinctions reflect a wide variety of names in common use across an extensive range of disciplines.

Categories of Functions

Despite the abundance of primitive functions in SHARP APL for UNIX, many have properties in common. By grouping together primitives with common attributes, several categories emerge. While many alternative groupings are possible, this guide uses the classification scheme given below:

- arithmetic functions, scalar
- arithmetic functions, non-scalar
- relational functions
- indexing functions
- structural functions
- miscellaneous functions

Default Argument Ranks

When a function is applied to an argument array of rank greater than that for which it is defined, the extra leading axes of the argument array are treated as *frame* axes and the function is applied *independently to each cell* within that frame. A function's argument rank is the number of axes devoted to an argument cell.

The ***rank*** operator (*f*⍤*n*—see "Chapter 5. Operators") can modify any function *f* to impose the argument rank specified by *n*. For most functions (but not all) there is also a *default argument rank*: the rank of an argument cell when the function is used without the ***rank*** operator. The default argument rank is shown for each function in this chapter.

In some cases, a function may have *infinite* argument rank: no matter how many axes the argument has, they are all treated as part of a single cell. A function with infinite default rank is indicated in the descriptions that follow by an ***infinity-sign*** (∞). A function with *undecided* argument rank is indicated by an asterisk (*).

Identity Elements

The derived functions called *reductions* (see "Reduce and Scan" in "Chapter 5. Operators") apply the same function between each of the major cells of their right arguments. For example, if an array *x* consists of three major cells `x0`, `x1`, and `x2`, then the *f*-reduction of *x*, written *f*⌿*x*, is computed by:

```
x0 f x1 f x2
```

To preserve identities when *x* is partitioned, it is useful to define reduction when the argument has zero cells. For that case, the result of the derived function is the function's identity element, a value which, when used as one argument, produces a result identical to the function's other argument. For example, the identity element for ***plus*** (⍺+⍵) is ***zero***: when added to any argument, zero does not change the result. Not all primitive functions have identity elements. In the following sections, the value of a function's identity element is noted where there is one.

Table 4.1. Ranks of primitive functions.

Function	Monadic	Dyadic	Function	Monadic	Dyadic
+	0	0 0	≡		∞ ∞
−	0	0 0	∈		0 ∞
×	0	0 0	≦		∞ ∞
÷	0	0 0	@	1	0 ∞
|	0	0 0	⍒	∞	∞ ∞
!	0	0 0	⍗	2	∞ 2
?	0	* *	⍴	∞	1 ∞
⌈	0	0 0	↑	∞	1 ∞
⌊	0	0 0	↓	∞	1 ∞
○	0	0 0	⊃	∞	∞ ∞
∧		0 0	⍪	∞	∞ ∞
∨		0 0	,		
⍲		0 0	⊣	∞	∞ ∞
⍱		0 0	⊢	∞	∞ ∞
~	0	∞ ∞	⍉	1	0 1
=		0 0	⊖	∞	∞ ∞
≠		0 0	⍉	∞	1 ∞
<		0 0	⊥		∞ ∞
≤		0 0	⊤		∞ ∞
≥		0 0	⍎	*	
>	0	0 0	⍕		* ∞
*	0	0 0	⍳	1	1 0
⍟	0	0 0	⍸		∞ ∞
⍲		∞ ∞			

∞ *Infinite rank*
* *Undecided rank*

Result Rank

Wherever the rank or shape of a result cell is different from that of the argument cell or cells, the descriptions note that fact. Otherwise, each result cell's rank and shape is the same as the argument cell's rank and shape.

Scalar Functions

A scalar function is a function that is defined on rank zero cells as arguments and returns rank zero cells as its result. Because its default operation is defined for scalars, applying a scalar function to an array causes it to treat all axes in the array as frame axes, and therefore to apply independently to each of the scalars throughout the frame.

When a scalar function is used dyadically, the following rules apply: either the two frames must have the same shape, or one of them must contain only a single element, to be paired with every cell in the other argument. For example, adding the numbers contained in frames of length three produces three parallel additions of corresponding cells:

```
     1 2 3+12 5 20
13 7 23
```

That is because `1+12` is `13`, `2+5` is `7`, and `3+20` is `23`.

Extending the Empty Frame

When one frame is empty, the single cell corresponding to that empty frame is paired with each cell in the other argument; the empty frame is said to be *extended* to match the frame of the other argument. The result frame has the same shape as the nonempty argument frame. This extension is illustrated in the following examples for addition:

```
     100+¯1 10 20
99 110 120
     ¯1 10 20+100
99 110 120
```

Singletons Sometimes Treated as Scalars

A ***singleton*** is an array containing exactly one element, and having one or more axes of length one. All primitive functions having either a monadic rank of zero, or a dyadic left and right rank of zero, treat singletons much like scalars. There are two cases to consider:

When ***one*** argument is a singleton and the other is not, the singleton is treated ***exactly*** like a scalar:

```
      x←1 2 3 4 5×1 1 1⍴8
      x
8 16 24 32 40
      ⍴x
5
```

When ***both*** arguments are singletons they are treated like scalars except that the rank of the result is the ***higher*** of the ranks of the two singleton arguments:

```
      x←(1 1 1⍴8)×1 1⍴2
      x
16
       ⍴x
1 1 1
```

Arithmetic Functions, Scalar

All the functions covered in this section have in common the following properties:

- They are defined only for unenclosed numeric arguments.
- They return unenclosed numeric results.
- The monadic functions have a rank of zero.
- The dyadic functions have both left and right ranks of zero.

+ω *Conjugate/Identity*

Rank: 0

The expression +ω gives the *conjugate* of ω, defined as `(|ω×ω)÷ω` (assuming that `0÷0` is `0`). The result is different from ω only where ω is complex. Since SHARP APL for UNIX does not at present support complex numbers, +ω is the same as ω, hence the alternative name ***identity***.

α+ω *Plus/Addition*

Ranks: 0 0; Identity Element: 0

The expression α+ω is defined as the arithmetic sum of α and ω.

−ω *Minus/Negate*

Rank: 0

The expression −ω is defined as `0`−ω, that is, the sign of every element in ω is reversed.

α−ω *Minus/Subtraction*

Ranks: 0 0; Identity Element: 0

The expression α−ω is defined as the arithmetic difference between α and ω.

×ω *Signum*

Rank: 0

The expression ×ω gives the *signum* of ω, defined as `ω÷|ω` (assuming that `0÷0` is `0`). Thus, the result ×ω is:

`1` where ω is positive
`0` where ω is 0
`¯1` where ω is negative

α×ω *Times/Multiplication*

Ranks: 0 0; Identity Element: 1

The expression α×ω is defined as the arithmetic product of α and ω.

÷ω *Reciprocal*

Rank: 0

The expression ÷ω is defined as 1÷ω, that is, the *reciprocal* of ω.

α÷ω *Divide/Division*

Ranks: 0 0; Identity Element: 1

The expression α÷ω is defined as the arithmetic quotient of α and ω. As in arithmetic, division by zero is undefined and is rejected as a *domain error*. However, 0÷0 ***is*** defined in SHARP APL to be ***one***.

⋆ω *Exponential*

Rank: 0

The exponential denoted by ⋆ω is equivalent to *e*⋆ω, where *e* is the base of the natural logarithms, given approximately by:

```
     *1
2.718281828
```

α⋆ω *Power*

Ranks: 0 0; Identity Element: 1

The expressions ω⋆2 and ω⋆3 and ω⋆.5 return, respectively, the square, cube, and square root of ω. The general definition of α⋆ω is ⋆ω×⊗α. For the simple case of an *integer right argument,* it is equivalent to ×/ωρα; in particular, ×/ applied to an empty vector yields 1, and α⋆0 is 1 for any α including the case where α is 0. The expression α⋆ω is often read as α *to the power* ω.

⊗ω *Natural Logarithm*

Rank: 0

The *natural logarithm* is inverse to ***exponential*** (⋆ω):

$$\omega \Longleftrightarrow \otimes\star\omega \Longleftrightarrow \star\otimes\omega \quad (\text{for } \omega>0)$$

Moreover,

$$\otimes\omega \Longleftrightarrow e\otimes\omega$$

where *e* is the base of the natural logarithms.

A *domain error* is signalled when ω≤0.

α⊗ω *Base-*α *Logarithm*

Ranks: 0 0; Identity Element: None

The ***base-***α ***logarithm*** (α⊗ω) is ***inverse*** to ***power,*** in the sense that

$$\omega \Longleftrightarrow \alpha\otimes\alpha\star\omega \Longleftrightarrow \alpha\star\alpha\otimes\omega \quad (\text{for } \omega>0)$$

A *domain error* is signalled when ω≤0.

|⍵ Magnitude/Absolute Value

Rank: 0

The result of |⍵ is defined as (⍵*2)*.5 giving the *magnitude* or *absolute value* of ⍵.

Residue/Remainder

Ranks: 0 0; Identity Element: 0; Implicit Arguments: ⎕ct

The most familiar use of the ***residue*** function (|⍵) is to determine the remainder resulting from dividing a nonnegative integer by a positive integer. For example:

```
     3|0 1 2 3 4 5 6 7 8
0 1 2 0 1 2 0 1 2
```

The definition

```
      α|⍵ ⇐⇒ ⍵-α×⌊⍵÷α+0=α
```

extends this notion to a zero left argument (in which case it returns the right argument unchanged), to non-integer right arguments, and to negative and fractional left arguments. When α is a negative integer, the result ranges between α and 0, just as it does when α is positive. For example:

```
      ¯3|¯4 ¯3 ¯2 ¯1 0 1 2 3 4
¯1 0 ¯2 ¯1 0 ¯2 ¯1 0 ¯2
```

When ⍵ contains non-integer values:

```
     1|2.5 3.64 2 ¯1.6
0.5 0.64 0 0.4
```

However, for cases such as (÷3)|2÷3, in order to produce a true zero (rather than a small fraction) the residue is made tolerant by the following definition:

Let $S \leftarrow \omega \div \alpha + \alpha = 0$

If $(\alpha \neq 0) \wedge (\lceil S) \neq \lfloor S$

then $\alpha | \omega \Leftrightarrow \omega - \alpha \times \lfloor S$

else $\alpha | \omega \Leftrightarrow \omega \times \alpha = 0$

For example:

```
      .1|2.5 3.64 2 ¯1.6
0 0.04 0 0
```

For α=0, the result of α|ω is ω unchanged.

Figure 4.1 illustrates the results obtained when the modulus is 3 or ¯3 and the right argument is 10 or ¯8:

```
      3 ¯3∘.|10 ¯8
 1  1
¯2 ¯2
```

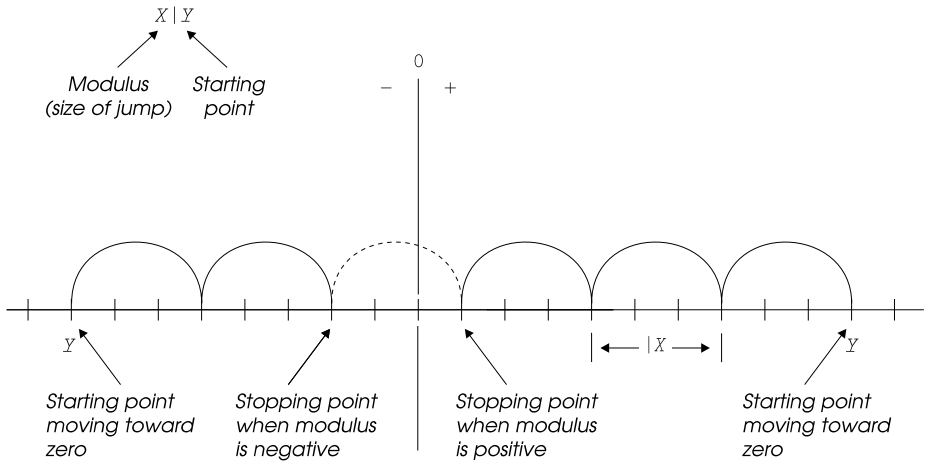

Figure 4.1. Schematic representation of residue.

!ω *Factorial*

Rank: 0

For a nonnegative integer argument, the ***factorial*** is defined by:

```
      !ω ⇐⇒ ×/1+⍳ω
```

For example:

```
      !4
24
      1+⍳4
1 2 3 4
      ×/1+⍳4
24
      !0
1
```

For an argument other than a nonnegative integer, `!⍵` is defined in terms of the *gamma* function, represented here by the (non-APL) symbol Γ:

```
!⍵ ⇐⇒ Γ (⍵+1)
```

α!ω *Out-Of/Combinations*

Ranks: 0 0; Identity Element: 1; Implicit Arguments: `⎕ct`

The dyadic function ***out-of*** is defined in terms of the *beta* function; subject to the interpretation given below, this definition is equivalent to

```
α!⍵ ⇐⇒ (!⍵)÷(!α)×(!⍵-α)
```

For nonnegative integer arguments, `α!⍵` yields the number of distinct ways of selecting `α` things from `⍵` things; this accounts for the function's name, and for its use in producing binomial coefficients:

```
      (⍳n+1)!n←3
1 3 3 1
```

For a negative integer i, the expression `!i` is not defined, because near a negative integer the magnitude of `!i` approaches infinity. Nevertheless, the definition of `α!⍵` can be understood by assuming that these infinite values occur in the expression for the dyadic definition in the following ways:

- When `α` and `⍵` are both nonnegative, but `α>⍵`, the term `!⍵-α` is infinite, yielding `0` for the result of `α!⍵`. This agrees with the notion that `α` things can be picked from a smaller collection in `0` ways.
- When infinite values occur in both numerator and denominator of the defining expression, they are assumed to cancel. This can be seen in the values in the following function matrix:

```
      i←¯3 ¯2 ¯1 0 1 2 3
      i∘.!i
  1  ¯2   1   0   0   0   0
  0   1  ¯1   0   0   0   0
  0   0   1   0   0   0   0
  1   1   1   1   1   1   1
 ¯3  ¯2  ¯1   0   1   2   3
  6   3   1   0   0   1   3
¯10  ¯4  ¯1   0   0   0   1
```

⌊ω *Floor*

Rank: 0; Implicit Arguments: ⎕ct

The expression ⌊ω returns the ***floor*** or *integer part* of ω. When ω is an integer, the result is the same integer, but when ω has a fractional part, the result is the next smaller integer. Thus, ⌊ω is the largest integer which is not greater than ω:

```
      ⌊3.1
3
```

α⌊ω *Minimum*

Ranks: 0 0; Identity Element: ca. 1.8e308

The expression α⌊ω returns the ***minimum*** of α and ω (whichever is smaller):

```
      3⌊4 ¯4
3 ¯4
```

This function is ***not*** subject to the value of ⎕ct (comparison tolerance).

⌈ω *Ceiling*

Rank: 0; Implicit Arguments: ⎕ct

The expression ⌈ω returns the ***ceiling*** of ω. When ω is an integer, the result is the same integer, but when ω has a fractional part, the result is the next larger integer. Thus, ⌈ω is the smallest integer which is not smaller than ω:

```
      ⌈3.1
4
```

The algorithm for ⌈ω can be defined in terms of ⌊ω:

```
      ⌈ω ⇐⇒ ⌊¨- ω ⇐⇒ -⌊-ω
```

Fuzzy Floor and Ceiling

The implied comparison of ω with the integers is *tolerant*. When ω is *sufficiently close* to an integer, both ⌈ω and ⌊ω are that integer, where *sufficiently close* is defined in terms of the value of the system variable ⎕ct (comparison tolerance):

- ⌊ω is the *largest integer* not greater than ω+⎕ct
- ⌈ω is the *smallest integer* not less than ω−⎕ct

For example, for the default value of ⎕ct:

```
      ⌈2+10*¯14
2
```

α⌈ω Maximum

Ranks: 0 0; Identity Element: ca. ¯1.8e308

The expression α⌈ω returns the maximum α and ω (whichever is greater):

```
      3⌈4 ¯4
4 3
```

This function is ***not*** subject to the value of ⎕ct (comparison tolerance).

○ω Pi Times

Rank: 0

The expression ○ω is equivalent to ω×π and is given approximately by:

```
      ○1
3.141592654
```

⍺○⍵ *Circle Functions*

Ranks: 0 0; Identity Element: None

In the expression ⍺○⍵, the left argument ⍺ must be an *integer* from the set ¯7 through 7, inclusive. The value of ⍺ selects one of the families of circular functions described in Tables 4.3 and 4.4, and illustrated in Figures 4.2 and 4.3.

The circle functions expect all angular arguments—and return all angular results—in ***radians, not degrees:*** 1 radian = 360÷○2 degrees (about 57.2958)

Table 4.2. Families of circular functions.

⍺∈ 1 2 3	Trigonometric
⍺∈¯1 ¯2 ¯3	Inverse Trigonometric
⍺∈ 5 6 7	Hyperbolic
⍺∈¯5 ¯6 ¯7	Inverse Hyperbolic
⍺∈ 0 4 ¯4	Pythagorean

Table 4.3. Individual circle functions.

0○⍵	(1−⍵*2)*0.5		
1○⍵	sine ⍵	¯1○⍵	arcsine ⍵
2○⍵	cosine ⍵	¯2○⍵	arccosine ⍵
3○⍵	tangent ⍵	¯3○⍵	arctangent ⍵
4○⍵	(1+⍵*2)*0.5	¯4○⍵	⍵×(1−⍵*¯2)*0.5
5○⍵	hyperbolic sine ⍵	¯5○⍵	hyperbolic arcsine ⍵
6○⍵	hyperbolic cosine ⍵	¯6○⍵	hyperbolic arccosine ⍵
7○⍵	hyperbolic tangent ⍵	¯7○⍵	hyperbolic arctangent ⍵

	PR: *sine A*
OP: 1	*OR*: *cosine A*
PR: 0○*OR*	*PT*: *tangent A*
OR: 0○*PR*	*PS*: *cotangent A*
OT: 4○*PT*	*OT*: *secant A*
PT: ¯4○*OT*	*OS*: *cosecant A*

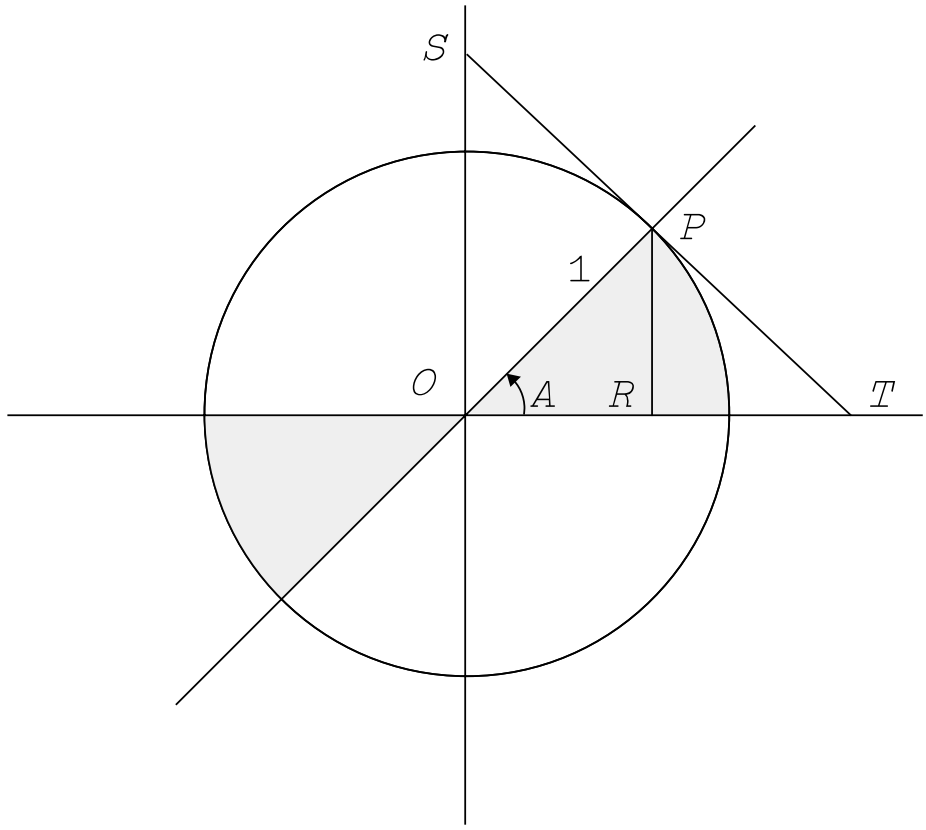

Figure 4.2. Pythagorean and circular functions in terms of the unit circle.

RP: 1○A	sine	TQ: 5○B	hyperbolic sine
OR: 2○A	cosine	OT: 6○B	hyperbolic cosine
PT: 3○A	tangent	VW: 7○B	hyperbolic tangent
OS: ÷1○A	cosecant	PS: ÷5○B	hyperbolic cosecant
OT: ÷2○A	secant	OR: ÷6○B	hyperbolic secant
PS: ÷3○A	cotangent	UV: ÷7○B	hyperbolic cotangent

(where *B* is ¯5○3○*A*)

.

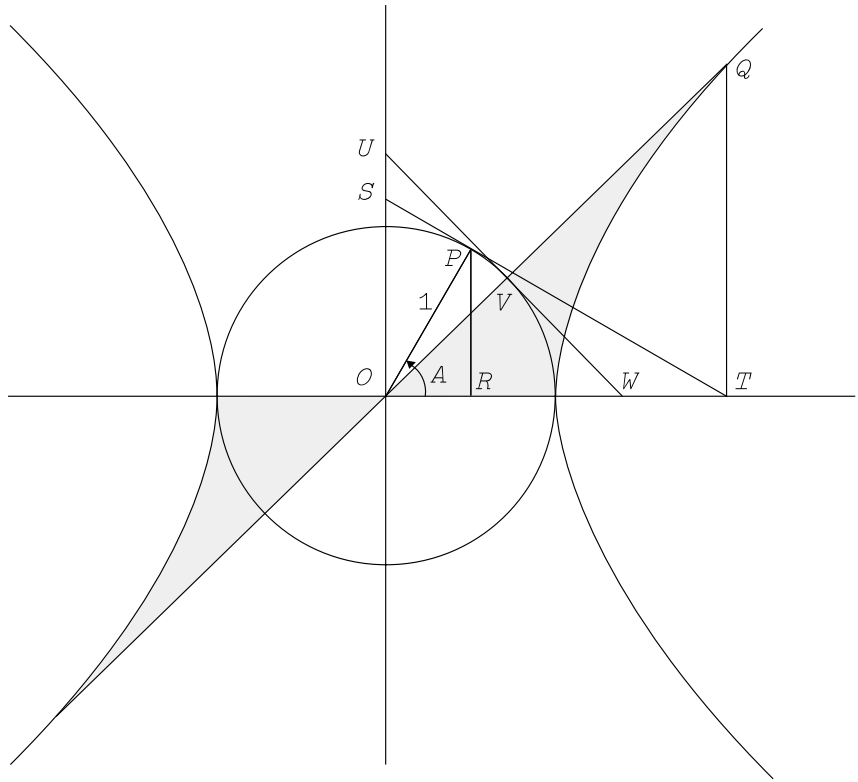

Figure 4.3. Hyperbolic functions in terms of the unit hyperbola.

α∧ω *And/Least Common Multiple (LCM)*

Ranks: 0 0; Identity Element: 1; Implicit Arguments: ⎕*ct*

For Boolean arguments, the result of α∧ω is the logical ***and*** of α and ω:

```
     p←0 1
     p∘.∧p
0 0
0 1
```

Logical ***and*** is related to logical ***or*** by the following identity:

```
α∧ω ⇐⇒ ~(~α)∨(~ω) ⇐⇒ α∨⍨~ω
```

More generally, for arbitrary numeric arguments, the expression α∧ω returns the ***least common multiple (LCM)*** of α and ω:

```
     12∧30
60
      3.6∧4
36
```

α∨ω *Or/Greatest Common Divisor (GCD)*

Ranks: 0 0; Identity Element: 0; Implicit Arguments: ⎕*ct*

For Boolean arguments, the result of α∨ω is the logical ***or*** of α and ω:

```
     p←0 1
     p∘.∨p
0 1
1 1
```

Logical ***or*** is related to logical ***and*** by the following identity:

```
α∨ω ⇐⇒ ~(~α)∧(~ω) ⇐⇒ α∧⍨~ω
```

More generally, for arbitrary numeric arguments, the expression α∨ω returns the ***greatest common divisor (GCD)*** of α and ω:

```
      12∨30
6
       3.6∨4
0.4
```

Fuzzy GCD and LCM

For non-Boolean arguments, α∧ω and α∨ω are subject to the value of ⎕⍎⍵t (comparison tolerance). Fuzzy computation permits "reasonable" results with fractions that cannot be precisely represented. For example, the least common multiple of `0.6` and `1.4` is `4.2`. This result is returned correctly with the default value of ⎕⍎⍵t. But if ⎕⍎⍵t is set to `0`, the interpreter fails to detect that each is an integer multiple of `0.2`, and hence returns an LCM as if they had no common factor.

α⍲ω *Nand*

Ranks: 0 0; Identity Element: None

The expression α⍲ω is defined only on Boolean arguments, and returns the logical ***nand*** of α and ω.

```
      p←0 1
      p∘.⍲p
1 1
1 0
```

Logical ***nand*** is related to logical ***and*** by the following identity:

```
      α⍲ω ⇐⇒ ~α∧ω
```

α⍱ω *Nor*

Ranks: 0 0; Identity Element: None

The expression α⍱ω is defined only on Boolean arguments, and returns the logical ***nor*** of α and ω.

```
      p←0 1
      p∘.⍱p
1 0
0 0
```

Logical ***nor*** is related to logical ***or*** by the following identity:

```
      α⍱ω ⇐⇒ ~α∨ω
```

~ω Not

Rank: 0

The expression ~ω is defined on Boolean arguments only and returns the logical ***not*** or logical negation of ω.

```
      p←0 1
      ~p
1 0

10
      ⍴10~20
1
```

?ω Roll

Rank: 0; Implicit Arguments: ⎕io ⎕rl

The function ***roll*** is named from the analogy with rolling a die to choose one of a set of numbers with equal probability. For example:

```
      ?6
4
```

The size of the population from which numbers are selected is determined by the value of the right argument. The population is the first ω integers; that is, a result cell is selected from ⍳ω. The population is therefore affected by the value of ⎕io (index origin).

See also "Pseudorandom Numbers Depend on Random Link" on page 4-23.

Arithmetic Functions, Non-scalar

All the functions covered in this section have in common the following properties:

- They are defined only for unenclosed numeric arguments.
- They return unenclosed numeric results.
- The monadic functions have a rank other than zero.
- The dyadic functions have a non-zero left rank and/or a non-zero right rank.

⌹ω *Matrix Inverse*

Rank: 2

This function is denoted in APL by the ***domino*** symbol (⌹). It is a generalization of the mathematical function of *matrix inverse* for which there is no special symbol, but which mathematicians indicate by M^{-1}.

An argument that does not have an inverse is rejected with a *domain error*.

The inverse of a nonsingular matrix *M* is found by ⌹*M*. It satisfies the identity

```
I ⇐⇒ (⌹M)+.×M
```

where *I* is an *identity matrix,* a square matrix having the same number of rows and columns as *M*, with 1s along the diagonal and 0s elsewhere. The monadic use ⌹*M* is equivalent to *I*⌹*M*. The shape of ⌹*M* is ⌽⍴*M*. More generally, the shape of each result cell is the reverse of the shape of an argument cell.

α⌹ω *Matrix Divide*

Ranks: ∞ 2

For a matrix *M* and column-matrix (that is, a one column matrix) *C* having the same number of rows as *M*,

```
C⌹M
```

is the solution of a *set of linear equations* in which the left argument *C* contains the *constant terms* while the right argument *M* is a matrix containing the *coefficients* for each of the unknowns. *M* has a column for each of the unknowns, and a row for each equation. Since you must have at least as many equations as unknowns, *M* must have at least as many rows as it has columns. The interpreter signals a *rank error* if the left argument has a rank greater than 2.

Multiple linear regression requires a vector of dependent observations in the column-matrix *C* and a matrix of independent or *predictor* observations *P*. The matrix *P* has a column for each of the independent variables, and a row for each row of *C*. Since the interpreter accepts a vector in the place of a column-matrix, there may be an element of *C* for each row of *P*.

The coefficients of the best-fitting linear combination of the independent variables are obtained from *C*⌹*P*. Finding a best-fitting polynomial is essentially the same procedure but with a single independent variable. The columns of a cell of ⍵ consist of the vector of observations on that variable, each raised to the successive powers by an expression such as *P*∘.*⍳*N*. Hence, to find the coefficients of the best-fitting polynomial of degree *D* to approximate the observations ⍺ from a predictor variable *P*, the expression is:

```
⍺⌹P∘.*⍳D+1
```

The expression contains *D*+1 because a polynomial of degree *D* has *D*+1 terms: it includes the constant term (represented by ⍵*0) as well as ⍵*1 ... ⍵*D.

When ⍺ has multiple columns, the result is a matrix containing an independent solution for each of the columns of ⍺. The result has a column corresponding to each of the columns of ⍺. In general, the matrix quotient ⍺⌹⍵ is linked to the inverse of a matrix by the identity:

```
⍺⌹⍵ ⌸⍸⌸ (⌹⍵)+.×⍺
```

When ⍺ is square (that is, has as many rows as columns), the solution is exact:

```
⍺ ⌸⍸⌸ ⍵+.×⍺⌹⍵
```

But when ⍵ has more rows than columns, the result is the least-squares approximation which minimizes the sums (down the columns) of the squares of the difference between ⍺ and ⍵+.×⍺⌹⍵.

Frame and Cells with Matrix Divide

Matrix division treats its left argument as a single cell, regardless of the number of axes; the left argument rank is therefore *infinite.* The right argument rank, however, is limited to two. When the right argument has more than two axes, any additional axes are *frame axes* in which the rank two cells are embedded. Each of the ⍵-cells is paired with the single cell formed by the entire left argument.

The length of the first axis of ⍺ must match the length of the first axis of each ⍵-cell. Those matching axes disappear from the result, as they do in reductions. The axes of ⍺ after the first—call them ⍺'s *trailing axes*—contain the *constant terms* for independent sets of equations to be solved, each set in a *column* of ⍺. There is a result vector corresponding to each of ⍺'s column vectors.

The various columns of ⍺ do not interact. When ⍺ has more than one trailing axis, the organization into trailing axes has no effect on the result, except that the result columns are grouped in the same way. Suppose each of the ⍵-cells has **r** rows. Then the first axis of ⍺ must have length *r* to match. Each of the columns in the trailing axes of ⍺ contains a separate set of constants for which a separate solution will be found. If there are (for example) six such columns, it makes no difference to the calculation whether they are arranged so that ⍺ has shape (*r*,6) or (*r*,6 1) or (*r*,2 3) or (*r*,3 2), and so on. The manner in which the axes of the arguments ⍺ and ⍵ contribute to the axes of the result can be summarized as follows:

- **Left argument axes.** First axis length must match number of rows of ⍵; all other axes *trailing*.
- **Right argument axes.** All but last two axes are *frame*; next-to-last axis is *rows*; last axis is *columns*.
- **Result axes.** Frame, columns, trailing.

⍺?⍵ Deal

Ranks: ∗ ∗; Implicit Arguments: ⎕*io* ⎕*rl*

The function ***deal*** is named by analogy with dealing from a pack of cards. The result of ⍺?⍵ for each of the corresponding cells of ⍺ and ⍵ is a vector of length ⍺, with all elements distinct. Each cell of the result has rank one.

```
     4?6
3 1 2 4
```

```
     3 4 (?⍤0) 100 200
 56 27  84   0
119 16 133 180
```

The zero at the end of the first row is *padding* provided by the *tolerant frame-builder* in SHARP APL to give each result cell the length of the longest (in this case, 4).

Pseudorandom Numbers Depend on Random Link

The values generated by the functions ***roll*** and ***deal*** are *pseudorandom*; they meet most tests of randomness, but are in fact generated by a determinate process. The algorithm refers to the value of the system variable ⎕*rl* (random link). Each use of ***roll*** or ***deal*** captures the current value of ⎕*rl*, and then sets ⎕*rl* to a new value. If at different times you set ⎕*rl* to the same initial value and then use ***roll*** or ***deal*** with the same arguments, you will get the same results. The algorithm by which pseudorandom numbers are generated is described in the *SHARP APL Reference Manual.*

⍺⊤⍵ *Encode/Representation*

Ranks: ∞ ∞

The ***encode*** function (⍺⊤⍵) constructs the ***representation*** of the numbers in ⍵ in the number system whose radix is ⍺. For example, to represent a time interval recorded as 3723 seconds as a triplet in radix 24 60 60 (hours, minutes, seconds):

```
     24 60 60⊤3723
1 2 3
```

There is a limit to the range of values that can be represented in a radix of a given length. The largest number representable in base ⍺ is ⍺⊥⍺−1. More generally, ⍺⊥⍺⊤⍵ equals (×/⍺)|⍵ rather than ⍵. For example, (3⍴10)⊤3247 is 2 4 7, but (3⍴10)⊥2 4 7 is 247, which is (×/3⍴10)|3247.

For higher rank right arguments, the extra axis representing the individual elements is placed *first*. For example:

```
      2 2 2⊤⍳2*3
0 0 0 0 1 1 1 1
0 0 1 1 0 0 1 1
0 1 0 1 0 1 0 1
```

The convention of placing the individual representations along the first axis was introduced early in the development of APL, before the generalization of frame/cell structures. A more consistent placement of the new axis (that is, at the end) can be obtained by using `⊤⍤1 0`:

```
      2 2 2⊤⍤1 0⊢⍳2*3
0 0 0
0 0 1
0 1 0
0 1 1
1 0 0
1 0 1
1 1 0
1 1 1
```

α⊥ω *Decode/Base-10 Value*

Ranks: ∞ ∞

The ***decode*** function (α⊥ω) is the inverse of the ***encode*** function, and returns the base-10 value of the representation ω in the number system whose radix is α.

```
      24 60 60⊥1 2 3
3723
```

For the vectors α and ω, the result of α⊥ω is `+/ω×`*`weights`*, where *`weights`* is the vector of weights corresponding to the successive positions in the representation:

```
      weights ⇐⇒ ×\⍨⌽1↓α,1
```

For example,

```
      α←24 60 60
      ω←1 2 3
      α⊥ω
3723
      weights←×\⍨⌽1↓α,1
```

```
3600 60 1
     +/⍵×weights
3723
```

Base values are produced for each vector of ⍺ and each column vector of ⍵. For example:

```
     10 10⊥2 3⍴⍳6
3 14 25
     (2 2⍴10 10 100 100)⊥2 3⍴⍳6
3  14  25
3 104 205
```

When ⍺ is a scalar, it is treated as `(1↑⍴⍵)⍴⍺`. Thus:

```
      2⊥1 0 1
5
```

Relational Functions

The relational functions are used to compare arrays. All of these functions return Boolean results.

⍺=⍵ Equals

Ranks: 0 0; Identity Element: 1; Implicit Arguments: ⎕*ct*

A result cell of ⍺=⍵ is `1` when the elements compared are equal, and `0` otherwise. For example:

```
     a←3 5⍴
```

```
abcdefghijklmno'
     a
abcde
fghij
klmno
```

```
      α=⌽α
0 0 1 0 0
0 0 1 0 0
0 0 1 0 0
```

Definition of Equality

An element in α is judged equal to an element in ω as follows:

- To be equal, the elements must be of the same type (number or character). Members of different types can never be equal. For example, the numeric scalar `6` is not equal to the character `'6'`. Enclosed elements are outside the domain of ***equals***; you can use `≡⍤0` to overcome this deficiency.
- Two characters are equal if and only if their positions in ⎕*av* (atomic vector) are equal; for example, `'a'` (position 97 in ⎕*av*) is not equal to `'A'` (position 65 in ⎕*av*).
- The interpreter judges two numbers to be equal if the absolute difference between them is not greater than the value of ⎕*ct* (comparison tolerance) times the greater of their magnitudes. Formally:

  $$\alpha=\omega \iff (|\alpha-\omega)\leq \square ct\times \alpha\lceil[0]|\omega$$

SHARP APL limits the maximum permissible value of ⎕*ct*. It may not be less than `0` nor more than `16*¯8`. If you attempt to set the value of ⎕*ct* to a value not within this range, the interpreter signals a *domain error*, and the value of ⎕*ct* remains unchanged.

There are two practical consequences of the combined effects of this limit and rule 3:

- Tolerant comparison can make a difference only to numbers whose internal representation requires floating point. Stating it another way, no permissible value of ⎕*ct* can permit a number represented internally in integer format to be judged equal to another unless they are exactly equal.
- Tolerance cannot affect comparisons with zero. That is, no comparison in which one of the elements is `0` can be judged equal to another unless they are exactly equal.

α≠ω *Not Equals*

Ranks: 0 0; Identity Element: 0; Implicit Arguments: ⎕ct

The function ***not equals*** is defined as the negation of ***equals***:

α≠ω ⟸⟹ ~α=ω

α<ω *Less Than*

Ranks: 0 0; Identity Element: 0; Implicit Arguments: ⎕ct

α≤ω *Less Than or Equal*

Ranks: 0 0; Identity Element: 1; Implicit Arguments: ⎕ct

α≥ω *Greater Than or Equal*

Ranks: 0 0; Identity Element: 1; Implicit Arguments: ⎕ct

α>ω *Greater Than*

Ranks: 0 0; Identity Element: 0; Implicit Arguments: ⎕ct

Dyadic use of these four functions is discussed together because they are closely related. There are no monadic uses of ≤ or ≥. Monadic use of the symbols < and > denote unrelated functions discussed in the section "Structural Functions."

A result cell returned by expressions containing these functions is `1` where the proposition is true, and `0` otherwise. The domain is restricted to numbers. APL has no default collating sequence for characters, and comparisons such as `'a'<'b'` are not permitted. Each of these functions makes tolerant comparisons, subject to the value of the system variable ⎕ct (comparison tolerance).

```
     1 2 3<2
1 0 0
     1 2 3≤2
1 1 0
```

```
     1 2 3≥2
0 1 1
     1 2 3>2
0 0 1
```

α≡ω *Match*

Ranks: ∞ ∞; Implicit Arguments: □*ct*

The ***match*** function (α≡ω) returns a single Boolean scalar result reporting whether α is identical to ω. To be judged identical, α and ω must have the same ***rank***, the same ***shape***, and each element in α must have the same ***value*** as the corresponding element in ω. For numeric elements, the magnitude of their difference must be no greater than the value of □*ct* (comparison tolerance) times the greater of their magnitudes.

The ***match*** function has infinite default argument rank. Unless explicitly modified by the ***rank*** operator, ***match*** treats its left and right arguments as single cells, and returns a scalar one or zero.

By explicitly modifying ***match*** with the ***rank*** operator, you can restrict the rank of the cells to which it applies:

v≡⍤1 *m*	Compare vector *v* with each vector (row) of *m*.
m≡⍤2 *h*	Compare matrix *m* with each matrix of *h*.

With imposed rank zero, *match* provides the facility that ***equals*** (α=ω) would provide if enclosed elements were in its domain.

α∈ω *Member*

Ranks: 0 ∞; Implicit Arguments: □*ct*

For each element (number, character, or enclosed) within α, the expression α∈ω reports whether a matching element can be found anywhere in ω (regardless of the rank and shape of ω). The result contains a 1 wherever an element in α is matched by an element somewhere in ω.

```
      v←2 3 5
      m←2 3⍴1 2 3 4 5 6
      m
1 2 3
4 5 6
      2∈m
1
      v∈m
1 1 1
       m∈v
0 1 1
0 1 0
```

Note that elements with different levels of enclosure do not match, so that:

```
      ((<'abc'),<'def')∈(<<'abc'),<<'def'
0 0
```

The ***rank*** operator may organize the frame in which membership in a cell of ⍵ is reported. For example, the following example returns a result frame of rank one, within which there are seven single-element cells. Each element of the result is a `1` or a `0`, showing whether that element in *v* is a ***member*** of *m*:

```
      v
7 2 8 5 9 2 5
      m
 1  3
 5  7
 9 11
13 15
17 19
      v∈m
1 0 0 1 1 0 1
```

By contrast:

```
      v∈⍤1 m
0 0 0 0 0 0 0
1 0 0 1 0 0 1
0 0 0 0 1 0 0
0 0 0 0 0 0 0
0 0 0 0 0 0 0
```

Here *v* is treated as a single rank one cell, contained in an empty frame, while *m* is treated as a set of rank one cells arranged in a rank one frame. As usual, the empty frame is paired with each of the cells in the frame that is not empty. That produces a result frame that has the same rank and shape as the nonempty argument frame. In the example, the frame has rank one and length five.

Each of the five result cells is a seven element vector, in which the presence of a `1` or a `0` indicates whether an element in *v* is a member of a particular cell of *m*. The result has five rows (the frame shape) and seven columns (the cell shape).

α⍷ω *In*

Ranks: ∞ ∞; Implicit Arguments: ⎕*ct*

The ***in*** function (α⍷ω) returns a Boolean result which has the same shape as ω. Each `1` in the result marks the position in ω at which the entire sequence of elements in α can be found in ω:

```
      a
The
      b
The theory of the theists' theme.
      a⍷b
0 0 0 0 1 0 0 0 0 0 0 0 0 1 0 0 0 1 0 0 0 0 0 0 0 1 0 0 0 0 0
      b ◇ 1 0⍤a⍷b
The theory of the theists' theme.
0000100000000010001000000000100000
```

The ***in*** function is defined to include overlapping matches:

```
     b←'bababababababababa'
     b ◇ 1 0⍤'baba'⍷b
bababababababababa
1010101010101010000
```

For arrays of other ranks and types, consider the example below.

```
     b←2 3⍴⍳6
0 1 2
3 4 5
     (2 2⍴1 2 4 5)⍷b
0 1 0
0 0 0
```

Indexing Functions

The indexing functions either produce a set of indices into, or apply a set of indices to, an argument array.

α@ω *From*

Ranks: 0 ∞

The *from* function selects elements from an array.

This function is unlike the indexing facilities available in traditional APL: it uses a single symbol to represent the function, not a pair of syntactically ambiguous ones (square brackets) combined with a special purpose delimiter (the semicolon). **The *from* function is not sensitive to** ⎕*io* (index origin)—indices you supply are treated as though in origin zero.

From selects elements from *ω* based on *indices* in α. Each element in α independently selects one or more cells from *ω* as expected from a function with left rank of zero. The shape of α@ω for simple α is `(⍴α),1↓⍴ω`.

```
      v
11 12 13 14 15 16
       1@v
12
       m
0 1 2
3 4 5
6 7 8
       1@m
3 4 5
       ⍴1@m
3
```

For an *integer scalar* left argument in the range `-1↑⍴ω` to `¯1+1↑⍴ω`, the ***from*** function selects a *major cell* of *ω* indexed by `(1↑⍴ω)|α`. Positive indices select from the *head* of *ω*, while negative indices select from the *tail* of *ω*.

```
      ¯1@v
16
       ¯1@m
6 7 8
```

```
      i
1 ¯1
0 ¯2
      i@'abcde'
be
ad
      a
 0  1  2
 3  4  5
 6  7  8
 9 10 11
12 13 14
      i@a
 3  4  5
12 13 14

 0  1  2
 9 10 11
```

Indices which lie outside the range of locations in an array are not acceptable and cause the interpreter to signal an *index error*.

Using Enclosed Indices

In the examples above, the left argument to ***from*** was unenclosed. When the left argument contains enclosed elements, ***from*** treats the contents of each enclosure as a complete list of indices, that is, one (and only one) index for each axis of the array being indexed. These enclosed elements are called ***selection expressions***.

By using an enclosed vectors of numbers as the left argument, arbitrary elements may be selected right argument. This is sometimes called ***scattered indexing***. The following example shows how unenclosed indices differ from enclosed indices:

```
      a
1 2
0 1
      a@m
3 4 5
6 7 8

0 1 2
3 4 5
      b←<⍤1 a
```

```
      b
|---| |---|
|1 2| |0 1|
|___| |___|
      b@m
5 1
```

When a selection expression is to contain *more than one index* for a particular axis of ω, the element is *two levels deep* (doubly enclosed). For example, to select row `1` and columns `0 2` from *m*, build a selection expression with two cells—one for rows and one for columns:

```
      a←<1⊃0 2
      a
|----------|
||¯| |¯¯¯||
||1| |0 2||
||_| |___||
|__________|
       a@m
3 5
```

Trailing axes from which you wish to select *all indices* may be omitted from a selection expression. For example, to select all columns from rows `2 1` of *m*, build a selection expression with one cell for the rows, and omit reference to the columns:

```
      a←<<2 1
      a
|-----|
||¯¯¯||
||2 1||
||___||
|_____|
      a@m
6 7 8
3 4 5
```

To select a set of columns *for all rows,* consider using the ***from*** function with the *rank* operator so that the indices you specify apply to a particular frame:

```
      2 0@⍤1 m
2 0
5 3
8 6
```

Also consider using *complementary indices* as described in the section *Excluded Index Selection.*

Excluded Index Selection

The ***from*** function also lets you indicate which indices along an axis of ⍵ you want to ***exclude***. This approach, which can shorten lists of indices, makes possible selection of all indices along an axis without referring explicitly to the length of the axis. To make such a *complementary selection,* enclose excluded indices *one level deeper* than for those to be included. For example, to select *all rows but the first,* and columns `1 2`, the selection expression is:

```
      α←<(<0)⊃1 2
      α
┌────────────┐
│┌───┐ ┌───┐│
│└¯┘│ │1 2│└┘
│0│ │ └───┘
│└─┘│
└────────────┘
      α@m
4 5
7 8
```

When you wish to include all elements along an axis, use a ***jot*** (∘) to *exclude none.* A *trailing* jot in an index expression is always superfluous, as is a trailing enclosure containing all indices for a trailing axis.

For example, the result obtained with the ***rank*** operator (to select all rows for a given set of columns) could also have been obtained with complementary indexing:

```
      ⊢⍺←<⍥⊃2 0
┌──────────┐
│┌──┐ ┌───┐│
│└──┘ │2 0│└┐
│     └───┘ │
└───────────┘
      ⍺@⍤
2 0
5 3
8 6
```

⍳⍵ *Count*

Rank: 1; Implicit Arguments: ⎕*io*

The ***count*** function (⍳⍵) counts in the sense that ⍳*n* yields a list of the first *n* integers, beginning with the value of ⎕*io* (index origin):

```
      ⍳3
0 1 2
```

Default Monadic Rank of Count

SHARP APL for UNIX departs from *A Dictionary of APL* in its treatment of the argument rank for ***count***. *A Dictionary of APL* prescribes rank zero, which would have some convenient properties. However, principally for consistency with existing applications, SHARP APL treats ***count*** as expecting its argument to be a one element vector, for which it returns a vector result. In this way, SHARP APL continues to support the widespread use of expressions such as ⍳⍴*v*. If ***count*** were a rank zero function, one would have to use ⍳0⍤⍴*v*.

Use the ***rank*** operator (*f⍤n*)where it is desirable to force ***count*** to take argument cells of rank zero. For example, several vectors of integers are generated by the following (with the short rows of the result automatically padded with zeros):

```
      (⍳⍤0) 2 3 4
0 1 0 0
0 1 2 0
0 1 2 3
```

Index Origin

The result produced by ***count*** is affected by the value of ⎕*io* (index origin).

```
      ⍳5
0 1 2 3 4
       ⎕io←1
       ⍳5
1 2 3 4 5
```

⍺⍳⍵ *Index Of*

Ranks: 1 0; Implicit Arguments: ⎕*ct* ⎕*io*

The expression ⍺⍳⍵ returns a result which indicates where each of the elements of ⍵ may be found in the vector ⍺. The left argument cell or cells must be of rank one so that a single number is sufficient to identify a location. The right argument cell or cells may be of any rank or shape. A result cell has the same rank and shape as a cell of the right argument ⍵.

You could think of ⍺⍳⍵ as a mapping that, for each element within ⍵, proceeds from a value in ⍵ to an ***index of*** ⍵ within ⍺:

- For an element in ⍵ that occurs *nowhere* in ⍺, ⍺⍳⍵ returns ⎕*io*+⍴⍺.
- For an element in ⍵ that occurs *more than once* in ⍺, ⍺⍳⍵ returns the index of the first occurrence.

The latter property is exploited in a common expression for returning the *nub* (unique elements) of a vector:

```
((v⍳v)=⍳⍴v)/v
```

This statement looks up where the elements of v occur in v itself. It compares these locations to those computed with the consecutive integers. For an element that is duplicated in v, the first occurrence has an index equal to its position. However, later occurrences also have that same index, which therefore must be different from their positions in v, resulting in their exclusion from the nub.

Tolerant Comparison in Index Finding

The expression $\alpha\,\underline{\iota}\,\omega$ asks, in effect, for the index of the element in α which matches an element in ω, where *match* is affected by $\square ct$ (comparison tolerance). Strictly speaking, $\alpha\,\underline{\iota}\,\omega$ returns the first location in α which is tolerantly equal to an element in ω. Comparison tolerance can have no effect on the result when the arguments are characters or integers in the range -2^{31} to $2^{31}-1$. However, it may indeed affect both the results and the time required to compute them for fractional values. When you are using fractional arguments and you know for certain that any elements in ω do indeed exist somewhere in α, it may speed computation to set $\square ct \leftarrow 0$.

$\alpha\,\underline{\iota}\,\omega$ *Index*

Ranks: $\infty\ \infty$; Implicit Arguments: $\square ct$ $\square io$

The ***index*** function ($\alpha\,\underline{\iota}\,\omega$) is defined in terms of the ***in*** function ($\alpha\underline{\in}\omega$), and $i \leftarrow \alpha\,\underline{\iota}\,\omega$ yields the index in $b \leftarrow \omega\underline{\in}\alpha$ of the first `1`, or the value ρb if ω does not occur in α.

```
      a
0 1 2 0 1 2
3 4 5 3 4 5
1 2 0 1 2 0
4 5 3 4 5 3
      b
4 5
2 0
      a∈b
0 0 0 0 0 0
0 1 0 0 1 0
0 0 0 0 0 0
0 0 0 0 0 0
      a⍳b
```

```
1 1
      a⌶-b
4 6
```

⍋ω Numeric Grade Up
⍒ω Numeric Grade Down

Rank: ∞; Implicit Arguments: ⎕io

Each of the various forms of the function *grade* returns a vector of integers which, if used to sort the *major cells* of ω, would arrange them in order. The result has the same length as the number of elements along the first axis of ω. The result is a *permutation vector* since it contains each of the first 1↑⍴ω integers exactly once. Since the result is formed from the first *n* integers, the result is sensitive to the value of the ⎕io (index origin). Grade is not affected by ⎕ct (comparison tolerance).

The two monadic ***grade*** functions can be used only on a unenclosed homogeneous numeric array, that is, an array that consists entirely of numeric elements.

The expression ⍋ω returns the vector of indices which would sort the major cells of ω in ascending order, while the expression ⍒ω returns the vector of indices which would sort the major cells of ω in descending order. Cells that are equal are left in the order in which they occur in ω.

When ω is a vector and its major cells are scalars, the result of ⍋ω is the vector of indices that would sort the elements. For example:,

```
      v←1.2 1.1 1.2 1.4 1.1
       ⍋v
1 4 0 2 3
       ⍒v
3 0 2 1 4
      (⍋v)@v
1.1 1.1 1.2 1.2 1.4
      (⍒v)@v
1.4 1.2 1.2 1.1 1.1
```

The right argument must have at least one axis; an attempt to grade a scalar is rejected with a *rank error*.

When you are grading an array that has more than one axis, the grade functions return a list of indices for the major cells. When the array has more axes, there are more axes in each cell. The cells are graded by assigning to each cell a numeric value equal to the base value of all the elements in the cell, taking them in row-major order (see "Shape, Reshape" on page 4-46).

The base value is evaluated using a radix larger than the value of any of the individual elements. That is equivalent to saying that the cells are graded so that the left-most element in the ravel of the cell has greatest weight.

Example:

```
      m
1 4 9 2.1
2 0 0 1.8
1 9 3 0
1 7 7 6
      ⍋m
0 3 2 1
      (⍋m)⌷m
1 4 9 2.1
1 7 7 6
1 9 3 0
2 0 0 1.8
      m
2 0 0 1.8
1 9 3 0
1 7 7 6
1 4 9 2.1
```

α⍋ω *Character Grade Up*
α⍒ω *Character Grade Down*

Ranks: ∞ ∞; Implicit Arguments: ⎕io

For grading character data, the grade functions are used with a left argument which describes the *collating sequence* for the ordering. To grade character data, you must provide a collating sequence in the left argument. APL has no default collating sequence for characters. The left argument is used to assign a numeric value to each of the characters in the right argument. With those values, the cells are graded in the same way as the major cells of a numeric array.

When the collating sequence ⍺ is a vector, the numeric value assigned to each character in ⍵ is the value that would be obtained by ⍺⍳⍵, and ⍺⍋⍵ is equivalent to ⍋⍺⍳⍵. When the right argument ⍵ contains a character that appears nowhere in the collating sequence ⍺, it is assigned a value greater than that assigned any of the characters that are present (just as ⍺⍳⍵ returns ⎕*io*+⍴⍺ for an element of ⍵ that is not present in ⍺).

```
      alf1←' abcdefghijklmnopqrstuvwxyz'
      m
jones
baker
jacobs
james
abel
      (alf1⍋m)@m
abel
baker
jacobs
james
jones
```

A one-axis alphabet such as *alf1* makes no provision for placing two characters at the same position in the collating sequence. For example, you might wish to place *b* and *B* after *a* or *A* and before *c* or *C*. To achieve that, the collating sequence ⍺ may have more than one axis:

```
      alf2
 abcdefghijklmnopqrstuvwxyz
 ABCDEFGHIJKLMNOPQRSTUVWXYZ
```

Where ⍺ has more than one axis, the dominant consideration is the position at which a character is found along the last axis of ⍺. For example, *b* and *B* are both at position two along the last axis. A character's location on the other axes contributes to the resulting sequence only to break a tie between two words which otherwise would have the same weight. The next-to-last axis breaks a tie on the last axis; the axis before that breaks a tie on the last two, and so on. To see the effect of this rule, compare the ordering produced by the tabular alphabet *alf2* with the ordering produced by the single axis alphabet *alf1* in Table 4.5.

Characters with Identical Weights

While the matrix *alf2* provides that *b* or *B* comes after *a* or *A* and before *c* or *C*, it also provides that, wherever there is a tie, a word containing *B* always comes later than a word that contains *b*. However, it is possible to construct the left argument so that certain letters have identical weights, and thus words that differ only in those letters are unaffected by grading.

Table 4.4. Comparing characters with identical weights.

m	(alf2⍋m)@m	(alf1⍋m)@m
Ama	acid	acid
YMCA	aluminum	aluminum
Trudgen	ama	ama
Tektite	Ama	ammonia
pi	AMA	embolism
stroke	ammonia	pavilion
pavilion	DPD	phosphate
piping	embolism	photosynthesis
respiration	NSPF	pi
pump	pavilion	piping
photosynthesis	pH	plug
underwater	Philodendron	pool
tsunami	phosphate	porosity
pool	photosynthesis	pump
NSPF	pi	pH
recovery	piping	recovery
aluminum	plug	respiration
embolism	pool	stroke
plug	porosity	trudgen
Tsunami	pump	tsunami
trudgen	recovery	underwater
pH	respiration	Tektite
porosity	stroke	Philodendron
phosphate	Tektite	Ama
DPD	trudgen	Trudgen
ammonia	Trudgen	Tsunami
AMA	tsunami	DPD
Philodendron	Tsunami	AMA
acid	underwater	YMCA
ama	YMCA	NSPF

For two characters to have identical weight, they must have the same effective position in ⍺. Clearly, they cannot both occupy the same position. However, when the same character occurs at two or more positions in ⍺, that character's *effective* position is at the *minimum* of its various coordinates. For example, if the letter *S* occurs at row 2, column 7 and also at row 5, column 3, its effective

position is `2 7⌊5 3` (i.e. row `2`, column `3`). The letter *S* has identical weight to a unique character located at row `2`, column `3`, or to ***any*** character that occurs twice, both in row `2` and a higher row, and column `3` and a higher column.

```
   0 1 2 3 4 5 6 7 8 9
0: ∘ ∘ ∘ ∘ ∘ ∘ ∘ ∘ ∘ ∘
1: ∘ ∘ ∘ ∘ ∘ ∘ ∘ ∘ ∘ ∘
2: ∘ ∘ ∘ ⊡ ∘ ∘ ∘ S ∘ ∘
3: ∘ ∘ ∘ ∘ ∘ ∘ ∘ ∘ ∘ ∘
4: ∘ ∘ ∘ ∘ ∘ ∘ ∘ ∘ ∘ ∘
5: ∘ ∘ ∘ S ∘ ∘ ∘ ∘ ∘ ∘
6: ∘ ∘ ∘ ∘ ∘ ∘ ∘ ∘ ∘ ∘
7: ∘ ∘ ∘ ∘ ∘ ∘ ∘ ∘ ∘ ∘
```

Figure 4.4. Effective weight ⊡ of `S` occurring at two positions in `⍺`.

To give capital and lowercase letters identical weight in sorting, one set (for example, the lowercase letters) is usually placed at the *top left corner* of the alphabet matrix, while the other set is duplicated, both to the right and below:

```
abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ
ABCDEFGHIJKLMNOPQRSTUVWXYZ
```

In such a matrix, the effective weight of *a* is exactly the same as the effective weight of *A*.

`⍺[⍵]` *Bracket Indexing*

Ranks: Not applicable; Implicit Arguments: `⎕io`

The expression or expressions enclosed in square brackets select elements from within the variable `⍺`, which may be any array having at least one axis. Bracket indexing is anomalous in several respects, and is considered obsolete. It is supported in this version of SHARP APL for upward compatibility with applications written in older dialects of APL. Unlike any other APL function, indexing was indicated not by a single symbol but by a pair of symbols. In some respects the left bracket (`[`) is the primary symbol, while the right bracket (`]`) simply delimits the argument, which implies an order of execution—like putting parentheses around the entire indexed expression.

```
a[i]+b ⟸⟹ (a[i])+b
```

The brackets surround not an expression but a set of independent expressions, one for each axis of α. To select from an array whose rank is greater than one requires more than one index expression; the various expressions are separated from each other by semicolons.

When α is a vector (with only one axis), one expression suffices to specify positions within it, and is written `α[ω]`. When α is an array of rank two or greater, selecting a subarray from within it requires a separate expression for each of the axes of α; the number of semicolons is always one less than the rank of α. For example, when α is a matrix use an expression of the form `α[i0;i1]`; when α has rank four, use an expression of the form `α[i0;i1;i2;i3]`.

Each array in the index expression `i0,i1,i2...`, may be of any rank or shape, provided only that every one of its elements is a valid position within α, subject to the value of the `⎕io` (index origin). If this requirement is not met, the interpreter signals an *index error*.

The semicolon is punctuation: a delimiter, not a function. Writing `i0;i1` does not produce a single array from `i0` and `i1`. You cannot write a general expression such as `α[ω]` that will be valid regardless of the rank of α. In fact, `α[ω]` is valid only when α is a vector (since it contains zero semicolons), and `α[i0;i1]` is valid only when α is a matrix (because it contains one semicolon), and so on. The semicolon acts as a complete separator in these expressions: you do not need to surround the expressions for each axis with parentheses.

To write a program that can index an array of any rank, you have either to generate a character expression containing the correct number of semicolons and then use ***execute*** (`⍎ω`), or else index it by referring to the ravel of the array, exploiting an identity mentioned below.

Keeping All Positions Along an Axis

You can retain all positions along a particular axis of α by not selecting along that axis. You do that by writing *no* expression for that axis at the appropriate position between brackets. That axis is then reproduced unchanged in the result. For example, when α is a vector, the result of `α[]` is α. When α has three axes, then `α[i0;i1;]` indicates that along the first axis, positions are to be selected by `i0`, along the second axis by `i1`, but the last axis (for which no index has been specified) is to be included completely in the result.

Similarly, `α[1 2;]` selects positions `1` and `2` from axis zero and all positions from axis one; and `α[;0 1]` selects all positions along axis zero, and positions `0` and `1` along axis one.

Shape of the Result from Bracket Indexing

The result of `α[⍺0;⍺1;...]` is an array formed by selecting from `α` elements at the positions indicated by the values in `⍺0`, `⍺1`, etc. The rank and shape of the result are determined solely by the ranks and shapes of the various arrays in the right argument. In the simple case when `α` is a vector, indexing is `α[ω]`, and

```
ρα[ω] ⇐⇒ ρω
```

More generally, this identity can be stated thus:

```
ρα[⍺0;⍺1;⍺2...] ⇐⇒ (ρ⍺0),(ρ⍺1),(ρ⍺2)...
```

For any axis `⍺` for which ***no*** indices are stated, the shape is `(ρα)[⍺]`.

The process of selection is based on a Cartesian product. The elements selected from `α` are those whose positions are described by all possible combinations of the first axis indices mentioned in `⍺0`, with the second axis indices mentioned in `⍺1`, and so on for all the axes.

If any of the selection arrays is present but empty (for example `α[⍺0;⍺1]` when `⍺0` is an empty vector), the result has length zero along that axis.

Where an axis is indexed by a scalar (which has no axes), the corresponding axis is omitted from the result. For example, when `α` is a vector, `α[2]` returns a scalar (with no axes) because `2` is a scalar with no axes. When `α` is a matrix, `α[1 3 5;2]` returns a `3`-element vector, and `α[1 3 5;,2]` returns a `3`-by-`1` matrix.

Indexed Assignment

A bracket expression `α[⍺0;⍺1;⍺2...]` may appear to the left of the ***assignment arrow*** (`←`); see "Chapter 2. Naming". Within the brackets, you must have a valid index expression, or a list of several such expressions, depending on the rank of `α`. The requirements for a valid list of indices are the same as for indexed selection.

The list of expressions inside the brackets implies the shape of a subarray within `α` whose elements are to be *respecified* in the same way as the indices of indexed selection determine the shape of the result. This is also the shape of the result received by any other function to the left of `α[ω]`. With one exception (explained below), the shape of the right argument must match the shape implied by the indices.

Effect of Indexed Assignment

The indicated positions within ⍺ are assigned the values of the corresponding elements of ⍵.

```
      b
 0  1  2  3
 4  5  6  7
 8  9 10 11

12 13 14 15
16 17 18 19
20 21 22 23
      c
101 102
103 104
      b[0 1;2;0 3]
 8 11
20 23
      b[0 1;2;0 3]←c
      b
  0   1   2   3
  4   5   6   7
101   9  10 102

 12  13  14  15
 16  17  18  19
103  21  22 104
```

When ⍺ is initially a homogeneous array, reassigning parts of it to elements of different type causes the entire array to become heterogeneous. This affects the space required to store it. Similarly, when ⍺ is initially a numeric array whose internal representation is Boolean or integer, but ⍵ contains elements that cannot be represented in that way, the entire array ⍺ is converted to the internal type required to accommodate the new elements, again affecting the space needed to store it.

Structural Functions

The functions discussed in this section are concerned primarily with the ***structure*** of arrays, rather than the values of the elements within the array.

ρω *Shape*

Rank: ∞

The ***shape*** function reports the length of each of the axes of ω. The result returned by ***shape*** is a list of numbers which are nonnegative integers. The vector contains one element for each axis of ω. Thus, the number of elements in the result is the number of axes in ω; that is, the *rank* of ω. The value of each element in the result is the length of the corresponding axis of ω. When ω has no axes, the result of ρω is a vector containing no elements (an empty vector). For example:

```
      ρω←6

      ρρω
0
      ρω←⍳5
5
      ρρω
1
      ρω←2 3ρ⍳6
2 3
      ρρω
2
```

αρω *Reshape*

Ranks: 1 ∞

The ***reshape*** function creates an array whose shape is given by α and whose elements are taken from ω. To reshape an array, the left argument α must be a list of nonnegative integers indicating the length of each of the axes of the array to be restructured. An scaler is accepted as equivalent to a one element vector.

The right argument ω may be an array of any type or shape. However, ω may be empty only when the left argument contains a zero. (An array is empty if one or more of its axes has length zero.) Thus, the right argument may be empty only if the result to be produced is also empty. If you attempt to create a nonempty array from an empty right argument, the expression is rejected with the message *length error*.

The expression ⍺⍴⍵ returns an array whose shape comes from ⍺ and whose elements come from ⍵. The elements in the result are filled by taking the elements of ⍵ in *row-major* order (see "Row-Major Order" below). When the total number of elements required for the result (that is, ×/⍺) is less than the number of elements in ⍵, elements are taken from ⍵ in row-major order until the result is complete, and the remaining elements of ⍵ are ignored. When the result requires more elements than ⍵ contains, the elements of ⍵ (always in row-major order) are repeated cyclically in the result. For example, suppose ⍵ is a 2-by-5 matrix:

```
      ⍵
0 1 2 3 4
5 6 7 8 9
```

Then a result containing fewer elements than ⍵ might be obtained by:

```
      7⍴⍵
0 1 2 3 4 5 6
```

or by

```
      2 4⍴⍵
0 1 2 3
4 5 6 7
```

A result requiring more elements than there are in ⍵ might be obtained by expressions such as:

```
      12⍴⍵
0 1 2 3 4 5 6 7 8 9 0 1
      3 7⍴⍵
0 1 2 3 4 5 6
7 8 9 0 1 2 3
4 5 6 7 8 9 0
```

If ⍺ is a vector longer than ⍴⍵ but containing ⍴⍵ as its last elements, the effect of reshaping is to replicate ⍵ completely at each of the positions along the new axes. For example:

```
      ⍴⍵
2 5
      3 2 5⍴⍵
0 1 2 3 4
5 6 7 8 9
```

```
0 1 2 3 4
5 6 7 8 9

0 1 2 3 4
5 6 7 8 9
```

If the shape of the result differs from the shape of the right argument, a systematic pattern may be generated. For example:

```
      4 4⍴5↑1
1 0 0 0
0 1 0 0
0 0 1 0
0 0 0 1
```

The shape of a scalar is empty, so to produce a scalar, specify an empty left argument for ***reshape***. Since the type of an empty vector does not matter, you can get a scalar from ⍵ equally well with `(⍳0)⍴⍵` or `''⍴⍵`

In almost every case, the shape of the result is ⍴⍺; the only exception is that a scalar ⍺ is accepted in place of a one element vector:

```
⍴⍺⍴⍵ ⇐⇒ ,⍺
```

and

```
(⍴⍵)⍴⍵ ⇐⇒ ⍵
```

Note on Row-Major Order

Row-major order is a term borrowed from matrix algebra. Referring to a matrix, it means that you take first the element located in the first row at the first column, then the one in the first row, second column, and so on, stepping through all the columns before moving to the next row. This can be applied to arrays of any rank: you count off through the positions along the last axis, then along the next to last, and so on. This is the same order as produced by the ***ravel*** function (`,`⍵):

```
⍺⍴⍵  ⇐⇒  ⍺⍴,⍵
```

α↑ω *Take*

Ranks: 1 ∞

The ***take*** function constructs an array by selecting a segment along each of the axes of ω. The left argument α is an integer vector which specifies for each axis of ω the length of the segment to be selected. The length of α is always less than or equal to the rank of ω. However, α may be a scalar, in which case it is treated as though it were a one element vector. The result always has the same rank as ω (except when ω is a scalar).

Each value in α specifies the length of a segment along the corresponding axis of ω to be included in the result. For example, when ω is a rank three array, the expression

```
2 3 4↑ω
```

returns a result constructed by taking the first two positions along the first axis of ω, the first three positions along the second axis of ω, and the first four positions along the third axis of ω. A negative value in α refers to the last positions along an axis. For example:

```
2 ¯3 4↑ω
```

returns a result constructed by taking the first two positions along the first axis of ω, the ***last*** three positions along the second axis of ω, and the first four positions along the third axis of ω.

Take when `(⍴α)<(⍴⍴ω)`

When α has fewer elements than ω has axes, the α provided is assumed to describe how much to take of the leading axes of ω, while the remaining axes are returned entire. For the same rank three ω mentioned before, the expression `2 3↑ω` selects the first two positions of the first axis of ω, the first three positions of the second axis of ω, and ***all*** positions of the remaining third axis. The expression `2↑ω` selects the first two positions of the first axis of ω, and all positions of the remaining axes. The expression `''↑ω` selects all positions along all the axes of ω.

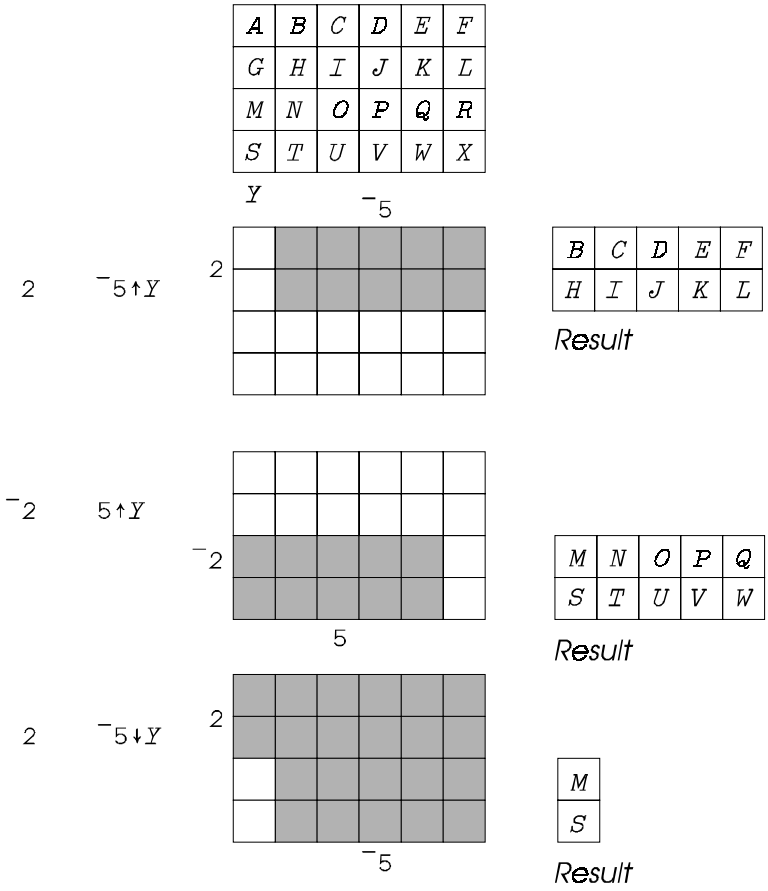

Figure 4.5. Take and drop.

Overtake and Padding

When the magnitude of an element of ⍺ is greater than the length of the corresponding axis of ⍵, ⍺↑⍵ requests that more elements than exist along that axis to be taken from ⍵. The result nevertheless contains as many positions as you request.

The additional positions contain a *fill* value, which is:

`0` when ⍵ is an unenclosed *numeric* array

`''` when ⍵ is an unenclosed *character* array

`⍬` when ⍵ is a *heterogeneous* array or contains *enclosed* elements

The ***jot*** symbol (`⍬`) denotes the *enclosed empty numeric vector,* a constant with a value of `⊂⍳0`.:

When such an oversized element of ⍺ is positive, the elements are taken from the front of the corresponding axis of ⍵, followed by the fill value for the remaining positions, so that the fill values appear at the end of that axis. When an oversized element of ⍺ is negative, elements are taken from the end of the corresponding axis of ⍵ and placed at the end of the corresponding axis in the result. The fill value appears in the remaining positions at the front of that axis.

Suppose ⍵ is the following `3`-by-`4` matrix of numbers:

```
      ⍵
1  2  3  4
5  6  7  8
9 10 11 12

      5 ¯7↑⍵
0  0  0  1  2  3  4
0  0  0  5  6  7  8
0  0  0  9 10 11 12
0  0  0  0  0  0  0
0  0  0  0  0  0  0
```

Take with Empty Right Argument

When ⍺ calls for creation of a nonempty array, but ⍵ is empty, the result contains only fill values with the same internal type as the empty right argument. Although in general the type of an empty array is not significant, information regarding its type is retained and affects the result produced by the ***take*** function and by the derived functions (*m*⍀⍵ and *m*⌿⍵), as described in "Chapter 5. Operators".

When the empty right argument has character type, the result consists of blanks. When it has numeric type, the result consists of zeros. An empty array formed by removing all cells from a heterogeneous array, or an array of enclosed elements, reverts to numeric type.

Take Applied to a Scalar

When ⍵ is a scalar, but ⍺ is not empty, ⍵ is treated as if it had the rank implied by the length of ⍺, and every axis had length one:

```
      3 ¯7↑6
0 0 0 0 0 0 6
0 0 0 0 0 0 0
0 0 0 0 0 0 0
```

⍺↓⍵ *Drop*

Ranks: 1 ∞

The ***drop*** function constructs an array by omitting a segment along each of the axes of ⍵. The left argument ⍺ is an integer vector which specifies for each axis of ⍵ the length of the segment to be omitted. The length of ⍺ is always less than or equal to the rank of ⍵. However, ⍺ may be a scalar, in which case it is treated as though it were a one element vector. The result always has the same rank as ⍵ (except when ⍵ is a scalar).

Drop when `(⍴⍺)<(⍴⍴⍵)`

When ⍺ has fewer elements than ⍵ has axes, the ⍺ provided is assumed to describe how much to drop from the leading axes of ⍵, while the remaining axes are returned without dropping anything. Given a right argument ⍵ of rank three, the expression `2 3↓⍵` selects all but the first two positions of the first axis of ⍵, all but the first three positions of the second axis of ⍵, and all positions of the remaining third axis. The expression `2↓⍵` selects all but the first two positions of the first axis of ⍵, and all positions of the remaining axes. The expression `''↓⍵` selects all of the axes of ⍵.

Dropping Beyond the Limits of an Axis

When ⍺ includes a scalar whose magnitude is greater than the length of the corresponding axis of ⍵, you are in effect requesting to drop beyond the length limit of that axis. In the result, the length of that axis is zero: the excess value in ⍺ has no effect.

Drop Applied to a Scalar

When ⍵ is a scalar, but ⍺ is not empty, ⍵ is treated as if it had the rank implied by the length of ⍺, and every axis had length one:

```
      ⍴3 ¯7↓6
0 0
```

↓⍵ *Raze*

Rank: ∞

The formal definition of the ***raze*** function is as follows:

If `1<1↑(⍴⍵),1` then `,⍤¯1↓⍵≡⍤¯1↓⍵`
If `0=1↑(⍴⍵),1` then `(1↓⍴⍵)⍴<⍬`
If `1=1↑(⍴⍵),1` then `1≡⍤¯1↓(1↓⍴⍵)⍴⍵`

Consider *v*, a list of enclosed words:

```
      v
|-----| |---|
|Every| |one|
|_____| |___|
      ↓v
|--------|
|Everyone|
|________|
```

The rank of the result of ***raze*** is one less than the rank of its argument:

```
      m
0 1 2
3 4 5
6 7 8
```

```
      ↓m
|-----| |-----| |-----|
|0 3 6| |1 4 7| |2 5 8|
|_____| |_____| |_____|
      ↓↓m
|------------------|
|0 3 6 1 4 7 2 5 8|
|__________________|
```

The raze function may be used in *row-tolerant catenation*, which is useful when two matrices with differing lengths along the last axis are to be joined along the first axis:

```
      a
abc
def
      b
gh
ij
kl
```

We can achieve the desired effect using the functions ***raze*** and ***link***:

```
      >>↓<⍥⍨n⍨>a⊃b ⊣ n←¯1+(⍴⍴a)⌈⍴⍴b
abc
def
gh
ij
kl
```

To understand this expression, examine the intermediate results at each stage of execution. First link the two arrays to form a two element vector:

```
      ⊢x←a⊃b
|---| |--|
|abc| |gh|
|def| |ij|
|___| |kl|
      |__|
```

Then use the function derived from enclose (<⍥¯1 ⍵) to enclose the vectors of each matrix:

```
      ⊢y←<⍤¯1⍤1>x
┌──────────┐ ┌──────────────┐
│┌──┐ ┌──┐│ │┌──┐ ┌──┐ ┌──┐│
│└──┘ └──┘│ │└──┘ └──┘ └──┘│
│abc  def │ │gh   ij   kl  │
└──────────┘ └──────────────┘
```

Now use ***raze*** to join the outer enclosures in *y*:

```
      ↓y
┌────────────────────────┐
│┌──┐ ┌──┐ ┌──┐ ┌──┐ ┌──┐│
│└──┘ └──┘ └──┘ └──┘ └──┘│
│abc  def  gh   ij   kl  │
└────────────────────────┘
```

Finally, remove the outer enclosure altogether, then disclose the resulting enclosed array to produce the simple matrix:

```
      >>↓y
abc
def
gh
ij
kl
```

Without ***raze***, we would first make the two matrices conform in length, then join them:

```
      (((×/1↑⍴a),m)↓a),[⍳2]((×/1↓⍴b),m)↑b ⊣ m←1↓(⍴a)⌈⍴b
abc
def
gh
ij
kl
```

If the expression for row-tolerant catenation is captured in a user function called *ovr*, then a last-axis tolerant catenation may be achieved through composition of *ovr* with ***transpose***:

```
      a ovr⍤⍳ b
abcgh
defij
   kl
```

<ω Enclose/Box

Rank: ∞

The expression <ω *encloses* or ***boxes*** ω, thereby encoding it as a scalar. It has no axes, and fits into an array structure in the same way as any other element (such as a single number or character). Enclosed arrays let you handle variables of arbitrary rank and shape within a simple rectangular framework.

The display of an nested array shows the contents. The display is conditioned by the value of the system variable ⎕*ps* (position and spacing) (see the *System Guide, Chapter 5*). Setting the value ⎕*ps*←⁻1 1 2 2 (not the default) causes a display which shows each enclosure surrounded by characters that mark its boundaries. To make explicit the effect of enclosure, that option has been used in the examples that follow.

Forming a Enclosed Array

Suppose *n* is a matrix of characters and *n* a matrix of numbers:

```
      c
Tom
Jane
Harry
      n
1203 118  234.56  27678.08
 409  22 1074.60  23641.20
8004 165  482     79530
   0   0    0    130849.28
      ⍴n
3 5
      ⍴<c ⍝ enclose produces a scalar

      ⍴quantities
4 4
      ⍴<n ⍝ enclose produces a scalar
```

```
      ⍴x←(<c),(<n)
2
      ⊢⎕ps←¯1 1 2 2
¯1 ¯1 ¯2 ¯2
      x
|-----| |-------------------------|
|Tom  | |1203 118  234.56  27678.08|
|Jane | | 409  22 1074.60  23641.20|
|Harry| |8004 165  482     79530   |
|_____| |   0   0    0    130849.28|
        |_________________________|
```

An enclosed element is never equal to its contents. In particular, an element (for example, a number or character) is never equal to the same element enclosed:

```
     2≡<2
0
```

Moreover, an element may be enclosed to any level. Elements with different levels of enclosure are not considered equal:

```
     (<2)≡(<<2)
0
```

When the display shows lines around an enclosed element, the number of lines reveals the level of enclosure. For example, here is a display of a two element heterogeneous vector formed by catenating the number `1023` (disclosed) to an enclosed element formed by enclosing the two element vector *x* shown earlier:

```
      1023,<x
1023 |-----------------------------------|
     ||-----||-------------------------||
     ||Tom  ||1203 118  234.56  27678.08||
     ||Jane || 409  22 1074.60  23641.20||
     ||Harry||8004 165  482     79530   ||
     ||_____||   0   0    0    130849.28||
     |       |_________________________||
     |___________________________________|
```

⊃ω *Conditional Enclose/Box*

Rank: ∞

The expression ⊃ω ***conditionally encloses*** ω provided ω is a simple array, that is, is not enclosed and does not already contain any enclosed cells. If ω is already enclosed, or contains any enclosed elements, ***conditional enclose*** returns ω unchanged.

You can use the expression ~ω≡⊃ω to determine whether array ω is simple or not.

α⊃ω *Link*

Ranks: ∞ ∞

The expression α⊃ω encloses α, encloses ω *if it is not already enclosed,* and then catenates the two to build a nested array. The formal definition of ***link*** is:

```
α⊃ω ⇐⇒ (<α),⊃ω
```

The *conditional enclosing* of the right argument lets you use a simple expression to enclose and join several objects without requiring parentheses:

```
      v←'Tom'⊃'Dick'⊃'Harry'
     ⍴v
3
      v
┌───┬────┬─────┐
│Tom│Dick│Harry│
└───┴────┴─────┘
```

To add another enclosed element to *v*—for example, to add `'Jane'` to the list of names—use either of the following two expressions:

```
     ⊢v←'Jane'⊃v
┌────┬───┬────┬─────┐
│Jane│Tom│Dick│Harry│
└────┴───┴────┴─────┘
     ⊢v←v,⊃'Jane'
┌───┬────┬─────┬────┐
│Tom│Dick│Harry│Jane│
└───┴────┴─────┴────┘
```

Notice that the second expression uses ***conditional enclose*** rather than ***link***. If the comma is omitted, the expression returns a rather different result:

```
      ⊢v←v⊃'Jane'
┌──────────────────┐
│┌───┐┌────┐┌─────┐│┌────┐
│└───┘└────┘└─────┘│└────┘
│Tom Dick Harry    │Jane
└──────────────────┘
```

>ω *Disclose/Open*

Rank: 0

The ***disclose*** or ***open*** function (>ω) is the inverse of ***enclose***. It discloses whatever array lies within ω. If the argument of ***disclose*** is not enclosed, the result preserves the argument unchanged.

There are two important points of contrast between ***enclose*** and ***disclose***:

- While ***enclose*** has unbounded rank and returns a scalar result regardless of the rank of its argument, ***disclose*** has a rank of zero.
- The result of ***enclose*** never matches the argument to ***enclose***, that is, ω≡<ω is **never** true**.** The ***enclose*** function is sometimes called ***strict enclose*** to emphasize this behavior and contrast it with ***conditional enclose***. ***Disclose***, on the other hand, **can** return its argument unchanged—if the argument is not enclosed. In other words, is for unenclosed ω, ω≡>ω.

Given these points, it can be seen that ***enclose*** and ***disclose*** are not true inverses in every case.

>><ω ⇐⇒ ω ⇐/⇒ <>ω

Tolerant Cell-Extension

When ω is an array of enclosed elements, >ω discloses all of them. Since ***disclose*** has an argument rank of zero, any axes in ω are frame axes. The result returned by >ω has the same frame axes as ω, followed by the axes needed to contain the arrays disclosed from the various enclosed elements.

When all the enclosed elements contain arrays of the same rank and shape, the cell axes returned by >ω are simply the axes common to all of them. For example, when ω is a 5-by-2 matrix of enclosures, and each enclosure contains a 3-by-4-by-7 array, then the shape of >ω is `5 2 3 4 7`.

When the arrays disclosed by opening the various elements differ in rank or shape, every result cell is promoted to the rank of the highest-rank argument cell and then each axis is extended to match the maximum length. As a first step, the rank of the result cell is increased by adding as many as necessary leading axes of length one. The length of each axis is increased to bring it up to the maximum length for that axis among all the argument cells. Cells whose arrays have axes shorter than the maximum are padded in the same way as for the ***take*** function. Unused positions in a numeric array are filled out with zeros, in a character array with blanks, and in an enclosed or heterogeneous array with ***jots*** (the enclosed empty numeric vector). For example, suppose the three element vector *x* consists of boxes, which when opened separately, contain the following:

`0@x:` a 5-element character vector
`1@x:` a 3-by-3 integer matrix
`2@x:` a box (enclosure)

The ranks of the arrays in these three boxes, when opened, are respectively one, two, and zero. Adding leading axes of length one where necessary, the shapes of the three cells become:

```
1 5
3 3
1 1
```

The maxima for the two axes are `3 5`. Hence >*x* has frame shape `3` and cell shape `3 5`. The three result cells are therefore computed as:

`0@>x:` `3 5↑1 5⍴` the 5-element character vector
`1@>x:` `3 5↑3 3⍴` the 3-by-3 integer matrix
`2@>x:` `3 5↑1 1⍴` the enclosure (box)

,ω Ravel

Rank: ∞

The expression ,ω creates a vector containing all the elements of ω in row major order; for example:

```
      m
One
Two
Three
      ,m
One  Two  Three
```

The ravel of a scalar is a one element vector.

⍪ω Table

Rank: ∞

The expression ⍪ω creates a matrix in which each row is the ravel of a major cell of ω. For example, suppose *x* is a 2-by-3-by-5 array:

```
      x
abcde
fghij
klmno

pqrst
uvwxy
zabcd
```

Then the ***table*** of *x* is a 2-by-15 array:

```
      ⍪x
abcdefghijklmno
pqrstuvwxyzabcd
```

When ω is a vector, ⍪ω is a 1-column matrix; when ω is a scalar, ⍪ω is a 1-by-1 matrix:

```
      ⍪⍳3
0
1
2
      ⍴⍪⍳3
3 1
      ⍪3
3
      ⍴⍪3
1 1
```

α,ω Catenate
α⍪ω Catenate-Down

Ranks: ∞ ∞

The expression α⍪ω catenates the major cells of α and ω along their first axes. The expression α,ω functions in similar fashion, except that it catenates the contra-major cells along their last axes. Applied to matrices, α⍪ω catenates them vertically, while α,ω catenates them horizontally. In general, the arrays thus joined must be of the same rank. However, their ranks may differ by 1, or one of them may be a scalar.

Result Type Produced by Catenate and Catenate-Down

The arrays being catenated may be of any type. When α and ω are homogeneous and both are of the same type (character, numeric, or enclosed), the result has that type also. When α and ω are of different types, the result is heterogeneous. When both arguments are entirely numeric but use different internal representations for numbers (Boolean, integer, or floating), the result is in the internal representation of whichever argument had the more extensive representation.

Shape of Arguments and Result

The arrays being catenated must have the same length in all their axes except the axis along which they will be catenated. Arrays catenated by α⍪ω must have the same length in all axes except the first, and arrays catenated by α,ω must have the same length in all axes except the last.

The result of $\alpha\overline{,}\omega$ is an array for which the length of the first axis is the sum of the lengths of the first axes of α and ω, and all other axes have the same lengths that they have in α and in ω. For example:

```
      a
abc
def
      b
uvw
xyz
      a⍪b
abc
def
uvw
xyz
      ⍴a⍪b
4 3
```

Similarly, the result of α,ω is an array for which the length of the last axis is the sum of the lengths of the last axes of α and ω, and all other axes have the same lengths that they have in α and in ω:

```
      a,b
abcuvw
defxyz
      ⍴a,b
2 6
```

When one of the arrays has rank one less than that of the other, it is treated as if it had an extra leading axis (in the case of ***catenate-down***) or an extra trailing axis (in the case of ***catenate***) of length one. For example:

```
      a⍪'xyz'
abc
def
xyz
      a,'xy'
abcx
defy
```

When one of the arrays is a scalar it is replicated to match the lengths of all the axes other than the one along which catenation will occur:

```
      a⍪'*'
abc
def
***
```

When both arguments are vectors, the distinction between ***catenate*** and ***catenate-down*** vanishes since, with only one axis, the first axis *is* the last axis. When both arguments are scalars, the result is a two element vector. This is the **only** case in which the rank of the result is greater that the rank of either of the arguments.

Catenating Two Arrays Along a New Axis

When you join ⍺ and ⍵ along a new axis, you are said to *laminate* them. ***Laminate*** does not refer to a new function, but to different way in which ***catenate*** and ***catenate-down*** are used.

To laminate two arrays along a new last axis, it is necessary that they have the same shape. Use ,⍤0 or ⍪⍤0 to ravel each element of ⍺ and ⍵. Catenating each element of ⍺ to the corresponding element in ⍵ produces an array with a trailing axis of length two:

```
      'abc',⍤0'xyz'
ax
by
cz
      ⍴'abc',⍤0'xyz'
3 2
```

Joining two arrays along a new first axis can be achieved in several ways. One is to compose ,⍤0 with ⍉, which reverses the order of axes of ⍺ and of ⍵, joins element by element, and then reverses the transposition:

```
      'abc',⍤0⍤⍉'xyz'
abc
xyz
```

Alternatively, joining the vector ⍺ to a 1-row matrix formed from the vector ⍵ has the same effect, which exploits the ***table*** function described above.

```
      'abc'⍪⍉⍪'xyz'
abc
xyz
```

Another possibility is: `⍬'abc',⍤0'xyz'`

Still another approach to creating a new first axis is to exploit the rules for opening an enclosed array:

```
      >'abc'⊃'xyz'
abc
xyz
```

Because of the flexibility of the tolerant frame-builder in SHARP APL for UNIX, the arrays being joined need not have a common shape:

```
      ⊢x←>'Long list of characters'⊃'Shorter list'
Long list of characters
Shorter list
      ⍴x
2 23
```

See the section on the ***raze*** function for examples of row-tolerant catenation.

Catenate Modified by Bracket-Axis Notation

Earlier APL systems permitted certain functions to be modified by a number in brackets placed after the function symbol. The functions ***catenate*** and ***catenate-down*** were among those that could be so modified. The value inside the brackets was a numeric scalar identifying the axis along which the arguments were to be joined. With the axis thus explicitly identified, `⍺,[i]⍵` had exactly the same effect as `⍺⍪[i]⍵`. As usual, the arguments were required to have the same length along the axes other than the axis along which they were chained to form the result. The value in brackets was an integer scalar identifying an axis, and therefore affected by the `⎕io` (index origin). For example, to join along the second axis, use `⍺,[1]⍵`.

To use bracket-axis notation to join along a new axis, identify the new axis by a non-integer number that falls somewhere between the existing axis numbers. For example, to join along a new axis to be placed between the first and second axes use `⍺,[0.5]⍵`. To join along a new axis to be created ahead of the first axis, `⍺,[-0.5]⍵`. The fraction does not have to be `0.5`; the requirement is that an interpolated axis differ from the number of an existing axis by an amount less than `1`. For compatibility with existing systems, bracket-axis modification of ***catenate*** and ***catenate-down*** is still permitted in SHARP APL for UNIX, but is considered obsolescent.

ϕω *Reverse*

Rank: 1

The ***reverse*** function (ϕω) reverses the order of elements along the last axis of its argument, that is, ϕω reverses the order of every vector of ω. For example:

```
      v
0 1 2 3 4
      ϕv
4 3 2 1 0
      m
0 1 2
3 4 5
      ϕm
2 1 0
5 4 3
```

⊖ω *Reverse-down*

Rank: ∞

The ***reverse-down*** function (⊖ω) reverses the order of elements along the first axis of its argument. ***Reverse-down*** can be defined in terms of the ***reverse*** function:

⊖ω ⟸⟹ ϕ⍤¯1⊖ω ⟸⟹ ⊖ϕ⊖ω

Applied to a vector, ***reverse*** and ***reverse-down*** give the same result. Applied to a matrix, independently within each column, the positions of elements are reversed:

```
      m
0 1 2
3 4 5
      ⊖m
3 4 5
0 1 2
```

You should not think of the preceding as *putting the last row first* because that suggests that what has been changed is the order of the vectors `0 1 2` and `3 4 5`. The effect is better understood as reversing the column `0 3`, the column `1 4`,

and the column `2 5`. Conceiving the effect as something that happens within each column independently is essential to understanding the effect of ***rotate-down*** (⍺⊖⍵).

⍺⌽⍵ *Rotate*

Ranks: 0 1

The expression ⍺⌽⍵ returns a cyclic rotation ⍵. The left argument is an integer indicating the amount and direction through which ⍵ is rotated. When ⍺ is positive ⍵ is rotated backwards; when ⍺ is negative, ⍵ is rotated forwards. For example:

```
      ⊢v←⍳11
0 1 2 3 4 5 6 7 8 9 10
      2⌽v
2 3 4 5 6 7 8 9 10 0 1

      ¯2⌽v
9 10 0 1 2 3 4 5 6 7 8

      11⌽v
0 1 2 3 4 5 6 7 8 9 10
```

Notice that rotation by `11` (the number of elements in *v*) has no net effect: the amount of rotation is treated as ⍺|⍴⍵.

Rotate Applies to Vectors

Since ***rotate*** has argument ranks `1 0 1`, it follows that when ⍵ has more than one axis, the additional leading axes are frame axes, and the last axis contains the cells. When *m* is rotated, since it contains four cells (each of which is a `5`-element vector), its left argument must be a `4`-element vector indicating by how much to rotate each cell:

```
      ⊢m←4 5⍴'One  Two  ThreeFour'
One
Two
Three
Four
      ¯1 0 2 3⌽m
```

```
 One
Two
reeTh
r Fou
```

Since ⍵ is composed of 1-cells while ⍺ is composed of 0-cells, in general the shape of ⍺ is ¯1↓⍴⍵. Rotate generalizes in the usual fashion to frames of higher rank. When ⍺ is a single element, it applies to each of the last-axis vectors in ⍵, thus:

```
      1⌽m
ne  O
wo  T
hreeT
our F
      ¯1⌽m
 One
 Two
eThre
 Four
```

⍺⊖⍵ *Rotate-down*

Ranks: ∞ ∞

The ***rotate-down*** function (⊖⍵) rotates the order of elements along the first axis of its argument. ***Rotate-down*** can be defined in terms of the ***rotate*** function:

```
⍺⊖⍵ ⟸⟹ ⍺⌽⍤⍉⍤⍉⍵ ⟸⟹ ⍉(⍉⍺)⌽⍉⍵
```

Applied to a vector right argument, ***rotate*** and ***rotate-down*** give the same result. Applied to a matrix right argument, independently within each column, the positions of elements are rotated:

```
      ⊢m←3 5⍴0 0 0 0 0 1 1 1 1 1 0 0 0 0 0
0 0 0 0 0
1 1 1 1 1
0 0 0 0 0
      ¯1 0 1 0 ¯1⊖m
1 0 0 0 1
0 1 0 1 0
0 0 1 0 0
```

In general, the left argument ⍺ must contain an element for each of the column vectors of ⍵, so the shape of ⍺ is 1↓⍴⍵. When ⍺ contains a single element, it is applied to all the column vectors in ⍵:

```
      1⊖m
0 0 0 0 0
0 0 0 0 0
1 1 1 1 1
```

⍉⍵ *Monadic Transpose*

Rank: ∞

The ***monadic transpose*** function *reverses the order of the axes* of its right argument. When ⍵ is a vector, there is only one axis, so no changes are possible in the ordering of axes. When ⍵ is a scalar, there are no axes to transpose, and the interpreter returns it unchanged.

When ⍵ is a matrix, the ***monadic transpose*** of ⍵ has rows where ⍵ has columns, and columns where ⍵ has rows:

```
      ⍵
Copenhagen
Libreville
Luxembourg
      ⍉⍵
CLL
oiu
pbx
ere
nem
hvb
aio
glu
elr
neg
```

The ***monadic transpose*** of a rank three array interchanges the first and last axes of the array. For example, when ⍵ is a 2-by-3-by-6 array, ⍉⍵ is a 6-by-3-by-2 array:

```
      ω
London
Athens
Ottawa

Peking
Moscow
Prague
      ⍒ω
LP
AM
OP

oe
to
tr

nk
hs
ta

di
ec
ag

on
no
wu

ng
sw
ae
```

The shape of the result is the reverse of the shape of the argument:

```
ρ⍒ω ⇐⇒ ⌽ρω
```

α⍉ω *Dyadic Transpose*

Ranks: 1 ∞

The ***dyadic transpose*** function *permutes the order of the axes* of its right argument. The left argument specifies how the axes are to be permuted, and may cause axes to be mapped together to take a *diagonal slice* through ω. The left argument α is a vector of integers which indicate what is to become of each of the axes of ω. The vector must contain one element for each axis of ω; that is, its length must equal ⍴⍴ω. When ω has rank one, the system also accepts a *scalar* for α where, strictly speaking, it should require a one element vector. However, transposition has no effect unless ω has at least two axes.

The positions within the left argument α correspond to the various axes of ω. That is, the first position within α describes what is to be done with the first axis of ω. The values within α refer to the axes of the result. Depending on the value of α, the result rank is always less than or equal to the right argument rank.

The elements of α must be integers such that (when the desired result has rank n) α includes each member of ⍳n at least once. For example, when you want a rank three result, `0 2 1 2` is a possible value for α, whereas `0 3 2 3` is not (because it omits `1`).

The result returned by α⍉ω is an array formed by permuting or mapping together the axes of ω. The result contains one axis for each unique number in α. When α contains no repetitions, the rank of the result is the same as the rank of the right argument ω.

Suppose z←α⍉ω. Then:

```
⍴ω ⇐⇒ (α-⎕io)@⍴z
```

When the numbers in α are distinct, all axes of ω are included in the result; only their order is changed. Suppose ω has four axes. Then the expression `2 0 3 1` ⍉ω says that:

- axis `0` of ω becomes axis `2` of the result
- axis `1` of ω becomes axis `0` of the result
- axis `2` of ω becomes axis `3` of the result
- axis 3 of ω becomes axis `1` of the result.

When the left argument is a vector of consecutive integers in ascending order, the axes are returned in the same positions they had before:

```
(⍳⍴⍴⍵)⍉⍵ ⟸⟹ ⍵
```

When the left argument contains consecutive integers in *reverse order,* the sequence of the axes of ⍵ is reversed, which is the same as the result from ***monadic transpose***:

```
(⌽⍳⍴⍴⍵)⍉⍵ ⟸⟹ ⍉⍵
```

When ⍺ contains repetitions, the number of distinct numbers in ⍺ is less than the rank of ⍵, so the rank of the result is also less than the rank of ⍵. Thus, although the rank of the result may not be greater than the rank of ⍵, it may be less.

Consider what happens when the left argument ⍺ contains repetitions, as it would if ⍺ is `0 2 1 2`. The value `2` occurs both at position `1` and at position `3`. That means that axis two of the result is to be formed from both axis one and axis three of ⍵. Two distinct axes of ⍵ are mapped into a single axis of the result. That is done by taking from ⍵ only those elements whose position is the same on both axes.

By mapping axes together, you are *selecting a diagonal.* For a four axis array ⍵, the expression

```
0 0 0 0⍉⍵
```

gives you a vector formed by

```
((<0 0 0 0)@⍵),((<1 1 1 1)@⍵),((<2 2 2 2))@⍵)
```

and so on. Similarly, for a matrix ⍵, `0 0`⍉⍵ is the vector forming the main diagonal.

When the axes mapped together differ in length, the number of positions common to both axes is only as great as the length of the shortest axis. For example, if ⍵ is a `3`- by-`10` matrix of characters, then `0 0`⍉⍵ picks out only the elements shown as ⊛ and ignores those shown as ∘:

```
⊛∘∘∘∘∘∘∘∘∘∘
∘⊛∘∘∘∘∘∘∘∘∘
∘∘⊛∘∘∘∘∘∘∘∘
```

Index Origin and ⍉

Since the elements of the left argument of ***dyadic transpose*** are drawn from consecutive integers up to the desired rank of the result, they depend on the ⎕*io*, index origin.

Miscellaneous Functions

⊢*ω* *Pass*

Rank: ∞

The function ***pass*** has an extremely simple definition:

```
⊢ω ⇐⇒ ω
```

Pass is used primarily to pass an argument cell into a composition without changing it. For example, the ***cut*** operator (*m*⍤*g*) modifies a monadic function to apply to each segment of an array argument. If you wish merely to select the segments without transforming them, use ***pass*** as the functional argument to ***cut***.

The following example selects each word (that is, each sequence of characters delimited by a blank) from a character vector and puts it into a frame, thereby forming a matrix:

```
      ¯1⍤⊢' ','A list of words'
A
list
of
words
```

Another common use of ***pass*** is to force the display of an assignment:

```
      ⊢x←⍳6
0 1 2 3 4 5
```

Pass can also be used to separate adjacent numeric constants:

```
      x←1 2 3,⍤0⊢4 5 6
```

⊣⍵ *Stop*

Rank: ∞

The ***stop*** function takes any array and returns `0 0⍴0`. ***Stop*** is used primarily to discard an unwanted result:

```
⊣⎕ex ⎕nl 2 3
```

⍺⊣⍵ *Left*

Ranks: ∞ ∞

The ***left*** function returns its left argument and ignores its right argument:

```
⍺⊣⍵ ⍝⍝ ⍺
```

Left is used primarily to assign a value that can not otherwise fit syntactically into an expression. For example, mathematicians are fond of statements in the form "*x* is a function of *y,* where *y* is . . . ." That is really two statements, linked by *where.* A similar construction in APL might be expressed as:

```
x←(1+y)*2⊣y←⍳10
```

The ***left*** function can often be used instead of the ***diamond*** symbol (⋄) to combine statements.

⍺⊢⍵ *Right*

Ranks: ∞ ∞

The ***right*** function returns its right argument and ignores its left argument:

```
⍺⊢⍵ ⍝⍝ ⍵
```

One use of the ***right*** function is to embed comment-like text within an APL statement. In the following example, the text in quotes has no effect upon the result:

```
a←⍵×⊢'count to n'⊢⍳n←⊢'length of r'⊢⍴r
```

≠ω Nubsieve

Rank: ∞; Implicit Arguments: □ct

The ***nubsieve*** function (≠ω) produces a Boolean vector that marks the ***nub*** (the unique major cells) of ω. For example, with a vector argument, the unique scalars are marked. With a matrix, the unique vectors of the matrix are marked, and so on:

```
     ≠'abacus'
1 1 0 1 1 1
     b
mabel
toto
lisa
fred
lisa
fred
dorothy
     ≠b
1 1 1 1 0 0 1

     (≠b)≠b
mabel
toto
lisa
fred
dorothy
```

The commonly used expression for the nub of a vector

```
((⍳⍴ω)=ω⍳ω)/ω
```

can be written more efficiently and compactly using nubsieve as

```
(≠ω)/ω
```

or simply as

```
↑ω
```

↑ω *Nub*

Rank: ∞; Implicit Arguments: □⎕ct

The expression ↑ω selects the *nub* (unique major cells) of ω. Formally:

```
      ↑ω ⇐⇒ (≠ω)⌿ω

      a
abc
abc
def
      ↑a
abc
def
      ↑3
3
      ⍴↑3
1
```

=ω *Nubin*

Rank: ∞; Implicit Arguments: □⎕ct

The expression =ω classifies the major cells of the nub of ω according to equality with the major cells of ω, producing an *m* by *n* boolean matrix, where *m* and *n* are the number of major cells for the nub of ω (that is, ↑ω) and of ω, respectively.

```
      a
allah
      ↑a
alh
      =a
1 0 0 1 0
0 1 1 0 0
0 0 0 0 1
      b
abc
def
abc
      ↑b
abc
```

```
def
abc
     =b
1 0 1
0 1 0
```

Formally

```
=ω ⇐⇒ (MC ↑ω)∘.(≡⍤0) MC ω
```

where `MC ω` is defined as `,<⍤¯1 ω`.

α~ω *Less*

Ranks: ∞ ∞; Implicit Arguments: `⎕ct`

The result of α~ω is the array whose cells are the major cells of α *less* the major cells of ω.

```
     (⍳6)~7 2 4
0 1 3 5
      a←3 4⍴'abcdefghijkl'
      b←2 4⍴'mnopabcd'
      a~b
efgh
ijkl
```

Formally,

```
α~ω ⇐⇒ (~(MC α)∈MC ω)≠α
```

where `MC ω` is defined as `,<⍤¯1 ω`.

The expression of α~α~ω is gives the ***intersection*** of α and ω:

```
     a←⍳6
     b←7 4 2
     a~a~b
2 4
```

The result of α~ω is always rank one or higher even if that result is a singleton:

```
      1 2 3~1 3
2
      ⍴1 2 3~1 3
1
      10~20
10
```

@⍵ *All*

Rank: 1

The monadic function ***all*** forms a catalogue of the elements of its argument. It returns a boxed cartesian product of the opened elements of the argument. For example:

```
      @ 'ab'⊃'cd'
|──||──|
|ac||ad|
|──||──|
|──||──|
|bc||bd|
|──||──|

      ⍴@⍳¨>⍳6
1 2 3 4 5 6

      ⍴@2 3⍴⍳¨>⍳6
2 4 5 6

      @ 'abc'
|───|
|abc|
|───|
```

⍎⍵ *Monadic Format*

Rank: ∞; Implicit Arguments: ⎕*pp* ⎕*ps*

Monadic format converts an array that contains elements that are other than characters to a representation constructed entirely of characters.

- When ⍵ is already a homogeneous array of characters, ***monadic format*** returns ⍵ unchanged; that is, ⍵≡⍕⍵.
- When ⍵ is a vector all of whose elements are numbers, the interpreter formats each number independently, following the same rules that it applies to a single numeric scalar. It separates adjacent elements in the vector by one blank. There is neither a leading blank before the vector nor a trailing blank at the end.
- When ⍵ is a numeric scalar (a single number with no axes), the interpreter converts it to characters as follows:

A number whose magnitude is less than `10*⎕pp` but not less than 10^{-6} is represented in decimal format with no more than ⎕*pp* significant digits; no trailing zeros to the right of the decimal point; no leading zeros to the left of the decimal point, *except* that a magnitude smaller than one has a leading zero in the units position. The decimal point is omitted unless there is a significant digit to its right. A negative value is preceded by a ***high minus*** (¯).

A number whose magnitude is less than 10^{-6}, or greater or equal to `10*⎕pp`, is represented in *exponential* format. It always has a decimal point, with one digit to the left of the point and `1⌈⎕pp-1` digits to the right (preceded by a ***high minus*** when the value is negative) and followed by the letter *e* and the *exponent* (which itself may include a ***high minus***), and up to three digits.

```
      ⎕pp←4
      ⊢x←-0 1e7,(÷3)×10*¯7+⍳15
0 ¯1.0e7 ¯3.333e¯8 ¯3.333e¯7  ¯0.000003333 ¯0.00003333
  ¯0.0003333 ¯0.003333 ¯0.03333 ¯0.3333 ¯3.333
  ¯33.33 ¯333.3 ¯3333 ¯3.333e4 ¯3.333e5 ¯3.333e6
```

Formatting Higher-Rank Arrays

For an array of rank two or greater, APL aligns all columns, then restores the leading axes (if there were any) so that (in general) the result of formatting has the same rank as the argument. The interpreter picks either decimal or exponential format. An array any of whose elements requires exponential format is shown in exponential format throughout. For arrays in decimal format, a field width is computed sufficient to accommodate the actual values in the array. A

decimal representation may occupy up to `1+⎕pp` positions to the left of the decimal point, and up to `5+⎕pp` positions to the right, so the maximum printed width of an array element is `7+(2×⎕pp)`.

Blanks Between Columns

One additional blank position is appended at the start of the representation of each column *except the first.*

Example of Monadic Format

The following example illustrates the points just mentioned. In particular, you can see that the same value (for example, the `0`s in the second plane) appear in either decimal or exponential format (depending on what other values are in the same column) and that the space allotted to a representation also depends on the other numbers represented in that column, for example `3333` in column `0` and in column `3`:

```
      t←÷3
      x←,1e4 1e¯5 1×t,t,¯1
      x←x,1 2 3
      x←x,1 2 3333
      x←x,t,t,¯1e9
      x←x,t,t, 1e199
      ⍉2↑(1,⍴x)⍴x⊣⎕pp←4
 3.333e3    1.000e0    1.000e0    3.333e¯1   3.333e¯1
 3.333e¯6   2.000e0    2.000e0    3.333e¯1   3.333e¯1
¯1.000e0    3.000e0    3.333e3   ¯1.000e9    1.000e199

 0.000e0    0.000e0    0.000e0    0.000e0    0.000e0
 0.000e0    0.000e0    0.000e0    0.000e0    0.000e0
 0.000e0    0.000e0    0.000e0    0.000e0    0.000e0
```

In displaying an array of rank greater than two, the session manager inserts an extra blank line to separate successive matrices, two extra blank lines to separate successive rank three arrays, and so on. Such a blank line is visible in the example above. Any *newline* characters required to produce the separation are supplied by the session manager and are **not** part of the result of ***monadic format*** itself.

Format of a Heterogeneous Array

When ⍵ is a heterogeneous array (that is, one containing elements from more than one of the types number, character, or enclosed), each of the elements is formatted independently without regard to any of the others. For each element, a character array is generated having the same appearance as ⍕⍵ for that element.

The representation of a single character is the character itself, while the representation of a number is a vector of one or more characters. The representation of an enclosed element may be a matrix (see below). Hence, the representations of numbers vary in width, and the representations of boxes vary in both height and width.

In the result, each of the representations thus produced, one for each element in the array, is arranged in a rectangular grid with a horizontal band of consecutive rows corresponding to each row of ⍵, and a vertical band of consecutive columns corresponding to each column of ⍵. Each vertical band is given a width sufficient to accommodate the widest representation in that column; each horizontal band is given a height sufficient to accommodate the tallest matrix in that row. This preserves the alignment of the boxes in the display of the frame of the array.

To separate the various representations, additional blanks may be inserted between adjacent horizontal or vertical bands. In addition, elements that are enclosed may be surrounded with characters that mark their horizontal or vertical boundaries (in effect, drawing a box around them). The amount of additional separation and the presence of delimiting markers around boxes are specified by the last two elements of the system variable ⎕*ps*, which controls spacing.

The ***magnitude*** of `2@⎕`*ps* controls the number of positions separating ***vertically*** adjacent elements. When `2@⎕`*ps* is ***negative***, lines are drawn at the top and bottom edges of adjacent elements. Similarly, the ***magnitude*** of `3@⎕`*ps* controls the number of positions separating ***horizontally*** adjacent elements, and when `3@⎕`*ps* is ***negative***, lines are drawn at the top and bottom edges.

To illustrate the effect of independent element-wise formatting in heterogeneous arrays, consider the various representations of the array *m*:

```
      m←((÷3)×10*1+⍳5)∘.×1 ¯1
      ⎕pp←5
      m
    3.3333      ¯3.3333
   33.333      ¯33.333
  333.33      ¯333.33
```

```
  3333.3      ¯3333.3
 33333       ¯33333
      □ps←2↑□ps,0 1   ⍝ One blank horizontal separation
      m,⍤1 'Cat'
3.3333 ¯3.3333 C a t
33.333 ¯33.333 C a t
333.33 ¯333.33 C a t
3333.3 ¯3333.3 C a t
33333   ¯33333  C a t
      □ps←2↑□ps,0 0   ⍝ No blank horizontal separation
     m,⍤1 'Cat'
3.3333¯3.3333Cat
33.333¯33.333Cat
333.33¯333.33Cat
3333.3¯3333.3Cat
33333 ¯33333 Cat
```

Position of Individual Representations

When the matrices representing the various representations within a row or column differ in size, a particular representation may be smaller than the rectangle available to it in the grid. The positioning of each element's representation within the available rectangle is controlled by the first two elements of □*ps*:

```
          ¯1  at the top of the rectangle
0@□ps      0  vertically centred in the rectangle
           1  at the bottom of the rectangle

          ¯1  at the left of the rectangle
1@□ps      0  horizontally centred in the rectangle
           1  at the right of the rectangle
```

For example, consider the heterogeneous array formed by joining characters to either side of a one-column array of numbers:

```
      n←⍪1 12 123 123e9 123456
      □ps←¯1 ¯1 0 1
      '□',n,'□'
□ 1         □
□ 12        □
□ 123       □
□ 1.23e11   □
```

```
□ 123456   □
      □ps←1 1 0 1
      '□',n,'□'
□       1 □
□      12 □
□     123 □
□ 1.23e11 □
□  123456 □
```

Format of an Enclosed Element

To display an enclosed element, SHARP APL first discloses it to reveal the contents. If need be, the process of disclosure is continued recursively until no more enclosures remain. Each of the elements thus disclosed is formatted independently in the manner already described for heterogeneous arrays.

When `2@□ps` or `3@□ps` is negative, the characters `¯` and `_` are used to draw borders above and below boxes, respectively, and `|` is used to draw borders at the sides. Where a box touches the edge of the rectangular space assigned to it in the grid, the border characters surrounding the box replace the first of the separating blanks on that edge. For that reason, border characters may be drawn around a box only when there is space for them (when the magnitude of the controlling element in `□ps` is at least two).

At the edges of the result of `⍕⍵` there are no separating blanks. There, the border characters are placed in an additional row or column. Thus, when `2@□ps` is `¯2` or less, the result has two or more extra rows, one at the top and one at the bottom. Similarly, when `3@□ps` is `¯2` or less, the result has two or more extra columns, one at the left and one at the right edge. The following example illustrates the effects of `□ps` on a nested array:

```
      b←2 2⍴(3 5⍴'AB')⊃1 2⊃'CD'⊃'EFGHIJ'
      □ps←¯1 1 3 3

|¯¯¯¯¯| |¯¯¯|
|ABABA| |1 2|
|BABAB| |___|
|ABABA|
|_____|

|¯¯|    |¯¯¯¯¯¯|
|CD|    |EFGHIJ|
|__|    |______|
```

```
      □ps←0 0 0 0
      b
ABABA
BABAB 1 2
ABABA
 CD  EFGHIJ
```

α⍕ω *Dyadic Format*

Ranks: ∗ ∞; Implicit Arguments: □*fc*

The expression α⍕ω yields a character array; the appearance of the result is controlled by α in two different ways, depending upon whether the type of the left argument α is empty, numeric or character.

Dyadic Format with Empty α

When α is empty, ***dyadic format*** returns exactly the same result as ***monadic format***.

Dyadic Format with Numeric α

When α is numeric, ***dyadic format f***ormats each of the columns of ω according to the specifications encoded in α. As with ***monadic format***, *column* refers to a position along the last axis. That is, the interpreter treats ω as if it were `((×/¯1↓⍴ω),¯1↑⍴ω)⍴ω`.

The left argument α requires two integers to specify the format of each column of ω. In SHARP APL for UNIX α is an integer vector whose length is twice that of the last axis of ω.

Each *pair of integers* in α controls the formatting of one column of ω. Alternatively, α may be a single pair, which then applies to all the columns of ω, or a single integer, in which case it is treated as if it were the pair `¯2↑α`.

Unlike SHARP APL for OS/390, which requires that the number of pairs be the same as the number of columns in ω, SHARP APL for UNIX *recycles* the pairs to accommodate all the columns of ω. Within each α-pair, the first element controls the *width* of the resulting representation. The second element controls the type of representation: the number of decimal places (*positive* for decimal representation); the number of significant digits (*negative* for exponential representation).

Both types of representation are *right-justified* within their fields. The only mechanism to insert blanks between the representations of consecutive columns is to assign additional width through the first element in a pair; the extra width is filled with blanks on the left side.

Decimal Representation

The first member of an ⍺-pair specifies the *total width* of the field. However, when the first member of the field is zero, the width is made wide enough to accommodate the widest representation in that column, plus one leading blank. An *additional position* in the field width is required for the decimal point itself and for ***high minus*** (¯) if required. Unused positions to the right of the decimal point are filled with trailing zeros.

The second member of an ⍺-pair specifies the number of places to the right of the decimal point.

If the value from ⍵ cannot be represented in the number of positions specified, the field is filled with the character ⋆ to indicate *overflow.* Consider the following example:

```
      ⎕pp←5
      ⍕m
     3.3333      ¯3.3333
    33.333      ¯33.333
   333.33      ¯333.33
  3333.3      ¯3333.3
 33333       ¯33333
      '[',(8 2⍕m),']'
[    3.33    ¯3.33]
[   33.33   ¯33.33]
[  333.33  ¯333.33]
[ 3333.33 ¯3333.33]
[33333.33********]
```

Exponential Representation

When the second member of an ⍺-pair is negative, the result is represented in *exponential format.* The value of `0@⍺` sets the *total width* of the field. The absolute value of `1@⍺` sets the *number of significant digits.*

In an exponential representation, positions are allocated as follows:

- The *last four* positions of each field are reserved for the exponent and its sign. They are set flush left within those four positions, and unused positions are filled with blanks.
- The *fifth position from the end* is reserved for the character *e*.
- In the remaining leftward positions, there must be space for the decimal point; the number of digits specified in `1⌾⍺`, the first of which appears to the left of the decimal point, and the others between the decimal point and the character *e*; and a ***high minus***, if required, for the value being represented.

In general, `0⌾⍺` (the total width) must be at least seven greater than the number of significant digits requested. However, if the value being represented is not negative, a width one less suffices. Any further unused positions at the left are filled with blanks. The following illustrates some of the points mentioned:

```
      t←⍕⍤¯1(10÷3)×1 ¯1 1e¯300
       t←t,t
       '[',(9 ¯3 16 ¯3⍕t),']'
[3.33e0           3.33e0   ]
[*********       ¯3.33e0   ]
[3.33e¯300         3.33e¯300]
```

Dyadic Format with Character ⍺

When ⍺ is a character vector, the expression ⍺⍕⍵ is called *format by example*: ⍺ provides a *pattern* for the result. The right argument ⍵ is a numeric array of any rank. The left argument ⍺ is a *character vector* of the *same length* as the *last axis* of the desired result, and shows how each column of ⍵ is to be represented.

```
      ⊢⍺←'Balance is $(55,510.50) on 05/55/55'
Balance is $(55,510.50) on 05/55/55
      ⍺⍕23456.714, 100⊥84 6 11
Balance is   23,456.71  on 84/06/11
      ⍺⍕456.78, 100⊥84 6 11
Balance is      456.78  on 84/06/11
      ⍺⍕¯23456.714 840611
Balance is $(23,456.71) on 84/06/11
      ⍺⍕¯456.78 840611
Balance is    $(456.78) on 84/06/11
      ⊢r←'555.55'⍕⍤¯1 0 10.1 100
   1
```

```
  10.1
 100
     ⍴r
4 6
```

You construct ⍺ in terms of a series of fields, each of which must contain at least one digit; a field is bounded by blanks or (on the right) by the first nondigit following a 6. The digits in a field are both *place–holders* and *controls* for that field. Any nondigits are *decorators*.

Simple decorator. A simple decorator may be embedded in a digit field or stand alone; it is reproduced in place.

Controlled decorator. A controlled decorator is a group of one or more characters immediately adjacent to a leading or trailing digit in a field that contains a 1, a 2, or a 3 which is nearer that decorator than a 4. For example:

```
      x←1 ¯1∘.×123 321
      'Ctrl55154Not  Not54321Ctrl'⍕x
    123Not  Not    321
Ctrl123Not  Not    321Ctrl
```

Conventional decorator. The dot and comma are conventional decorators; they specify *decimal digits* or *group separators* according to common conventions (as in 23,456.78). A dot is a conventional decorator if it either precedes the first digit or is surrounded by digits, and if it is the only such dot in the field. A comma is a conventional decorator if it is surrounded by digits. For example:

```
      '5...55.55'⍕⍪1 12.3 123.45
 ... 1
 ...12.3
1...23.45

      'LF,55,525,,5,RT'⍕⍪12345 ¯12345 123 ¯123
LF,1,234,,5,RT
   1,234,,5
   LF,12,,3,RT
      12,,3
```

The characters employed to display these conventional dots and commas are controlled by certain elements of the system variable ⎕*fc* (format control) as follows:

`0@⍺⍵fc` Separate integer from fractional part.

`1@⍺⍵fc` Separate digits where `⍺` (in `⍺⍕⍵`) contains "`,`" (comma) surrounded by digits.

The default value (in a clear workspace) of `0 1@⍺⍵fc` is `'.,';` reversal of these values provides printing according to a common European convention.

A decimal point not followed by a fractional value is a part of the trailing decorator, and as such is suppressed according to the rules that apply to any trailing decorator. For example:

```
     '-5125.POS'⍕¯123 ¯123
 123.POS
-123
```

Control Digits for ⍕

`'0'⍕⍵` **Leading/trailing zeros.** Unused positions are filled with zero from here toward the decimal point:

```
     '55,550.099'⍕1 0 1000.1
    1.000
    0.0
 1,000.100
```

`'1'⍕⍵` **Negative values.** The decorator text on this side (of the decimal point) is included and floats (to touch the digits of the number); otherwise, it is replaced by blanks. The effects of the digits `1`, `2`, and `3` do not depend in any way on the decimal point. For example:

```
     'THIS1555.55HERE'⍕12 ¯12
   12
THIS12HERE
```

`'2'⍕⍵` **Nonnegative values.** The decorator text on this side is included and floats; otherwise, it is replaced by blanks.

`'3'⍠⍵` **Float.** For all values, the decorator text on this side floats.

`'4'⍠⍵` **No float.** For all values, the decorator text on this side does not float; nullifies any floating specified by a `1`, `2`, or `3` further from this side:

```
      ⍺
→→55125.5←← →→54124.5←← →→54125.5←← →→55124.5←←
      ⍺⍠1 ¯1∘.×4⍴1234
   1234←←   →→ 1234  ←← →→ 1234←←      1234  ←←
 →→1234     ********** →→ 1234      →→1234  ←←
```

`'5'⍠⍵` **Normal digit.** The position is available for a digit, or for a sign decorator (when a sign decorator has been supplied and has been selected by the presence of a `1`, `2`, or `3`).

`'6'⍠⍵` **Field delimiter.** The next rightward noncontrol digit (any character other than `0 1 2 3 4 5 6 7 8 9` or `.`) starts a new field. This is needed where fields are not separated by blanks. The digit `6` provides an additional field break at the next noncontrol digit. For example, this allows you to include slashes rather than blanks as field delimiters in a date:

```
     '06/06/05'⍠7 11 84
07/11/84
```

`'7'⍠⍵` **Exponential format.** The value is represented in exponential notation; next rightward noncontrol digit separates exponent from mantissa. Engineering notation differs from scientific notation in requiring all exponents to be multiples of three. For example:

```
      '¯1.75E5 VS ¯155.75E5'⍠(10*⍳8)∘.×1 1
1.00E0 VS   1.00E0
1.00E1 VS  10.00E0
1.00E2 VS 100.00E0
1.00E3 VS   1.00E3
1.00E4 VS  10.00E3
```

```
1.00E5 VS 100.00E3
1.00E6 VS   1.00E6
1.00E7 VS  10.00E6
```

`'8' ⍕ ⍵` **Check protection.** Unused positions from here to the decimal point are filled with the character specified in `2@⎕fc` (default is `*`).

```
      ⊢a←'$558,555,535.50'
$558,555,535.50
      a⍕¯123 12345 1234567 123456789
  $******123.00
  $***12,345.00
  $1,234,567.00
$123,456,789.00
```

`'9'⍕⍵` **Conditional zero-fill.** Just like positions marked by `0`, these positions are filled with zeros when not otherwise used in the representation of a number. However, when the value being represented is zero, these positions are left blank.

Format Control (⎕fc)

The system variable ⎕fc (format control) is an implicit argument to ***dyadic format***. In a clear workspace, the six character elements in ⎕fc are `'.,**_¯'`

The meaning of the character in each position of ⎕fc is:

`0@⎕fc` Character to separate the *integer part* from the *fractional part* of a number. Where the left argument ⍺ contains the character "`.`" used as a decimal point, `0@⎕fc` is substituted for "`.`" in the result of `⍺⍕⍵`. The value of `0@⎕fc` is "`.`" in a clear workspace.

`1@⎕fc` Character to separate *groups of digits* in the representation of a number. (The groupings are usually but not necessarily triplets.) Where the left argument ⍺ contains `,` surrounded on *both sides* by digits, `1@⎕fc` is substituted in the result of `⍺⍕⍵`. The value of `1@⎕fc` is "`,`" in a clear workspace.

`2@⎕fc` Character to fill *insignificant positions* in the representation of a number. Where the left argument ⍺ contains `8` and the corresponding digit in the result is not significant, the result substitutes `2@⎕fc`. The value of `2@⎕fc` is "`*`" in a clear workspace.

`3@⎕fc` Character to indicate that the representation requires more space that has been specified in the left argument ⍺. In the result, the entire field is replaced by `3@⎕fc`. However, when `0@⎕fc` is `0`, and any of the numbers in ⍵ cannot be represented in the space provided, the entire expression is rejected with a *`domain error`*. The value of `3@⎕fc` is "`*`" in a clear workspace.

`4@⎕fc` Character to indicate where a blank should be substituted in the result. Where the left argument ⍺ contains this character as a *controlled character* (embedded within the positions devoted to the number's representation), the result replaces `4@⎕fc` by a blank. The value of `4@⎕fc` is "`_`" in a clear workspace.

`5@⎕fc` Character to indicate a *negative value.* However, at present the only permitted value is "`¯`", and that is the value of `5@⎕fc` in a clear workspace.

⍎⍵ *Execute*

Rank: ∗

The ***execute*** function allows a character vector (or scalar) to be evaluated like an APL expression. The interpreter treats ⍵ as a *line to be executed* in the same way as a line in the definition of a function or a line entered from the keyboard is executed. The result of ⍎⍵ is the result of executing the character vector ⍵. For example:

```
      ⍎'2+2'
4
      ⍎'A←3+4'
```

assigns the value seven to the name *A*.

The ***execute*** function permits an APL program to construct statements to be executed. Some uses of ***execute*** are conditional execution, conversion of numeric constants, and passing of unevaluated arguments to user functions.

To show *conditional execution,* consider processing a user's response to the question `Work?` with a statement using ***execute***:

```
⍎('Yes'≡3↑⍞⊣⎕←'Work?')/'analyze data'
```

Execute can convert character vectors representing numeric constants to their numeric values. For example:

```
      a←⍎'1 2 3'
      +/a
6
```

In this rather limited sense, ***execute*** is inverse to ***monadic format***. See also the system functions `⎕fi` (Input Format Conversion) and `⎕vi` (Verify Input Format) in the *System Guide, Chapter 5*.

Since system commands are not APL statements, they *cannot* be executed by this function. Nor may ***execute*** be used to invoke the ***del*** (∇) editor.

Result Returned by Execute

Whether ***execute*** returns a result depends upon whether the statement represented in its argument, when evaluated, returns a result. When it does, the result of evaluating the statement becomes the result of ***execute***. Otherwise, ***execute*** has no result. For example:

```
⍎'¯1+2×⍳⌊0.5×⍴⍵'
```

returns a result. However, if a defined function is written so that it does not return an explicit result, the function ⍎ does not.

```
⍎'foo x'
```

Nor does `⍎''`.

Consequently, the first statement in the preceding example can be embedded in a larger statement:

```
v[⍎'¯1+2×⍳⌊0.5×⍴⍵']
```

whereas the second statement cannot:

```
      a←⍎'foo x'
result error
      a←⍎'foo x'
      ^
```

Display of Result

When the ***execute*** function returns a result, the result is displayed if the result would normally be displayed. For example, `⍎'⍳5'` displays a value, but `⍎'a←⍳5'` does not.

Evaluation of Compound Statements

Several statements can be evaluated in one call to ⍎ if they are included in the character vector argument separated by diamonds. For example:

```
      x←⍎'a←b/⍳⍴b ⋄ ra←⍴a'
```

In such a case, the value (if any) returned is the value resulting from the ***last*** statement evaluated. In the above example, `ra←⍴a` is the last statement evaluated, and the value of `ra` becomes the value of `x`.

Occurrence in the State Indicator

If ***execute*** has been invoked but has not completed execution, it appears in the state indicator as a *separate line* in the same way as a user function. For example, if `fn` is a function invoked by `⍎'fn'` and its execution is suspended on line three, the state indicator appears as:

```
      )si
fn[3]*
⍎
      ⎕lc
3 0
```

The contents of the system variable `⎕lc` (line counter) for the pending execute is either zero, which indicates that ***execute*** was invoked from immediate execution, or the line number of the function that invoked ***execute*** so that expressions like `→⎕lc` will work as expected even from within `⎕trap` expressions.

Branch Within Execute

When a statement evaluated contains a ***branch arrow*** (→) the effect is the same as if the statement had been encountered as part of a defined function being executed, or entered from the keyboard for immediate execution. That is, if the value to the right of the ***branch arrow*** identifies a line number in the defined function currently highest on the execution stack, the branch is taken (and the remainder of the argument of ***execute***, if any, is abandoned).

Errors Encountered During Execution

Error conditions encountered during execution of the represented statement immediately yield the usual display, including an error message, the statement in error, and the caret. The statement containing the error is displayed, rather than the one at the level of the calling environment of ***execute***:

```
      ⍎'a←⎕⎕'
syntax error
⍎     a←⎕⎕
      ^
```

The ***pawn*** symbol (⍎) is displayed in the state indicate to mark a statement originated from a call to ***execute***, as opposed to a statement in immediate execution mode or on a line of a user-defined function. Further, as an aid in determining how the erroneous statement was created, the error display is preceded by one ***pawn*** symbol for each occurrence of ***execute*** in the state indicator between the level where the error occurred and the level where the suspension occurred.

For example:

```
      ⍎'⍎''?0'''
domain error
⍎⍎    ?0
      ^
```

5
Operators

An *operator* produces a new function called a *derived function.* Each of the symbols ***slash*** (⌿), ***slash-bar*** (⍀), ***back-slash*** (⍂), ***back-slash-bar*** (⍁), ***left shoe*** (⊂), and ***ampersand*** (&) denotes a *monadic* operator, which applies to a single function argument or a single array argument. Each of the ***symbols*** dot (.), ***dieresis*** (¨), ***paw*** (⍤) and ***hoof*** (⍥) denotes a *dyadic* operator, which applies to either two function arguments or one array argument and one function argument.

Tables 5.1 and 5.2 list the operators supported by SHARP APL for UNIX. In these tables, and in all tables and examples, the letters f, g, m, and n represent the operator's arguments. The letters f and g are, respectively, left and right function arguments. The letters m and n are, respectively, left and right array arguments.

In tables where ranks are provided, the infinity symbol (∞) indicates infinite rank; the symbol ⋆ indicates that the case is undefined; the abbreviations mf, ιf, and rf indicate the monadic, left, and right ranks, respectively, of function f, and the abbreviations mg, ιg, and rg indicate the monadic, left, and right ranks, respectively, of function g.

Table 5.1. SHARP APL monadic operators.

Form	Name	Monadic Rank	Dyadic Rank
f/ ω	f-Reduce	∞	
f⌿ ω	f-Reduce-down	∞	
f\ ω	f-Scan	∞	
f⍀ ω	f-Scan-down	∞	
m/ ω	Copy/Compress	∞	
m⌿ ω	Copy-down/Compress-down	∞	
m\ ω	Expand	∞	
m⍀ ω	Expand-down	∞	
f⊂ ω	Swap	∞	
α f⊂ ω	Swap		rf lf
f& ω	Select	∞	
α f& ω	Merge		∞ ∞

Table 5.2. SHARP APL dyadic operators.

Form	Name	Monadic Rank	Dyadic Rank
f⍤g ω	f-On-g	mg	
α f⍤g ω	f-On-g		mg mg
f⍤n ω	f-Rank-n	n	
α f⍤n ω	f-Rank-n		n n
m⍤g ω	m-Cut-g	mg	
α m⍤g ω	m-Cut-g		mg mg
f⍥g ω	f-Upon-g	mg	
α f⍥g ω	f-Upon-g		mg mg
f⍨g ω	f-Under-g	mg	
α f⍨g ω	f-Under-g		mg mg
m⍨g ω	m-With-g	mg	
f⍨n ω	f-With-n	mf	
f.g ω	Alternant	2	
α f.g ω	Inner-product		∞ ∞
α m.g ω	Tie/Outer-product		∞ ∞
f.m ω	ply	mf	

Any function (primitive, system, user, or derived) may be used as the argument of an operator.

This chapter opens with a summary of general characteristics of operators, and then describes each of the primitive operators individually.

Precedence of Operators

APL has one rule of precedence: execute operators before functions. All operators have equal precedence (just as among functions, all functions have equal precedence). The interpreter must know what derived function an operator produces before it can parse the rest of the statement. Therefore, as you read a statement from left to right, whenever you come to an operator, you evaluate it at once, and replace it and its arguments by the resulting derived function. The difference between the way you treat a function and the way you treat an operator is this:

- When you come to an operator, evaluate it at once.
- When you come to an unmodified function, delay evaluation until you know the value of its right argument. You do not have to delay for the left argument because, since you are reading from left to right, you already know what that is.

Valence of Operators

Monadic operators apply to the array or function that they follow. Dyadic operators apply to the functions (or array and function) on each side.

Syntactic Class of Arguments to Operators

The effect of an operator depends upon whether its arguments are arrays or functions; that is, on the *syntactic class* of its arguments.

- For each monadic operator, there are two possible meanings, depending upon whether the single argument is an array or a function.

- For each dyadic operator, there are four possible meanings, depending upon whether each of the arguments is an array or a function.

For most operators, there is a generic name for the symbol (applicable to all the cases), and specific names for the separate meanings linked to the syntactic class of the arguments.

The derived function produced by an operator (like all functions) may be used either dyadically or monadically. The dyadic and monadic uses of the derived function are also sometimes given distinct names, depending on the effects they produce. The various names are noted as part of the description of each monadic or dyadic operator.

Consecutive Operators

It is possible to have a sequence of several operators in succession in a statement. Read them from left to right. Operators appear next to each other when the derived function produced by the first is in turn modified by the next. For example, in the statement

```
A +.×⍤2 B
```

the first function ***plus*** (α+ω) is modified by the dyadic operator ***dot*** (*f*`.`*g*), which takes ***plus*** and ***times*** (α×ω) as its arguments and produces the derived function ***plus-times inner product*** `(+.×)`. That derived function is then modified by the ***rank*** (*f*⍤*n*) operator to produce a new derived function which calculates the inner product of each of the two-dimensional matrices within `A` and `B`.

Here is an example that mixes some functions that are not modified by operators with some that are. Assume that the names `A`, `B`, and `C` refer to arrays.

```
C−A+.×⍤2−B÷2
```

Reading from left to right, the first function is ***subtract*** (α−ω). Since the function is not modified by an operator, it is an ordinary function. Set it aside until you know what its right argument is.

The next function is ***plus*** (α+ω) which is modified by the dyadic operator ***dot*** (*f*`.`*g*). Evaluate the operator at once: ***dot*** takes arguments on both sides, so you get the derived function ***plus-times inner product*** `(+.×)`.

Following the derived function is the ***rank*** (*f⍤n*) operator. Evaluate it at once. Its arguments are the derived function `(+.×)` on its left and the number two on its right. That gives you the derived function `((+.×)⍤2)`.

The next function is ***minus*** (−*ω*). Since it is not modified by an operator, it is an ordinary function. Set it aside until you know what its right argument is.

The next function is ***divide*** (*α*÷*ω*). It is not modified. You know the value of its right argument. Evaluate it. Replace *B*÷2 by the resulting array.

Now you know the value of the right argument of the ***minus*** function that you most recently set aside. Go back and evaluate it. It has no left argument, since to its left there is a derived function. Replace the segment −*B*÷2 by the resulting array.

Now you know the value of the right argument of the derived function `((+.×)⍤2)`. Evaluate the derived function. Replace *A*+.×⍤2−*B*÷2 by the resulting array.

Now you know the value of the right argument of the ***subtract*** function you set aside earlier. Go back and evaluate it. It is the root function. You have now evaluated the entire statement.

Look-ahead

Reading from left to right requires you to look ahead. Each time you come to a array or function, you have to check whether it is modified by an operator. You do not have to look ahead far: if there is a modifier, it must be the next thing to the right. Whenever you discover that what is to the right is a array or function, that is sufficient to show that your current array or function is not modified.

Separating Side-by-Side Constants

Because operators have higher precedence than functions, it often happens that two constants appear side by side in a statement. One of them is the right argument of the dyadic operator just to its left, and the other the left argument of a function just to its right.

For example, if you want to apply the function *analyze* to the rank-*i* cells of *X*, you write

```
analyze⍤i X
```

That statement is perfectly acceptable. But suppose you know that *i* is `1` and that *X* is simply the numbers `4` and `5`. The blank which serves to separate the name *i* from the name *X* will not do to separate the number `1` from the numbers `4` and `5`. If you were to write

```
analyze⍤1 4 5
```

or even

```
analyze⍤1    4 5
```

the numbers would run together to form the single constant.

```
1 4 5
```

There are several ways to make plain what you intend:

`(analyze⍤1)4 5`	Put parentheses around the derived function to separate it from its argument.
`analyze⍤1(4 5)`	Put parentheses around the right argument of the derived function.
`analyze⍤1⊢4 5`	Use the ***pass*** function (⊢ω) to separate the two numeric constants. The ***pass*** function returns the argument to the right without changing it, but its presence makes clear that `1` and `4 5` are separate constants.

Monadic Operators

The operators represented by the ***slash*** (/), ***slash-bar*** (⌿), ***back-slash*** (\) and ***back-slash-bar*** (⍀) symbols form a family with related properties. The operators represented by ***left shoe*** (⊂) and ***ampersand*** (&) are in separate classes.

These monadic operators produce a family of derived functions with the following in common:

- All have unlimited right rank: they apply to the entire right argument at once, rather than independently to cells within a right-argument frame.
- All partition their right argument into major cells. The operators whose symbols have a crossbar (⌿ and ⍀) produce derived functions which split the argument array into major cells in the standard way, along the first axis. The operators whose symbols lack the crossbar (/ and \) produce derived functions that split the right argument into contra-major cells, splitting the argument along its last axis rather than the first.[1]

Summary of Forms

The single argument of these monadic operators may be either a function or array.

- When the argument is a function, the function's dyadic use is understood.
- When the argument is an array, it must be a scalar or vector of integers.

Derived Functions Modified by Axis Bracket Notation

All the derived functions produced by ***slash*** (/), ***slash-bar*** (⌿), ***back-slash*** (\) and ***back-slash-bar*** (⍀) work by partitioning the right array argument into cells. The derived functions produced by ***slash-bar*** and ***back-slash-bar*** partition along the first axis, while the derived functions produced by ***slash*** and ***back-slash*** partition along the last axis. You can override the axis along which partitioning takes place by modifying the derived function with an *axis specifier.*

For example, the expressions $n⌿[i]\ \omega$ and $n/[i]\ \omega$ both mean that the derived function partitions the array ω along its i^{th} axis rather than the first or last.

Such axis notation is considered an obsolescent holdover from early APL implementations. In SHARP APL for UNIX, it is supported only for those constructions for which it works in ISO standard APL. The functions that can be used with the axis specifier are ***catenate*** (α,ω), ***catenate-down*** ($\alpha⍪\omega$), ***reverse*** ($\phi\omega$), ***reverse-down*** ($\ominus\omega$), ***rotate*** ($\alpha\phi\omega$), ***rotate-down*** ($\alpha\ominus\omega$), f***−reduce*** (f/ω), f***−reduce-down*** ($f⌿\omega$), ***compress/copy*** (n/ω), ***compress-down/copy-down*** ($n⌿\omega$), f***−scan*** ($f\backslash\omega$), f***−scan-down*** ($f⍀\omega$), ***expand*** ($n\backslash\omega$), and ***expand-down*** ($n⍀\omega$)

1. Historically, the last-axis split was introduced first, which is why it has the simpler symbols.

See the ***rank*** (*f⍤n*) operator on page 5-25 for a mechanism to control the relationship between frame and cells that is more general than the axis specifier.

f/ ω f-Reduce
f⌿ ω f-Reduce-down
f\ ω f-Scan
f⍀ ω f-Scan-down

When the argument to ***slash*** (*/*), ***slash-bar*** (*⌿*), ***back-slash*** (**) or ***back-slash-bar*** (*⍀*) is a function, the derived function splits the array *ω* into cells, and applies *f* dyadically between each of them.

Reduce	The result of reduction is found by evaluating *f* between all the cells thus formed.
Scan	The result of a scan is built up as a series of cells, by evaluating *f* for the first argument cell, the first two argument cells, the first three argument cells, and so on.

For example:

```
      ⊢a←2 3 4⍴⍳24
 0  1  2  3
 4  5  6  7
 8  9 10 11

12 13 14 15
16 17 18 19
20 21 22 23
      +⌿a
12 14 16 18
20 22 24 26
28 30 32 34
      ,⌿a
 0  1  2  3 12 13 14 15
 4  5  6  7 16 17 18 19
 8  9 10 11 20 21 22 23
      +⍀a
 0  1  2  3
```

```
 4  5  6  7
 8  9 10 11

12 14 16 18
20 22 24 26
28 30 32 34
      ,⍀⍺
 0  1  2  3  0  0  0  0
 4  5  6  7  0  0  0  0
 8  9 10 11  0  0  0  0

 0  1  2  3 12 13 14 15
 4  5  6  7 16 17 18 19
 8  9 10 11 20 21 22 23
```

Effect of Order of Grouping on f-Scan and f-Reduce

Because a scan's n^{th} result cell is constructed from the first n cells of the argument, but APL's leaf-to-root execution implies evaluating the last function first, the value of the $(n+1)^{th}$ cell is obtainable directly from the n^{th} cell only when the function is associative. Consider a simple example with the nonassociative function ***subtract*** (⍺ −⍵).

Suppose *X* is a numeric array having three axes, with lengths `4 5 6`. A scan or reduce splits *X* along the first axis forming four cells, each of which is a `5`-by-`6` table. Call those tables *X*`1`, *X*`2`, *X*`3` and *X*`4`. Thus

```
      −⍀X
```

returns a `4`-by-`5`-by-`6` result, whose four tables (each `5`-by-`6`) are computed respectively by:

```
X1                  ⟺   X1

X1 - X2             ⟺   X1 - X2

X1 - X2 - X3        ⟺   X1 - (X2 - X3)

X1 - X2 - X3 - X4   ⟺   X1 - (X2 - (X3 - X4))
```

The left-to-right grouping for the cells produced by a scan was deliberately adopted because, in this order, the scans produced by several nonassociative functions produce valuable results. For instance, `-⍀X` is the *alternating sum* of cells of `X`, and `÷⍀X` is the *alternating product* of cells of `X`.

In the result of any scan, the first cell is the same as the first cell in the argument, and the last cell is the same as the result of *reduce.* Thus the single cell returned by `f⌿X` is the same as the last of the cells returned by `f⍀X`.

Where scan and reduce are used with *scalar functions,* the result of a scan applied to an array `X` has the *same rank and shape* as the argument `X`. The result of a reduction by an scalar function has a rank one less than the rank of `X`. When the argument to reduce already has rank zero, the result also has rank zero.

f-Reduce or f-Scan with Derived or User Functions

The functional argument to these four operators can be any dyadic function, including the function produced by another operator or a user function (identified by a name rather than a symbol).

For example, using the notation introduced earlier for the various cells of an array `X`, `foo⌿X` is executed as

```
X0 foo X1 foo X2
```

and `+.×⌿X` is executed as

```
X0 +.× X1 +.× X2
```

f-Reduce when Argument Has One Cell or None

By convention, when the array argument of the derived function has only one cell, the result of reduction by any function is that cell, unchanged. That holds true even when the values in the argument cell would otherwise be outside the function's domain, as for example `∧⌿'a'`.

When the array argument of the derived function has no cells and the reduction is by a primitive function with a known identity element, the result is a cell whose value is that identity element. An identity element for a dyadic function is a value that causes the function to return a result that is the other argument

unchanged. Thus, the identity element for addition is zero because adding zero to another number does not change it; the identify element for multiplication is one, and so on.

For the dyadic function ***maximum*** (⍺⌈⍵), the identity is that value compared to which all other numbers are greater. In principle, it is minus infinity; in practice, it is the smallest representable number. Similarly, for the function ***minimum*** (⍺⌊⍵), the identity is that value compared to which all other numbers are smaller. In principle, it is infinity; in practice, it is the largest representable number.

For some functions, there is a right identity, but not a left. For example, anything divided by `1` remains unchanged, but that is not true for `1` divided by anything. For some functions, there is an identity element when the function's domain is restricted. For example, as long as you stick to Boolean values, one is an identity element for the function ***equals*** (⍺=⍵).

When reduction is applied to an empty array, SHARP APL returns whatever identity element is available. The values of the identity elements are shown in tables accompanying the presentation of each primitive function (See "Table 4.1. Ranks of primitive functions." on page 4-3).

Reduction by a user function (such as *foo*≠*X*) is undefined for an *X* that has no cells. Even though the defined function may in fact have an identity element, at present SHARP APL lacks a way of recording what its identity element is.

m⌿ ⍵ *Compress/Copy*
m≠ ⍵ *Compress-down/Copy-down*
m⍀ ⍵ *Expand*
m⍒ ⍵ *Expand-down*

When the argument to ***slash*** (⌿), ***slash-bar*** (≠), ***back-slash*** (⍀) or ***back-slash-bar*** (⍒) is an array, the derived function first splits the array *X* into cells, and then *m*⌿*X* copies the cells a specified number of times, while *m*⍀*X* inserts additional cells between them.

The array argument *m* in *m*⌿⍵ or *m*≠⍵ must be a scalar or vector of nonnegative integers. The array argument *m* in *m*⍀⍵ or *m*⍒⍵ must be a Boolean scalar or vector.

Derived Function Copy: Replicating Cells

Use this derived function monadically to select from an array; restrict *m* to nonnegative values.

```
      ⊢x←>'First'⊃'Second'⊃'Third'
First
Second
Third
      2 0 2⌿x
First
First
Third
Third
      2 0 2⌿⍳3
0 0 2 2
```

The number of cells in the result is the sum of the values in *m*. The result consists of replications of the cells of the right argument.

When the values in the vector *m* are all either zero or one, a cell from the derived function's right (and only) argument is either retained in the result (where *m* has a one) or discarded (where *m* has a zero). In that case, the derived function is called *compress,* since the result differs from the argument only in having certain cells removed. This often occurs in an expression in the form:

```
      (Proposition_on_X)⌿X
```

Here *Proposition_on_X* stands for any APL expression that takes *X* as an argument and returns a Boolean result with a zero or one for each major cell in *X*; for example, `0 (∊⍤¯1) X` ("accept any major cell that contains a zero"). The result retains from *X* those major cells for which the proposition is true (represented by one), and omits those for which the proposition is false (represented by zero).

Derived Function Expand: Introducing Cells

The array argument to ⍀ or \ is required to be a Boolean scalar or vector in which the number of occurrences of a one agrees with the number of cells along the axis of the array you are about to expand. Occurrences of a zero produce *fill elements* appropriate to the type of the array you are expanding.

Fill elements are:

`''` for character arrays
`⊘` for enclosed arrays
`0` for all other arrays (heterogeneous or numeric)

For example, using the array *x* defined above:

```
      1 0 1 0 1↖x
First

Second

Third
```

f⊂ ⍵ Swap

When the derived function *f*⊂ is used monadically, the value of ⍵ is used for both the left and right arguments to function *f*:

f⊂⍵ ⟺ ⍵*f*⍵

For example:

```
      ⊘.=⊂⍳5
1 0 0 0 0
0 1 0 0 0
0 0 1 0 0
0 0 0 1 0
0 0 0 0 1
```

⍺ *f*⊂ ⍵ Swap

When the derived function *f*⊂ is used dyadically, the order of the arguments function *f* is reversed.

⍺*f*⊂⍵ ⟺ ⍵*f*⍺

For example:

```
      3÷⊂6
2
```

f& ω Select

When you use the derived function *f&* monadically, you can *select* from an array using any **monadic** selection function to specify the desired subarray. Monadic selection functions include ***copy*** (*m/*), ***copy-down*** (*m⌿*), ***take*** (⍺⍤↑), ***drop*** (⍺⍤↓), and ***from*** (⍺⍤@).

```
      2 2⍤↑&3 3⍴⍳9
0 1
3 4
```

α f& ω Merge

When you use the derived function *f&* dyadically, you can *merge* two arrays without the pitfalls of indexed assignment. The explicit result of applying such a merge is an array which does not overwrite the source array, as was the case for indexed assignment. The resulting merged array may become the argument of any appropriate function.

If the selection function you use is already monadic (as for any function derived from ***slash*** (*/*) or ***slash-bar*** (*⌿*), then the stage is set for merging. The subarray to be merged occurs as the left array argument of the derived function.

```
      ⊢v←⍳5
0 1 2 3 4
      b←0 1 1 0 0
      (55 66) b/&v
0 55 66 3 4
```

If your selection function is not monadic, use the ***with*** operator (*m⍤g*) to derive a monadic function to use as the argument to ***merge***:

```
      (55 66) ¯2⍤↑&v
0 1 2 55 66
      (7) 3⍤@&v
0 1 2 7 4
```

You can use the ***rank*** operator (*f⍤n*) with the function derived from merge to replace a single subarray into multiple locations in a source array:

```
      ⊢m←3 3⍴⍳9
0 1 2
3 4 5
6 7 8
       (10 11 12) 1 2⍤@&⍤0 1 m
 0 10 10
 3 11 11
 6 12 12
```

Dyadic Operators

Cases of Dyadic Operators

For every dyadic operator, four cases are possible based on the syntactic class of the operator's arguments:

1. Left and right function arguments
2. Left array argument and a right function argument
3. Left function argument and right array argument
4. Left and right array arguments

In theory, each of these cases may produce an ambivalent derived function which may be used either monadically or dyadically. In practice, some of the cases are not yet implemented or defined.

α *m*.*g* ω Tie

The tie operator reduces the number of frame axes that must agree, leaving the remaining frame axes free to interact as a Cartesian product. The argument *m* specifies the number of *leading frame axes* which must still agree; those axes are said to be *tied*. *m* must be a nonnegative integer scalar not greater than the maximum rank of the arguments of the derived function.

The interesting effect of tie is not on the axes that are tied, but on those that are not. The remaining frame axes—those that are not tied—are said to be *free.* A result cell is generated for every possible pairing of an α-cell with an ω-cell. As a consequence, there is no need for the free axes of α to match the free axes of ω in either number or length, and both the free axes of α and the free axes of ω are represented in the result. Suppose the shapes of *A* and B are both `2 3 4`. Then, for a function *g* whose dyadic argument ranks are `0 0`, the statement

```
A g B
```

requires that the frame shapes of *A* and *B* agree. Since *A* and *B* have rank `3`, the statement

```
A 3 .g B
```

has the same effect, although it makes explicit that the first three frame axes of *A* must match the first three frame axes of *B*.

Using the tie operator with a value of *m* less than the rank of *A* or the rank of *B* leaves the remaining frame axes free. For example,

```
A 2 .g B
```

requires the first two frame axes of *A* and *B* to agree (and to be reproduced in the result), while the remaining frame axes of *A* form a Cartesian product with the remaining frame axes of *B*. In the example, the remaining frame axes are `4` for *A* and also `4` for *B*, resulting in a `4`-by-`4` table for each of the `2`-by-`3` tied frame positions. The frame shape of the result is therefore `2 3 4 4`.

For example, using ***rotate***:

```
      (⍳3) (0 .⌽)>'abcdefghij'⊃'0123456789'
abcdefghij
0123456789

bcdefghija
1234567890

cdefghijab
2345678901
```

For example, using ***catenate***:

```
      ⊢M←2 4⍴⍳8
0 1 2 3
4 5 6 7
      ⊢N←3 2⍴10+⍳6
10 11
12 13
14 15
      M (0 .(,⍤1)) N
0 1 2 3 10 11
0 1 2 3 12 13
0 1 2 3 14 15

4 5 6 7 10 11
4 5 6 7 12 13
4 5 6 7 14 15
```

Outer Product

When none of the axes is tied, all frame axes of ⍺ and ⍵ appear in the result. This is obtained by:

```
⍺ 0 .g ⍵
```

and is called the *outer product* of ⍺ and ⍵. Historically, before the tie operator was defined, the outer product was introduced as:

```
⍺ ∘.g ⍵
```

For consistency with past applications, this version of SHARP APL recognizes `⍺ ∘.g ⍵` as an alternative way of eliciting `⍺ 0 .g ⍵`.

Note: Since the left argument of the tie operator is an integer scalar, it frequently appears as a constant. To prevent the dot from being treated as a decimal point, it is essential to separate such a constant from the dot that denotes the operator. The separator may be a blank or parentheses surrounding either the dot or its argument. Perhaps the easiest to read is:

```
⍺ (0 .+) ⍵
```

f.m ω Ply

If *m* is a non-negative scalar integer, the ***ply*** operator applies the monadic function *f* *m* times. If *m* is a negative scalar integer, the inverse of *f*, if it exists, is applied absolute *m* times. For example:

```
      ⍴⍴ x ← 2 3⍴ 23
2
      ⍴. 2 x
2
      ⍟.¯1 ⍟ 1
1
      1⍤⍤+. 0 ( 2 )
2
```

α *f.g ω Inner Product*

The expression α+.×ω is equivalent to the *inner* or *matrix* product as defined in mathematics for vectors (+/α×ω) and matrices (where the element in row *i* and column *j* of α+.×ω is the *dot product* of row *i* of α and column *j* of ω). For example:

```
      1 2 3 +.× 3 4 5
26
       a←2 3⍴⍳6
       b←3 4⍴⍳12
       a+.×b
20 23 26 29
56 68 80 92
      a+.×1 2 3
8 26
```

For scalar functions other than ***plus*** and ***times***, the same definition holds for argument ranks not greater than two. For example:

```
      a⌊.+b
0 1 2 3
3 4 5 6
      a+.⌊b
3  3  3  3
9 10 11 12
```

The general definition for arbitrary functions and arguments of arbitrary rank is:

```
      α f.g ω ⟸⟹ f⌿((1⌽⍳⍴⍴)⍪α) 1 .g ω
```

In other words, the result is reduction by *f* over the result of applying (`1 .`*g*) to the major cells of `(1⌽⍳⍴⍴)⍪α` and *ω*.

The general case follows:

```
      ⊢c←(1⌽⍳⍴⍴)⍪a
0 3
1 4
2 5
      b
0  1  2  3
4  5  6  7
8  9 10 11
      (0@c) 0 .×0@b
0 0 0 0
0 3 6 9
      (1@c) 0 .×1@b
 4  5  6  7
16 20 24 28
      (2@c) 0 .×2@b
16 18 20 22
40 45 50 55
```

The final result of *a*`+.×`*b* is the sum over these tables, which is the same as the example of *a*`+.×`*b* given earlier.

The overall effect of the definition of *f*`.`*g* is to split *α* along its last axis, just as *ω* is split along its leading axis. This asymmetric treatment of the arguments was originally due to the desire to make the simple case of `+.×` on matrix arguments agree with the *matrix product* of mathematics, whose definition exhibits the same sort of asymmetry.

f.g ω Alternant

The expressions −.×ω and +.×ω are for square matrix arguments ω, the *determinant* and the *permanent* of mathematics. The generalization to arguments other than ***plus***, ***minus*** and ***times*** is based on construing the determinant as an *alternating sum* (−⌿) over products over the diagonals of tables obtained by permuting the major cells of ω.

The model for the determinant described here is illuminating for what it reveals about the close association between the monadic and dyadic forms of the dot operator (*f.g*). The determinant is generated by the expression −.×ω, where ω is a matrix of integers (usually square, but possibly with more rows than columns). Only the top-left square portion of an array with more columns than rows is considered; that is, ω is truncated to be (⌊/⍴ω)↑ω.

A model supplied by McDonnell and Hui for monadic inner product is available in the recursive function *DOT* shown below; the two functions *f* and *g* are assumed to be ***subtract*** (α−ω) and ***times*** (α×ω) in the following example.

```
    ∇z←DOT ω;M
[1]  →(1≠¯1↑⍴ω)/⍳0
[2]  →0⊣z←f⌿ω ⍝ one-column case
[3]  ⍳0:M←⍉⍤1 0 ⊂ ⍳ 0@ ⍴ω ⍝ minors
[4]  z←ω[;□io] f.g DOT⍤2 ω[M;1↓⍳(⍴ω)[1+□io]]
    ∇
```

For example, given the array *M*

```
      ⊢M←3 3⍴6 7 2  1 5 9  8 3 4
6 7 2
1 5 9
8 3 4
```

the expression which *DOT* evaluates to produce the determinant uses each scalar in the first column of *M* with subarrays called *minors* derived from the remaining columns of the remaining rows:

```
      (6×(5×4)−3×9) − (1×(7×4)−3×2) − 8×(7×9)−5×2
360
```

```
      f⍤⍤g ⍵ Upon
α     f⍤⍤g ⍵ Upon
      f⍤∘g ⍵ On
α     f⍤∘g ⍵ On
      f⍤⍤g ⍵ Under
α     f⍤⍤g ⍵ Under
```

When one of the symbols ***hoof*** (⍤), ***paw*** (⍤∘), or ***dieresis*** (⍤⍤) is used with a pair of functions, it *composes* the two functions to form a new one. Functions derived from the composition operators are always ambivalent.[2]

In all cases, composition is regarded as *close*: applied independently to each cell of the array argument (when used monadically) or to each of the pairs of corresponding cells (when used dyadically) as opposed to applying to the entire array. Refer to the section "Cellwise Composition" on page 5-23.

Default Rank of the Derived Function

When the function derived from a composition operator is not explicitly modified by the rank operator, the derived function divides its argument(s) into cells in the same way as would *g*, the operator's right argument. In particular, the rank of a function derived from *f*⍤*g* is the same as the rank of *g*, whereas the rank of functions derived from either *f*⍤⍤*g* or *f*⍤∘*g* is the same as the monadic rank of *g*.

Three Forms of Composition

Each of the three symbols ***hoof*** (⍤), ***paw*** (⍤∘), or ***dieresis*** (⍤⍤) may take two functions as its arguments: one function which it uses monadically and another which it uses ambivalently. Each of the three operators produces an ambivalent derived function. The three operators can be summarized as follows:

f⍤*g* (***upon***) The right argument *g* is the ambivalent one. It acts on corresponding cells of the derived function's array arguments ⍺ and ⍵ (or just on cells of ⍵ when ⍺ does not exist); the left argument *f* acts on each result cell.

2. You could also consider the ***dot*** operator (**.**) to be a form of composition, but its properties are sufficiently different that it is discussed separately.

$f\ddot{\circ}g$ (***on***) The left argument f is the ambivalent one. It acts on the corresponding cells produced by applying g to cells of α and to cells of ω (or just to cells of ω when α does not exist).

$f\ddot{\circ}g$ (***under***) Like $f\ddot{\circ}g$ but to each of the cells produced, it applies the inverse of g.

The three forms of composition are illustrated in Figure 5.1

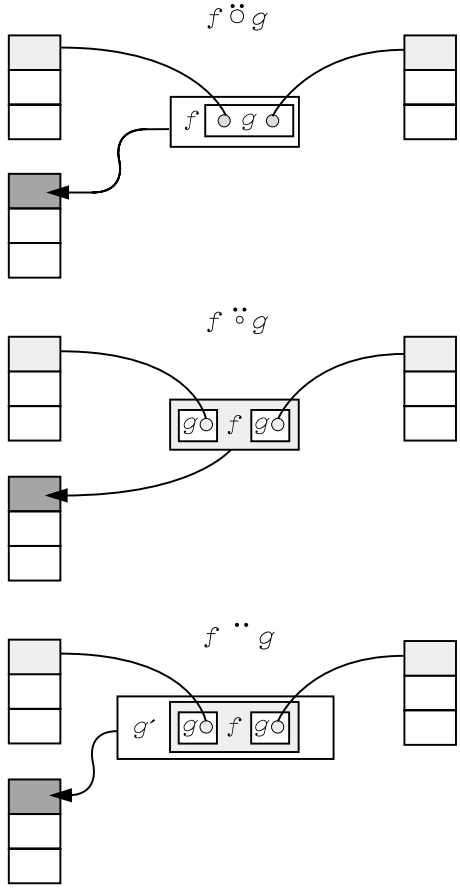

Figure 5.1. Three forms of composition.

Functions Derived from Composition Operators

Table 5.7 provides a template for the functions derived from the composition operators. In this table, ⍺ denotes a cell within the left array argument of the derived function (*LA*), and ⍵ denotes a cell within the right array argument of the derived function (*RA*).

Table 5.7. Composition templates.

upon	*LA f⍤g RA* *f⍤g RA*	*f ⍺ g ⍵* *f g ⍵*
on	*LA f⍥g RA* *f⍥g RA*	*(g ⍺) f g ⍵* *f g ⍵*
under	*LA f⍣g RA* *f⍣g RA*	*g' (g ⍺) f g ⍵* *g' f g ⍵*

g' ***is the inverse of*** *g.*

f⍣g Requires a Known Inverse

Because *f⍣g* automatically applies not only the monadic function *g* but also its inverse, the choice of *g* in compositions such as *f⍣g* is limited to functions for which the interpreter already knows the inverse. At present, the APL interpreter does not provide a way to indicate what inverse should be used with a user function *g*. So *g* must be taken from the following list of primitive functions:

Monadic Function	⋆	⍟	>	<	−	÷	⊖	⌽	−	~	+	⊢	⍉
Inverse	⍟	⋆	<	>	−	÷	⊖	⌽	−	~	+	⊢	⍉

Notice that several of these functions are self-inverses. The function ***pass*** (⊢⍵) is a self-inverse because it is an identity function, returning as its result the same value it received as an argument.

Cellwise Composition

In keeping with the general rules for partitioning arrays into frame and cells, all of the composition operators produce *close* or *cellwise* compositions. Figure 5.2 illustrates close or cellwise composition.

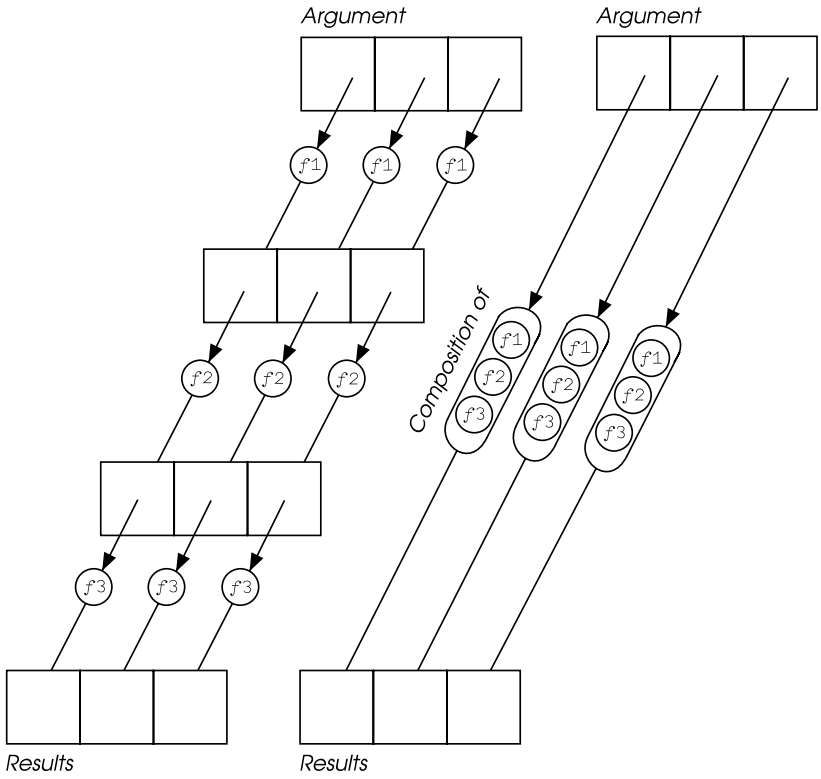

Figure 5.2. Close or cellwise composition.

The derived function is evaluated independently for each argument cell, without interaction with the other cells. Consider the difference between the following expressions:

`f1 f2 f3 x` and `f1¨∘f2¨∘f3 x`

The first applies the function `f1` to the array that results from applying `f2` to the array that results from applying `f3` to `x`. An array that contains the entire result of `f3 x` is computed, and that entire array becomes the argument of `f2`. The second statement evaluates `f1 f2 f3` for each cell of `x`, but does not assemble the resulting cell into the result frame until the sequence `f1 f2 f3` has been evaluated. Depending what the functions are and the rank of `x`, the two statements may produce the same result. However, the statement `f1 f2 f3 x`

requires building a frame after evaluation of each function, whereas the composition `f1⍤f2⍤f3` does not. For some functions, building the intermediate frame would produce a shape inappropriate to the next function, or would require conversion or padding of some result cells to fit the intermediate frame. Close composition avoids intermediate frames altogether.

Example:

```
      ⊢x←'this'⊃'is'⊃'a test'
|────| |──| |──────|
|this| |is| |a test|
|────| |──| |──────|
      ⍴>x
3 6
      ⍴⍤>x
4
2
6
      (<'is')∈⍤>x
1 1
1 1
0 1
      (<'is')⊆⍤>x
0 0 1 0 0 0
1 0 0 0 0 0
0 0 0 0 0 0
      x∈<'is'
0 1 0
       x∈⍤><'is'
0 0 1 1 0 0
1 1 0 0 0 0
0 0 0 0 1 0
```

f⍤*n* ⍵ *Rank*
α *f*⍤*n* ⍵ *Rank*

When the left argument of ***paw*** (⍤) is a function and the right argument is an array, the array specifies the argument rank of the function. For example, the expression

```
      α foo⍤2 1 3 ⍵
```

indicates that the function *foo* is to be applied to rank-1 cells of α and rank-3 cells of ω. It also says that when *foo*⍤2 1 3 is used monadically, it applies to the rank-2 cells of ω.

The array argument of *f*⍤*n* must be an integer scalar or a vector of one, two, or three integers. The formal definition of the treatment of the variable argument is as follows:

```
⍵3⍴⍵n ⇐⇒ 3⍴⍤⍵n
```

Calling the scalars within the vector *a*, *b*, and *c*, respectively (when the variable has one, two, or three scalars), they specify the argument ranks as shown in Table 5.9.

Table 5.9. Specifying the rank of a function.

n	*Monadic rank*	*Dyadic left rank*	*Dyadic right rank*
a b c	*a*	*b*	*c*
a b	*b*	*a*	*b*
a	*a*	*a*	*a*

Argument Rank versus Frame Rank

When an scalar in the right argument of *f*⍤*n* is positive, it specifies the number of trailing axes in an argument cell. Any remaining axes are frame axes. When an scalar in the right argument of *f*⍤*n* is negative, its magnitude specifies the number of leading axes in the argument frame. Any remaining axes are cell axes. For each of the cells within the frame thus defined, the function argument is evaluated with its ordinary rank.

For example,

```
foo⍤¯1 ⍵
```

specifies that the frame rank is one and that *foo* is applied independently to each of the major cells of ω.

Ranks Can Nest

The operator *f*⍤*n* controls the initial partitioning of the arguments into frame and cells. To each of those cells, the argument function's definition applies. This definition may itself have a finite rank, in which case the cell can be partitioned

again. For example, to add the nine element vector α to each of the rows of the nine column table ω, you write:

```
α+⍤1⊢ω
```

This pairs the vector α with the vector `0⊃ω`, then with the vector `1⊃ω`, and so on. For the first result cell, the interpreter evaluates

```
α+0⊃ω
```

The function ***plus*** has default argument rank zero. ***Plus*** is applied to the arguments α and `0⊃ω` with rank zero. That causes a new level of partitioning between frame and cell. In this example, each argument cell is partitioned into a nine element frame consisting of rank zero cells. `0⊃α` is added to `0⊃0⊃ω`, `1⊃α` is added to `1⊃0⊃ω`, and so on.

In executing $f⍤i⍤j\ \omega$, the function $f⍤i$ is applied within cells determined with reference to j. The default rank is necessarily the most deeply nested.

Default Rank Not Defined for Some Functions

A function's default rank is the argument rank it assumes when the rank is not explicitly stated with the rank operator. For example, the ***matrix inverse*** function (⌹ω) has default rank of two. If x is a rank three array and you write ⌹x, it is assumed that you mean to invert to each of the matrices of x. However, user functions do not have a default argument rank, and there are a few primitive functions for which as yet no default rank has been defined.

When you use a composition operator to partition the arguments of a function which has no default rank you must create cells that the function can accept. Only a function that has a defined rank can repartition the argument cells that the rank operator produces. Stating it the other way, for any function that lacks a default rank, you may use the rank operator only in a way that produces cells of a rank satisfactory to the function's ordinary definition.

Default rank is not defined for any user function. Thus, a user function can be usefully modified by the rank operator only in a way that produces argument cells directly acceptable to the unmodified function. Similarly with primitive functions (those denoted by symbols), when you use a function for which default rank has not been defined, the rank operator must first be used to assign a rank acceptable in the ordinary case.

α *m*⍤⍤*g* ω Cut

When the left argument of ***paw*** (⍤) is an array and the right argument is a function, the derived function *m*⍤⍤*g* cuts the argument array ω into segments. The content of the segments is determined by the right (function) argument.

The left argument *m* is an integer scalar specifying the type of cut; *m* must be one of ¯3, ¯2, ¯1, 0, 1, 2, or 3. For all types except for 0, the partitioning occurs along the first axis.

- When the derived function is used monadically, the first (or last) of the major cells of ω is treated as a delimiter which marks positions at which the first axis of ω is to be cut.
- When the derived function is used dyadically, the left argument α specifies positions in ω at which ω is to be cut.

To each of the cells produced, *g* is applied monadically.

Types of Cuts

For every type of cut, the derived function has infinite right argument rank. When used dyadically, the derived function's left argument rank is either 1 or 2, depending on the type of cut, which is controlled by the value of *m*:

m=0	Each cell is selected by an integer table specifying the offset and length along each axis.
m=1 or 2	Cells are selected by a Boolean vector in α which delimits segments along the first axis of ω or by a major cell that marks partitions within ω.
m=3	***Tessellation***.

Cells Selected by Table

When *m* is 0, the derived function has left argument rank of two. Each matrix in α describes the *position and shape* of a cell passed to the function *g*. The description is a table having two rows; it has one column for each axis of ω. Within each matrix of α:

- Row 0 contains the position along ω's various axes from which a block is to start. A negative position is treated modulo the length of the corresponding axis of ω, so ¯1 indicates the last position, ¯2 the position before last, and so on, just as for indices used as arguments to the function ***from*** (⌷ω).

- Row `1` contains the length of the selected segment along each axis of ⍵. For example, an ⍺ containing the numbers

```
2 3 5
2 1 3
```

  indicates that, from an array ⍵ of rank three, a cell is to be extracted as follows:

- Start `2` positions after the start of axis `0` and run along that axis for `2` positions.
- Start at the third position from the end of axis `1` and run (forward) along that axis for `1` position.
- Start `5` positions after the start of axis `2` and run for `3` positions along that axis, but (because the value is negative) place them in reverse order in the cell passed to function *g*.

The `0`-cut requires a considerably more elaborate left argument than the `1`-cut or `2`-cut with a distinct ⍺-cell for each cell passed to *g*. However, the `0`-cut lets you specify the lengths a segment has on all the axes, not just the first. It also permits overlapping and empty segments.

Cells Selected by Boolean Partition

When *m* is other than zero, the derived function *m*⍤⍤*g* has a left rank of one. Each element `1` in an ⍺-cell marks the start (or end) of a cell passed to the function *g*. Each Boolean vector in ⍺ generates as many cells as it contains `1`s. The frame shape is `(¯1↓⍴⍺),⌈/+/⍺`. The maximum is relevant where there are several ⍺-cells and some have more `1`s than others.

Each cell has the same rank as ⍵ and the same length as ⍵ along all its axes except the first along which a cell passed to *g* has a length that depends on the distance between consecutive `1`s in the ⍺ cell. The positions selected for each cell are controlled by the value of *m* as follows:

`¯2` Each `1` in an ⍺-cell marks the *end* of a segment. The last segment ends at the last `1`; positions after it are *not* represented in the result. The position corresponding to each `1` in an ⍺-cell is *excluded* from the result cell.

`¯1`	Each `1` in an `α`-cell marks the *start* of a segment. The first segment starts at the first `1`; positions before it are *not* represented in the result. The position corresponding to each `1` in an `α`-cell is *excluded* from the result cell.
`1`	Each `1` in an `α`-cell marks the *start* of a segment. The first segment starts at the first `1`; positions before it are *not* represented in the result. The position corresponding to each `1` in an `α`-cell is *included* in the result cell.
`2`	Each `1` in an `α`-cell marks the *end* of a segment. The last segment ends at the last `1`; positions after it are *not* represented in the result. The position corresponding to each `1` in an `α`-cell is *included* in the result cell.
`¯3` or `3`	***Tessellation***. The `3`−cut has left rank `2`, and `α 3⍤g ω` applies *g* to each element produced by tessellation of `ω` using a size `1@α` and beginning points that are multiples of the "shift" `0@α`.

Selections produced by the cut operator are made readily visible by `α 1⍤< ω`, and the resulting boxes displayed with `□ps←¯1 1 3 3`.

Examples:

```
      ω←' Worlds on worlds'
      (ω=' ') 1⍤<ω
|¯¯¯¯¯¯¯| |¯¯¯| |¯¯¯¯¯¯¯|
| Worlds| | on| | worlds|
|_______| |___| |_______|
      α←1 0 0 0 1 1 0 0 0 0
      ω←0 1 2 3 4 5 6 7 8 9
      α 1⍤(+≠) ω
6 4 35
      ⊢Table←5 4⍴⍳6
0 1 2 3
4 5 0 1
2 3 4 5
0 1 2 3
4 5 0 1
      (1 0 0 1 0) ¯1⍤< Table
|¯¯¯¯¯¯¯| |¯¯¯¯¯¯¯|
|4 5 0 1| |4 5 0 1|
|2 3 4 5| |_______|
|_______|
```

```
      (1 0 0 1 0) ¯1⍤(+⌿) Table
6 8 4 6
4 5 0 1

      (>2 2⊃2 2)¯3⍤< 5 5 ⍴⍳25

|___|   |___|
|0 1|   |2 3|
|5 6|   |7 8|
|___|   |___|

|_____| |_____|
|10 11| |12 13|
|15 16| |17 18|
|_____| |_____|

      (>2 2⊃2 2) 3⍤< 5 5 ⍴⍳25

|___|   |___|   |¯|
|0 1|   |2 3|   |4|
|5 6|   |7 8|   |9|
|___|   |___|   |_|

|_____| |_____| |__|
|10 11| |12 13| |14|
|15 16| |17 18| |19|
|_____| |_____| |__|

|_____| |_____| |__|
|20 21| |22 23| |24|
|_____| |_____| |__|
```

The table at the beginning points of the elements of the tessellation is given by *s*`x⍤>@⍳⍤>0⌈⌊ ((⍴ω)-1@ω)÷`*s*`←0@⍺.` The case `¯3⍤`*g* is equivalent to `3⍤`*g* except that each cut of shape less than `1@⍺` are omitted.

The following example encloses each word in the character vector *ω* based on a set of common punctuation characters:

```
      ω←'While this (and that) are true, try these.'
      (ω∈Punct) ¯2⍤<ω⊣Punct←' ,.;()'

|_____| |____| || |___| |____| || |___| |____| || |___| |_____|
|While| |this| || |and| |that| || |are| |true| || |try| |these|
|_____| |____| || |___| |____| || |___| |____| || |___| |_____|
```

The result produced by dyadic use of a function derived from the ***cut*** operator has the frame shape of ⍺ followed by the number of partitions, followed by maximum cell-shape resulting when the function *g* is applied to the cut cells. As usual, result cells that are of lower rank or shorter length are padded to match the largest result cell.

m⍤⍤*g* ⍵ Cut

When the derived function *m*⍤⍤*g* is applied monadically, the result is what would be obtained if ⍺ were the Boolean vector `⍺←⍵ (≡⍤¯1) 0@⍵` when *m*`∈1 ¯1`; or `⍺←⍵ (≡⍤¯1) ¯1@⍵` when *m*`∈2 ¯2`. *m* cannot be zero when *m*⍤⍤*g* is applied monadically. When *m*`=3` or `¯3` is specified, a *domain error* occurs.

The partition vector ⍺ is formed by comparing the major cells of ⍵ with the first cell or the last cell of ⍵. The following expression arranges each of the elements between slashes into a table:

```
      ¯1⍤⊢'/The/first/book'
The
first
book
```

m⍤⍤*g* ⍵ With
f⍤⍤*n* ⍵ With

When one of the arguments of the ***with*** operator is an array and the other is a function (understood in its dyadic sense), the derived function is the monadic function formed by taking the supplied array (*m* or *n*) as one of the arguments of the supplied function (*f* or *g*). For example, `1⍤⍤+`*X* computes one more than *X*; `÷⍤⍤6` *X* computes a sixth of *X*.

Functions derived from the ***with*** operator are very useful in the context of the monadic operator ***merge*** (*f*`&`) which requires a monadic function as its argument. For example, you can bind a set of indices to the ***from*** function (⍺`@`⍵) through the ***with*** operator to derive a monadic function for use with ***merge***:

```
      (50 100) 1 3⍤⍤@&⍳4
0 50 2 100
```

6
Syntax

APL symbols are arranged to form expressions, statements and lines according to the rules of APL syntax (see "Chapter 1. Language Overview"). The symbols discussed in this chapter are the ones used for grouping characters meant to be interpreted as a unit or for separating characters meant to be treated as distinct:

⋄	separates statements within a line
''	in pairs, surround a character constant
()	in pairs, surround an expression to be treated together, and hence to be evaluated before other parts of a statement
[]	in pairs, surround a set of expressions used as an index
;	within [ ], separates successive expressions for indexing along the axes of an array
→	alters the sequence in which lines of a function are executed
:	separates a label from the rest of a line
⍝	separates a comment from the rest of a line.

Line: The Unit of Execution

To the APL system, the basic unit of execution is the *line.* In immediate execution mode, the signal to the session manager to pass characters to the APL interpreter is the *newline* character. The newline character is not itself part of what is transmitted; rather, newline signals the session manager that you release what

you have typed for passage to the interpreter. A line in the APL system can be any arbitrary length, but only lines of up to `1024` characters can be entered directly in function definition mode.

When the definition of a function is stored, it is arranged in lines. In the display of a function's definition, the lines are separated by the newline character. Here also, newline is not part of what is executed, but only what separates the successive segments. Executing that line may entail almost any level of work. The simplest possible line is empty, with literally nothing to do. A single line may contain a function that invokes other functions, or loops back to repeat the lines within its own definition, so that its execution goes on, in an immense or even endless calculation. When the system has completed work on that line, it returns to the source of control (that is, the user at the keyboard, or the controlling shared variable of an S-task) for another line to execute.

Lines, Statements, and the Diamond Separator

A *statement* is very much like a line: a sequence of characters passed to the APL interpreter for execution. In the first implementations of APL, there was no distinction between *line* and *statement.* In SHARP APL (and in ISO standard APL) a line may contain several statements. Each statement could be executed as a separate line. However, they are passed to the APL interpreter together.

The statements in a line are separated by the ***diamond*** symbol (⋄). Like newline, diamond is not itself part of the statement; it serves only to separate one statement from the next. The diamond is used in situations such as the following:

- ***To assure that a set of statements, if interrupted, is restarted only from the beginning.*** When the interpreter is executing a user function, it normally executes the lines in succession. When you signal that you wish to interrupt execution, the interpreter suspends work. You may subsequently resume execution. In some contexts, it is important to execute certain statements together. For example, you do not want an interruption to separate a preparatory test or signal from what follows. The ***branch arrow*** (→) causes the interpreter to execute the line having the same number as the value to the right of the arrow. Because the arrow can point only to a *line* (and not to a statement within the line), you can make it impossible to separate the execution of several statements by putting them on the same line. Thus, statements separated by diamonds on a single line form a *block of code* from which escape is possible by branching out, but to which re-entry is allowed only at the first statement. (Interruption and restarting are discussed in the *System Guide,* Chapter 7.)

- ***To adjust the appearance of a program.*** Some authors like to adjust the appearance of a function by putting several statements on the same line, separated by diamonds. This is essentially a matter of personal taste, and has no effect on the content or execution of the definitions.

During execution, if you signal a *weak interrupt* (defined in the *System Guide, Chapter 7*), the APL interpreter completes work on the line it is executing and then halts.

Grouping and Separators Within a Statement

The symbols that follow are concerned with grouping and separating items within a single statement.

Quotes Delimit a Character Constant

APL uses a single quote mark (') which indicates both the beginning and ending of a character constant. Within a statement, a character constant is a sequence of characters delimited by quotes. When there is just one character between the delimiting quotes, the constant is a scalar, otherwise it is a vector.

For example:

```
x←'⍋'
```

When there is any other number of characters (including zero characters) between the delimiting quotes, the constant is a *vector*.

For example:

```
x←'There were 3 ravens.'
```

When a blank occurs between quotes, it is a character like any other (even though it has no visible symbol). The quote character may itself appear in a character constant by typing it twice. Each of the following makes *x* a vector of *seven elements*:

```
x←'She can'
x←'      '
x←'I can''t'
```

Scanning a statement from left to right, the first occurrence of the quote character marks the beginning of a character constant. Until the quote that marks the end of the character constant occurs, any other characters whatever are treated simply as elements within the vector, and have no grammatical significance.

Quotes can surround any character you can type except cursor controls such as backspace, newline, or linefeed. A quoted vector can never span two lines of input. The newline can occur as an element in a variable, but you cannot write it by typing an open quote and then newline; the usual way is to select it from ⎕*av* (atomic vector):

```
NewLine←10@⎕av
```

Parentheses

Parentheses serve to group expressions within an APL statement just as they do in conventional mathematics. That is, they surround items that must be treated together. The items enclosed within parentheses must be evaluated before they can interact with items outside the enclosure. When evaluated, the expression inside the parentheses may produce an array, function or operator.

Parentheses are often required to surround the *left argument* of a dyadic function. For example, in the expression

```
(a-b)÷c
```

the parentheses show that the root function is ***divide*** (⍺÷⍵) rather than ***subtract*** (⍺−⍵), and that *a−b* is the left argument of ***divide***.

In general it is not necessary to enclose the right argument of a function: the function applies to the whole expression to the right of it anyway. The following are equivalent:

```
(a-b)÷(a+b)
(a-b)÷a+b
```

However, when the function is a *derived* function (i.e. the functional result of the application of an operator), you may need parentheses to distinguish the arguments of the operator from the arguments of the function. One way to do that is to put parentheses around the derived function; that is, around the operator and what it modifies. For example, the derived function ⍳⍤0 generates successive integers for each of the items in its argument. When the argument is

itself a constant (for example, `1 2 3`), it becomes essential to separate the argument of the dyadic operator from the argument of the derived function. You could bind the `0` to ⍤ as follows:

```
(⍳⍤0)1 2 3
```

Alternatively, it would have the same effect (but perhaps make the intent less clear) to write

```
⍳⍤0(1 2 3)
```

or

```
⍳⍤(0)1 2 3
```

When several operators appear in succession, it may be necessary to enclose the right arguments in parentheses. For example, the following expression finds the product of one more than each set of integers, in order to compute an array of factorials:

```
×≠⍤(1¨+⍤(⍳⍤0))n
```

Within the outermost parentheses, the root operator is ***on*** (*f*⍤*g*). Its left argument is the derived function `1¨+` and its right argument is the derived function `⍳⍤0`. The operator's right argument needs parentheses but its left argument does not. Similarly, the root operator of the expression as a whole is ***on*** (*f*⍤*g*) with the derived function `×≠` as its left argument and the derived function `1¨+⍤(⍳⍤0)` as its right argument. Here again, the operator's right argument requires parentheses, but its left does not.

Redundant Parentheses

It is not wrong to put redundant parentheses around an expression that produces either a variable or a function–although perhaps faintly wasteful of computer memory or interpreter time. The following are equivalent:

```
a×b
(a×b)
((((a×b))))
(a)×(b)
a(×)b
(a)(×)(b)
```

Brackets Delimit an Index

In an expression such as `X[i]` or `⊖[i]X`, the brackets (`[]`) surround an expression which describes an *index.* The index is the entire expression within the brackets. The bracketed expression applies to the variable or function which it follows. In the statement

```
X[i]
```

the expression within the brackets identifies the subarray to be selected from the array `X`. This form of selection is described in "Chapter 4. Functions".

The selection can also appear to the left of and ***assignment arrow*** (`←`), indicating that within an existing array certain positions but not others receive a new value, as in:

```
X[i]←a
```

Brackets may also be used to surround an expression which modifies the axis along which certain functions operate, as in the following:

```
⌽[i]X
+/[i]X
A,[i]X
```

Obsolescence of Bracket Notation

These uses of brackets are regarded as obsolescent; they are inconsistent with the general syntax of APL in several ways:

- Brackets require a *pair of symbols* to denote a single function (when used for indexing), or operator (when used to specify an axis); this is required of no other APL function or operator.
- Brackets may denote either a function or an axis specification, depending on context. No other APL primitive has this syntactic inconsistency.
- In order to index an array having more than one axis, bracket notation requires a special separator within the brackets, neither used nor required in any other APL construction.
- Indexed assignment is a *root* expression which both forces a new statement in order to reference the new array, and precludes merging a subarray into an existing array without overwriting the existing array. These objections

to bracket notation and a presentation of superior alternatives are developed further in *A Dictionary of APL*. For consistency with past implementations, SHARP APL continues to support bracket notation, with the warning that brackets are acknowledged to be an unsatisfactory aspect of APL for which proposals for replacement have already been implemented and published. See the section on the function ***from*** (⍺⌷⍵) in "Chapter 4. Functions", and the section devoted to rank in "Chapter 5. Operators"

Semicolon Separates Index Expressions

When brackets are used to denote selection by indexing from within an array whose rank is greater than `1`, the bracket notation requires a separate expression for the index on each axis. These expressions are separated by semicolons. For example, when *X* is a rank three array and *i*`0` denotes the index expression for axis `0`, *i*`1` the index expression for axis `1`, *i*`2` the index expression for axis `2`, and so on, the selection is written

```
X[i0;i1;i2]
```

The semicolons act as *complete separators*: that is, even when the values for the various axes are written as complex expressions, it is *not* necessary to surround each of the expressions with parentheses; the semicolons themselves separate the expressions without parentheses. For example, to select from *X* the matrices whose indexes are *a*`⍳`*b*, the rows whose indexes are `⌊(`*n*`+`*r*`)÷2`, and the columns whose indexes are *l*`↓⍳`*w*, you have only to write the following, with no parentheses surrounding the three expressions:

```
X[a⍳b;⌊(n+r)÷2;l↓⍳w]
```

Semicolons serve only to separate the index expressions; they may appear *only* within brackets. The semicolon does *not* denote a function. There is no way to assign the set of expressions within the brackets to a single variable. Indeed, if you write

```
X[index←a⍳b;⌊(n+r)÷2;l↓⍳w]
```

the name *index* receives the value of *a*`⍳`*b* and is unaffected by anything outside the boundary imposed by the semicolon.

Former Use of Semicolons in Vectors

SHARP APL for OS/390 and some earlier implementations of APL permitted two other uses of semicolon, although they were for a long time considered obsolescent. The following are ***not*** supported in this version of SHARP APL:

- ***Output vectors.*** A series of expressions separated by semicolons and entirely surrounded by parentheses caused display of the resulting expressions. The effect is now obtained either by catenating the results of separate uses of format, for example `(⍕a),⍕b`, or by formatting an enclosed array, for example `⍕a⊃b`.
- ***Vector argument to*** `⎕fmt`. The formatting primitive `⎕fmt`formerly accepted a right argument consisting of a series of expressions separated by semicolons and entirely surrounded by parentheses. This appeared as `⎕fmt (a;b;c)`. The same effect is now produced by `⎕fmt a⊃b⊃c` without semicolons and requiring no surrounding parentheses.

The Branch Arrow as a Grouping Device

Within the definition of a single user function, the interpreter executes successive lines in order. It starts with the first and, if nothing indicates otherwise, executes that and then each of the others in succession until it reaches the last. After the interpreter has executed the last line, control returns to the statement that invoked the function.

The branch arrow supplants the normal sequence of execution. It directs explicitly the line to which the interpreter should go next. The destination may be any of the lines in the definition of the function now being executed. Lines are identified by their line numbers, and the destination of a branch may be the number of any of the function's existing lines. *Labels* provide a convenient way of attaching a name to a line. Labels are mentioned below in the discussion of the colon separator, and are described in the *System Guide, Chapter 7*.

A statement containing a branch symbol consists of the branch arrow followed by a expression which, when evaluated, returns the number of the line to be executed next. There is no restriction on how the computation is performed. However calculated, that value determines which line comes next. A value that is *empty* indicates *no branch*: execution continues in the normal order. A value of zero or any integer which is not in fact the number of a line in the function now being executed causes an exit from the defined function.

A line which is *not* pointed to by an arrow can be executed only when the interpreter comes to it in its normal sequence. In that sense, such a line is linked to the one before it: the interpreter can execute it only by first executing the line before. By contrast, a line which is the destination of a branch (and usually has a label) is the start of a block of lines. A branch arrow can point only to a line within the same definition, or to exit. It cannot point to a line in the definition of some other function.

A branch arrow with nothing following it–a so-called naked branch–abandons execution of the program that contains it, and also of all programs that invoked it. The use of the branch arrow alone on a line returns control to the keyboard (or the controlling shared variable) at the point where external control started execution. The branch arrow and expressions for constructing branches are discussed in the *System Guide, Chapter 7*.

Colon and Label

A ***label*** is a name assigned to a line within the definition of a user function. (See the discussion of user-defined functions in *System Guide, Chapter 7*.)

The colon (`:`) serves to identify a *label.* A label is a name formed following the same lexical rules for names of variables or functions. If a line is labelled, the line must begin with the label and be followed immediately by a colon. That is the only use of the colon. Since there may be at most one label on a line, there may be at most one colon. Note that a label is attached to a *line,* not to a statement, so a line containing several statements separated by diamonds may nevertheless contain only one label, and the label must appear first on the line.

Lamp and Comment

The ***lamp*** symbol (`⍝`) marks the start of a *comment,* which is a portion of a line that contains text intended for the illumination of the human reader, but not for execution by the APL interpreter. As it scans a line from left to right, when the interpreter reaches a ***lamp*** symbol (other than one embedded within quotes) it *ignores* the lamp and everything to the right of the lamp on that line.

Note that a comment is part of a *line* rather than part of a *statement.* When several statements appear on the same line (separated by diamonds) each statement may *not* have its own comment, since the first comment causes the interpreter to disregard the remaining characters on that line.

Since ***lamp*** ensures that the remainder of the line will not be executed, it does not matter whether the comment contains unmatched quotes, parentheses, brackets, or additional comment symbols. Nor does anything written in the comment affect the function editor, which might otherwise respond to bracketed line numbers or the symbol ***del*** (∇).

Priority of Quotes and Comment

Any characters enclosed in quotes are simply characters; in particular, they are not names, numbers, symbols, or any executable instruction.

The ***lamp*** symbol has a similar effect: any characters that follow it are simply part of the comment text. When both occur in the same line, which dominates? It depends on position. The following rules resolve which characters are part of a character constant and which are part of a comment.

- Scanning from left to right, mark the zones enclosed by matching quotes.
- Scanning from left to right, note the first ***lamp*** that is *not* inside quotes. That ***lamp*** starts a comment, and all characters to the right of it including quotes are text.

In the body of the line (that is, in the zone that precedes the comment), the number of quotes must be even. However, within a comment, quotes need not be balanced.

SHARP APL *for UNIX*

SVP Manual

JUMP TO ...

OVERVIEW
CONTENTS
PREFACE
MASTER INDEX
USING THIS DOCUMENTATION

HANDBOOK
AUXILIARY PROCESSORS MANUAL
FILE SYSTEM MANUAL
INTRINSIC FUNCTIONS MANUAL
LANGUAGE GUIDE
SVP MANUAL
SYSTEM GUIDE

© 2000 by Soliton Associates Limited
44 Victoria Street, Suite 2100, Toronto, Ontario, Canada M5C 1Y2

Pubcode: UW-000-3001 (0007)

SHARP APL *for UNIX*

DEVELOPER LIBRARY

SVP Manual

SOLITON
ASSOCIATES

© 2000 by Soliton Associates Limited–Printed in Canada
44 Victoria Street, Suite 2100, Toronto, Ontario, Canada M5C 1Y2

All rights reserved. No part of this work may be reproduced in any form or by any means without prior written permission of the copyright owner.

Soliton Associates Limited shall not be liable for errors contained herein or for incidental consequential damages in connection with the use of this material. The information contained within this document is subject to change without notice.

Acknowledgments

SHARP APL is a registered trademark of Soliton Associates Limited.

SVP and NSVP are trademarks of Soliton Associates Limited.

IBM and AIX are registered trademarks of International Business Machines Corporation.

RISC System/6000 is a trademark of International Business Machines Corporation.

UNIX is a registered trademark of UNIX System Laboratories, Inc.

Sun Workstation is a registered trademark of Sun Microsystems, Inc.

Sun, Solaris, SunOS, Open Windows, and SunView are trademarks of Sun Microsystems, Inc.

All other products referred to in this document are trademarks or registered trademarks of their respective trademark holders.

Publication code: UW-001-0501 (0009)

Contents

Tables and Figures vii

Preface

Introduction ix
Chapter Outlines ix
Conventions x
Documentation Summary x
Contacting Soliton Associates xi

2. Shared Variable Communication

Overview 1-1
The Shared Variable Processor 1-1
Sharing with Non-APL Programs 1-2
Auxiliary Processors 1-2
The Network Shared Variable Processor 1-2

2. Discipline of Communication

How Name Sharing Works 2-2
Sending and Receiving 2-3
Establishing a Shared Name from an APL Task 2-4
Saving a Workspace Containing Shared Variables 2-5

3. SVP Variables and Functions

⎕*svo* Shared Variable Offer 3-1
Names and Surrogates 3-1
Processor ID 3-2
Result from an Offer to Share 3-4

Monadic ⎕svo: Inquire About Degree of Coupling 3-5
Duration of Sharing 3-5
⎕svr Shared Variable Retraction 3-5
Implicit Retraction of a Shared Name 3-6
Localization and Shared Names 3-6
⎕svn Clone ID and Sign-on Index 3-6
Clone ID 3-7
Default Clone ID 3-8
Clone ID of an Auxiliary Processor 3-9
Sign-on Index 3-9
Remote Sign-Off 3-11
⎕svq Shared Variable Query 3-12
State and Controls of a Shared Name 3-13
Set and Use 3-13
⎕svs Shared Variable State 3-13
First Pair: Which Sets Are Not Used? 3-14
Second Pair: Who Is Aware? 3-14
Possible Values of ⎕svs 3-15
⎕svc Shared Variable Access Control 3-16
Effective Access Control 3-18
Orderly Communication 3-19
Representation of Access Control 3-19
Shared Names and ⎕pack, 4 ⎕ws, and 6 ⎕ws 3-21
Which Actions Will Not Be Delayed? 3-21
Meaningful Settings for Access Control 3-22
⎕sc State Change 3-22
What Counts as a State Change? 3-23
Sequence for Using ⎕sc 3-23
Setting ⎕sc to Time Out 3-24

4. NSVP

Introduction 4-1
The NSVP Daemon 4-2

Starting the NSVP Daemon 4-2
Using NSVP 4-4
System Identifiers 4-4
Monadic □*svn* 4-4
Dyadic □*svn* 4-5
Dyadic □*svo* 4-5
Monadic □*svq* 4-6

5. Diagnostics/Maintenance
NSVP Log Files 5-1
SVP Tools 5-2
`svpdump` 5-2
`svpmaint` 5-4

Index IX-1

Tables and Figures

Figure 2.1. Shared variable interactions between tasks. 2-2
Table 3.1. Set and use. 3-13
Table 3.2. Shared variable states. 3-15
Figure 3.1. Transition diagram for state changes. 3-16
Table 3.3. Access control. 3-17
Figure 4.1. Example connection of NSVP under UNIX. 4-1
Table 4.1. Optional flags for nsvp.script. 4-3
Table 5.2. Options for svpmaint. 5-4

Preface

Introduction

This document discusses shared variables, the Shared Variable Processor (and Network Shared Variable Processor), and how shared variables permit one APL task to communicate with another APL task or with processes outside of APL. Shared variables, or shared names, are the means of communication between the active workspace of an APL task and other APL tasks or auxiliary processors provided with SHARP APL.

Rather than reproduce existing material, references to other SHARP APL for UNIX publications are supplied where applicable.

Chapter Outlines

This document is organized into the chapters described below.

Chapter 1, “Shared Name Communication,” provides an overview of SHARP APL shared name communication.

Chapter 2, “Discipline of Communication,” discusses name sharing protocol.

Chapter 3, “SVP Functions and Variables,” lists and describes the system functions and variables used in shared name communication.

Chapter 4, “NSVP,” focuses on the Network Shared Variable Processor.

Conventions

The following conventions are used throughout this documentation:

`⎕io←0`	Although the default value for `⎕io` in a clear workspace is *one*, all examples in this manual assume that the index origin is *zero*.
α and ω	Two APL symbols are used in syntax descriptions of monadic and dyadic functions; *alpha* (α) identifies the left argument and *omega* (ω) identifies the right argument.
`constant width`	Examples of non-APL input and output are represented in a `constant width` typeface.
default (%) prompt	Examples in this book assume you are using the C shell. The default prompt for the C shell is the *percent sign* (%).
passno	Argument variables appear in *sans-serif italic* type.
[ ]	Optional arguments are shown between square brackets. For example, in the following APL statement, the argument *passno* is optional:

'filename' ⎕stie tieno [,passno]

These default environment variables represent frequently used pathnames in SHARP APL for UNIX documentation and scripts:

`$SAXLIB`→`/usr/sax`
`$SAXDIR`→`/usr/sax/rel`
`$SAXCNF`→`/usr/sax/local`
`$HOME`→home directory of the current user.

Documentation Summary

The SHARP APL for UNIX reference documentation is published in several volumes. One or more of the following may be referenced in this manual.

SHARP APL for UNIX,

- *Handbook,* publication code UW-000-0401

- *Language Guide,* publication code UW-000-0802
- *System Guide,* publication code UW-000-0902
- *Auxiliary Processors Manual,* publication code UW-033-0501
- *File System Manual,* publication code UW-007-0501
- *Intrinsic Functions Manual,* publication code UW-033-0501

For a complete list of SHARP APL publications, please refer to the ***Support*** link on the Soliton Associates website: `http://www.soliton.com.`

Contacting Soliton Associates

Problems or questions regarding your SHARP APL system or this guide should be directed to the Soliton Associates Technical Support group:

`support@soliton.com`

Requests for additional technical documentation should be directed to the Soliton Associates Sales group:

`sales@soliton.com`

1
Shared Variable Communication

Overview

A name that two *programs* share is like a window between them; indeed, you can say that a shared name is a variable that exists in two different environments at the same time. Shared names (variables) are the means of communication between the active workspace of an APL task and other APL tasks or auxiliary processors that are provided with the SHARP APL system. A shared variable connection requires:

- managing all name sharing for your task and for all other tasks running under the same UNIX system at the same time
- maintaining tables of shared names and the tasks (auxiliary processors or other APL tasks) that are sharing them
- providing intermediate storage for the values set by one partner but not yet used by the other
- enforcing rules that prevent conflict over whose turn it is to set or use a shared name.

For a detailed description of the name-sharing process go to "Chapter 2. Discipline of Communication".

The Shared Variable Processor

This name-sharing work is made possible by a pseudo-device installed in the SHARP APL kernel called the ***Shared Variable Processor***, SVP: it's function is to manage the transfer of data and arbitrate access control between cooperating

programs. "Chapter 3. SVP Variables and Functions", lists and describes the SHARP APL system variables and functions used in shared name communication.

Sharing with Non-APL Programs

Sharing can be established between any pair of tasks that abide by the shared name discipline understood by the SVP. The active workspace of an APL task can share with the active workspace of any other APL task, and with almost any sort of program capable of running concurrently, even one written in a completely different language. In that manner, your APL workspace can send commands to the operating system via AP11 to execute, or make use of, non-APL files. AP11 is documented in the *Auxiliary Processors Manual, Chapter 3*.

Auxiliary Processors

An auxiliary processor (AP) is a program that provides a service to an APL task but which may not itself be part of APL; it uses the SVP to commuicate with the APL task. Auxiliary processors can be used to communicate with programs to perform specialized calculations, manage external communications, or provide special services for managing the session, controlling graphics displays, etc. For more information on auxiliary processors, consult the *Auxiliary Processors Manual*.

The Network Shared Variable Processor

The ***Network Shared Variable Processor***, NSVP, enables program-to-program communication between connected systems using an extended version of the SVP protocol. Information is transferred along a communications line between NSVP processors residing on each of the systems. Communication between the NSVP and the local task is handled by the SVP. The network shared variable processor is fully documented in "Chapter 4. NSVP".

2
Discipline of Communication

The original concept for name sharing was introduced by IBM in APL/SV. They were called ***shared variables*** then, which explains why the processors responsible for sharing names —SVP and NSVP— are called shared ***variable*** processors. It also accounts for the acronyms given to the system variables and functions (⎕`svo`, ⎕`svc`, etc.) used to control the mechanism of sharing. The term *shared name* originates from the realization that it is also possible to share functions and other named objects between users.

A shared name identifies a variable whose value can be set and used concurrently in different contexts. In APL terms, such a context is a workspace; but a name can be shared with a program outside APL where the term ***workspace*** may not apply. Indeed, any two programs can share a name if they abide by the conventions for using shared names.

When two otherwise independent programs (each in its own context) interact, the interaction is kept extremely simple so that the mechanics of communicating have minimal influence on either the form or the content of the dialogue. The mechanics of sharing between two processors active on a single host apply almost unchanged when the two processors are active on separate hosts interconnected through one or more networks.

Each of the cooperating processors begins a dialogue by offering to share a name with the other. When both offers are matched by the SVP, the name is fully shared. When one processor gives a new value to the variable whose name they share (sets it), the other has only to use it to become aware of its new value.

For example, suppose Task 1 and Task 2 represent the active workspaces of two APL users who have agreed to share the name `x`. The left side of Figure 2.1 shows events in Task 1's workspace, and the right side events in Task 2's

workspace. The fact that *x* is shared leads to the following simple example . Task 2 assigns to *x* the value *abc def*. Almost simultaneosly, Task 2 evaluates 2×*x* and gets the answer 10.

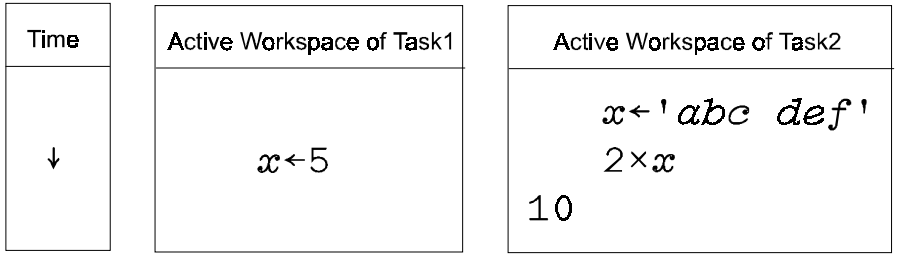

Figure 2.1. Shared variable interactions between tasks.

Of course the reason for this result is that after Task 2 gave *x* the value *abc def*, Task 1 gave *x* a new value. Since *x* is shared, Task 2 now sees the ***new*** value.

How Name Sharing Works

The partners that share a name may be a pair of APL tasks, an APL task and a non-APL task, or two non-APL tasks. Each of the partners keeps its own copy of the variable that the name identifies, but marks the name as shared. Whenever a partner refers to the shared name, that partner checks first with the SVP, which also maintains a place where the value of the shared name can be stored; that is, three places for one name's value. At any particular moment, two of those places contain duplicate copies of the value of the shared name, and the third does not matter.

For example, suppose you have arranged to share the name *x* in your APL workspace with the Host AP (the processor that communicates with the UNIX operating system); and you wish to assign a value to *x* so that the Host AP will see it and treat it as a command to be executed under UNIX, and then get back, in *x*, the UNIX response to executing your command in the same shared name. Typically, this routine appears as follows:

1. You establish that *x* is shared between your APL workspace and the Host AP. (The procedure to do this is described in the next section.) Inside your workspace, the interpreter marks the name *x* to show that it refers to a shared name.

2. You assign a value to x. When the APL interpreter executes that assignment, it notices that x is a shared name, and notifies the SVP that a new value has been received. The SVP stores its own copy of the value, which now exists both in your workspace and in the SVP's own storage. The SVP notes that you know the latest value of shared name x (since you put it there) but that the Host AP does not yet know the value.
3. Like most APs, the Host AP does not initiate work, but waits for orders from an APL task. It remains inactive until notified by the SVP that something has happened to affect a name it is sharing. As soon as the Host AP uses the value of the name that you call x, the SVP transfers its copy of the value to the Host AP. The SVP notes that both partners are now aware of the value of x, and discards its copy. Now the shared name exists with the same value in your APL workspace and in the storage space available to the Host AP.
4. In due course, the Host AP has a reply for you and it sets the value of the shared name. Whenever a shared name is set, the SVP is notified. It saves a copy of the value that the Host AP has set and notes that the Host AP knows the value (because it set it) but that you do not (because you have not yet used the name).
5. When you attempt to use x again, the SVP stores its copy in your workspace, notes that both partners are now aware of the value of x, and discards its copy.

That's how the SVP makes the value of x seem to exist in both workspaces. The value is not in fact transferred until one partner attempts to use it; but that is not apparent because any time a partner asks, ``What is the value of x?'' the SVP slips in and delivers the copy it has been holding.

Sending and Receiving

When two programs communicate by sharing a name, there are no special actions called ***send*** or ***receive.***

- To send some data to a partner, use the regular specification arrow ← to assign a value to the name that you share. That name is visible to the partner as well, and so you have sent the value just by giving a shared name that value in your own workspace.
- To receive data from a partner, you have only to use the value of the shared name.

Set and Use: References

The variable to which a name refers is ***set*** when the name occurs immediately to the left of ← (the assignment arrow). Any other occurrence is a ***use*** of the name. As will be apparent later, it is useful to have a term that includes both actions: the term ***reference*** is used to mean any occurrence of either set or use of a shared name.

Establishing a Shared Name from an APL Task

Following are the steps that each program must go through when it uses shared names. Even if an application embeds them in a cover function so that the steps are no longer apparent each time, the corresponding steps are always the same.

1. Using the function `⎕svo` (shared variable offer), you extend an offer to share a name with another task. The SVP notes your offer; it returns a result of `1` which means that it has received one offer to share this name. Initially, the offer is unreciprocated.
2. The SVP makes your offer available to your proposed partner. When the partner inquires, ``Has anyone offered to share with me?'' the SVP informs it that your task has made an offer. Since an auxiliary processor usually has nothing else to do but serve tasks that offer to share with it, it is usually written so that it is permanently vigilant for new offers to share, and hence notices at once when you extend an offer to it. You could write an APL task so that it behaves in the same way, and is therefore always ready to serve. Alternatively, an APL task expects to share only with the task that created it, and therefore expects no more than one offer.
3. The prospective partner responds with a counter offer. You can tell it has done so because the result of `⎕svo` then becomes `2`, when before it was `1`. When both sides have made matching offers to share the same name, full sharing is in effect: the degree of coupling between the two processors is `2`.
4. Using the function `⎕svc`, you or the AP set the ***shared variable control.*** This interlocks references to the shared name so the partners do not get out of step; it prevents misunderstandings that would arise if one partner were to rush ahead instead of waiting for the other.
5. Now begins a cycle that may be repeated many times as you work with the AP:

   You set the shared name.

Your partner uses the value you have set and does whatever is called for.

Your partner sets the shared name, either with an acknowledgment of your request or with a response to it.

You use the shared name to find out what your partner's response was.

6. At any point in the process, you can use the function `⎕svs` to inquire about the shared variable state. This gives you the answer to the question, ``Who does not yet know the shared name's latest value?'' In this way, you can tell:

   for a ***local partner***, whether your partner has used what you set, or has set a reply for you to use

   for a ***remote partner***, whether your most recent ***set*** of the name has been transmitted to the remote environment.

7. If you are waiting for a reply from several different shared names, but you do not know which of them will reply first, you can have your program wait for a ***state change*** by using the system variable `⎕sc`. When you ***use*** `⎕sc`, your program waits until there is a change in the state of one or another of the names you have shared. Only then does `⎕sc` return a value, whereupon your program can check the state of the name and take appropriate action.

Saving a Workspace Containing Shared Variables

When you save a workspace, SHARP APL requires that the value of each of its variables must be physically present within it. When you give the command `)save`, it is possible that you are sharing a variable whose value you have not yet used (your partner knows the value, but you do not). If you were to use the variable, the SVP would supply the value. But when you save the workspace, the value must be physically present in the workspace (not merely available from the SVP if needed). So, before it saves the workspace, SHARP APL moves into the workspace the current value of each of the shared variables to which you have not yet referred.

If it proves impossible to transfer the value of a shared variable because of insufficient space, the system displays the error message `ws full`, and saves the workspace without a value for the shared variable.

3
SVP Variables and Functions

The system functions ⎕*svo*, ⎕*svr*, ⎕*svn*, ⎕*svq*, and ⎕*svc*, together with the system variable ⎕*sc*, appear in the following sections. The convention for describing the rank of each function is the same as that used in the *System Guide, Chapter 5*: monadic, dyadic left, and dyadic right ranks appear in that order. For functions that are strictly monadic or dyadic, only the first element or the last two elements in the rank array are shown. Two additional symbols are used: ∞ indicates infinite rank, and ⋆ indicates undefined rank.

⎕*svo* Shared Variable Offer

Rank: `1 ¯1 1`

Used dyadically, the system function ⎕*svo* offers to share a name in its right argument with a processor identified in its left argument. The `¯1` indicated for the left dyadic rank indicates that the function treats major cells of the left argument array. (See the section "Making an Offer," below, for a summary of the options available for the left argument.)

Used monadically, ⎕*svo* reports the degree of coupling of variables named in its right argument.

Names and Surrogates

You may find it inconvenient to call your shared variable by the same name that your partner requires. For example, a server task may expect to share the name *command*, but for some reason your application cannot use the name *command* for that purpose. In such a situation, you can make your offer with a surrogate name.

For example, the statement

```
11 □svo 'wish command'
```

offers to the task whose processor ID is specified as `11` a name that in your workspace appears as *wish* but appears to the SVP as *command*. Processor `11` will see an offer to share *command* (not knowing that you privately call it *wish*). Processor `11` may, in fact, take advantage of the same mechanism and use a surrogate name of its own to refer to the variable that the SVP calls *command*.

On each row of the right argument to □*svo* there may be either one name or two names separated by blanks. Where there are two names,

- the one on the left is the name in your workspace
- the one on the right is the name visible to the SVP.

When the right argument ⍵ contains more than one row, ⍵ is a character matrix, with one row per name offered. A single name is offered as a character array or as a character element (for a name that has only one letter).

The name that you offer

- must be well formed (from the characters usually permitted in the names of APL user objects, and subject to the same maximum length). For more information, see *"Names for User Objects"* in the *Language Guide, Chapter 2*.
- must not have a conflicting use (that is, cannot already have a visible use as a label or as the name of a user-defined function). On the other hand, it is not necessary that it already exist as a variable.

Processor ID

Using a right argument ⍵ that is a matrix with several rows and a left argument ⍺ that has a single major cell, you can offer several names to the same processor. However, you cannot offer the same name to more than one processor. When you identify more than one processor (by giving ⍺ more than one major cell), you must give ⍵ the same number of rows as ⍺ has major cells.

The left argument to □*svo* is an array of one or more ***processor IDs.*** The general pattern for this array is an integer matrix with one row per processor ID, which allows a processor ID to have more than one element.

SHARP APL recognizes two classes of processor IDs:

- ***local,*** with two elements to identify a task active in the same environment
- ***remote,*** with three elements to identify a task active in a distinct environment.

SHARP APL requires a processor ID to be unique so that the SVP can unambiguously distinguish all processors using its services.

Local Processor ID

For two processors active in a single UNIX environment to cooperate, each is required to have a processor ID with two elements; the second element is often omitted, and in that case is assumed to be `0`. The following two elements are part of a local processor ID (*PID*):

`0@ PID` The ***processor number***. For an APL user task, this is the account number assigned by the UNIX administrator. For an auxiliary processor, this is a unique number assigned by the author of the AP.

`1@ PID` The ***clone ID***. This number is used to distinguish between different tasks active under the same account number. Each task must set its own clone ID (refer to the discussion of `⎕svn`) so that the combination of account number and clone ID is unique.

Making a Local Offer

When offering to share with a local processor, the three options for the integer left argument are:

`⍺` **as scalar** A single local processor (another APL user or an AP) is assumed to have clone ID `0`; a scalar `⍺` is treated as though it were `1 2⍴2↑⍺`; for example:

```
11 ⎕svo 'ctl'
```

⍺ as vector One or more local processors are all assumed to have clone ID `0`; a vector ⍺ is treated as though it were `((⍴⍺),2)⍴⍺`. The number of elements in ⍺ must agree with the number of names in ⍵; for example:

```
11 124 ⎕svo ⍪'cd'
```

⍺ as matrix One or more local processors have clone IDs as indicated in the second column of the matrix; for example:

```
(1 2⍴119 1) ⎕svo 'dat'
```

Remote Processor ID

For two processors active in separate host environments to cooperate, each is required to have a processor ID with three elements. Of these, the first, called the ***sign-on index,*** is placed at the front of the elements of a local processor ID. A sign-on index identifies a host and its environment.

See the sections on ⎕*svn* for the procedure used to obtain a current value for a sign-on index, and ⎕*svq* for the procedure used to determine if any remote processors have outstanding offers to share a name with you.

Making a Remote Offer

When you use ⎕*svo* to make an offer to a remote processor, the remote processor ID in the left argument must contain all three elements and be presented as an integer matrix. For example:

```
(1 3⍴600 11 0) ⎕svo 'ctl'
```

Result from an Offer to Share

The result of ⎕*svo* is an integer array showing the degree of coupling of each of the names in ⍵. It reflects the degree of coupling immediately following your offer. If you are the first to make a valid offer, the result is always `1`, even when the processor to which you make the offer accepts it immediately upon receipt.

Monadic ⎕svo: Inquire About Degree of Coupling

The number of partners sharing a name is its degree of coupling. The degree of coupling may be 0, 1, or 2. (Multi-way communication among multiple processors is built up from channels each of which has only two ends.)

Monadic use of ⎕svo tells you the degree of coupling of the names in its right argument. When the level is 2, full sharing is in effect. A level of 1 means that the SVP has recorded one offer, but has not yet received a matching offer. A level of 0 means that no valid offer has been made. To verify the degree of coupling of a name that you offered with a surrogate, specify the surrogate in an inquiry with monadic ⎕svo. For example:

```
      11 ⎕svo 'wish command'
1
      ⎕svo 'wish'
2
```

Duration of Sharing

Sharing goes into effect when both you and your partner have made matching offers to share a name. You can confirm this by checking that the result of ⎕svo is 2. Sharing remains in effect until either side retracts the offer. Once you have made a valid offer (and received a result of at least 1), your offer remains in effect until you retract it, implicitly or explicitly.

⎕svr Shared Variable Retraction

Rank: 1

You retract a shared name explicitly by executing the monadic system function ⎕svr. It takes as its argument the name or names that you wish to retract. For example, if *a* has been shared,

```
      ⎕svr 'a'
2
```

retracts the offer. When you are retracting your offer to share a name that you earlier offered with a surrogate, it is permissible (but not necessary) to include both names in the retraction.
For example:

```
      ⎕svr 'wish command'
2
```

The result returned by ⎕*svr* is an integer array with one number for each name in *ω*, indicating the level of sharing that was in effect immediately before you retracted. Thus, a result of *2* indicates that your retraction terminated a fully-coupled share then in effect.

Implicit Retraction of a Shared Name

An offer to share is retracted implicitly whenever the name that you offered ceases to exist, which it does if you execute the system commands *)clear*, *)xload*, or *)off*; execute the system functions ⎕*load* or ⎕*qload*; expunge the name; or exit from a function to which it is local.

It is standard practice to write an AP so that when it finds that you have withdrawn your offer to share, it assumes a normal end to its transactions. It withdraws its counteroffer and tidies up whatever it was doing. If you have shared a name with the Host AP, for instance, and used it to open a file, then erased your shared name (or did anything else that might terminate the share), the Host AP closes the file for you.

Localization and Shared Names

It is permissible to share a name that is local to a user-defined function, but when you do, the sharing ceases as soon as you exit from that function. So you probably do not want to do that unless the function to which your shared name is local is going to become active rarely enough, or remain active long enough, to make it worthwhile to reinitiate sharing each time the function is invoked.

When you have shared a global name *a*, and then, while that share is in effect, you invoke a sub-function that uses *a* as a local name, the local name is the only one visible. Temporarily, your global *a* is shadowed by the local meaning. This does not constitute an implicit retraction, and you regain access to the shared name *a* as soon as the shadowing sub-function completes.

□svn Clone ID and Sign-on Index

Rank: `1 * *`

The system function □svn used monadically enables you to set a ***clone ID*** for use within a local environment, and to set a ***sign-on index*** for use in connection with a remote environment. To distinguish these two uses of the system function, the right argument may be of two types:

- When the right argument is an integer, then the traditional use of □svn applies: it is used to inquire about or to set a clone ID destined to become the last element in a processor ID.
- When the right argument is a character array or an enclosed array, then the extended use of □svn applies: it is used to establish a sign-on index on a remote host.

Used dyadically, □svn lets you remove your presence from shared processing on a remote host: in effect, it lets you sign off by nullifying the sign-on index you specify.

Clone ID

When you refer to a local AP (active on the same host), you identify it by its local user ID. Similarly, when you refer to another APL task active on the same host, you refer to it by its account number (the number assigned by the UNIX administrator and reported in the first element of the system variable □ai).

However, it is quite possible to have more than one task with the same account number. Indeed, AP1 exists to create and manage additional APL tasks active under a single account number. Whenever two or more APL tasks are running under the same account number at the same time, and they both make use of the SVP, they need some further identification to differentiate them. They do it by catenating a unique nonnegative integer to the account number to complete their local processor ID. That additional number is called the ***clone identification*** or, informally, ***clone ID.***

In SHARP APL for UNIX, the value of your clone ID persists across)load and)clear. This is not the case in SHARP APL for OS/390.

For example, if you want your clone ID to be 99, you might execute:

```
id←⎕svn 99
```

The result of ⎕*svn* is an integer indicating the clone ID in effect after execution of ⎕*svn*. If your proposed clone ID is successfully adopted, the result is the same as the argument.

However, if you earlier set some other clone ID for a task and you still have shares (or outstanding offers to share) with others, you cannot adopt a new clone ID; also, you can never set an invalid clone ID. In either of those cases, your clone ID is unchanged and the result of ⎕*svn* is your former clone ID.

You can use ⎕*svn* to inquire what your clone ID is by executing ⎕*svn* `¯1`. Since `¯1` is not a legal clone ID, the SVP does not set your clone ID to `¯1`, but returns your current clone ID setting.

When there is no valid clone ID (usually, because the task has established none of its own, but another task has made us of the default clone ID 0), the result is `¯1`.

Default Clone ID

When you use a function that requires a processor ID but fails to provide a full one, the SVP assumes that you intend clone ID `0`. Notice that an offer to share a name requires complete identification of two processors: the task with which you propose to share and your own. Each must have a valid clone ID. Therefore, the default clone ID (`0`) may apply to either processor:

- Making an offer to another processor without stating its clone ID is assumed by default to be an offer to that account with clone ID `0`. Thus:

  ```
  114 ⎕svo 'x'   has the effect of
  (1 2⍴114 0) ⎕svo 'x'
  ```

- Making an offer to another processor when you have not yet established your own clone ID implicitly sets your own clone ID to `0`. Thus:

  ```
  1 2⍴114 0) ⎕svo 'x'   has the effect of
  (1 2⍴114 0) ⎕svo 'x' ⊣ ⎕svn 0
  ```

To assume that your own task has clone ID `0` is reasonable only for the first of your tasks. If your primary task runs another APL task, and the second task also fails to set its clone ID explicitly, the SVP will not permit it to take clone ID `0` by default since clone ID `0` is already taken by the first of your tasks.

Suppose your task has no valid clone ID, and the default clone ID `0` is already being used by another task under your account number. Your attempt to execute any of the shared name system functions results in the message *`identification in use`*, meaning that another task already has clone ID `0`. Until you establish a clone ID of your own, any attempt to execute a shared name function is ambiguous, and the SVP rejects it. Not only will you see *`identification in use`*, but if you use `⎕`*`svo`* to test the degree of coupling of the variable in question, the result is `0`. Therefore, tasks that use shared names should set their clone IDs explicitly.

Clone ID of an Auxiliary Processor

Auxiliary processors have `2`-element processor IDs, just as APL tasks do. However, each APL task sees a separate version of the AP, and that version always has clone ID `0`. For that reason, it is standard practice to omit the clone ID of an auxiliary processor when you make an offer to share with it.

Whenever you make an offer to share, your offer implicitly contains your processor ID and clone ID.

Sign-on Index

If you wish to share a name with a remote processor, you must determine if the host itself and the resource you wish to use on that host are available to you. There is a close analogy between obtaining a clone ID to distinguish local tasks and obtaining a sign-on index to distinguish remote tasks. If you do not establish a clone ID when you share a name with a local processor, the SVP provides you with the default clone ID, in effect identifying you fully. This act of initial identification is mirrored across the host or network boundary, and again it is `⎕`*`svn`* that enables you to establish a full identification on the remote host.

When `⎕`*`svn`* returns its integer result, it in effect says that you are identified and thereby ready to use the resources to which you have access.

To obtain a remote sign-on index, the right argument to `⎕`*`svn`* must be a character array, or an enclosed array of up to `4` elements. SHARP APL treats a noninteger argument `⍵` to `⎕`*`svn`* as `4↑⊃⍵`. The significance of the elements is as follows:

`0@⍵` ***Name*** of the remote host.

`1@⍵` ***Sign-on*** identifier. For UNIX hosts, the ***user name***; for non-UNIX hosts, the ***account number*** or other identifier appropriate to the particular host system.

`2@⍵` ***Password***. For UNIX hosts, you can (in the spirit of `rlogin`) omit the password, but if you do, password checking is based on the convention that your password in the remote environment is the same as your password in the environment from which you are making the offer. For non-UNIX hosts, you must follow the conventions appropriate to the particular host system.

`3@⍵` Clone ID as seen on remote host.

The user information (sign-on identifier and password) need not be for the user associated with the processor making the offer. On the remote machine, the sign-on index you obtain inherits the access privileges of the user identified in your request to *⎕svn*.

If cells in the argument are empty, their default values are:

- Host—this host
- User—user making the offer
- Password—password of user making offer
- Clone ID—shared names processor on remote system selects a clone ID for you.

Submit the required information to *⎕svn* as a character array (if you supply only the host name), or as an enclosed array of as many as `4` elements.

Example:

```
      ⊢n←⎕svn 'saxony'
100
      ⊢n←⎕svn 'intern'⊃'dje'⊃'passwd'
42
```

The sign-on index, which *⎕svn* returns as its explicit result, is a positive integer suitable for use as the first element in a remote processor ID. A remote offer presents a `3`-element processor ID to *⎕svo* as an integer matrix. For example:

```
      (1 3 ⍴ n,106 0) ⎕svo 'x'
```

The host name may refer to an OS/390 mainframe rather than a UNIX workstation. In this case, if the user name supplied is an empty array, the connection is a normal one—that is, you can make offers to and receive from any shared variable process running on the identified host. If the user name is nonempty, then your connection establishes an OS/390 address space on the remote host, and all subsequent transactions with the remote host are carried out as if you are signed on directly to that host.

Such a connection requires a special configuration. If this type of connection has been configured, then the password will be validated on the remote host, and if correct, an OS/390 address space will be started that is capable of supplying the services of several mainframe APs. The configuration selected for the connection by the system administrator determines which APs are available.

Remote Sign-Off

Dyadic use of ⎕*svn* allows you to remove your presence on a remote host, thereby nullifying a sign-on index you established with an earlier monadic use of ⎕*svn*. The left argument is a single sign-on index currently in use, presented as an integer. The right argument is an empty array. The result upon successful removal of the sign-on index is `¯1`. If the left argument is `0`, then all sign-on indices currently in use are removed. If the left argument is not a currently active sign-on index, ⎕*svn* signals *`remote no shares`*. If the right argument is not an empty rank-`1` array, ⎕*svn* signals *`rank error`*.

Any variables shared with processors that had the nullified sign-on index as the first element of their processor ID are no longer shared; auxiliary processors with no variables shared with their active tasks cease to exist on the remote host.

Example:

```
     42 ⎕svn ''
¯1
```

Implications of Network Sharing

When the processors that share a name are active on distinct hosts, there are three areas of shared name usage that differ from local conventions:

- ***Reoffer of a remotely shared name.*** It is not possible to recouple a remotely shared name. Once the degree of coupling of a remotely shared name goes down to `1`, it must then go to `0`, and a new offer must be made.

- ***Simultaneous set.*** When two processors sharing a remote name each set the name at what appears to be the same time, the processor that actually sees its value appear in the shared name is the one that executed □*svn*.
- ***Timing for multiple-name APs.*** If you have applications that take advantage of auxiliary processors that use a pair of names (such as *ctl* and *dat*) to pass control information and data separately, the arrival of values for the two names cannot be assumed to occur in the order you expect unless the interlocks you set with □*svc* are sufficiently strict to guarantee the proper sequence.

□*svq* Shared Variable Query

Rank: 1

The monadic system function □*svq* inquires about incoming offers addressed to you. You are most likely to need this when designing a program that runs in the workspace of a server task, such as an S-task launched by a parent APL task, or one that serves several users concurrently.

The function □*svq* returns different results depending on the form of its argument:

□*svq* '' — ***Empty argument.*** Result is a 2-column matrix of local processor IDs of tasks that have outstanding offers to you, one processor ID per row. An offer is outstanding when the processor making the offer sees a degree of coupling of 1.

□*svq* ⍬ — ***Boxed empty argument.*** Result is a 3-column matrix of remote and local processor IDs of tasks that have outstanding offers to you, one processor ID per row. An offer is outstanding when the processor making the offer sees a degree of coupling of 1.

□*svq ID* — ***Processor ID argument.*** When the argument is 2-element or 3-element integer array containing the processor ID of a processor that has an offer outstanding to you, the result is a character matrix of the name(s) which that processor has offered to you.

Ordinarily, you use □*svq* '' or □*svq* ⍬ to find out who is making offers. Then, if the resulting matrix of processor IDs is not empty, you inquire individually about each processor ID in turn.

State and Controls of a Shared Name

When two processors share a name, the sharing is symmetric. Neither takes precedence over the other. It does not matter which made its offer first and which second. (Remote sharing is an exception; this is described earlier in the section "Implications of Network Sharing.")

- The monadic system function □*svs* reports the state of a set of shared names.
- The function □*svc* reports on or sets the access controls of a set of shared names.

To preserve the symmetry of sharing between partners, each processor sees itself (in both the arguments and results of these two functions) as the first member of the pair, and its partner as the second. We refer in the descriptions that follow to

- ***you***—the task that executed □*svs* or □*svc*
- ***your partner***—the task with which each name is shared.

Set and Use

Within the result of either □*svs* or □*svc*, the columns are organized this way:

- ***first pair***—setting the shared name
- ***second pair***—using the shared name.

The four columns employed by □*svs* and □*svc*, which refer to use and set by the two partners, are shown in Table 3.1.

Table 3.1. Set and use.

Set the Name		*Use the Name*	
You	Partner	You	Partner

⎕svs Shared Variable State

Rank: 1

The monadic system function ⎕svs reports the ***shared name state*** of each of the names in its right argument. It has no dyadic use. The argument is a character matrix with one name per row. For a single name, the argument may be a character vector or character scalar.

The result of ⎕svs is a Boolean matrix with four columns or a 4-element array (when the argument had rank 1 or 0). There is a row in the result for each name in ⍵. The result indicates which partner has set a value of which the other is still ignorant. It also shows which partner is aware of the shared name's current value.

First Pair: Which Sets Are Not Used?

A 1 appears in one of the first two columns when the latest value of the shared variable has been set by one partner, but not yet used by the other. The position of the 1 indicates which partner set the new value (and therefore knows what it is).

For a particular name (that is, on a particular row of the result), there is no way that 1 could appear both in column 0 and in column 1. However, the first two columns could both be 0. That occurs when both partners are aware of the shared name's value, and neither of them has set a value of which the other is unaware. The first two columns also contain 0 0 when sharing has just been established, and neither partner has yet set the shared name.

Second Pair: Who Is Aware?

The second pair of columns in the result of ⎕svs indicates which of the partners is aware of the shared name's current value.

There are three possibilities:

- only you know
- only your partner knows
- you both know.

The partner who last set a value for the shared name must be aware of it (having just set it). The other partner can be aware of the shared name's value only if it has used the name since the name was set. Thus, the last two columns may contain either one `1` or two `1`s. When sharing has just been established, it is presumed that neither processor has yet set a value for the shared name (but that each knew what value the name had in their respective workspaces before it was shared).

Possible Values of □*svs*

Shared names exist in four possible states, summarized by the state settings shown in Table 3.2. Once sharing has started, there are three distinct states for □*svs* that correspond to the three circles in Figure 3.1. The fourth possible state is all zeros, which is reported when the name is not shared.

Table 3.2. Shared variable states.

Who has set a value the other has not used?		*Who knows the current value?*		*Interpretation*
You	*Partner*	*You*	*Partner*	
`0`	`0`	`1`	`1`	You and your partner are both aware of the current value, and neither has since set the variable. This is the initial value when sharing is first established.
`1`	`0`	`1`	`0`	You have set the shared variable, but your partner has not yet used it.
`0`	`1`	`0`	`1`	Your partner has set the shared variable, but you have not used it.
`0`	`0`	`0`	`0`	The name is not the name of a shared variable.

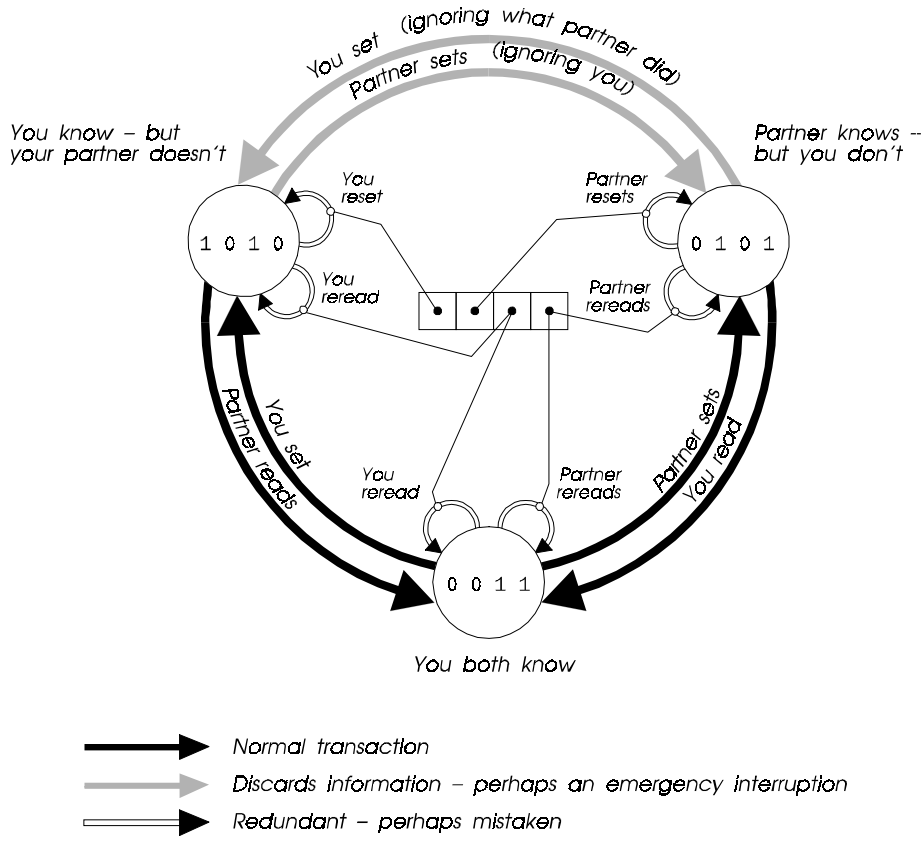

Figure 3.1. Transition diagram for state changes.

□svc Shared Variable Access Control

Rank: `1 1 1`

If you were to set your shared variable with a message for your partner at the same time that your partner is setting it with a message for you, at least one of those messages would be lost. If you were impatiently waiting to use the value of

your shared variable before your partner has had time to put a reply there, you would be rereading your own earlier request, or your partner's reply to an earlier request. `⎕svc` provides the mechanism for avoiding that kind of confusion.

Either partner can assure effective communication by setting the appropriate elements of the access control vector for the name that they share. Access control works by inhibiting one or both partners from doing something too soon. For example, it can inhibit you from setting a new value until your partner has referred to your previously set value. Or it can inhibit you from using the variable before your partner has set it for you.

When the access control is set to inhibit an action, the system does not report an error, but simply delays that action until the necessary conditions have been fulfilled.

Consider the expression that routinely appears in a program to send a command to another task and receive its reply:

```
r←shared ⊣ shared←cmd
```

When the access controls have been set properly, there is no risk of using the result too soon, nor of delaying longer than necessary.

Like its state vector, a shared name's access control is a `4`-element Boolean array. By making the appropriate elements `1`, either partner can specify that certain actions are inhibited. You set the access control by dyadic use of the function `⎕svc`. (Before either partner has set it, the access control is `0 0 0 0`; you cannot set a name's access control until after you have offered to share the name.)

The right argument `⍵` of `⎕svc` is a single name you have offered to share, or a matrix with one name per row. The left argument is the `4`-element array indicating which actions you propose to inhibit, or a `4`-column matrix of actions when you propose different controls for the various names in `⍵`.

The significance of the four elements of the access control is shown in Table 3.3.

Table 3.3. Access control.

Inhibit reset		*Inhibit reuse*	
You	Partner	You	Partner

Wherever a name's access control contains a `1`, the corresponding action is inhibited. For a control vector `c`, the effect of a `1` in each of the four positions is as follows:

`0@ c` ***You*** are inhibited from again setting the value of the shared variable until your partner knows what the value is (because the partner has referred to the variable; i.e., used it or set it).

`1@ c` ***Your partner*** is inhibited from again setting the value of the shared name until you know what the value is (because you have referred to the variable; i.e., used it or set it).

`2@ c` ***You*** are inhibited from reusing the value of the shared variable until your partner has set it.

`3@ c` ***Your partner*** is inhibited from reusing the value of the shared variable until you have set it.

Effective Access Control

Effective access control is control that is actually in effect (to distinguish it from the access control either one of the partners may have requested separately). Effective access control for a shared variable is set by or-ing the access control vectors proposed by each of the partners. The SVP keeps track of which partner requested what. The effective access control vector has a `1` set wherever either partner proposes it.

This rule means that either partner is free to set a `1` in any of the four positions (that is, to inhibit any of the four possible actions). However, one partner cannot change to `0` any element that the other has set to `1`. You can always increase the level of inhibition for yourself or your partner, but you can remove only those inhibitions that you set for yourself.

The APs distributed with SHARP APL set their access control vectors themselves as soon as they can after matching an offer to share a name with you. They set control either to all `1`s or so that it appears to you as `0 1 1 1`. However, the fact that the AP sets access control does not mean that you should not set it also. When you have just offered to share a variable with an AP, you do not want your first instruction to be lost. Since you cannot be certain when the AP will respond and will set its side of the access control, you should not assume that access

control has been taken care of: the AP may not yet have had time to do so when you are ready to send your first command. It is always prudent to start by setting the access control yourself.

The three circles in Figure 3.1 represent the three states of the shared variable. Inside each circle you see the 4-element array by which that state is identified in the result of □*svs*.

The curved paths show the various ways you could get from one state to another. For example, if you are in the state shown by the lower circle (``You both know''), and then you set the shared variable, the path labelled ``You set'' takes you to the state labelled ``You know—but your partner doesn't.''

Orderly Communication

Normal communication is represented by the pairs of solid lines in Figure 3.1 linking the lower state (``You both know'') to one or the other of the states at the top. A normal sequence takes place as follows:

- ***You*** start in the lower circle.
- ***You*** set the shared variable. That takes you to the upper left.
- ***Your partner*** uses the shared variable. That brings you back to the lower circle.
- ***Your partner*** sets the shared variable. That takes you to the upper right.
- ***You*** use the shared variable. That brings you back to the lower circle again; and so on indefinitely.

The other paths represent various exceptional transitions. They arise from programming error or from emergencies. Avoiding the erroneous paths is what □*svc* is for.

Representation of Access Control

At the center of Figure 3.1 you see the effective access control, established by the way you and your partner have both used □*svc*. For each element in the access control array, a line shows which paths are blocked by a value of 1.

Exceptional Routes

Since there are two of you, and either of you could set or use the shared variable, four arrows come out from each circle. Some of the arrows lead back to the same circle. This happens when either of you repeats an action—for example, if you already know the value of the shared variable but use it again. These paths are shown by outline arrows.

There are six of these paths. At best, they are redundant. At worst, they may arise by mistake (for example, you think you are about to read new information, but in fact your partner has not yet sent any). All six redundant paths can be blocked by an appropriately placed `1` in the access control vector.

The two paths across the top are shown by dotted arrows. They represent interruptions: one of you ignores what the other has just set, and sets a new value into the shared variable. When it is your turn to use the shared variable, you have the right to set its value instead (ignoring the value presumably set there by your partner); this is not something you want to do all the time, but it must be possible as a way of signalling to a partner that you have to break out of the ordinary pattern of exchange. Neither of these paths can be inhibited. Sending an interrupting message to your partner does not guarantee the partner will notice or heed it.

Attempting to Carry Out an Action That Is Inhibited

Suppose you have a variable whose state at present is described by `1 0 1 0` (``You have set'') and whose control has a `1` as its first element (to inhibit you from resetting until your partner knows what the value is). You attempt to give the shared variable a new value anyway. What happens? You experience a delay. There is no interrupt or error message, but the interpreter waits to make your specification until your partner uses (or sets) the variable you share. The delay lasts ***indefinitely***—as long as it takes.

If you decide you would rather not wait, signal a ***strong interrupt*** (by transmitting two successive interrupt signals). (You may want to do this if you suspect something is wrong with the AP or that your attempt to reset the shared variable was ill-conceived.) This aborts your attempt to use the shared variable. The interpreter then normally displays the error message *`sv interrupt`* and the statement it was attempting to execute.

A similar delay occurs when a variable's state has a `1` in the third position (meaning that you know what the variable's value is), and there is a `1` in the third position of the variable's access control (to inhibit you from rereading until after

your partner has set). There is an indefinite delay until your partner has again set a new value for the shared variable. As before, you can interrupt the attempt to use the interlocked variable by signalling a strong interrupt.

Warning: While you are trying to discover what went wrong with a program interrupted by an error, you may want to display the current value of a variable that is shared to see what is going on—such a display counts as a reference to the variable and might be inhibited.

To make repeated use of the value of an interlocked shared variable, try assigning its value immediately to another (unshared) variable, so you can refer to it as often as need be. Another technique is to use `6 ⎕ws 'shv'` to examine the value of the shared variable *shv* without counting the examination as reference of the shared name (see below).

Shared Names and ⎕pack, 4 ⎕ws, and 6 ⎕ws

Using the name of a shared variable in the argument to any of these three system functions does not constitute a reference to the shared variable as far as the SVP is concerned. With monadic ⎕*pack* and `6 ⎕ws`, you get the last value visible in your workspace; with `4 ⎕ws` you get the size of the last value visible in your workspace. The only value you can get is a value that you already know (either because you already made some other use of it, or because you are the one who set it). Using any of these three functions has no effect on the shared name's state.

Which Actions Will Not Be Delayed?

By examining a name's state and its access control, you can tell which actions will not be subject to delay. Recall the four possible actions: you set, your partner sets, you use, and your partner uses. For a variable `⍺`, an action you are free to take immediately (that is, with no delay) is any of those indicated by a `1` in the expression:

```
(⎕svs '⍺') ⍲ ⎕svc '⍺'
```

Meaningful Settings for Access Control

`0 0 0 0` ***No interlock.*** Such a setting is reasonable only when you know that other constraints will prevent the two processors from getting out of step. This might happen if the programs that use this shared name without interlocks also make use of another name that is interlocked.

`1 1 0 0` ***Waiting for read.*** Neither of you can set a new value until the other has referred to it. This is sometimes called ***half duplex.***

`0 0 1 1` ***Waiting for write.*** Neither of you can reread the shared variable until the other has set it. This is sometimes called ***half duplex.***

`1 0 0 1` ***Waiting when you send to partner.*** You can't reset until your partner has referred to it, and your partner can't reuse it until you have set. This is sometimes called ***simplex transmission from you to your partner***; appropriate if you are having a rather one-way sort of conversation, as you might while sending instructions to a printer or plotter.

`0 1 1 0` ***Waiting when partner sends to you.*** You can't reread until your partner has set, and your partner can't set until you have referred to it. This is sometimes called ***simplex transmission from your partner to you.***

`1 1 1 1` ***Full interlock.*** This is a two-way conversation; each must receive and then reply. Sometimes this is called ***reversing half duplex***.

□sc State Change

Suppose you just want to wait until an AP accepts an offer. How can you say to APL, "Notify me when it answers?" Or, suppose you have a program that uses several shared names (for example, to communicate with two different auxiliary processors). The program may, on occasion, reach a point where it has nothing to do immediately, but is waiting for a reply to come in from one of its shared variables—but can't tell which one will be set next. How can you have the program wait until something happens that may require its attention?

You wait for a change in the state of shared names by using the system variable □sc, which is automatically shared with SHARP APL. Whenever you use □sc, you get the value that the interpreter has set. That value is either `0` or `1`. The

interesting property of □*sc*, however, is not so much what its value is, but when you get it. While you are using shared names, the SVP must have detected a ***state change*** between any two successive uses of □*sc*. If there has not been a state change, execution is delayed until there is one. So when you receive the value of □*sc*, you are assured that some processor other than your current task has done something that affects the shared name state.

What Counts as a State Change?

A ***state change*** is an action (or several actions) occurring outside your workspace but affecting your use of shared names. A state change is any of the following:

- any new offer-to-share extended to you
- retraction of any offer-to-share extended to you
- any set of a variable you are already sharing
- any set of the access control vector of a shared variable
- any change in the degree of coupling of a name you have already offered.
- the first time your partner uses a value that you set

The system variable □*sc* has the value `0` when you have no valid clone ID. Once your task has established a valid clone ID, the SVP sets the value of □*sc* to `1` each time a state change occurs. You cannot use the value of □*sc* more than once; when you next use □*sc*, the value you get is a value that the SVP sets after the last time you used it. The fact that you get a value indicates that at least one state change has occurred since you last used □*sc*. Thus, the value of □*sc*—when you succeed in using it—is always `1` once sharing is in progress.

When a state change has occurred—either since the share was started or since the last reference to □*sc*—the interpreter returns a value immediately when *sc* is referenced. When there has been no state change, the interpreter waits and does not return the value of □*sc* until a state change occurs. (See the section "Setting □*sc* to Time Out," below.)

Sequence for Using □*sc*

A reasonable sequence for using □*sc* is as follows:

- With □*svs*, verify that there is nothing you can do at present with the relevant shared names.

- Use □*sc* in an expression such as →□*sc*↑*label*. Provided you have a valid clone ID, this usage will delay the branch to *label* until some state change occurs. But if you lack a valid clone ID (and you cannot make effective use of shared names without one), control passes immediately to the next line of your program, where you can deal with this situation.
- When you have received a `1` from □*sc*, again use □*svs* or □*svq* to review the situation and decide what action is possible or appropriate.

Setting □sc to Time Out

Your application can control the amount of time it waits for a state change instead of waiting indefinitely. If you assign any positive number to □*sc*, it is regarded as the number of seconds the system waits before it releases control of □*sc* and returns a value just as though a state change had occurred. Your application can then take appropriate action. For example:

```
      □sc←10
      □sc
1                 ⍝  (10 seconds later)
```

The number you assign to □*sc* must be a numeric scalar. If you assign to □*sc* a variable of rank-`1` or higher, the interpreter does not recognize the value you have set as a valid timing control; the next time you use □*sc*, you will wait indefinitely, just as though you had not set a value at all.

***Note*:** The time-out delay begins counting down as soon as you set □*sc*, not when you attempt to use □*sc*.

If you set a □*dl* delay after you set □*sc* but before you use it, the interpreter keeps the two entirely separate: if the □*dl* completes before the time you set in □*sc*, you wait for the rest of the time you set in □*sc*; otherwise the interpreter lets □*sc* return its value without a wait when you next use it.

Example:

```
      □sc←20      ⍝  set
      ⍝□dl 5      ⍝  delay 5 seconds
      □sc         ⍝  use
1                 ⍝  (15 seconds later)
```

4
NSVP

Introduction

NSVP, the SHARP APL Network Shared Variable Processor, is a multi-user auxiliary processor that enables program-to-program communication between connected systems, using an extended version of SHARP APL's SVP protocol. Information is transferred along a communications line between NSVP subtasks residing on each of the systems. NSVP subtasks and the communications line make up an NSVP *channel*, which represents a unique link between two systems.

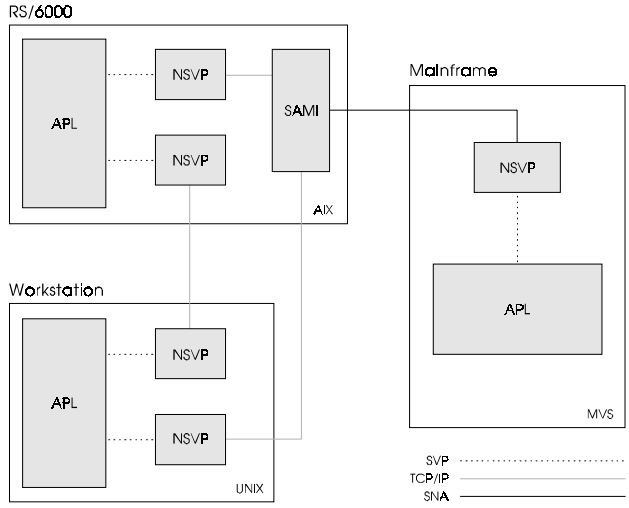

Figure 4.1. Example connection of NSVP under UNIX.

Communication between NSVP and the local task is handled by the host's SVP. Figures 4.1 illustrates one of the connections possible through NSVP. Using the optional mainframe interface SAMI, NSVP can also be used to link UNIX systems with OS/390 systems, with or without APL.

The NSVP Daemon

Each machine that is part of the NSVP network must run an NSVP daemon process, normally started each time the machine is booted. This task communicates directly with the SVP. When an APL task uses the system function □*svn*, the daemon starts a pair of processes, one on the local machine and one on the remote machine. These processes communicate with one another, using TCP/IP, and with the SVPs on the machines where the processes are running.

Starting the NSVP Daemon

The script for starting and stopping NSVP is `$SAXDIR/etc/nsvp.script`. It requires the environment variable `SAXDIR` be set to the appropriate SHARP APL directory (for example, `/usr/sax/rel`) and takes one argument, either `start` or `stop`. The default values specified in the script may be modified using the optional flags shown in the following example and described in Table 4.1.

```
nsvp.script start [-a path/nsvpd] [-c path/nsvp] [-d number]
  [-e | +e] [-f path/nsvp.hosts] [-g alternate:host@port -g ...]
  [-i svpid] [-p iport] [-s path/nsvp.shells] [-w wssize]
  [-x path/apd] [-z path/apd.tab] [-1 path/stdout]
  [-2 path/stderr]
```

Table 4.1. Optional flags for **`nsvp.script.`**

Flag	*Description*	*Default value*
`-a`	Path name of the executable binary for the NSVP daemon process.	`$SAXDIR/bin/nsvpd`
`-c`	Path name of the executable binary for the NSVP child process.	`$SAXDIR/bin/nsvp`
`-d`	Debug flags; produce messages to `stdout` (`2147483647`=full debug).	`0`
`-e` `+e`	If switched on (`-e`), reference `/etc/hosts.equiv` when passwords omitted. Off switch is `+e`.	`On`
`-f`	Path name for list of alternate route names for hosts in `/etc/hosts`	`$SAXDIR/etc/nsvp.hosts`
`-g`	Alias route name to a single host `"alias_name:real_name@ip_port"`; multiple arguments may be supplied.	`None`
`-i`	NSVP shared variable processor ID.	`13`
`-p`	NSVP Internet port number.	`405`
`-s`	Path name for list of shell programs recognized as valid.	`$SAXDIR/etc/nsvp.shells`
`-w`	Workspace size for APL (S-task).	`$WSSIZE, or 500000 if not defined.`
`-x`	Path name for the APL daemon executable binary process.	`$SAXDIR/bin/apd`
`-z`	Path name for list of valid APs and their parameters.	`$SAXDIR/etc/apd_nsvp.tab`
`-1`	Path name for debugging output (`stdout`).	`/dev/console`
`-2`	Path name for status/error log (`stderr`).	`/usr/adm/nsvperrs`

Flags are only applicable to the `start` option, but are tolerated for the `stop` option. Any invalid flags will cause either option to fail. All flags are optional; if not specified, the default values listed in Table 4.1 are used. For flags requiring additional arguments, spaces are allowed but not required; for example:

`-c`*path/name* ***is the same as*** `-c` *path/name*

With the exception of the `-g` flag, only the last occurrence of a flag or option that is specified more than once is retained.

Using NSVP

This section documents those instances where NSVP diverges from SVP protocol. User's should be aware of these differences from SVP, even though in the majority of cases they are transparent. A full discussion of the underlying SVP protocol appears in earlier chapters of this manual.

System Identifiers

Remote systems connected to your system via NSVP have a character identifier known as the `host id`. When you connect to a remote system by supplying the `host id` to *⎕svn*, *⎕svn* returns an integer result known as the neighborhood ID or sign-on index (`SOX`).

To accommodate NSVP, system functions *⎕svn*, *⎕svo*, and *⎕svq* are extended to recognize the remote system's neighborhood ID.

Monadic *⎕svn*

For NSVP, monadic *⎕svn* accepts a character vector argument. The character vector supplied must match the `host id` of a remote system. The result of *⎕svn* `'`*host*`'` (where `'`*host*`'` is the `host id` of a remote system) is the neighborhood ID. If *⎕svn* is used with a character argument that does not identify a remote system with NSVP support, APL error `96`, *`remote no shares`*, is signalled.

The neighborhood value returned by the character vector form of *⎕svn* can be used with the extended version of *⎕svo* to offer variables to tasks running on the selected remote system.

The following examples illustrate some uses of ⎕*svn* using NSVP.

```
      ⎕svn 'hostname'
40
```

This type of ⎕*svn* actually causes a remote signon. The neighborhood ID returned is not predictable, thus ⎕*svn* must be issued at least once, and if issued more than once, will cause multiple remote signons (with different neighborhood IDs).

```
      ⎕svn'hostname'⊃'username'⊃'password'⊃cloneid
41
```

This is the general ⎕*svn*, of which ⎕*svn*'*hostname*' is a special case.

Dyadic ⎕*svn*

The following variant of ⎕*svn* causes an NSVP signoff in UNIX, and also in OS/390 when there is an MVSLink sign-on.

```
      neighborhood ⎕svn ' '
  ¯1
```

The following form may also be used.

```
      0 ⎕svn ' '
 ¯1
```

where your task is signed off from any neighborhood it was signed on to.

Dyadic ⎕*svo*

To accommodate NSVP, dyadic ⎕*svo* accepts a three-column matrix as a left argument. The three columns of the left argument must contain the neighborhood, processor ID, and clone ID of the task or tasks with which you are offering to share. The right argument specifies the variable(s) to be shared. The result of this function is the degree of coupling for the specified variable(s).

Example:

```
      (1 3P3040 2214567 99)⎕svo 'v1'
2
```

Note: When a three-column left argument is used with ⎕*svo*, a neighborhood specification of ¯1 will offer the variable locally.

Monadic ⎕*svq*

For NSVP, monadic ⎕*svq* accepts an argument of an enclosed null value or a three-element vector.

If the argument is an enclosed null value, the result is a three-column matrix containing the neighborhood, processor ID, and clone ID of all incoming offers to share.

5
Diagnostics/Maintenance

This chapter are intended primarily for system administrators. With assistance from Soliton support personnel, the following tools can be used in several situations to examine, troubleshoot, and help maintain the NSVP and the SVP on multi-user systems.

If you have encountered problems with the NSVP and SVP, consult the Soliton Associates Technical Support group for the appropriate diagnostic procedure:

support@soliton.com

NSVP Log Files

SHARP APL can be set up to produce log files for tracing NSVP activity. This information can then be used (if requested by Soliton support personnel) for diagnosis of NSVP problems. The log is established via the environment variable `APL_TRACE` and by editing the `trc task` option in the trace control files `nsvp.trc.ctl` and `nsvpd.trc.ctl`. Trace control files are supplied by Soliton with default settings — only the third column can be changed. The following example shows the first three options in a trace control file:

```
internal             maxfilesize                    5000000
internal             paging                         on
============================================================
trc                 task                           on
```

`maxfilesize` sets the size of the `trc` file in kilobytes. When `paging` is on, a new `trc` file will be started whenever the current `trc` file reaches `maxfilesize`; when `paging` is off, output is accumulated in a single file that is overwritten once it reaches `maxfilesize`.

SVP Tools

SHARP APL includes two methods for monitoring the SVP. The command `svpdump` allows you to display the current state of the SVP. The command `svpmaint` allows you to modify the current state of the SVP.

svpdump

```
svpdump [-ver] [-site] [-proc] [-post] [-shr] [-svs] [-perf]
  [-glob] [-trace] [-loop]
```

The `svpdump` command gathers information about the SVP that can help support personnel isolate problems. When used without arguments, `svpdump` displays all the available information. One or more of the following options may be used to customize the output for specific details:

Table 5.2. Options for `svpdump`.

Option	*Meaning*
`-ver`	`svpdump` version number and SVP version number
`-site`	SVP site-specific information
`-proc`	processor table
`-post`	perpost table
`-shr`	share table
`-svs`	dump of share variable storage
`-perf`	performance analysis variables
`-glob`	dump of various globals
`-trace`	list of SVP events
`-loop`	10 second repeat of `svpdump`

Examples:

```
% svpdump -ver
SVP version 6.0.0


% svpdump -perf
Performance Analysis variables:
-------------------------------
```

```
SVP calls: 3637763, Posts: 151666, WFS: 0
Gcols: 0, # Bytes xfer'd: 4135976, SVP up time(sec): 4144285
Pick cid requests: 0, Extra iterations required: 0


Cmd     #Calls  Histogram
------------------------
svra         1:
on      151042: ....
off     151040: ....
svq    1868203: ..........................................
svo      50455: .
svr        218:
svc      50373: .
svs         12:
uis     251900: ........
sis      50391: .
cis          0:
udt        152:
sdt      50386: .
cdt          0:
ndt      50232: .
cid          2:
inq          0:
sig         75:
wait    963169: ............................
clup         0:
chn          0:
svos         0:
pwait       96:
non          2:
nons         0:
noff         0:
offall       0:
grid         0:
nsvpon       0:
getnon       0:
snrc         0:
srid         0:
ncis         0:
ncdt         0:
dec          0:
peek        14:
poke         0:
```

```
% svpdump -trace

SVP Trace

6104 TTE009  svfwait: pid 20483, mask=ffffffff, pending events=0
6105 TTE012  _sv_waitpost: pid 20646 resume, rc 34
6106 TTE003  _sv_ioctl: pid 20646, command end return code SVZINT
6107 TTE002  _sv_ioctl: pid 20646, starting command SVPWAIT
6108 TTE009  svfwait: pid 20646, mask=ffffffff, pending events=0
6109 TTE002  _sv_ioctl: pid 20646, starting command SVPOFF
6110 TTE053  svfoff: sign off ap 1 0, ppi 3
                                  .
                                  .
                                  .
```

svpmaint

```
svpmaint -respa |-sweep | -avail n | -ttypes n |
  -init [update|shutdown][users n] [shares n] [size n]
     [pin|nopin][ttypes n]
```

The command `svpmaint` allows `root` (superuser) to modify the state of a running SVP. The options available for the use of `svpmaint` are listed in Table 5.2.

Table 5.2. Options for `svpmaint`.

Option	*Meaning*
`-respa`	Reset performance analysis variables
`-sweep`	Sweep system for dead processes
`-avail` *n*	Set availability flag to *n* (testing only) `0` disable SVP `1` enable SVP
`-ttypes` *n*	Previous traced tte types
`-init [update|shutdown] [users` *n*`] [shares` *n*`] [size` *n*`] [pin|nopin] [ttsize` *n*`] [ttypes` *n*`]`	Initializes the SVP.

SHARP APL *for UNIX*

System Guide

JUMP TO ...

SYSTEM OVERVIEW
CONTENTS
PREFACE
MASTER INDEX
USING THIS DOCUMENTATION

HANDBOOK
AUXILIARY PROCESSORS MANUAL
FILE SYSTEM MANUAL
INTRINSIC FUNCTIONS MANUAL
LANGUAGE GUIDE
SVP MANUAL
SYSTEM GUIDE

© 2000 by Soliton Associates Limited
44 Victoria Street, Suite 2100, Toronto, Ontario, Canada M5C 1Y2

Pubcode: UW-000-0902 (0007)

SHARP APL *for UNIX*

USER LIBRARY

System Guide

SOLITON ASSOCIATES

© 2000 by Soliton Associates Limited–Printed in Canada
44 Victoria Street, Suite 2100, Toronto, Ontario, Canada M5C 1Y2

All rights reserved. No part of this work may be reproduced in any form or by any means without prior written permission of the copyright owner.

Soliton Associates Limited shall not be liable for errors contained herein or for incidental consequential damages in connection with the use of this material. The information contained within this document is subject to change without notice.

Acknowledgments

SHARP APL is a registered trademark of Soliton Associates Limited.

SVP and NSVP are trademarks of Soliton Associates Limited.

IBM and AIX are registered trademarks of International Business Machines Corporation.

RISC System/6000 is a trademark of International Business Machines Corporation.

UNIX is a registered trademark of UNIX System Laboratories, Inc.

Sun Workstation is a registered trademark of Sun Microsystems, Inc.

Sun, Solaris, SunOS, Open Windows, and SunView are trademarks of Sun Microsystems, Inc.

All other products referred to in this document are trademarks or registered trademarks of their respective trademark holders.

Publication code: UW-000-0902 (0009)

Contents

Tables and Figures xi

Preface

Introduction xiii
Chapter Outlines xiii
Conventions xiv
Documentation Summary xv
Contacting Soliton Associates xv

1. SHARP APL for UNIX

Introduction 1-1
System Overview 1-2
System Variables and Functions 1-2
APL Tasks 1-3
APL Workspaces 1-3
APL Files 1-4
Shared Variable Communications 1-5
Auxiliary Processors 1-6
Associated Processors 1-6
System Commands 1-6
System Facilities 1-7
Event Handling 1-8

2. Tasks and Workspaces

Tasks 2-1
T-tasks and S-tasks 2-2

N-tasks 2-3
The Workspace 2-5
The Active Workspace 2-5
Saved Workspaces 2-6

3. Files, Libraries, and Directories
UNIX Files 3-1
Pathnames 3-2
Directories 3-2
APL Files 3-3
Libraries 3-4
Naming Conventions 3-6
Qualifying Suffix 3-7
UNIX File Permissions 3-8
`fcreatesw=1` 3-8
`fcreatesw=0` 3-9

4. Shared Variable Processing
SVP and NSVP 4-1
System Variables and Functions for Shared Processing 4-2
Using NSVP 4-2
Auxiliary Processors 4-3
The AP Daemon 4-3

5. System Variables and Functions
System Variables 5-1
System Functions 5-2
Results of System Functions 5-2
Argument Ranks of System Functions 5-4
Alphabetical List of System Variables and Functions 5-4
□ Evaluated Input/Output,
□ Character Input/Output 5-4
□*ai* Accounting Information 5-10
□*append* Append Variable to Tied File,
□*appendr* Append Variable to Tied File with Result 5-10

)arbin Arbitrary Input,
)arbout Arbitrary Output 5-11
)av Atomic Vector 5-12
)avm Mainframe Atomic Vector 5-12
)avail File System Availability 5-13
)bind Intrinsic Function Bind 5-13
)bounce Terminate Task 5-14
)cr Canonical Representation 5-15
)create Create a File 5-15
)ct Comparison Tolerance 5-15
)dl Delay Execution 5-18
)drop Drop Components from a Tied File 5-18
)ec Environment Condition 5-18
)EDIT System Editor 5-19
)er Event Report 5-19
)erase Erase a Tied File 5-19
)ex Expunge Objects 5-19
)fc Format Control 5-20
)fd Function Definition 5-20
)fhold Hold Tied Files
)hold Hold Tied Files 5-23
)fi Input Format Conversion 5-23
)fmt Format Output 5-24
)fsm Full-Screen Manager 5-32
)fx Fix a Definition 5-32
)hold Hold Tied Files 5-33
)ht Horizontal Tab 5-33
)io Index Origin 5-34
)lc Line Counter 5-35
)lib File Library 5-35
)load Load a Workspace,
)qload Quiet Load a Workspace 5-36
)lx Latent Expression. 5-37

)MF Monitoring Facility 5-37
)NA Name Association 5-39
)names Library Names of Tied Files 5-41
)nc Name Class of Identifiers 5-41
)nl Name List 5-42
)nums Tie Numbers of Tied Files 5-42
)pack Build a Package 5-47
)PARSE Pre-parse a User-defined Function 5-48
)paths Pathnames of Tied Files 5-49
)pdef Package Define 5-49
)pex Package Expunge 5-50
)pins Package Insert 5-50
)plock Package Lock 5-50
)pnames Package Names 5-51
)pnc Package Name Class 5-51
)pp Print Precision 5-52
)ppdef Protective Package Define 5-53
)PR Prompt Replacement 5-53
)ps Position and Spacing in Display 5-54
)psel Package Select 5-56
)pval Package Value 5-56
)pw Printing Width 5-56
)qload Quiet Load a Workspace 5-57
)rdac Read Access Matrix of a Tied File 5-57
)rdci Read Component Information 5-57
)rdfi Read File Information 5-57
)read Read Component of a Tied File 5-58
)rename Rename a Tied File 5-58
)replace Replace Variable Stored in a Component 5-58
)resize Reset File Size Limit 5-58
)rl Random Link 5-58
)run Start N-task 5-59
)runs Report Running Tasks 5-63

□*sc* Shared Variable State Change 5-63
□*signal* Signal Event 5-63
□*size* Size of Tied File 5-64
□*sp* Session Parameter 5-64
□*stac* Set Access Matrix of a Tied File 5-65
□*stie* Share-tie 5-65
□*stop* Stop Execution of User-defined Function 5-65
□*svc* Shared Variable Access Control 5-66
□*svn* Shared Variable Clone ID 5-66
□*svo* Shared Variable Offer 5-67
□*svq* Shared Variable Query 5-67
□*svr* Shared Variable Retraction 5-67
□*svs* Shared Variable State 5-67
□*SYSVER* System Version 5-67
□*tie* Exclusive-tie 5-68
□*trace* Trace Execution of a User-defined Function 5-68
□*trap* Event Trap 5-69
□*ts* Time Stamp 5-69
□*twsid* Termination Workspace Identification 5-70
□*ul* User Load 5-72
□*untie* Untie Tied Files 5-72
□*USERID* UNIX User Name 5-72
□*vi* Verify Input 5-72
□*wa* Work Area Available 5-72
□*ws* Workspace Information 5-73
Summary of System Variables and Functions 5-76

6. System Commands

Workspace Names 6-2
Alphabetical List of System Commands 6-4
) Recall Previous Line 6-4
)*binds* Display Bound Intrinsic Functions 6-5
)*clear* Clear Active Workspace 6-5

)copy Copy Objects into Workspace,
)pcopy Protected Copy 6-7
)drop Delete a Workspace 6-8
)edit Full-Screen Editor 6-8
)erase Erase Global Objects 6-8
)fns Display Function Names 6-9
)ifdef Display IF Definition/Redefine IF 6-9
)iferase Remove IF Definitions 6-10
)iflist List IF Groups and Their Contents 6-11
)lib Display Saved Workspace Names 6-12
)libs Display Library Numbers and Directory Paths 6-13
)load Load a Workspace
)xload Load a Workspace 6-13
)off Terminate APL Session 6-14
)save Save Active Workspace 6-15
)sesm Session Manager Conditions 6-16
)sic State Indicator Clear 6-18
)si State Indicator,
)sinl State Indicator with Name List 6-18
)symbols Set Symbol Table Size 6-19
)vars Display Variable Names 6-20
)wsid Set Workspace Name 6-20

7. System Facilities

Session Manager 7-2
The Session Log 7-2
Scrolling 7-3
Support of the APL Character Set 7-4
Signalling an Interrupt 7-6
Programming and Immediate Execution 7-7
Defining Variables and Functions 7-8
User-Defined Functions 7-8
The ∇ editor 7-11
Locking a Function's Definition 7-13

The Full-Screen Editor 7-13
Activating the Editor 7-14
Multiple Screens 7-14
APL Session versus Editor Screen 7-15
Editor Commands 7-15
Errors 7-20

8. Public Workspaces
Library `1` 8-1

9. Event Handling
Events 9-2
Errors 9-2
Interrupts 9-2
Signalled Errors 9-3
Unreportable or Untrappable Events 9-3
`⎕er` Event Report 9-5
`⎕ec` Environment Condition 9-6
`⎕trap` Trap Definition 9-6
Event Numbers Field 9-7
Qualifier Field 9-8
Action Field 9-8
Recovery Expression Field 9-8
`⎕signal` Signal Event 9-12

Index IX-1

Tables and Figures

Figure 3.1. Sample file system tree. 3-2
Table 5.1. System variables. 5-2
Table 5.2. System functions and their argument ranks. 5-3
Table 5.3. Effects of comparison tolerance. 5-16
Table 5.4. Diagnostic codes for `3 ⎕fd ⍵` 5-22
Table 5.5. Result of `⎕fsm`. 5-32
Table 5.6. `⎕NA` classification codes. 5-40
Table 5.7. `⎕nc` classification codes. 5-42
Table 5.8. `⎕pnc` classification codes. 5-51
Table 5.9. Result codes for `⎕run`. 5-62
Table 5.10. Columns of `⎕runs`. 5-63
Table 5.11. Elements of `⎕ts`. 5-69
Table 5.12. Actions for `⎕twsid`. 5-71
Table 5.13. Right argument of `1 ⎕ws ⍵`. 5-73
Table 5.14. Result of `2 ⎕ws 3` 5-74
Table 5.15. Result of `2 ⎕ws 4` 5-75
Table 5.16. Result of `5 ⎕ws ⍵`. 5-76
Table 5.17. Categories of system variables and functions. 5-77
Table 6.1. System commands. 6-2
Table 6.2. Conditions in a clear workspace. 6-6
Table 7.1. Single key commands. 7-16
Table 9.1. Event numbers. 9-4

Preface

Introduction

This document describes the broad range of SHARP APL for UNIX system support that is not itself part of the interpreter. Although it is the principal reference for system functions, variables, and commands, this document presents an excellent overview of system features such as APL tasks, workspaces, shared variables, auxiliary processors, APL component files, event handling, public libraries, and system facilities.

The file system, SVP/NSVP, and auxiliary processors are fully documented in separate SHARP APL for UNIX publications. Rather than reproduce existing material, references to these publications are supplied where applicable.

Chapter Outlines

The SHARP APL for UNIX System Guide is organized into the chapters described below.

Chapter 1, “SHARP APL for UNIX”, provides an overview of the SHARP APL system.

Chapter 2, “Tasks and Workspaces”, describes different types of APL tasks and workspaces, explains how they are used, and defines the relationship between the two.

Chapter 3, “Files, Libraries, and Directories”, discusses the framework of the SHARP APL file system (account numbers, workspaces, component files, and libraries) under UNIX.

Chapter 4, “Shared Variable Processing”, explains, in general terms, SVP technology and how it is applied to auxiliary processors such as AP1, AP11, and AP124.

Chapter 5, “System Variables and Functions”, is the general reference for all SHARP APL system variables and functions.

Chapter 6, “System Commands”, is the general reference for all SHARP APL system commands.

Chapter 7, “System Facilities”, discusses various system components including the session manager, session log and tools for editing user-defined functions.

Chapter 8, “Public Workspaces”, presents a list of the public workspaces that are currently distributed with SHARP APL for UNIX.

Chapter 9, “Event Handling”, covers the SHARP APL event-handling facility.

Conventions

The following conventions are used throughout this documentation:

`⎕io←0`	Although the default value for `⎕io` in a clear workspace is *one*, all examples in this manual assume that the index origin is *zero*.
`⍺` and `⍵`	Two APL symbols are used in syntax descriptions of monadic and dyadic functions; *alpha* (`⍺`) identifies the left argument and *omega* (`⍵`) identifies the right argument.
`constant width`	Examples of non-APL input and output are represented in a `constant width` typeface.
default (%) prompt	Examples in this book assume you are using the C shell. The default prompt for the C shell is the *percent sign* (%).
passno	Argument variables appear in *sans-serif italic* type.
[ ]	Optional arguments are shown between square brackets. For example, in the following APL statement, the argument *passno* is optional:

```
'filename' ⎕stie tieno [,passno]
```

These default environment variables represent frequently used pathnames in SHARP APL for UNIX documentation and scripts:

```
$SAXLIB    →    /usr/sax
$SAXDIR    →    /usr/sax/rel
$SAXCNF    →    /usr/sax/local
$HOME      →    home directory of the current user.
```

Documentation Summary

The SHARP APL for UNIX reference documentation is published in several volumes. One or more of the following may be referenced in this guide.

SHARP APL for UNIX,

- *Handbook,* publication code UW-000-0401
- *Language Guide,* publication code UW-000-0802
- *SVP Manual,* publication code UW-001-0501
- *File System Manual,* publication code UW-037-0501
- *Auxiliary Processors Manual,* publication code UW-033-0501
- *Intrinsic Functions Manual,* publication code UW-007-0501

For a complete list of SHARP APL publications, please refer to the ***Support*** link on the Soliton Associates website: *http://www.soliton.com.*

Contacting Soliton Associates

Problems or questions regarding your SHARP APL system or this guide should be directed to the Soliton Associates Technical Support group:

support@soliton.com

Requests for additional technical documentation should be directed to the Soliton Associates Sales group:

sales@soliton.com

1
SHARP APL for UNIX

Introduction

This chapter provides an overview of the following standard system features of SHARP APL for UNIX:

- system variables and functions
- system commands
- public workspaces
- shared variable processors, SVP and NSVP
- auxiliary processors
- an APL file system
- system facilities such as editors and the session manager.

Four different APL language interpreters are distributed with the UNIX version of SHARP APL. The default, `apl`, provides a session manager and full-screen editor. `apl-8` is a complete interpreter that is invoked through a UNIX session manager. The run-time interpreters, `rtapl` and `rtapl-8`, support the execution of closed applications, but do not allow the user to enter APL immediate execution mode.

For more information on selecting APL interpreters, refer to the instructions on customizing your work environment in the *Handbook*, *Chapter 3*.

System Overview

At the most basic level, SHARP APL is a system that executes instructions written in APL[1]. Instructions may be typed one statement at a time from a keyboard for *immediate execution* or contained within a *program* (which, in APL, is also called a user-defined function). Anyone who is enrolled in the SHARP APL system, and thus entitled to enter instructions for execution, is said to be a *user.*

The UNIX operating system maintains the list of enrolled users and identifies each user both by a name and by a unique account number. Anyone who can log on to the UNIX host can choose to run SHARP APL—SHARP APL does not impose additional passwords or requirements for enrollment.

System Variables and Functions

Instructions written in APL interact with the SHARP APL system and the UNIX host by way of a special class of variables and functions each of which is denoted by a fixed *distinguished name*—one that starts with ⎕ (quad).

- *system variables* are variables that are always shared with the system; you can refer to their values, but so can the system. The system's behavior may be changed by setting the value of a system variable. For example, the system variable ⎕*pw* is used to change the page widths for displays.
- *system functions* are functions related to the use of the system. For example, several system functions (⎕`create` , ⎕`erase`, ⎕`read` , ⎕`append` , etc.) are used to manage transactions with the file system.

Like APL *primitives*, system variables and functions may be used in a program, are subject to the normal rules of syntax, and are available throughout the system. Each distinguished name is reserved for the system and cannot be applied to used-defined objects.

These are described in "Chapter 5. System Variables and Functions".

1.The SHARP APL language (formal APL) is described in detail in the SHARP APL for UNIX *Language Manual.*

APL Tasks

A session of work executed for a user is considered an APL *task*[2] A task is actually a UNIX process (via the APL interpreter); each is associated with an account number to identify the task's owner. A user may have several tasks running at the same time. Once a task has been started, it may itself start other tasks.

In general, tasks are independent. Tasks owned by different users are completely invisible to each other. There is only one situation in which a task has direct knowledge of what another task is doing and is directly able to control that task: when the first task has initiated the second one.

SHARP APL recognizes three kinds of tasks: T-tasks, S-tasks, and N-tasks.

- A **T-task** is an APL task controlled from a keyboard and monitor, either directly from the UNIX machine's console or from a terminal connected to an external port on the UNIX machine.
- An **S-task** is a subordinate APL task with no monitor or keyboard. It is started and controlled by another APL task through a variable shared with the auxiliary processor AP1. Such a task receives instructions from the controlling task through values contained in a shared variable; results are returned by setting the same shared variable.
- An **N-task** is a non-interactive task with no keyboard or monitor. It is started by the system function □*run* for immediate execution.

See "Chapter 2. Tasks and Workspaces", for further information on APL tasks.

APL Workspaces

Each active task has associated with it a working environment called an *active workspace*. There is a one-to-one relationship of task to active workspace.

An APL workspace serves both as the environment in which names are understood and the working storage in which calculations take place. A *saved workspace* is a snapshot of the active workspace saved for later use.

2.In SHARP APL documentation, the only tasks discussed are those that involve APL, so we refer to any of them just as a "task."

Certain criteria must be met while a program is being executed: the program's definition must be present in the active workspace; any data used by a calculation must be present in the active workspace; and, there must be enough room in the active workspace to produce results or partial results during calculations. For more information on workspaces, see "Chapter 2. Tasks and Workspaces".

Public Workspaces

Commonly used programs are often placed in workspaces that are saved in one or more *public libraries*. For more information on public workspaces, including those supplied with SHARP APL for UNIX, refer to "Chapter 8. Public Workspaces".

APL Files

APL files are collections of data outside the APL workspace that are accessible to statements executed from within the workspace. These files are made up of components that are identified by component numbers. SHARP APL includes system functions that allow you to transfer these components between the file and the active workspace.

The APL component file system establishes a common administration for all use of APL component files by any APL task. It not only manages the separate use of private files, but permits files to be shared and manages concurrent read/write sharing of the same file by several concurrent tasks.

SHARP APL can also access non-APL files, such as standard UNIX files. A non-APL file appears to APL as a continuous character vector. Users access these files by sharing a variable with an auxiliary processor. For example, users can access UNIX files by sharing a variable with the Host AP, AP11.

APL Libraries and UNIX Directories

Historically, APL systems have had naming conventions different from those of UNIX: the only unit of storage was the saved workspace, and workspaces were identified solely by the owner's account number and a single-level name. Sets of workspaces having common account numbers are called workspace libraries. The same naming convention is extended to APL files: each has a two-part name consisting of a numeric account number and a single-level name. A set of files having a common account number is called a file library.

SHARP APL for UNIX offers two alternative styles of naming files and workspaces. One is consistent with the naming conventions of earlier APL systems, and the other is consistent with current UNIX practice:

APL library style	The name is written in the form `123 data`
UNIX path style	The name is written in the form `/usr/sue/data`

These are different ways of referring to the same file or workspace. The APL system maps its numeric account and library information to the UNIX system's directory of names.

See "Chapter 3. Files, Libraries, and Directories", for complete details of the SHARP APL file system (account numbers, workspaces, component files, and libraries).

Shared Variable Communications

A *shared variable* can be defined as a variable whose value is set and used in different environments at the same time. In SHARP APL, shared name communications are handled via multi-user processors called the *shared variable processor* (SVP) and the *network shared variable processor* (NSVP).

The SVP permits one APL task to share a variable (and thereby communicate) with other APL tasks, or with an auxiliary processor. It provides the link through which an APL workspace is able to start and then communicate with subordinate APL tasks, or to communicate with the independent APL tasks of other users. Shared variable communication can be between any tasks which follow SVP protocol, including those outside of APL.

The NSVP enables program-to-program communication between connected systems using an extended version of the SVP protocol. Information is transferred along a communications line between NSVP processors residing on each of the systems. Communication between the NSVP and the local task is handled by the SVP.

For more information on the SVP and NSVP, see "Chapter 4. Shared Variable Processing". Also, refer to the SHARP APL for UNIX *SVP Manual*.

Auxiliary Processors

An *auxiliary processor* (AP) is a task that accepts commands and returns results via the SVP. Typically, an AP offers a specialized service to an APL task, doing work that the APL task is not itself equipped to do. The AP is not usually written in APL and often operates under a command language unrelated to APL. SHARP APL for UNIX is supplied with three auxiliary processors: AP1, AP11, and AP124.

AP1 runs another APL task which behaves in exactly the same way as an APL task controlled from the keyboard, except that its input and output is via the shared variable rather than by way of the keyboard and terminal.

AP11 communicates with the UNIX operating system, letting you interact directly with UNIX files and issue UNIX commands directly from your active workspace. Every time the shared variable is set, a UNIX subshell is started to perform the command. The subshell terminates after the command is executed, and its result is returned via the shared variable.

AP124 provides full-screen capabilities to applications running in APL that require full-screen interaction with the user.

For more information refer to the SHARP APL for UNIX *Auxiliary Processors Manual*

Associated Processors

You can prepare C routines in UNIX as *associated processors* to be used as user-defined functions in your APL applications. The mechanism for doing this is provided by the system function `⎕NA` (explained in "Chapter 5. System Variables and Functions").

System Commands

In the management of a session, there are certain commands that are not part of your task. Rather, they are instructions to the system to tell it how to manage the session. These are called *system commands* (see "Chapter 6. System Commands").

The system commands constitute a language completely outside of APL. They are used to communicate directly with SHARP APL for a number of housekeeping tasks, such as saving or loading a system workspace, ending an

APL session, and so on. Each system command starts with a right parenthesis (for example, the)*clear* command); the) symbol distinguishes a system command from a statement in APL.

A system command must come to the system through standard user input. For a task using SHARP APL directly (T-task), the system command must be typed from the keyboard; for a task controlled by a variable shared with AP1 (S-task), it must come as input from the controlling shared name. A system command cannot appear as part of the definition of a user-defined function. However, certain system functions—such as □*load*— parallel the system commands, and thus provide a way of performing the same action from within a user-defined function.

System Facilities

SHARP APL provides some basic facilities to integrate the APL interpreter with its environment, both within the active workspace and beyond it. (These are further described in ("Chapter 7. System Facilities".)

The Session Manager

The session manager is responsible for the management of an interactive APL task. It is responsible for transferring output from the APL system to the device driver, for receiving and interpreting input from the device drivers, and for providing basic services to the user at a terminal (such as scrolling, editing, copying, and moving text). SHARP APL provides a built-in session manager with its default interpreter `apl`. It also gives you the option of using a UNIX session manager with the interpreter `apl-8`.

The Session Log

The session manager maintains a log in which all session input and output are recorded. This is a record of each side of the dialogue between you and the APL system, with a newline character at the end of each line. A useful feature of the session log is the ability to scan the log and reuse earlier input or output as part of your new input.

The ∇-editor

This traditional APL editor is used for entry and revision of user-defined functions. You invoke it by typing the symbol ∇ followed by the name of the function to be edited. The ∇-editor was designed for the sole purpose of creating definitions of user-defined functions and is ill-suited for other kinds of editing.

The Full-Screen Editor

The session manager provides a full-screen editor for your APL session, and also makes it available for editing *user-defined functions* and *variables of rank* `1` or `2` ***character type***.

Event Handling

The flow of an APL session may be diverted or interrupted by a number of circumstances, some internal and some external. Collectively, these are called *events*—a deliberately neutral term intended to embrace a variety of circumstances.

For each class of events, the system defines a standard response. However, an APL task may substitute its own list of actions to be taken when an event of a particular class occurs. The facility for supplying alternate actions when events are encountered is called *event trapping.* Through this facility, recovery from some conditions can be automated.

More detailed information is provided in "Chapter 9. Event Handling".

2
Tasks and Workspaces

The basic organization of the SHARP APL system can be defined as a relationship between tasks and workspaces. A task is a sequence of work executed in an active workspace. A workspace is both an environment in which variables and functions are defined and the working storage in which calculations take place. This chapter describes different types of tasks and workspaces, explains how they are used, and defines their relationship in the SHARP APL system.

Tasks

All work done during the interval from the start to the termination of an APL session is called a *task.* For each task there is an active workspace where all computation takes place.

There are three types of APL tasks, referred to as T-tasks, S-tasks, and N-tasks. These three task types are defined as follows:

T-task A T-task is started from the command processor of the host operating system. A T-task is presumed to have a human controller at a standard input/output device, specifically a monitor and keyboard. (The name "T-task" is derived from "terminal task.")

S-task An S-task is started by the auxiliary processor AP1 in response to an offer to share a name with it. The offer usually comes from an already-running APL task (but in principle could come from a non-APL task, provided it communicated correctly with the SVP).

The task that made the offer to AP1 controls the S-task by setting the name that the initiating task shares with AP1. Instead of typing lines from a keyboard, the controlling task assigns a vector of characters to the controlling shared name. Instead of receiving output on a screen or printer, the controlling task refers to the shared name to see values that the interpreter has set there. (The name "S-task" is derived from "shared variable task".)

N-task An N-task is started by an already-running task as a side-effect of the system function □*run*. Such a task starts as soon as you execute □*run*. The argument to □*run* supplies what the task needs to know. (The name "N-task" is derived from "non-interactive task.")

T-tasks and S-tasks

In T-tasks and S-tasks, the APL interpreter executes APL statements. The process is *interactive* and *cyclic:* the controller (a human for a T-task or the controlling task for an S-task) provides one line to be executed, the interpreter executes it, the controller provides another, and so on. The cycle is as follows:

1. The task's active workspace is initially in immediate execution mode.
2. The interpreter receives a line, either typed from the keyboard or transmitted by setting the shared variable that controls an S-task.
3. The interpreter executes the instructions contained in the line received. The line may be a system command or one or more APL statements. For example:
   - When the command is)*load*, the interpreter loads the active workspace as instructed and then executes the line it finds stored in the new workspace's □*lx*.
   - When the statement being executed calls for the execution of other programs, they are executed in turn.
   - When the statement being executed calls for input or output, the interpreter collects input or transmits output, as indicated. During that time it is operating in *input mode*; what it receives is treated as input and not as a fresh command.

     This continues until the interpreter has completed execution of the controlling line last given to it (and all its consequences), or an interrupting event occurs, signalled either from the controlling task, or from an error condition the interpreter encounters.

4. Execution halts. If execution has been completed, the state indicator is as it was before the controller sent the last controlling line. If execution has not been completed (because work was halted by an error), the uncompleted work may remain on the state indicator.
5. The interpreter examines □*trap* for an instruction regarding what should be done when execution halts. This instruction may cover various possible errors, interrupts, or halts for any reason, including normal completion. If the interpreter finds there a trap instruction that covers the present case, it executes it. In this case, the halt is imperceptible to the user or controlling task.
6. When execution has halted and □*trap* provides no instruction covering the halt, the interpreter returns to immediate execution, issues the default input prompt (the newline and six blanks) and waits for further input. The line next entered may call for a new execution, or may resume execution of the suspended work left at the top of the state indicator, however the user or controlling task decides.

The default value of □*trap* is empty. When □*trap* has not been set, the interpreter returns to immediate execution and awaits new instructions following any halt (whether from error, interrupt, or normal completion). However, when □*trap* has been given an instruction that covers the case at hand, the interpreter immediately executes the instruction it finds there. The effect of a □*trap* instruction is to do automatically what you might otherwise do by typing a new instruction from the keyboard. The system variables □*trap* and □*er* (event report), together with the system variable □*ec* (environment condition) and the system function □*signal*, are described in "Chapter 9. Event Handling".

N-tasks

In addition to interactive tasks (which are directly controlled from a keyboard or through a shared variable), SHARP APL also provides non-interactive tasks called N-tasks.

An N-task has no ongoing interactive control. It receives its instructions at the moment it is launched. After that, it progresses on its way until it completes what it started to do, or fails due to an error or interruption. Everything an N-task does must proceed from its initial instruction. There are two ways to assign the starting points: *autostart* and *split workspace*.

Autostart. You provide the name of a saved workspace. When the system starts your N-task, it loads the workspace you have indicated then executes the workspace's latent expression, □*lx*. When SHARP APL completes execution of the latent expression, the task ends.

Split workspace. The system starts your N-task by cloning your present active workspace. The new task continues executing exactly the same program as your current active workspace. There is no loading of a saved workspace, and hence no latent expression. However, the spawned N-task will be in a somewhat different state from its parent, as it will not have any files tied, nor will it have any shared variables outstanding. It too runs only until execution of that program is completed, and it ends when the program's execution halts.

Starting an N-task

An N-task is launched by □*run*. The argument to □*run* is a character vector consisting of five fields separated by blanks. These fields contain the following information:

- the account number that will own the task
- the name of the workspace to be loaded to start the task
- the name under which the task's active workspace will be saved when the task ends
- the maximum number of CPU units permitted for this task
- the maximum elapsed time, in seconds, permitted for this task.

Note: Currently the CPU and elapsed time limits are not honored; however, they are reported in the result of □*runs*.

For more information on the arguments to □*run*, see "Chapter 5. System Variables and Functions".

Result of □*run*

The result of the system function □*run* is a two-element numeric vector indicating the outcome of your instruction to start an N-task. When the first element is `0`, the N-task was successfully started. The second element is usually the task number of the new task. In the case of a split-workspace □*run*, the task number reported by the spawned N-task is `0`. This allows the parent task and the spawned task to distinguish themselves from one another.

When the first element is not `0`, the N-task was not started, and the first element of the result is an error code indicating why.

For an explanation of □*run* return codes, see the description of □*run* in "Chapter 5. System Variables and Functions".

The Workspace

The workspace serves both as the environment in which names are understood and the working storage where calculations take place. This area of working storage is called the *active workspace*. Using the system command `)save`, you can make a copy of all the variables and functions present in the active workspace, creating a duplicate *saved workspace.*

Programs are developed and executed from within the active workspace. You can establish new definitions for variables and functions or make use of definitions that are already written. During execution, a user-defined function can refer to any function or variable whose definition is present in the same workspace; however, these functions can establish their own *local names*, which are hidden from general view.

Keeping user-defined functions and variables together in a workspace makes it easy for them to refer to each other. You can combine several functions in a single statement. The fact that they are all stored in the same workspace makes it unnecessary to *link* the definition of a function with the names that it uses (although in some other languages, that is required when separate programs work together). For more information on the tools used for defining variables and functions, see "Chapter 7. System Facilities".

The Active Workspace

When a task starts, SHARP APL allocates a block of the system's storage to serve as the task's active workspace. Depending on the task's level of activity, the active workspace may reside in the system's main storage, or it may reside on auxiliary storage, such as disk. The active workspace is the place in which all the task's computation takes place. Any program or data involved in a calculation must be present in the workspace at the moment the calculation is performed.

Size of the Active Workspace

The amount of memory available in the initial workspace is limited by the parameters used to start the task. After your APL task is started, the amount is determined by the size with which a loaded workspace was saved, or by specifying other limits using `)clear` or `)load`. Additionally, the operating system may impose limits on the amount of memory available and these limits

may depend on the demands that other tasks are making on the UNIX system's real or virtual memory. APL includes a system function, $\square wa$, that returns the amount of free space remaining in the active workspace.

Saved Workspaces

The SHARP APL system permits you to set aside a copy of everything in the active workspace. Such a snapshot of the active workspace is called a *saved workspace*. Development of an APL application usually involves bringing together the various definitions it will use and then saving them together as a saved workspace.

The command `)save` records the contents of the active workspace to permanent (but inactive) storage. Also, at the end of a non-interactive task or when an interactive task is interrupted by a communications or system failure, the SHARP APL system automatically saves the task's active workspace.

Note: It is also possible to save a workspace while sharing a name (via the SVP) whose value you have not yet used. This situation is explained in more detail in the *SVP Manual, Chapter 2*.

Saving the Execution Stack

A saved workspace contains the *execution stack* used by the interpreter while executing a program. A workspace may be saved in such a way that its work starts automatically whenever it is reactivated, as a convenience to its users. It is possible to save a workspace with work in progress; that is, with execution of some function started but not completed. That may happen while you are debugging a program. When you subsequently reactivate the saved workspace, the entire working environment is restored to exactly the same state that existed when it was saved, and you can resume work where you left off; however, shared variables are no longer shared and tied files are no longer tied. Restarting with $\rightarrow x$ starts your session at the beginning of line number x of the suspended function.

3
Files, Libraries, and Directories

The APL workspace storage facility and the APL component file system evolved in an environment in which the APL system managed the storage of APL files entirely independent of its host's native file systems. APL workspaces and APL component files are named in a two-level hierarchy: the first level is a user account identified by a number; the second level is an arbitrary name assigned by the user when creating the workspace or file.

Under UNIX, SHARP APL is not restricted to that two-level representation of names. It is entirely your choice: you can access APL workspaces and component files by their library-style names, or you can use the UNIX path format. Similarly, file functions may also take arguments that are UNIX-style names.

This chapter describes the framework of the SHARP APL file system (account numbers, workspaces, component files, and libraries) under UNIX.

UNIX Files

The UNIX file system is organized in a hierarchial *tree* structure where files are the *leaves*; directories, which list groups of files, are the *branches*; and the initial directory, the starting point of the tree, is the *root*. A directory contained within another directory is also called a *subdirectory.*

A full reference to a file traces a *path* that starts from the root directory, passes through each subdirectory, and finishes at the location of the file itself. A listing of this path, which is another way to identify a file, is called a *pathname*. Figure 3.1 shows part of a hierarchial file system shaped in the form of an inverted tree.

You can expand the file system by adding any number of subdirectories (branches) to the tree, and subdirectories to those subdirectories. You may also use identical filenames in different locations of the file system at the same time.

For example, in Figure 3.1, there are two files named `beans`: one is in the `soup` directory, and one is in the `chili` directory. Although these files share the same name, they are identified with different pathnames.

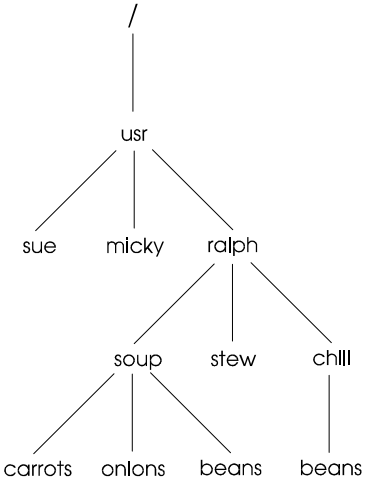

Figure 3.1. Sample file system tree.

Pathnames

Within a pathname, each directory is separated from the following level by a slash (`/`). The full pathname of a file starts with a slash (the root directory), and includes as many subdirectory names as necessary before indicating the filename. For example, pathnames for the two `beans` files look like the following:

/usr/ralph/soup/beans
/usr/ralph/chili/beans

Directories

The UNIX administrator by convention sets aside the directory `/u` or `/usr` for the work of individual users. For example, in a system whose users are Smith, Dupont, and Nakamura, the administrator creates the following directories:

/usr/Smith
/usr/Dupont
/usr/Nakamura

Each of those is the private *home* directory of a user. Normally, users may read but not modify files in the directories of other users.

Names of files use your current directory as their relative point of departure. Using the simple filename `beans` is understood to mean the file of that name within the current directory. Referring to a file `chili/beans` is understood to mean the file reached by going from the current directory down one level to the subdirectory `chili` before selecting the file called `beans`.

You use double dots (`..`) to indicate a directory one level closer to the root than the current directory. The name `../data/xxx` indicates a file called `xxx` within the directory `data` which is a *sibling* (has the same parent) of the current directory.

APL Files

The SHARP APL system comprises account numbers, workspaces, component files, and libraries:

Account number	UNIX user identification number that also serves as your APL library number.
Active workspace	Environment where variables, functions, and packages are understood, and the working storage area where calculations take place.
Saved workspace	Copy of the active workspace saved in storage.
Component file	Collection of data stored independently but accessed by an active workspace.
Library	Group of saved workspaces or component files belonging to one account number.

For more information on APL workspaces see "Chapter 2. Tasks and Workspaces". For more information on component files, see the *File System Manual, Chapter 2*.

Under UNIX, SHARP APL workspaces and component files are written to or read from UNIX files saved with an APL-specific internal format. Files created from within your active workspace but for which you do not state a path are, by default, written in your *home* UNIX directory.

It is also possible to read and write a UNIX file from APL using AP11 (see the *Auxiliary Processors Manual*, Chapter 3). In this case, the contents of a UNIX file appears to APL as a continuous character vector. You may read or write the entire file as a single object. Alternatively, you may read or write a block of consecutive bytes. A block is described by its length and its offset from the beginning of the file. Characters are passed back and forth between the UNIX file and the active workspace without translation or modification.

Libraries

Under OS/390, SHARP APL enforces two types of libraries: public and private.

By convention a *public* library is accessible to all users of the system, whereas a *private* (home) library is accessible only to a single owner. However, because UNIX controls file access–and hence workspace access–this distinction is not as clear under UNIX as it is under OS/390. SHARP APL cannot enforce private ownership of files beyond what UNIX enforces.

Home Libraries

SHARP APL for UNIX provides the conceptual equivalent of private libraries, called *home libraries*. The system creates these libraries by associating the user IDs of all enrolled users with the full paths of their respective home directories. Thus, in general, an APL user's home directory and user ID will share the same number.

When you start an APL session, the system builds a table of these associations. You can display this table by using the system command `)libs`. The table displayed looks similar to the following:

```
1 /usr/sax/rel/lib/wss      100 /usr/guest
109 /home/gsd               108 /home/msy
104 /home/sax               101 /home/users
116 /home/gbe               125 /home/dlf
111 /home/ema               148 /home/ran
118 /home/aba               105 /home/bca
102 /home/mnu               110 /home/radhak
```

Although the system automatically associates a user's home directory with his library number, any user may override this association using the session startup parameter `-L` (see the *Handbook, Chapter 3*).

Public Libraries

You can also establish public libraries. One such library, ***library*** `1`, is established automatically at the start of each session. As distributed, SHARP APL for UNIX provides one public library, which associates the library number `1` with the directory `$SAXDIR/lib/wss`. This library contains several public workspaces and files that provide utilities for the SHARP APL system. For more information, see "Chapter 8. Public Workspaces".

Library/Path Mapping

The `-L` startup parameter can be used to add or modify the available public libraries; you can associate any available directory with your user number (see the *Handbook, Chapter 3*). In such an APL session, any workspace or component file you access in its simple form (without library number or UNIX special characters such as slash, dot, dollar sign, or tilde) refers to this newly defined home library. For example, if you are user `154`, you could start a session requesting the directory `/usr2/mohan/wss` as your home library:

```
% sax -L154/usr2/mohan/wss
```

The following example names a workspace and saves it in your home library, regardless of the current directory.

```
     )wsid drivers
was clear ws
     )save
drivers saved 2000-02-19 14:24:36
```

Using UNIX backquotes as part of the startup parameter, SHARP APL can be set up to make your current directory your home library. This type of special setting can be stored in the UNIX environment variable `APLPARMS`, or in a private system profile called `$HOME/.saxrc`. The following UNIX command starts SHARP APL making the current directory the new home directory of user `154`.

```
% sax -L154`pwd`
```

Naming Conventions

Under UNIX, SHARP APL allows you to access APL workspaces and component files in two ways: by library or by directory.

Library Format

You are allowed to organize as APL libraries those UNIX files in which APL workspaces or component files are stored. Each library is identified by a number. If the number is 100 or greater, it is associated with the home directory of a particular user (private library). Workspaces or component files intended for communal use are organized in public libraries, which are normally associated with an arbitrary number between 1 and 99 that is set up using the `-L` start-up parameter.

Library numbers `1` to `99` are reserved for public libraries. Should a user account be assigned a user number in this range, that user will be unable to save a termination workspace in his home directory using library format, as the APL interpreter prohibits saving termination workspaces in public libraries.

Furthermore, such a user will not have a default termination workspace name set for him when his process starts, so will never save a termination workspace without explicitly using *⎕twsid* to set a termination workspace name.

When you provide the name of a file in directory format, SHARP APL attempts to carry out the requested action on the name as given. When you provide an argument in library format, the name of the file or workspace is either a simple name (with no slash, tilde, dots, or dollar sign), or begins with a library number. This name is converted into a complete UNIX pathname (from the root directory) before SHARP APL attempts to deal with the request you have made. For this to work, the name you give must indicate a file in your home directory, a file in a private library to which you have access, or a file in an existing public library. The system command *)libs* presents a vector of all public and private libraries available to your session (see "Chapter 6. System Commands").

Directory Format

Since APL workspaces and APL component files reside in UNIX files within directories, they have UNIX pathnames associated with them.

As explained in the section, "UNIX files," files can be identified either by their full pathnames, starting from the root directory (`/`), or by a relative point starting from the current directory.

Double dots (`..`) refer to levels above the current directory, the tilde (~) refers to your own home directory, and the dollar sign ($) refers to UNIX environment variables which contain pathnames or portions of pathnames.

File Name Examples

To load the workspace called *utils* from `fred`'s home directory, you write a command of the form

```
)load /usr/fred/utils
```

If `fred` is user `101`, then the expression

```
)load 101 utils
```

loads the same workspace.

If the directory `/usr/fred/utils` is the current directory when you start your session, then you can type

```
)load ./utils
```

from SHARP APL, which has the same effect as the above two commands.

As distributed, the UNIX command `sax` defines only one public library which associates the library number `1` with the directory `$SAXDIR/lib/wss`, which contains public utility workspaces and component files (see "Chapter 8. Public Workspaces").

Qualifying Suffix

SHARP APL identifies UNIX files that contain APL files by their suffixes: the suffix `.sf` is attached to the name of a UNIX file which is structured as an APL component file; and, the suffix `.sw` is attached to the name of a UNIX file that contains a saved APL workspace. These suffixes are visible from the UNIX shell, or from commands executed for you through AP11. However, when you refer to an APL file with the APL file functions or to a workspace using APL system functions or system commands, you neither see the suffix nor write it.

The system command)*lib*, for example, when used with no argument, shows you a list of saved APL workspaces in your home directory. SHARP APL carries out the command by finding all files in your home directory that have names ending in `.sw`. When it shows you the names, it removes the suffix. When you ask to load a workspace, you ask for it by a name that does not include the `.sw` suffix; that suffix is automatically supplied when your request is interpreted.

Similarly, the function □*lib* shows you the names of APL component files in a particular directory. SHARP APL executes that statement by looking for files in a specified directory whose names end in `.sf`. When it returns the names to you, it removes the suffix. When you use an APL file, you give its name without the suffix; the system automatically supplies the suffix when it interprets your request.

UNIX File Permissions

The owner of an APL file may specify access controls for it. Every file has an access table that identifies which user accounts may tie it and which file functions each user account may employ with respect to it. At creation, a file has an empty access table that permits its owner to do anything to the file and denies any form of access to any other user. The access table may require the task that ties it to supply a *passnumber* (a numeric password) in order to tie the file.

However, the underlying UNIX files are always accessible outside the APL interpreter under the rules of the native UNIX file system. Since the standard UNIX file permission bits give relatively little detailed control over access to a UNIX file, it is not possible to reflect the access rules implicit in APL.

Two different schemes for handling this issue are made available with SHARP APL for UNIX. The preferred scheme can be chosen for each APL session, using the "suit yourself" session startup option `fcreatesw` (see the *Handbook, Chapter 3*). The two schemes are as follows:

`fcreatesw=1`

When an APL file is created, the APL file system will attempt to create the underlying UNIX file with permission bits 666 (read/write access for owner, group and user). The actual bits set will be modified by the rules of the user's `umask` setting at the time. Thus, by adjusting the `umask` setting, the user can completely control the permissions set on the UNIX file at the time the APL file is created (though all files created during a single APL session must use the same

`umask` setting). Any subsequent change in the APL file access table is not accompanied by any change in the UNIX permissions for the underlying UNIX file. These UNIX file permissions can, of course, be changed outside of the APL interpreter by normal UNIX administration procedures.

`fcreatesw=0`

When an APL file is created, the APL file system attempts to create the underlying UNIX file with permission bits 600 (read/write access for owner only). The actual bits set are modified by the rules of the user's `umask` setting at the time. Any subsequent change in the APL file access table may be accompanied by a change in the UNIX permissions for the underlying UNIX file, according to the following rules:

- If only the owner appears in the access table, the UNIX permissions are set to 600 (read/write access for owner only).
- If any user other than the owner appears in the access table, the UNIX permissions are set to 666 (read/write access for owner, group and user).

The general intention of these two schemes is that, since it is impossible in general to map the APL rules exactly onto the UNIX rules, one has to choose between the following:

- sometimes giving too much UNIX access but never too little (`fcreatesw=0`).
- sometimes giving too little UNIX access but never too much (`fcreatesw=1`).

If you prefer to ensure that the APL file access table rules not fail because of UNIX rules, at the expense of lax UNIX protection, you would choose `fcreatesw=0`. If you prefer to control UNIX permissions outside of APL, you would choose `fcreatesw=1`.

4
Shared Variable Processing

This chapter is an introduction to SHARP APL's SVP technology and how it is applied to the auxiliary processors AP1, AP11, and AP124. For more information on shared variable communications, please refer to the *SVP Manual*. SHARP APL Auxiliary Processors are described in the *Auxiliary Processors Manual*.

SVP and NSVP

The SVP makes it possible for the same variable to appear in two otherwise separate contexts by:

- managing all name sharing for your task and for all other tasks running under the same UNIX system at the same time
- maintaining tables of shared names and the tasks (auxiliary processors or other APL tasks) that are sharing them
- providing intermediate storage for the values set by one partner but not yet used by the other
- enforcing rules that prevent conflict over whose turn it is to set or use a shared name.

The Network Shared Variable Processor, NSVP, uses an extended version of the SVP protocol to enable the same type of communication between connected systems.

System Variables and Functions for Shared Processing

The system functions □*svn*, □*svo*, □*svc*, □*svs*, □*svq*, and □*svr*, together with the system variable □*sc*, can be applied during the various stages that each program must go through when it uses shared names:

- □*svn* (shared variable name) is used to establish the identifier which distinguishes multiple tasks for the same user number.
- □*svo* (shared variable offer) is used to extend an offer to share a name with another task.
- After both tasks make matching offers to share the same name, □*svc* (shared variable control) is used to interlock references to the name they share.
- At any point during the sharing process, □*svs* (shared variable state) may be used to determine which side is unaware of a shared name's latest value, □*svq* (shared variable query) may be used to inquire about incoming offers, and □*svr* (shared variable retraction) may be used to retract an offer to share a name.
- While waiting for a reply from several different shared names, the system variable □*sc* (state change) is used to determine a change of state in one or another shared names (and, to delay execution until there is one).

Using NSVP

NSVP, the Network Shared Variable Processor, lets an APL task share variables with other APL tasks or APs on separate machines linked through a TCP/IP network. Using the optional mainframe interface SAMI, NSVP can also be used to link UNIX systems with OS/390 systems, with or without APL.

The NSVP Daemon

Each machine that is part of the NSVP network must run an NSVP daemon process, normally started each time the machine is booted. This task communicates directly with the SVP. When an APL task uses the system function □*svn*, the daemon starts a pair of processes, one on the local machine and one on the remote machine. These processes communicate with one another, using TCP/IP, and with the SVPs on the machines where the processes are running.

The SVP, the NSVP, and the system variables and functions that facilitate the use of shared names are fully described in the *SVP Manual, Chapter 3*.

Auxiliary Processors

A SHARP APL auxiliary processor, AP, is a program that provides a service to an APL task but is not itself part of the APL interpreter. An AP communicates with an APL task through the use of *shared names*. Each AP runs in the UNIX environment at the same time as SHARP APL and communicates via the Shared Variable Processor, SVP, which is installed as a pseudo-device driver in the UNIX kernel.

The following auxiliary processors are distributed with SHARP APL:

AP1	**The APL auxiliary processor**. By sharing a name with AP1, one APL task can start and control another APL task. Input and output is handled via a shared name.
AP11	**The Host auxiliary processor (Host AP)**. By sharing a name with AP11, an APL task can pass commands to and receive responses from a UNIX shell. Commands may be *simple,* for transmitting a character vector, or *enclosed,* for transmitting a nested array.
AP124	**The full-screen auxiliary processor**. By sharing a name with AP124, an APL task can pass commands and data to a full-screen manager and capture data from fields on the screen. These full-screen capabilities are similar to those provided by SHARP APL for OS/390.

Refer to the *Auxiliary Processors Manual* for more information on SHARP APL auxiliary processors.

The AP Daemon

When you start an APL session, a special process (called the AP daemon) is intialized on your behalf. As soon as you offer to share a name with an AP, the special process starts the AP so that it will reciprocate your offer to share. If you do not require any of the services APs provide, no AP tasks are started for you.

The AP daemon makes initial startup of an APL session faster than it would be if all APs were started for every session. However, because the AP daemon must start the AP before the AP can reciprocate offers to share with your task, your first transaction with an AP may be slightly slower than any that follow. The AP daemon is fully described in the *Auxiliary Processors Manual, Chapter 1*.

5
System Variables and Functions

An APL program can interact with the environment provided by SHARP APL and its UNIX host. Most interactions are handled using a special class of variables and functions each of which is denoted by a *distinguished name*—one that starts with ⎕ (quad). Such names are called distinguished because they are set apart by the leading ⎕. System variables are generally spelled in the *first alphabet* and must begin with an alphabetic character. Distinguished names that are spelled in the *second alphabet* (for example ⎕*PR* or ⎕*MF*) represent experimental facilities. They may be provided in different forms (or not at all) in future releases of SHARP APL for UNIX.

The symbol ⍞ is another distinguished name, but there are no system names combining it with other letters. There is a fixed list of distinguished names, described later in this chapter; you are not free to coin your own.

System Variables

System variables are automatically shared with the system: you can both point to and set each of their values (so can the interpreter). This produces behavior unlike that of other variables:

The system can reject your attempt to set a value that is outside the acceptable type, rank, shape, or range of values for that system variable. Although the interpreter does not verify that the statement you store in ⎕*lx* (latent expression) is in fact executable, it does not permit you to make ⎕*lx* a numeric array of any rank, or a character array of rank greater than 1.

The system can set the value of a system variable. You can set a value and subsequently refer to it only to find that the value differs from what you set (although only in minor ways). Although the interpreter accepts both a character vector and a character scalar when you set □*lx*, it is always a vector (even if a 1-element vector) when you refer to it.

Any system variable (except □ or ⍞) can be localized in a user-defined function. Table 5.1 provides a complete list of the available system variables.

Table 5.1. System variables.

Variable	*Meaning*	*Variable*	*Meaning*
□	Evaluated input/output	□*pp*	Print precision
⍞	Character input/output	□*PR*	Prompt replacement
□*ct*	Comparison tolerance	□*ps*	Position and spacing
□*ec*	Environment condition	□*pw*	Printing width
□*er*	Event report	□*rl*	Random link
□*fc*	Format control	□*sc*	State change
□*ht*	Horizontal tab	□*sp*	Session parameter
□*io*	Index origin	□*trap*	Trap expressions
□*lx*	Latent expression		

System Functions

System functions are primitive functions without symbols in the sense that they are built-in. For example, □*cr* returns the canonical representation of a user-defined function, and □*svo* offers to share a name with another process. Table 5.2 provides a complete list of the available system functions and their ranks.

Results of System Functions

All system functions return results. Those system functions that return no result under OS/390 return an empty matrix under UNIX. This allows all system functions to be modified by an monadic operator or dyadic operator without requiring definitions of monadic operators or dyadic operators to take into account the lack of a result. Since an empty matrix is displayed in zero lines, those functions that formerly returned no result continue to evoke no display.

Table 5.2. System functions and their argument ranks.

System Function	Monad Rank	Dyad Ranks	System Function	Monad Rank	Dyad Ranks	System Function	Monad Rank	Dyad Ranks
□*ai*			□*load*	*		□*resize*		0 1
□*append*		∞ 1	□*MF*	1	0 1	□*run*		
□*appendr*		∞ 1	□*NA*	*	* *	□*runs*		
□*arbin*	1		□*names*			□*signal*	*	* *
□*arbout*	1		□*nc*†	1		□*size*	1	
□*av*			□*nl*	1	1 1	□*stac*		2 1
□*avail*	*		□*nums*			□*stie*		1 1
□*avm*			□*pack*	*	* *	□*stop*	1	1 1
□*bind*	1	1 1	□*PARSE*	*		□*svc*	1	1 1
□*bounce*†	0		□*paths*			□*svn*	1	* *
□*cr*	1		□*pdef*	*	* *	□*svo*†	1	¯1 1
□*create*		1 1	□*pex*		* *	□*svq*	1	
□*dl*	0		□*pins*		* *	□*svr*†	1	
□*drop*	1		□*plock*	*	* *	□*svs*	1	
□*EDIT*			□*pnames*	*		□*SYSVER*		
□*erase*		1 1	□*pnc*	*	* *	□*tie*		1 1
□*ex*†	1		□*ppdef*	*	* *	□*trace*	1	1 1
□*fd*		∞ ∞	□*psel*		* *	□*ts*		
□*fhold*	2	1 2	□*pval*		* *	□*twsid*	*	
□*fi*	1		□*qload*	*		□*ul*		
□*fmt*		1 2	□*rdac*	1		□*untie*	0	
□*fsm*			□*rdci*	1		□*USERID*		
□*fx*	2		□*rdfi*	1		□*vi*	1	
□*hold*	1		□*read*	1		□*wa*		
□*lc*			□*rename*		1 1	□*ws*		∞ ∞
□*lib*	1		□*replace*		∞ 1			

† *Rank-1 result* ∞ *Infinite rank* * *Undefined rank*

Argument Ranks of System Functions

System functions are listed in this chapter with information regarding default argument rank: if one value appears, a function is monadic; if two values appear, a function is dyadic; if three values appear, a function is ambivalent; if there is no rank, the function is niladic. Rank information appears in the order monadic, dyadic left, and dyadic right. The symbols listed in the rank information (other than numbers) are the same as those used for primitive functions: ∞ for infinite rank and ⋆ for undefined rank. (This is further explained in the *Language Guide*, *Chapter 4*.)

For those functions for which a default argument rank has been defined, you can use the rank operator ⍤ to supply arguments of any rank:

function ⍤ *rank*

You can modify any system function ***function*** to impose the argument rank specified by ***rank***. When you do not use the rank dyadic operator explicitly, the interpreter treats the function as if it were used with the default rank.

The three system functions ⎕*bounce*, ⎕*ex*, and ⎕*nc* produce a rank-`1` result even when invoked with rank `0` imposed:

```
      ⍴⎕nc⍤0 'a'
1
```

Alphabetical List of System Variables and Functions

This section provides a an alphabetically arranged summary of distinguished system names (functions and variables). Each entry presents a definition, points to a related definition, or points to a discussion in an other chapter. Descriptions of system functions include default argument ranks (explained in the previous section "Argument Ranks of System Functions").

⎕ Evaluated Input/Output, ⍞ Character Input/Output

The symbols ⎕ (quad) and ⍞ (quote-quad) denote system variables shared with the console screen and keyboard. The action of these two system variables depends on whether they are used or set.

A variable is *set* when it occurs to the left of the ← (assignment arrow):

```
x←foo □←analyze x
```

It is *used* when it occurs anywhere else:

```
x←foo analyze □
```

Setting a value into □ or ⍞ causes the value assigned to □ or ⍞ to be displayed.

***Note*:** If any characters are present before `apl-8` encounters end-of-file (`EOF`), those characters are processed normally, and the `EOF` is ignored. No characters before the `EOF` (i.e., a naked `EOF`) causes the signalling of event `1004`, input interrupt.

z←□ Evaluated Input

In a statement to be executed, when the symbol □ occurs anywhere except to the left of the assignment arrow, it represents the value of an expression to be obtained from the keyboard.

When in the course of evaluating a statement, the interpreter encounters the □ symbol used in this way, it displays the prompt `□:` and awaits the entry of one line from the keyboard. It evaluates what you type in response to the `□:` in the same way that it would evaluate a line typed during immediate execution. That done, it substitutes the resulting value for the □ symbol and resumes execution of the statement that contained □.

Example:

```
      Years←20
     (1+□)*Years×12
□:
     .08÷12
4.926802771
```

The presence of □ in the expression `(1+□)` causes the interpreter to issue the prompt `□:`. In reply, you type `.08÷12`. The interpreter evaluates `.08÷12` and substitutes the result `0.006666666667` for □ so that it can evaluate `1+□`. That done, it raises `1.006666666667` to the power `20×12` and prints the result `4.926802771`.

In response to the □: prompt, you can enter from the keyboard any executable expression, invoking any primitive function or any user-defined function present in the workspace. Whenever evaluation of that statement is complete, the interpreter passes its value back to the expression that contained the □. The result (for this example) could be anything in the domain of +.

When, in response to the □: prompt, you type a line containing several statements separated by diamonds (see the *Language Guide*, *Chapter 6*), the interpreter evaluates all of them, but the value substituted for □ is the value returned by the last of those statements.

When, in response to the □: prompt, you do something that produces no statement to evaluate (for example, you simply press ENTER without writing a statement), the interpreter repeats the □: prompt for evaluated input. It is acceptable to type a system command during □-input. The session manager carries out the command but, because no system command can satisfy a request for an input statement, the interpreter reissues the □: prompt. You cannot exploit this to copy into the workspace something that was not initially present because)*copy* does not work under □.

When, in response to the □: prompt, you enter the naked right arrow →, the interpreter abandons execution of □ and the statement that contained it, as well as any pending functions back to the previous immediate execution entry from the keyboard. (See "Programming and Immediate Execution" in "Chapter 7. System Facilities").

z←⍞ Character Input

In a statement to be executed, when the symbol ⍞ occurs anywhere except to the left of the left arrow, it represents a character vector to be obtained from the keyboard. The interpreter displays no special prompt (except as produced by ⍞-output, described below). The interpreter simply awaits the entry of one line from the keyboard. It treats that line literally as characters (even if what you then type looks like a statement or a system command). It substitutes the resulting characters for the ⍞ symbol and resumes execution of the statement containing ⍞.

Example:

```
      'Please enter characters.'⋄  (⍕⍴c←⍞),' characters
received.'
Please enter characters.
2 4⍴'A B C D'
13 characters received.
```

Escape from ⍞-input

While the interpreter is waiting for one line of characters from the keyboard, anything you type is accepted as characters and not as an instruction to be executed. However, you can still signal an input interrupt by transmitting CTRL-G at any point in the line before you have pressed ENTER.

CTRL-G sent during input mode signals an *input interrupt*, which normally causes the interpreter to abandon execution of the statement that contains ⍞ and halt execution of the program that contains it. However, if ⎕*trap* has been set to respond differently to this event, the interpreter can do something else, making it impossible to escape (see "Chapter 9. Event Handling").

Explicit Output from ⎕ and ⍞

When the system variables ⎕ or ⍞ occur to the left of the assignment arrow ←, the value to the right is displayed. The displays produced are identical to those produced by default output or by the monadic use of ⍕, except that output produced by ⍞ is not followed by a newline character, whereas the session manager automatically supplies a newline character following default output or output elicited by ⎕.

The effect of this difference between ⍞ and ⎕ is that a series of statements produced by ⎕ appear on successive lines, whereas a series of statements produced by ⍞-output may accumulate on the same line. This allows a line of output to be built up from the results of several expressions. For example, a program might contain the following sequence:

```
[ ]    year←1982
[ ]    miles←56.1
[ ]    todaymiles←12
[ ]    ⍞←year
[ ]    ⍞←''''s record was '
[ ]    ⍞←miles×1.6
[ ]    ⎕←' km.'
[ ]    ⍞←'Today''s distance: '
[ ]    ⍞←todaymiles×1.6
[ ]    ⎕←' km.'
[ ]    ...
```

When that fragment of the program is executed, the separate pieces of ⍞-output fit together thus:

```
1982's record was 89.76 km.
Today's distance: 19.2 km.
```

When the value to the right of the assignment arrow is an array of rank 2 or higher, ⍞-output always displays that value at the start of a new line:

```
      ⍞←'abc' ⋄ ⍞← 2 2⍴1 2 3 4
abc
1 2
3 4
```

But when the value to the right of the assignment arrow has a rank less than 2, ⍞-output displays the value immediately following the previous display:

```
      ⍞←2 2 ⍴1 2 3 4 ⋄ ⍞←'abc'
1 2
3 4abc
```

Input on the Same Line as the Character Prompt

You may want to accept input on the same line as a prompt message supplied by the program. You can do that by placing ⍞-output immediately before a request for ⍞-input.

Example:

```
      name←⍞ ⍝⍞←'Name: '
Name: John
```

The variable *name* receives both the reply typed from the keyboard (*John*) and as much of the prompt as appeared on the same line as the reply:

```
      name
Name: John
```

The example could also be written with the prompt and reply on separate lines. However, in writing a program, it is prudent to keep the prompt and the reply on the same line so that they remain together even if the program is halted and restarted.

When a ***weak interrupt***, CTRL-C, is signalled in ⍞-input, any characters entered are discarded and the prompt is reissued. Where the prompt takes more than one line, only the part of the prompt that occurs on the same line as the input is represented in the result. When the session manager has finished displaying the prompt and is ready to receive your reply, the cursor is at the next position to the right, just beyond the end of the prompt, as follows:

```
Name: _
```

If you simply type the four letters *John* without first moving the cursor somewhere else, the result contains ten characters. The first six are the prompt and the last four are the letters you typed: *John*. However, if you move the cursor leftward into the prompt zone and there write some other characters (overtyping the characters used in the prompt), the characters you substitute are reproduced in the result.

While in the prompt zone, you may elect APL overstrike mode to form a new character by merging one of the characters emitted in the prompt with one that you supply from the keyboard (for example, you can overstrike / with − to form ≠). See ⎕*PR*, prompt replacement, for further information on controlling characters in the prompt zone.

Discarding the Prompt

When you have no interest in entries that replace or overstrike the prompt, you simply discard them from the result. Expressions such as the following are common:

```
prompt←'Supply name: '
name←(⍴prompt)↓⍞ ⊣⍞←prompt
```

Alternatively, you can exploit a quirk of ⎕*arbout* to prevent the cursor from moving back into the prompt area and at the same time eliminate the prompt from the result, as follows:

```
prompt←'Supply name: '
name←⍞ ⊣⎕arbout '' ⊣⍞←prompt
```

Interpolating any use of ⎕*arbout* between the input and output uses of ⍞ discards the output from the ⍞-input buffer and prevents the prompt zone from being treated as part of the input line.

⎕*ai* Accounting Information

The niladic system function ⎕*ai* returns a result containing current accounting information, which in this version of SHARP APL for UNIX is a 4-element integer vector containing the following information:

`0@ ⎕ai`	your account number (the number assigned by the UNIX administrator, which also serves as your APL library number)
`1@ ⎕ai`	cumulative central processing unit (CPU) time used by this APL session, in milliseconds
`2@ ⎕ai`	elapsed time since the start of the APL session, in milliseconds
`3@ ⎕ai`	cumulative system time (CPU time logged by your activities) since the start of this UNIX session, in milliseconds.

⎕*ai* relies on the UNIX operating system's clock for time measurement, which typically limits its resolution to 1/60 or 1/100 of a second. The definition of ⎕*ai* is different from that used in SHARP APL for OS/390, both in length and in the interpretation of elements other than the first two. User-defined functions should not depend on the result of ⎕*ai* having length 4, since in future releases more elements may be included.

⎕*append* Append Variable to Tied File, ⎕*appendr* Append Variable to Tied File with Result

Rank: ∞ 1

The dyadic system functions ⎕*append* and ⎕*appendr* are used to place an array in a new component which is appended at the end of a tied file. These functions are fully documented in the *File System Manual*, *Chapter 2*.

⎕*arbin* Arbitrary Input, ⎕*arbout* Arbitrary Output

Rank: `1`

The monadic system functions ⎕*arbin* and ⎕*arbout* cause transmission to the device from which your APL session is being managed (the console screen, the screen of a terminal, or the controlling shared variable of an S-task). They take as arguments either a numeric vector whose elements are all integers between `0` and `255`, inclusive, or a character vector. The transmission consists of a sequence of bytes identified by the characters themselves, or characters in ⎕*av* (atomic vector) indexed by the numbers in the right argument. The effect of transmitting those characters depends on the physical device to which they are sent.

This output is not subject to default folding imposed by ⎕*pw* (printing width). Indeed, ⎕*pw* ignores use of ⎕*arbin* or ⎕*arbout*, and continues to fold any ⎕ output which may surround the use of ⎕*arbin* or ⎕*arbout* as though arbitrary I/O had not taken place, possibly leading to misplaced folds. SHARP APL does not supply a newline character following arbitrary output. When the arbitrary output produces a series of characters on the screen, the cursor is left at the position following the last character.

The function ⎕*arbout* produces a transmission to the output device. The function ⎕*arbin* also causes a transmission to the output device in exactly the same way as ⎕*arbout*, and then waits for the next transmission from the input device. It returns that transmission as its result, encoded as an integer vector in which each byte is identified by a number in the range `0` to `255`.

For each task connected to a screen and keyboard, the UNIX system maintains an input buffer in which transmissions from the keyboard are recorded and passed to the APL session manager. The result of ⎕*arbin* is the transmission accumulated in that buffer since the time it was last emptied up to and including the first occurrence of the carriage return character (ASCII `13`) or the line feed character (ASCII `10`). Unless your device is operating in block mode (so that it resolves cursor movements or overstrikes before transmitting them), the buffer (and thus the result of ⎕*arbin*) contains a record of each keystroke. If you type ahead, it is possible for the result of ⎕*arbin* to reflect keystrokes you made before the prompt was displayed on your screen.

Note: If any characters are present before `apl-8` encounters end-of-file (`EOF`), those characters are processed normally, and the `EOF` is ignored. If there are no characters before the `EOF` (a naked `EOF`) ⎕*arbin* returns an empty vector.

The transmission produced by ⎕*arbout* or by the ⎕*arbin* prompt bypasses output formatting imposed by SHARP APL and to some extent that imposed by UNIX device drivers. It is quite possible to transmit codes that disrupt the display, or leave the keyboard in a state from which you can no longer control the APL session. Since the prompt can solicit a response directly from the device without interaction from the keyboard, the reply reported in the result of ⎕*arbin* may contain characters the device supplied without your involvement.

⎕*av* Atomic Vector

The niladic system variable ⎕*av* returns a vector of all possible character values in SHARP APL, in numeric order of their internal encoding, from hex `00` to hex *`FF`* (decimal `0` to `255`). The first `128` characters are an exact mapping of the ASCII characters; the last `128` are all the APL characters, including composite characters.

⎕*av* is commonly used as a reference for translating characters to their numeric equivalents:

```
⎕av⍳⍵
```

It can also be used to translate from one character code to another, by an expression such as the following:

```
(⎕av⍳old)@ NewAlf
```

To avoid loss of information, *`NewAlf`* should be a permutation of ⎕*av* : a list in which each of the `256` possible characters occurs exactly once.

⎕*avm* Mainframe Atomic Vector

The niladic system function ⎕*avm* returns a `256`-element character vector that is a permutation of ⎕*av*. Each element of the result, ⎕*avm*`[`*i*`]`, is defined as follows:

- If ⎕*av*`[`*i*`]` is an APL primitive function or operator under ***OS/390***, then ⎕*avm*`[`*i*`]` is the corresponding primitive function or operator.
- If ⎕*av*`[`*i*`]` is a first alphabetic character under ***OS/390***, then ⎕*avm*`[`*i*`]` is the corresponding lowercase alphabetic character.
- If ⎕*av*`[`*i*`]` is a second alphabetic character under ***OS/390***, then ⎕*avm*`[`*i*`]` is the corresponding uppercase alphabetic character.

- If ⎕*av*[*i*] is a third alphabetic character under ***OS/390***, ⎕*avm*[*i*] is ⎕*av*[129+*i*−166]. That is, ⎕*avm* "maps" the mainframe third alphabet into positions 129 through 154 (0-origin) of the UNIX ⎕*av*. All other elements of mainframe ⎕*av* are mapped to "convenient" locations in ⎕*avm*. The precise mapping is readily available via the expression ⎕*av* ⍳ ⎕*avm*.

⎕*avail* File System Availability

Rank: ⍟

The monadic system function ⎕*avail* reports whether the APL component file system is running. This system function is documented in the *File System Manual*, *Chapter 2*.

⎕*bind* Intrinsic Function Bind

Rank: 1 1 1

The ambivalent system function ⎕*bind* is used to create an IF bind or to query the association level. This functionality is also provided by the system intrinsic function *system.bind*.

Right Argument: The right argument is a character vector or scalar containing the name of a single APL user identifier (reserved names are not permitted). The dyadic use of *system.bindrep* requires that the name be locally undefined, and creates an IF bind from the bind display. The monadic use is a query only.

Left Argument: The left argument, if provided, is a character vector or matrix containing a valid IF bind display (see below). Note that only line 0 (*group.function*) of the display form is required by *system.bindrep*; that is, the vector '*system.bindrep*' is a valid left argument.

Result and Side Effects: The result of both monadic and dyadic ⎕*bind* is a single integer reporting the association state of the named symbol. The association state describes the usability of an IF bind, as described below:

0 Monadic use only. The symbol is not an IF bind.

`1` The symbol is a valid IF bind, but it refers to an IF that is not present in this system.

`2` The symbol is a valid IF bind to a known and available IF.

Association state `1` is provided so that a system can construct and process IF binds for or from other systems with a different set of available IF routines. An attempt to use an IF bind whose association state is not `2` will result in an *`association error`* being signalled (event number `10`).

For more information on intrinsic functions and IF binds refer to the *Intrinsic Functions Manual*, *Chapter 3*.

⎕*bounce* Terminate Task

Rank: `0`

The monadic system function ⎕*`bounce`* terminates (bounces) any UNIX process(es) identified by task number in its integer argument. A task number is assigned whenever a task is started. The banner displayed at the start of a T-task begins with the task number of that APL session. That task number is also available as `10@ 2 ⎕ws 3`.

The result contains a `1` for each existing task that is bounced and a `0` elsewhere; `1` means that the task exists but is not necessarily terminated. The UNIX system closes any files that the terminated tasks had opened, but does not otherwise save their work. However, an APL task with ⎕*`twsid`* set does save the active workspace.

***Note*:** It is possible to terminate your own task if you include its number in the argument of ⎕*`bounce`*, or if you use the value `0` as its argument. This is because APL task IDs and UNIX process IDs are equivalent: executing ⎕*`bounce`* `n` is the same as using the UNIX `kill` command for process ID `n`. It is also possible to ⎕*`bounce`* any UNIX task that could be terminated from UNIX using the `kill` command.

⎕cr Canonical Representation

Rank: 1

The monadic system function ⎕cr returns the *canonical representation* of the most local definition of the user-defined function named in its argument. The argument can be a character vector or a character scalar for a one-character name.

The result is a character matrix containing the definition of the user-defined function in canonical form; that is, without leading or trailing ∇ symbols or bracketed line numbers. Each line (including the header) appears flush left. Lines shorter than the longest are right-padded with blanks.

Example:

```
      ⎕cr 'plus'
z←a plus ω
z←a+ω
      ⍴⎕cr 'plus'
2 10
```

When the argument contains a well-formed name but the visible referent of the name is not that of an unlocked user-defined function, the result is an empty matrix.

⎕create Create a File

Rank: 1 1

The dyadic system function ⎕create is used to create (and leave share-tied) a new APL component file. This system function is documented in the *File System Manual*, Chapter 2.

⎕ct Comparison Tolerance

The system variable ⎕ct is the maximum relative difference allowed between two numbers if they are to be considered equal. It is used to overcome problems of inexact internal representation and cumulative rounding errors inherent in computer arithmetic on noninteger values. Comparison tolerance permits the interpreter to ignore small differences between numbers likely to arise from either of these two problems.

Inexact internal representations of numbers entered from the keyboard arise because input and display are normally in base `10`, but internal representations are in base `2` (binary). For example, because the representation of `0.1` in base `2` is a nonterminating sequence of binary digits, `0.1` cannot be represented exactly, no matter how many bits are devoted to its binary representation.

A rounding error can occur when the result of a calculation cannot be represented exactly, or cannot be represented exactly within the available precision (which uses `52` bits for the mantissa, `1` for the sign, and `11` for the exponent).

In a clear workspace, the default value for `⎕ct` is `5.684341886e¯14`, or `2*¯44`. You can assign the visible `⎕ct` to be any numeric scalar not less than `0` nor greater than `10*-10`. See Table 5.3 before setting `⎕ct` to `0`. This table is based on three variables:

```
a←0 0 1 1 ⊣ b←(0+eps),(0-eps),(1+eps),1-eps ⊣ eps←1e¯14
```

Table 5.3. Effects of comparison tolerance.

`eps←1e¯14`	*a* *b*	0 0+*eps*	0 0−*eps*	1 1+*eps*	1 1−*eps*
`⎕ct←0`	`⌊b`	0	¯1	1	0
`⎕ct←10×eps`	`⌊b`	0	0	1	1
`⎕ct←0`	`⌈b`	1	0	2	1
`⎕ct←10×eps`	`⌈b`	0	0	1	1
`⎕ct←0`	`a=b`	0	0	0	0
`⎕ct←10×eps`	`a=b`	0	0	1	1
`⎕ct←0`	`a<b`	1	0	1	0
`⎕ct←10×eps`	`a<b`	1	0	0	0
`⎕ct←0`	`a⍳b`	5	5	5	5
`⎕ct←10×eps`	`a⍳b`	5	5	3	3
`⎕ct←0`	`a∈b`	0	0	0	0
`⎕ct←10×eps`	`a∈b`	0	0	1	1

Primitives Dependent on Comparison Tolerance

The value of ⎕*ct* is used when computing the result of any of the following primitive functions with floating-point arguments:

Dyadic use | ∧ ∨ ∈ ⍳ < ≤ = ≡ ≠ ≥ > ⍷

Monadic use ⌊ ⌈

The value in ⎕*ct* is a relative comparison tolerance. When two numbers are compared and at least one of them is a floating-point number, their relative difference (which depends on their magnitudes) is compared against ⎕*ct*. Two numbers are considered equal if their relative difference is less than or equal to ⎕*ct*. That is, ⍺ and ⍵ are judged equal when

```
(|⍺-⍵) ≤ ⎕ct × ⍺⌈⍤¨|⍵
```

When one of the arguments (say, ⍵) is `0` but the other is not, that inequality reduces to

```
(|⍺) ≤ ⎕ct × |⍺
```

Since ⎕*ct* must be less than `1`, the inequality can never be satisfied. Thus, the value of ⎕*ct* can never affect a comparison with `0`.

When ⎕*ct* is `0`, all comparisons are exact (and the judgment not equal may be produced by bits otherwise considered insignificant).

The following illustrate the effects of ⎕*ct* in a few simple examples:

```
      ⎕ct
1e¯14
      ⊢a←(○11)÷○1
11
      a=11
1
      ⌊a
11

      ⎕ct←0
      a=11
0
```

```
      ⌊⍺
10
     ⍺-11
¯1.776356839e¯15
```

⎕dl Delay Execution

Rank: 0

The monadic system function ⎕dl delays execution of the statement in which it appears. The argument to ⎕dl is a numeric scalar, specifying the duration of the pause in seconds. The result of ⎕dl is a numeric scalar whose value is the actual delay in seconds. When the delay completes normally, the result is not less than the argument.

The delay can be aborted by a *weak interrupt*, CTRL-C. The interruption cuts short the delay, but causes the function to return its result in a normal manner. The result is the actual delay. The line containing ⎕dl completes, and execution halts before starting a new line. Thus, an interrupt issued during execution of a line containing multiple instances of ⎕dl cancels the delays for all ⎕dls that have not yet been processed. However, results are nonetheless returned for each ⎕dl.

⎕drop Drop Components from a Tied File

Rank: 1

The monadic system function ⎕drop lets you drop components from the head or the tail of a tied file. This system function is documented in the *File System Manual*, *Chapter 2*.

⎕ec Environment Condition

The system variable ⎕ec, which is part of the event trapping mechanism, is a Boolean scalar. When set to 1, it causes an event (error, interrupt, or return to immediate execution) in a user-defined function to force an exit rather than a stop. It has the special property that, in a function which localizes it but has not set its value, the interpreter sets its value to 1. This system variable is documented in "Chapter 9. Event Handling".

⎕*EDIT* System Editor

Rank: ∞

The monadic system function ⎕*EDIT* invokes the full-screen editor for the object named in the right argument. This experimental system function may not be available in future releases of SHARP APL for UNIX.

The right argument to ⎕*EDIT* is a character scalar, vector, or 1-row matrix containing a single, valid name. The result of ⎕*EDIT* is an empty matrix. The full-screen editor is fully documented in "Chapter 7. System Facilities".

⎕*er* Event Report

The system variable ⎕*er* is set by the interpreter when an error, interrupt, or other event occurs. This system variable is documented in "Chapter 9. Event Handling".

⎕*erase* Erase a Tied File

Rank: `1 1`

The dyadic system function ⎕*erase* lets you remove a tied file (to which you have proper access) from its file library. This system function is documented in the *File System Manual, Chapter 2*.

⎕*ex* Expunge Objects

Rank: `1`

The monadic system function ⎕*ex* erases, if possible, the most local referent of each of the objects named in its argument. The argument is a name list containing zero or more names. The names can be arranged in a character table with one name per row, or in a character vector with successive names separated by blanks.

The result of ⎕*ex* is a Boolean vector with an element for each name in the argument. Even when used with imposed rank-`0`, the result is a `1`-element vector.

The result is:

`1` when the corresponding name is now free (either because it had no referent or because its former referent has been erased)

`0` when the corresponding name is not free (either because it has a referent that cannot be erased, or because it is not a well-formed name).

□*ex* does not erase a name whose referent is a *label* or a *system function*. A function that is executing, pendent, suspended, or waiting can be erased, and its name thereby freed; but the interpreter's working copy of its definition does not disappear from the workspace until it has been cleared from the state indicator.

If □*ex* produces a *ws full* or *domain error*, nothing has been erased.

□*fc* Format Control

The system variable □*fc* is in effect an additional argument to ⍕, serving to specify how it deals with a number of special cases. It is a character vector of `6` elements. For more information on ⍕, refer to the *Language Guide*, *Chapter 4*.

□*fd* Function Definition

Rank: ∞ ∞

The dyadic system function □*fd* provides services for conversion between a function and the character representation of its definition. Historically, these precede and overlap the system functions □*cr*, □*fx*, and □*ex*. The left argument, a numeric element from the set `1 2 3 6 7`, selects one of the services described below.

`1 ⎕fd ω` **Vector representation**. The right argument ω is a character vector containing a name whose visible referent is a user-defined function. The result is a character vector (including newline characters) that shows the definition in the form displayed by the ∇ editor (with line numbers enclosed in square brackets and a leading and closing ∇). Applied to an argument that is not the name of a function, the function `1 ⎕fd ω` signals *domain error*.

```
    1 ⎕fd 'plus'
    ∇ z←α plus ω
[1]   z←α+ω
    ∇
```

`2 ⎕fd ω` **Matrix representation**. The right argument ω is a character vector containing a name whose visible referent is a user-defined function. The result is a character matrix containing the canonical representation of the function's definition, with one row for each line of the definition. The lines are shown in the same form as produced by the system function ⎕cr: flush left, without line numbers or their enclosing brackets, and without the opening and closing ∇. Applied to an argument that is not the name of a function, the function `2 ⎕fd ω` signals *domain error*.

```
    2 ⎕fd 'plus'
z←α plus ω
z←α+ω
```

`3 ⎕fd ω` **Fix a function from its character representation**. The right argument ω is the character representation of the definition of a user-defined function. The character representation can be either the matrix form (returned by ⎕cr or `2 ⎕fd`), or the vector form with successive lines separated by a newline character, with or without line numbers, and with or without the opening and closing ∇. The result, as for the system function ⎕fx, is the name (as a character vector) of the function just fixed.

```
    3 ⎕fd 1 ⎕fd 'plus'
plus
```

As a side effect, use of `3 ⎕fd ω` causes the function defined in *ω* to be fixed (materialized) in the workspace, in the same manner as `⎕fx`. A definition can be fixed provided the name shown in *ω* is available. A name is available if it has no visible use or its visible referent is a function (in which case the new definition replaces the old). When for any reason the definition cannot be carried out, the result is a 2-element integer vector. The first element is a diagnostic code; the second element is the number of the line in error for matrix form *ω*, or it is the 0-origin index of the character in error for vector form *ω*. The diagnostic codes are listed in Table 5.4.

Table 5.4. Diagnostic codes for `3 ⎕fd ω`***.***

Code	*Meaning*
1	Workspace full
2	Definition error
3	Character error
4	Symbol table full
5	Unused
6	Unused
7	Unused

`6 ⎕fd ω` Expunge names. The right argument *ω* is a character name list (either a vector in which successive names are separated by blanks, or a table with one name per row). Each of the names in *ω* is freed: its visible referent is expunged so that the name is free for other uses. The result is a table containing those names from *ω* that were not freed (because their referents could not be expunged). If a name in the argument is not well-formed, the system function `6 ⎕fd ω` signals a *domain error*. See also `⎕ex` in this chapter.

`7 ⎕fd ω` Lock functions. The right argument ω is a name list (either a character vector of names in which successive names are separated by blanks, or a character table with one name per row). Each of the functions named in ω is locked. The result is a table containing names that were not locked (because their visible referents were not functions).

⎕fhold Hold Tied Files
⎕hold Hold Tied Files

Rank: `2 1 2`

The system functions ⎕fhold and ⎕hold provide a way for two or more tasks running at the same time to coordinate the use of files they share. ⎕fhold is the newer and more general form. ⎕hold is unable to accept a left argument, and is unable to include passnumbers or component ranges in its right argument. These system functions are documented in the *File System Manual*, *Chapter 2*.

⎕fi Input Format Conversion

Rank: `1`

The monadic system function ⎕fi converts images of numbers within a character vector to their numeric form. The function takes as its argument a character vector which is treated as a series of fields delimited by blanks.

⎕fi returns a numeric vector containing the value of each well-formed numeric field within the vector, and `0` for each other field. (The system function ⎕vi reports whether each field is well-formed.) It recognizes as valid those numeric forms appearing in the result of monadic ⍕, including integers, decimal fractions, *e*-formats and numbers preceded by ¯ (negative sign). It does not recognize as well-formed any other fields; in particular, it does not consider well-formed a field that contains a currency symbol or embedded commas, or one that is enclosed in parentheses.

Example:

```
     ⎕fi '666 ¯1.20      .1 314159e¯5'
666 ¯1.2 0.1 3.14159
```

```
      ⎕fi 'Take 3 spoons of 150 proof rum'
0 3 0 0 150 0 0

      ⎕fi 'Sent {$}3.00 for 6 kilos on 87/4/1'
0 0 0 6 0 0 0
```

Techniques, such as the following, may help ⎕fi convert some common but otherwise unacceptable formats:

```
      ext←' 0123456789e.,-¯'
      int←' 0123456789e. ¯¯*'
      ⎕fi ((-⎕io)+ext⍳' .25 -6.25 8,9,10')@ int
0.25 ¯6.25 8 9 10
```

Use of ⎕fi with ⎕vi

Since the function ⎕vi returns a Boolean result indicating which fields in its character vector argument are well-formed, the expression

```
(⎕vi ⍵)/⎕fi ⍵
```

returns the values of the well-formed numbers from a character vector ⍵.

⎕fmt Format Output

Rank: 1 2

The dyadic system function ⎕fmt formats character and numeric data into a character array for display. Its features largely overlap those provided by dyadic use of ⍕ (thorn), which was introduced after ⎕fmt. The conventions for the left argument of ⎕fmt resemble some of those used in Fortran or Cobol.

⎕fmt Format Phrases

The left argument of ⎕fmt is a character vector divided into *format phrases* separated by commas. Each format phrase governs the appearance of a column of data from the right argument. For example, the following contains three phrases:

```
      f←'i6,f8.2,g<999.99>'
```

The number of phrases need not match the number of columns in the right argument; the phrases are repeated cyclically as needed.

Each phrase has a *type,* indicated by a letter (such as *i*, *f*, or *g* in the preceding example). In addition, a *constant text* is reproduced in the output, but does not correspond to any column of the data to be formatted. In the following example, constant phrases are shown enclosed by the characters <>. Other pairs of text-delimiting characters are listed in the section "⎕fmt Summary of Format Phrases" on page 5-28 .

```
f←'<Part: >,i6,< Cost: >,f8.2'
```

A tab phrase does not itself control a field of data, but controls the position in the result where the next field starts. The position of the next field is specified either by the distance forward or back from the end of the previous field, or as an absolute location in the result. A summary of all the various phrase types appears later in this section.

⎕fmt Repetition Factors

A phrase or group of phrases can be preceded by a repetition factor:

```
g←'2i6,3f8.2,g< 999.99>'
```

This indicates that the first phrase is to be used for two columns, the next for three, the last for one. The repetition factor can apply to a set of consecutive phrases enclosed in parentheses:

```
g←'2(i6,3f8.2),g< 999.99>'
```

This indicates that the sequence of one column formatted by `i6` and three columns formatted by `f8.2` is to be repeated twice.

⎕fmt Type and Shape of Right Argument

The right argument, containing the data to be formatted, can be one of the following:

- A numeric array or character array of rank no greater than `2`. When ⍵ is a simple vector or scalar, it is treated as if it were the `1`-column matrix ⍪⍵.

- A vector of one or more boxes containing such arrays. When ⍵ is formed as a succession of boxes, each is disclosed and treated as a table, and their successive columns are formatted in order from left to right. Columns need not have the same height; short columns in the result are padded with blank rows at the bottom to match the length of the tallest. The boxed form of the right argument takes the place of a format allowed in SHARP APL for OS/390 that used ; (semicolon) to separate segments of the right argument. See "Semicolon Format Not Supported" below.
- A mixed array. Only in the restricted case where the right argument is a vector can it be a mixed array.

⎕*fmt* Semicolon Format Not Supported

Earlier versions of SHARP APL permitted a special syntax for ⎕*fmt* in which the right argument was formed by placing required parentheses around two or more expressions separated by *semicolons*. This form is no longer supported. The effect formerly obtained by

```
f ⎕fmt (a;b;c)
```

is now obtained by

```
f ⎕fmt a⊃b⊃c
```

⎕*fmt* Shape of Result

The result returned by ⎕*fmt* is always a character matrix. The number of rows is sufficient to accommodate the tallest column in ⍵.

In the result, successive bands of adjacent print positions form *fields* corresponding to the phrases in the left argument ⍺. Each field has the width specified by the corresponding phrase in ⍺. There is no additional space between fields, so the visual separation of fields in the result is produced either by constant phrases in ⍺, or by specifying formats that are wider than needed to represent the data being formatted. If one or both arguments of ⎕*fmt* are empty, the result is an empty character matrix.

```
      ⍴'' ⎕fmt ⍳0
0 0
      ⍴'' ⎕fmt ⍬⍬⍬
0 0
      ⍴'' ⎕fmt ''⊃''
```

```
0 0
      ⍴'⍳5' ⎕fmt ⍳0
0 5
```

⎕fmt Qualifiers and Decorators

Within each format phrase, various qualifiers or decorators can be inserted to the left of the letter that identifies its type.

- A *qualifier* provides additional rules, for example ``insert commas between triplets of digits'' or ``leave this field entirely blank when the value being represented is zero.''
- A *decorator* is a piece of text attached to the representation of a number, usually to indicate its sign. When used, decoration text appears to the right of the right most digit used in the representation, or to the left of the left most. For example, to enclose negative numbers in parentheses, you need to include in the format phrase the decorators *m* (the left text of a negative number) and *n* (the right text of a negative number):

```
      'm<(>n<)>f8.2' ⎕fmt ¯123.45
(123.45)
```

To preserve right-alignment, when you include decorators that appear to the right, you need to specify equal-length decorators for negative and nonnegative numbers. In the following example, *q* places text to the right of a non-negative number:

```
      'm<(>n<)>q< > f8.2' ⎕fmt 123.45 ¯123.45
 123.45
(123.45)
```

Alternatively, you can exploit decorators of different size to give different alignment to negative and nonnegative values:

```
      'm<>q<       >f16.2' ⎕fmt 123.45×1 ¯1 1 ¯1
123.45
        123.45
123.45
        123.45
```

The *background decorator* `r` first fills the field by repeating as necessary the character(s) specified, and then overwrites the field with the digits of the representation and the sign decorators. The effect is to leave the background text visible only in the positions unused by the representation. This is commonly used for check protection:

```
      '⍵r<$*********>12.2' ⎕⍵fmt⍵ 12345.67 5.67
$***12345.67
$*******5.67
```

⎕⍵fmt⍵ Summary of Format Phrases

In the following definitions, the five symbols *n*, *p*, *w*, *d*, and *q* appear with the principal format phrase types `a`, `e`, `f`, `i`, `x`, `t`, and `g`. The meanings of the symbols are as follows:

n number of repetitions (optional)

p number of positions to displace the next field

w field width

d number of digits after decimal point (`f` format) or number of significant digits (`e` format)

q qualifier or decorator (see below)

Text-delimiting characters

```
< >
≤ ≥
⊂ ⊃
•• ••
/ /
⌈ ⌈
⌊ ⌊
```

***Note*:** Each of these pairs of characters is used to delimit text within a format phrase.

Phrase Types

n a *w*	**Character data.** Print each character in a field *w* positions wide. (Note that when *w*>1, blanks are inserted between adjacent letters).
n q e *w.d*	**Exponential fraction.** Print in a field *w* positions wide, with *d* significant digits.
n q f *w.d*	**Fixed-point fraction.** Print in a field *w* positions wide with *d* decimal digits. Trailing zeros always print.
n q i *w*	**Integer.** Leading zeros are represented by blanks unless the z qualifier is used.
n x *p*	**Relative Position.** Start the next field at the position displaced *p* positions from the end of the preceding field. The displacement may be negative, in which case the next field may overwrite earlier ones.
n t *p*	**Absolute Reposition.** Start the next field at position *p* (possibly over-writing fields).
n<*text*>	**Constant.** Insert *text* in every row of the result.
n q g <*text*>	**Picture.** Insert digits that represent the value (rounded to the nearest integer) into the ``picture'' text. Copy *text* to the field. Overlay successive digits of the number's representation. Place the digits at the positions marked in the picture by 9 or z, starting with the low-order digit in the right-most 9 or z. In positions marked z represent leading or trailing zeros by blanks. Where other text characters are embedded between digits used in the result, keep them. Overlay right decorators to the right of the rightmost digit, and left decorators to the left of the leftmost digit.

Qualifiers (optional)

b	Make the entire field blank when the value is zero.
c	Insert commas between successive triplets of digits. Cannot be used with g format.
kn	Scale the result by displaying a value 10^n times the value in ω.

`l` Left-justify the representation within its field. Cannot be used with `g` format.

`z` Print leading zeros or trailing zeros. Note that in `g` format, `z` within the picture makes leading or trailing zeros blank.

Decorators (optional)

`m<text>` Insert ***text*** to the *left* of the representation of a *negative* value.

`n<text>` Insert ***text*** to the *right* of the representation of a *negative* value.

`p<text>` Insert ***text*** to the *left* of the representation of a *nonnegative* value.

`q<text>` Insert ***text*** to the *right* of the representation of a *nonnegative* value.

`r<text>` **Background text.** First fill the field by cyclically repeating ***text*** across it. Then overwrite the digits and decorators needed to represent the value. Leave the background text visible at positions not used for digits, sign, or decorators.

`s<text>` **Symbol substitution.** In each pair of characters within ***text***, the first indicates a character used by default and the second the desired substitute. For example, `s<, .,>` displays blank where `⎕fmt` would put comma, and comma where `⎕fmt` would put dot.

The only symbols that may be substituted are:

`*` field overflow

`0` leading zero fill character

`,` character inserted with `c`-qualifier

`.` decimal point

`e` character to mark exponent

`¯` character to indicate negative exponent; does not apply to the sign of the number being represented, which is governed only by the `m` decorator.

⎕fmt Examples

```
      n
4163649355 0 ¯4163649355 3649355 ¯14163649355

     'bcm<->k-2f14.2' ⎕fmt n
41,636,493.55

-41,636,493.55
 36,493.55
**************

     'm<(>n<)>q< >bck-2f17.2' ⎕fmt n
   41,636,493.55

  (41,636,493.55)
       36,493.55
 (141,636,493.55)

      'q<            >m<>i26' ⎕fmt n
   4163649355
            0
               4163649355
      3649355
             14163649355

     '<Phone >,g<(zzz) zzz-zzzz>' ⎕fmt n
Phone (416) 364-9355
Phone (
Phone ¯416) 364-9355
Phone (     364-9355
Phone **************
```

A similar result can be obtained from dyadic ⍠ with character left argument (format by example). The dyadic function ⍠ is discussed in the *Language Guide, Chapter 4*.

```
     'Phone (342)_111-1111' ⍠¯1⊣n
Phone (416) 364-9355
Phone     )    -
Phone (416) 364-9355
Phone    () 364-9355
Phone **************
```

⎕fsm Full-Screen Manager

The niladic system function ⎕fsm determines if AP124 is available and provides device information when it is. If AP124 is not available, ⎕fsm results in ⍳0, a numeric empty vector; otherwise, it returns a 68-element integer vector (defined in Table 5.5). This behavior is similar to the behavior of ⎕arbin 24⍴0 in IDSH, the session manager for SHARP APL under OS/390.

Table 5.5. Result of ⎕fsm.

Position	*Description*
1–27	Zero (reserved)
28	Rows of AP124 screen
29	Columns of AP124 screen
30–31	Zero (reserved)
32	Extended capability
33–34	Zero (reserved)
35	Color(=2) + highlighting(=1)
36–52	Zero (reserved)
53–54	Maximum number of format rows (256 base)
55–56	Zero (reserved)
57–60	Size of AP124 shared variable buffer (256 base)
61–68	Zero (reserved)

⎕fx Fix a Definition

Rank: 2

The monadic function ⎕fx *fixes* (that is, makes available for execution) a user-defined function from the characters of its *canonical representation*. The argument to ⎕fx is a character matrix containing the function's definition. The form of the character representation is similar to that found in the result of ⎕cr. Each row of the matrix represents a line of the definition. In the canonical form, the rows do not start with bracketed line numbers.

The principal effect of ⎕*fx* is to create in the active workspace the function whose name appears in the first line of the argument to ⎕*fx*. The spacing of the characters in the argument is preserved in the definition thus produced; this means that extra blanks, while not semantically significant, are preserved in the definition as a feature of SHARP APL for UNIX. This was not true of the traditional behavior of ⎕*fx*.

Producing a new function is a *side effect*, and not the explicit result of ⎕*fx*. When the definition is successfully fixed, the result is a character vector containing the name of the function that has been defined. When the definition cannot be fixed, the result is a numeric scalar containing the line number where the first fault was found. The result may indicate a row that does not exist in the argument if the argument is empty or if the error is detected after the last row of the argument.

When the name of the newly defined function corresponds to a name that has been localized in a function that is currently executing, pendent, or suspended, the newly defined function is local, and ceases to exist when the function to which its name is local completes execution.

When the name of the function thus fixed is the same as the name of an existing function, the existing function is replaced by the new one; all internal code for the former definition is discarded and any ⎕*stop*, ⎕*trace*, or ⎕*MF* settings are removed. The system functions ⎕*fx* and `3` ⎕*fd* provide similar capabilities.

⎕*hold* Hold Tied Files

Rank: `1`

See the description of ⎕*fhold*.

⎕*ht* Horizontal Tab

In those APL systems where the system itself manages transmission between remote terminal and APL, the system variable ⎕*ht* is used to describe tab stop positions on the terminal so the system can interpret the *tab* character in input or make use of it in output. However, these actions are handled by UNIX and the value of ⎕*ht* is ignored. In SHARP APL it can be assigned any value.

⎕io Index Origin

The system variable ⎕io establishes the *index origin* used by several primitive functions. The domain for assignment to ⎕io is a numeric scalar or 1-element vector whose value is either 0 or 1. In a clear workspace, the default value for ⎕io is 1.

When generating or using index values, the interpreter assumes that indices are numbered starting at ⎕io. Compare the different effects produced when ⎕io is 1 and ⎕io is 0:

```
      ⊢⎕io←1
1
      ⍳5
1 2 3 4 5
      ⊢x←5+⍳5
6 7 8 9 10
      x[3]
8
      x[5]
10



      x[0]
index error
      x[0]
      ^
      3?3
3 1 2
      v←6 23 11 4 ¯6
      ⍋v
5 4 1 3 2
      x,[0.5] v
6  7  8  9 10
6 23 11  4 ¯6
```

```
      ⊢⎕io←0
0
      ⍳5
0 1 2 3 4
      ⊢x←5+⍳5
5 6 7 8 9
      x[3]
8
      x[5]
index error
      x[5]
      ^
      x[0]
5



      3?3
2 0 1
      v←6 23 11 4 ¯6
      ⍋v
4 3 0 2 1
      x,[0.5] v
5  6
6 23
7 11
8  4
9 ¯6
```

Functions Dependent on Index Origin

The value of ⎕io is used by both monadic and dyadic ⍳, ?, ⍒, and ⍋: in interpreting the left argument of dyadic ⍉; in selection by indexing using square brackets; and, when the axis qualifier is used, in selecting the axis of application of rotate (⌽), rotate-down (⊖), copy (⌿), expand (⍀), and catenate (,).

⎕lc Line Counter

The niladic system function ⎕lc is an integer vector with the current value of the execution *line counter,* which contains one number for each row in the *state indicator* beginning with the most recently invoked function.

Like the state indicator itself, ⎕lc includes an entry for each pending use of ⍎ or ⎕-input. The value reported for ⍎ or ⎕ is the number of the line whose statement invoked ⍎ or ⎕ (usually, the value which appears next in ⎕lc).

When a user-defined function has been interrupted, the expression →⎕lc is commonly typed from the keyboard to mean, ``Resume execution with the next line in sequence.'' This is a shorthand for →0⍴⎕lc, since → considers only the first of the values in its argument.

```
     )si
tri[2]*
⍎
example[3]

     ⎕lc
2 3 3
```

⎕lib File Library

Rank: 1

The monadic system function ⎕lib returns a matrix containing the names of files in the indicated library or directory. This system function is documented in the *File System Manual, Chapter 2*.

⎕load Load a Workspace, ⎕qload Quiet Load a Workspace

Rank: ⋆

The monadic system functions ⎕load and ⎕qload have the same effect as the system command)load (see "Chapter 6. System Commands"). Each replaces the present contents of the active workspace with the contents of a saved workspace.

The argument is a character vector or scalar identifying the workspace to be loaded (optionally including the pathname or library number to which the workspace belongs). The workspace can be in any directory. When no path or library is specified, the home directory is assumed.

It is a moot question whether ⎕load returns a result: when the interpreter has executed ⎕load, the active workspace in which ⎕load was executed no longer exists.

The two forms ⎕load and ⎕qload differ only in whether they display the message *saved* followed by the date and time where the workspace was saved. The function ⎕load causes the display, while ⎕qload (for ``quiet load'') does not.

The functions ⎕load and ⎕qload attempt to adjust the size of the active workspace to the size of the saved workspace if the sizes are different. If resources are not available to accommodate the new size, ⎕load fails and reports *ws too large* (event 70). The command)load may be used to bring the contents of a saved workspace into a workspace of a different size.

Following either ⎕load or ⎕qload, the interpreter automatically executes the line it finds stored as the visible value of ⎕lx in the workspace it has just loaded. Applications that make use of ⎕load require some mechanism to pass control from the workspace that invoked ⎕load to the workspace loaded in consequence. Usually the invoking workspace depends upon the ⎕lx of the new workspace to initiate appropriate action. Information can be passed between the two workspaces by way of ⎕sp (session parameter), or by a convention that the new workspace looks in a particular file for information previously stored there by the preceding workspace.

⎕*lx* Latent Expression.

The system variable ⎕*lx* can be set before saving a workspace to control the way processing starts each time it is loaded. At each load, the interpreter automatically executes the line it finds stored as the visible value of ⎕*lx*. The system variable ⎕*lx* thus provides a mechanism for making an immediate and automatic start on such things as conditioning the workspace environment, starting or restarting an application, or validating a user's access to an application. When you set ⎕*lx*, it must be a character scalar or vector; when you refer to it, it is always a character vector. In a clear workspace, the default value for ⎕*lx* is `''`.

Whenever you load a workspace (*except* when you load it by using the *)xload* command), the interpreter executes the expression ⍎⎕*lx*. When a statement stored in ⎕*lx* is invalid, the interpreter reports an error and suspends execution just as it would for a statement entered in immediate execution.

The system variable ⎕*lx* can be used to deny workspace access to an unauthorized user; for example:

```
⎕lx←'⎕trap←(~(0@ ⎕ai) ∈ approved )/'' ∇2001 d clear'''
```

Or, it might be used to start processing in an application workspace; for example:

```
      ⎕lx←'autostart'
      )save
ws saved 1988-05-18 11:00:45
      1 ⎕fd 'autostart'
    ∇ z←autostart;t;p
[1]   CodeFile ⎕stie t←999          ⍝ tie code file
[2]   p←⎕read t,1                   ⍝ obtain bootstrap package
[3]   z←run p                       ⍝ start up the application
    ∇
```

Lines input from the keyboard are not stored in the recall buffer by state indicator level (as they are in SHARP APL for OS/390). The expression `'⍎⎕lx'` appears if you recall the last input line immediately after loading a workspace. See the system command) (recall) in "Chapter 6. System Commands"

⎕*MF* Monitoring Facility

Rank: `1 0 1`

The ambivalent system function ⎕*MF* lets you monitor resources consumed during execution of each line of a user-defined function. This experimental system function may not be available in future releases of SHARP APL for UNIX.

The resource units that ⎕*MF* returns are unmodified operating system clock ticks. The granularity of the ticks must be considered carefully when estimating absolute usage, as in any timing calculations. ⎕*MF* generates statistics indicating total resource units for the line, resource units for the line excluding calls to user-defined functions, and visits to the line.

Like the system functions ⎕*stop* and ⎕*trace*, you

- set monitoring on or off when you use it dyadically
- report current values when you use it monadically.

You also turn monitoring off if you use the full-screen editor to modify the definition, or if you use ⎕*fx* or ⎕*fd* to refix the definition of a function for which monitoring is turned on.

The right argument *ω* is a name list of user objects; if any entry in this vector is not a properly formed name, ⎕*MF* signals a *domain error*. The name list may consist of a character scalar or vector containing a single name, a character vector in which multiple names are separated by blanks, or a character matrix with one name per row.

Dyadic ⎕*MF*

For dyadic use, the left argument of ⎕*MF* must be an scalar or vector of Boolean values in which a `1` means turn on and a `0` means turn off monitoring for the corresponding function in *ω*. When the left argument is a vector, it must contain the same number of elements as the right argument contains names.

The result of dyadic ⎕*MF* is a Boolean vector in which a `1` indicates success and a `0` failure in modifying the monitoring status of the corresponding name. A failure occurs if a name in *ω* is not that of an unlocked user-defined function that is neither pending nor suspended.

For example, if both *host* and *ls* are the names of two unlocked user-defined functions, then the following expression turns monitoring on:

```
     1 ⎕MF 'host ls'
1 1
```

Each dyadic use of ⎕*MF* to turn monitoring on resets all values being stored for resource usage.

Monadic ⎕MF

When you use ⎕*MF* monadically, the right argument must be a single name. The result is an integer matrix with as many rows as there are lines in the function named in ⍵, and three columns. The values in this matrix contain the total resources used by the function named in ⍵ since monitoring was turned on. The information generated is provided in three columns:

First column: Total resource units consumed for the line.

Second column: Resource units for the line excluding calls to user-defined functions (which you can monitor as well if you wish).

Third column: Number of times the line is visited.

The statistics reported for the function header (the first line of the matrix) show resource totals for all lines in the function being monitored. The result contains all zeros if the function has not been executed since monitoring was turned on.

⎕*NA* Name Association

Rank: ⋆ ⋆ ⋆

The ambivalent system function ⎕*NA* is used dyadically for the definition of hybrid user-defined functions, and used monadically for the examination of the name of a user-defined function to determine its hybrid status. This experimental system function may not be available in future releases of SHARP APL for UNIX.

From the perspective of the active workspace, a hybrid function is an APL user-defined function: it takes array arguments and returns array results. The definition of the function is not in APL, however, nor is the definition stored in the active workspace. The definition exists in a separate associated address space that the name association mechanism handles for you automatically. In SHARP APL for UNIX, the definition of the hybrid function is compiled object code. APL applications may thereby take advantage of compiled programs.

Subsequent versions of ⎕*NA* may relax the restriction of the separate address space and allow compiled definitions to access APL objects directly in the active workspace. Even then, however, the development of compiled routines in a

separate address space will remain a powerful modelling tool to which the added speed of execution gained from combining the APL and non-APL definitions in the same address space can be viewed as a final production step.

An application developer, having selected a set of object programs that an APL application requires, builds a private associated processor that contains the compiled definitions. Hybrid functions are defined to the workspace using the system function ⎕*NA*.

This modular approach provides security at two levels. Execution problems encountered in the compiled programs in their own address space do not affect the contents of the active workspace. The communications channel between the APL workspace and the associated processor is managed by the shared variable processor; such a channel is secure, even when the associated processor is located on a remote UNIX host.

Defining Hybrid Functions

Dyadic use of ⎕*NA* is similar to dyadic use of ⎕*svo* (refer to the *SVP Manual, Chapter 3*). The left argument is the processor ID of the associated processor that contains the compiled definitions to which you are going to link. The right argument (for both dyadic and monadic use) is the name by which you refer to a particular compiled routine. This name, as with ⎕*svo*, can be a single name or two names separated by a blank. When two names are used, a surrogate name is used in the same fashion as for ⎕*svo*.

The current version of ⎕*NA* is restricted to a rank-1 right argument. Future versions will support a list of names. The result for both monadic and dyadic use is an integer indicating the name class of the hybrid object named in the right argument. Possible name classes for the returned integer values are listed in Table 5.6.

Table 5.6. ⎕NA classification codes.

Value	*Class*
¯3	⎕*NA* function, no link available
0	not a ⎕*NA* object
3	⎕*NA* function

For example, if your associated processor has processor ID `99`, and the hybrid function you wish to define is to be named *format*, then the expression

```
      99 ⎕NA 'format'
3
```

defines a user-defined function *format* with the syntax established for that function when the associated processor was compiled. Like other user-defined functions, a hybrid function is saved when you save your active workspace. When you next load the workspace, your hybrid functions are still defined. They are active as long as the associated processor containing their definitions is active. If this associated processor is not running, then a monadic use of ⎕*NA* to inquire into the status of the name of the hybrid function returns the value `¯3`. When the associated processor is reactivated, you can again make use of the hybrid function.

A hybrid function cannot be displayed using ⎕*fx*, ⎕*fd*, or the ∇ editor.

⎕*names* Library Names of Tied Files

The niladic system function ⎕*names* returns a character matrix each row of which contains the library name (in APL library style) of a currently tied file. This system function is documented in the *File System Manual*, Chapter 2.

⎕*nc* Name Class of Identifiers

Rank: `1`

The monadic system function ⎕*nc* describes the name class to which the visible use of each name in its argument refers. The argument is a name list, containing zero or more names. Multiple names must be arranged in a matrix with one name per row; a single name can be presented as a character vector or character scalar.

The result of ⎕*nc* is an integer vector of classification codes, one for each name in the argument. Even when used with imposed rank-`0`, the result is a `1`-element vector. Table 5.7 lists the values that may be returned.

Table 5.7. ⎕nc classification codes.

Value	*Class*
0	Not in use
1	Label
2	Variable
3	Function
4	Other

A value of 4 indicates that the identifier is a *distinguished name* (begins with ⎕), or that the argument is *not a well-formed name*. 4 ⎕ws can be used to distinguish between these cases (for system variables).

⎕nl Name List

Rank: 1 1 1

The ambivalent system function ⎕nl takes as its right argument an integer scalar or vector indicating classes of use (see ⎕nc), and returns an alphabetized character matrix of names whose visible (most local) use is to refer to a member of those classes. The result matrix is sorted in ⎕av order.

The classes that appear in the right argument are identified by the numbers 1 through 3, with the same meanings as in the result of ⎕nc. For example, ⎕nl 1 2 returns a matrix of names whose visible use refers to a label or a variable.

When ⎕nl is used dyadically, the left argument ⍺ is a character scalar or vector. It restricts the result to names whose first letter is a member of ⍺. For example,

```
'XYZ' ⎕nl 2
```

returns a matrix of the names of variables that begin with X, Y, or Z.

⎕nums Tie Numbers of Tied Files

The niladic system function ⎕nums returns the tie numbers of currently tied files. This system function is documented in the *File System Manual*, Chapter 2.

⎕*out* Record Session I/O to File

Monad ⎕*out* permits you to control the destination of the APL session's display. You can divert to a file what would otherwise appear at the terminal (of a T-task) or at the controlling shared variable (of an S-task). Or you can both capture the output in a file and display it in the usual way.

Use of ⎕*out* also permits you to capture normal output generated by a batch task, which would otherwise be lost.

ω is an integer vector of either zero or two elements. When *ω* is empty, ⎕*out* reports the current setting. When *ω* is a two-element vector, the value provided in *ω* takes effect, and the result is the former setting. When *ω* is not empty, its elements have significance as follows:

`0@ω` This value determines whether output is displayed, either to the controlling shared variable of an S-task or to the display device of a T-task. Possible values are:

`0`	Do not display output.
`1`	Display output.

`1@ω` This value determines whether a record of output is appended to a file. Possible values are:

`0`	Do not file the output data.
tn	Append the output data to the file tied to *tn*.

When `1@ω` is not `0`, the system appends the output data as components of the APL file tied to `1@ω`. (You must have tied the file, without a passnumber, beforehand.)

If SHARP APL for UNIX finds it is unable to redirect output as requested (e.g., because the file is not tied, or is tied but is full, or due to *ws full*), it executes ⎕*out* `1 0` on your behalf, and signals the event in the usual way (for example, as *file tie error*, *file full*, or *ws full*).

Timing of ⎕*out Result*

The change in output routing produced by ⎕*out* takes effect before the result of ⎕*out* is returned. Thus, if your program permits the result of ⎕*out* itself to be displayed, so that your program includes statements like this,

```
[ ]   ⎕out 0 45      ⍝ Start sending output to file 45
[ ]   Output
[ ]   ⎕out 1 0       ⍝ Resume sending output to terminal
```

you will find that the first component written to the file contains `1 0` (the result of your initial instruction ⎕*out* `0 45`, indicating what the condition was before you executed ⎕*out* `0 45`), and when your program executes ⎕*out* `1 0` to return output to the terminal, you will see its result `0 45` (the former setting) displayed at the terminal, but not in the file.

Control Messages

While copying output to a file, the system may also generate *control messages* to describe the circumstances of input or output. When you subsequently examine the contents of the file, the control messages have to be interpreted (rather than printed). The control messages generated by ⎕*out* are a subset of those recognized by the SHARP APL for OS/390 facility HSPRINT.

A control message is a `22`-character vector whose positions have significance as follows:

1	The character ⎕*av*`[`⎕*io*`+255]`.
2–5	Message number. Four-digit number, with leading zeros. The numbers refer to the list of standard control messages (see below).
6	Blank.
7–10	Component count. Four-digit number, with leading zeros. The number of following components to which this control message applies.
11	Blank.
12–22	Mnemonic. In output generated in response to ⎕*out*, the system writes here a standard mnemonic to help human readers interpret the control message.

The list that follows shows all the control messages recognized by HSPRINT. The public workspace `1 hsp` is supplied in SHARP APL for UNIX to interpret HSPRINT files. The control messages recognized by `1 hsp` are marked by a dagger (†). The control messages that may arise in a file recorded by `⎕out` are marked by an asterisk (*).

List of Control Messages

† `0000 000n format`	Arguments to `⎕fmt` follow. The following components are not the result of formatting but the arguments. Formatting is left to the program that interprets the `⎕out` file. Where the "number of components" field indicates more than two, those after the first arose from the paren/semicolon convention for the right argument of `⎕fmt`. (See ***Note*** at the end of this section.)
† `0001 000n mixed`	Mixed output follows. The following components are not the result of formatting but the arguments.
† `0002 0001 pagen`	New page number follows.
†* `0003 000n quadp`	`⍞-output` follows.
† `0004 0001 translate`	Translate table follows. The table maps APL characters to the output device, and is specific to the device employed.
† `0005 0001 digits`	New value for `⎕pp` follows. The new value affects the display of subsequent components containing arrays that were stored unformatted; for example, for mixed output.
† `0006 0001 width`	New value for `⎕pw` follows. The new value affects the display of subsequent components containing arrays that were stored unformatted; for example, for mixed output.
† `0007 0001 title`	Page title follows.
† `0008 0001 subtitle`	Page subtitle follows.

	0010 0001 *prtuctm*	Control for printing user control messages follows. When the following component contains 1, control messages are printed; when 0 (the default), they are not.
†	0011 0000 *page*	Skip to beginning of new page.
†	0012 0001 *carriage*	Vertical forms control follows. The control is a Boolean vector with as many elements as lines on the output page, and 1 for each line on which printing is permitted.
†	0013 0001 *prtarbout*	Control for printing □*arbout* data follows. When the following component contains 1, arbitrary output is printed; when 0 (the default), it is not.
†*	0014 000*n* *arbout*	□*arbout* argument follows.
*	0015 000*n* *arbin*	□*arbin p*rompt and result follows.
	0016 0001 *prtarbin*	Control for printing □*arbin* prompt and data follows. When the following component contains 1, the arbitrary prompt and resulting input are printed; when 0 (the default), they are not.
	0018 000*n* *hfile*	Identifies a file created by HSPRINT under its print-to-file option.
†*	0019 0001 *quad er*	□*er* follows.
†*	0020 000*n* *input*	Normal input follows.
*	0021 000*n* *quad*	□-input follows.
*	0022 000*n* *quad prime*	⍞-input follows.
	0023 000*n* *fn defn*	Function definition input follows.
*	0024 000*n* *ws full*	Workspace full while preparing input.
*	0025 000*n* *st full*	Symbol table full while preparing input.

*	0026 000*n defn err*	Definition error while preparing input.
	0027 000*n not in defn*	Obsolete: Formerly "Input not allowed in definition mode."
	0028 000*n open quote*	Obsolete: Formerly "Input contained unmatched quote."
†	0029 0001 *prtinput*	Control for printing input follows. When the following component contains 1, input is printed; when 0 (the default), it is not.
	0030 0001 *quadps*	New value for □*ps* follows.
	0031 0001 *char error*	Obsolete: Formerly "Character error while preparing input."

***Note*:** SHARP APL for OS/390 uses the format control message as described above. In SHARP APL for UNIX, this has been abandoned and the result of a □*fmt* statement is appended to file.

□*pack* Build a Package

Rank: ⋆ ⋆ ⋆

The system function □*pack* builds a package, which is a variable containing a set of names, a set of name classes, and the referent of each name. Each referent can be a user-defined function, a variable (itself possibly a package), a bound intrinsic function, or undefined. See the *Language Guide, Chapter 3* for more information on SHARP APL packages.

Used dyadically, □*pack* forms a one-member package whose value is that of the variable in the right argument, and whose name (inside the package) is that indicated by the character vector or scalar in the left argument.

Used monadically, the right argument is a name list (either a character vector of names in which successive names are separated by blanks, or a character matrix with one name per row). The result of monadic □*pack* is a package containing the names in ⍵ together with the visible referent of each name.

⎕*PARSE* Pre-parse a User-defined Function

Rank: ⋆

The monadic system function ⎕*PARSE* pre-parses (prepares and preserves the internal form of) all lines of a user-defined function with the guarantee that all lines have been visited. Such internal preparation can streamline the performance of a production application which, for example, always executes the same code when the application workspace is started. This experimental system function may not be available in future releases of SHARP APL for UNIX.

The SHARP APL interpreter goes through a two-step process when it first executes a line of APL statements. In the first step, a line is examined for information about the name class and valence of each word (or token) it contains, and its syntax is analyzed and translated to an internal form.

In the second step, the parsed form of the line is executed to produce the expected results. When the executing line is from the definition of a user-defined function, the parsed information is preserved with the function's definition so that the next and all subsequent executions of that line are processed faster than the first. If you save your workspace after you have executed a user-defined function, you save not only the lines you write, but also the pre-parsed internal form for those lines that have completed execution.

⎕*PARSE* Usage

The right argument of ⎕*PARSE* contains names of user-defined functions, either as a character vector with the names separated by spaces, or as a character matrix with one name per row. The result is a character matrix containing the names of functions (one per row) which could not be completely pre-parsed. This usually happens because they refer to objects that are not in the workspace (though syntax errors may be responsible).

Successful execution of ⎕*PARSE* on all functions named in its argument produces an empty character matrix. A future extension may provide additional information about the functions whose names appear in the result.

For example, to preparse all the functions in a workspace, you might try the following:

```
      ⎕PARSE ⎕nl 3
      )save
ws saved 2000-02-19 17.43.26
```

Functions prepared in this fashion must be saved in your workspace library in order to take advantage of the streamlining. Functions stored in packages in components of APL component files do not preserve the parse information.

⎕*paths* Pathnames of Tied Files

The niladic system function ⎕*paths* returns the names of tied APL files in UNIX directory-style form. The names appear as a vector of enclosed character vectors, the elements of which correspond to the tie numbers in ⎕*nums*. This system function is documented in the *File System Manual*, *Chapter 2*.

⎕*pdef* Package Define

Rank: ⋆ ⋆ ⋆

The system function ⎕*pdef* can be used monadically or dyadically. The right argument ⍵ must be a package. The optional left argument is a name list (a character vector of names separated by blanks, or a character matrix with one name per row). ⎕*pdef* defines objects from the package to become the visible referent of those names in the active workspace.

- When ⎕*pdef* is used dyadically, only the names specified in ⍺ are defined. If one or more names in ⍺ are not names of objects in the package ⍵, a *domain error* results.
- When ⎕*pdef* is used monadically, all names in the package are defined.

The new objects are defined as a side effect of using ⎕*pdef*. The explicit result is an empty character matrix.

The objects defined replace the visible (most local) referent of those names. When the name of an object to be defined corresponds to a name that has been localized in a function currently executing, pendent, or suspended, the newly defined object is local and ceases to exist once the function completes execution. See also ⎕*ppdef* in this chapter.

Settings for ⎕*stop*, ⎕*trace*, or ⎕*MF* in effect when a function is packaged are not maintained through ⎕*pdef*.

⎕pex Package Expunge

Rank: ⋆ ⋆

The result of α ⎕pex ω is a package containing a subset of the names in the package ω, together with their referents.

α is a name list (a character vector in which successive names are separated by blanks, or a matrix with one name per row) containing names to be excluded from the result. It is not necessary that the names in α actually exist in ω.

⎕pins Package Insert

Rank: ⋆ ⋆

Both arguments of the dyadic function ⎕pins are packages. The result is a package in which the names of the objects is the union of the names found in the two argument packages.

Where the same name occurs in both α and ω, the referent of that name in the result is the referent from ω. All names and their referents in ω are always represented in the result, whereas α contributes only those names and referents not contained in ω.

⎕plock Package Lock

Rank: ⋆ ⋆ ⋆

The ambivalent system function ⎕plock takes a package as its right argument. The optional left argument α is a name list (a character vector of names separated by blanks, or a matrix with one name per row).

The result is a package identical to ω except for the fact that the definitions of the functions named in α are locked. When ⎕plock is used monadically, the definitions of all functions in ω are locked. If one or more names in α do not represent functions in ω, a *domain error* results.

⎕*pnames* Package Names

Rank: ⋆

The monadic system function ⎕*pnames* reports the names contained in package ⍵. The result is a character matrix with one row for each name in ⍵.

When ⍵ is an empty package, the result is a 0-by-0 matrix. However, when ⍵ is not a package, the result is an empty vector. The expression ⍴⍴⎕*pnames* *n* thus provides a test of whether a variable *n* is or is not a package.

⎕*pnc* Package Name Class

Rank: ⋆ ⋆ ⋆

⎕*pnc* reports the name class of names in a package. The result is a numeric vector, one element per name, representing the class of each name by codes similar to those returned by ⎕*nc*. Table 5.8 lists the values that may be returned.

Table 5.8. ⎕pnc classification codes.

Value	*Class*
¯1	Undefined
0	Referent not in package
1	[unused]
¯2	IF Bind
2	Referent is variable
3	Referent is function
4	[unused]

Monad ⎕*pnc* returns the name class of every name in ⍵, in the same order as ⎕*pnames* ⍵.

For dyad ⎕*pnc*, ⍵ is a name list (a character vector in which successive names are separated by blanks, or a matrix with one name per row). The result contains the name class of each of the names in ⍺, in the order that they appear in ⍺.

If a package contains an object that has no definition, the name class returned by ⎕*pnc* is ¯1.

```
      )erase a
      ⎕pnames p←⎕pack 'a'
 a
      ⎕pnc p
¯1
```

⎕pp Print Precision

The system variable ⎕pp specifies the maximum number of significant digits, or *print precision,* provided by the system when it displays numeric floating-point data using either default output or monadic ⍕ (thorn).

The domain for assignment to ⎕pp is a scalar whose value is an integer between 1 and 17, inclusive. In a clear workspace, the default value for ⎕pp is 10.

Functions Dependent on Print Precision

The value of ⎕pp affects the result of monadic use of the function ⍕ (thorn) applied to a fractional value. Since system output displayed by default uses the same rules as monadic ⍕, it is also affected.

The system variable ⎕pp has no effect on the display of any value that can be represented internally as an integer, or on a number which is not different from an integer, even when represented internally in floating-point form

Example:

```
      ⎕pp←3
     ÷3
0.333

     1234.
1234

     1234.0000000000000001
1.23e3
```

When ⎕pp←17, the interpreter displays output with full internal precision. Every internal floating-point value is distinguishable from its nearest neighbours. The final digit cannot be otherwise significant.

⎕*ppdef* Protective Package Define

Rank: ⋆ ⋆ ⋆

This ambivalent system function takes a package as its right argument. The optional left argument ⍺ is a name list (a character vector of names separated by blanks, or a matrix with one name per row).

As a side effect, ⎕*ppdef* causes those objects contained in the package ⍵ (and named in ⍺, when ⎕*ppdef* is used dyadically) and whose names have no visible referent in the active workspace, to be defined in the active workspace. The result is a character matrix of names *not defined* because of visible use in the workspace.

If one or more names in ⍺ do not represent objects in ⍵, a *domain error* results.

⎕*PR* Prompt Replacement

The system variable ⎕*PR* governs how characters obtained from ⍞-input are affected by the prompt provided by the preceding use of ⍞-output. The usage described here applies only to T-tasks, and is not available for S-tasks. This experimental system function may not be available in future releases of SHARP APL for UNIX.

The domain for assignment to ⎕*PR* is a character scalar or an empty character vector. When referenced, ⎕*PR* is an empty character vector or a 1-element character vector.

The value of ⎕*PR* determines how an input prompt, if there is one, is merged with the result of ⍞-input. There are three useful cases:

⎕*PR* is an empty vector (default in a clear workspace) User is permitted to modify the prompt. The result of ⍞-input retains all of the input prompt that is on the same line as the input (as far back as, but not including, the last newline character). Where the user moves the cursor back into the prompt and substitutes characters, or forms new characters by overstriking those already displayed, the new characters thus formed merge with those remaining from the prompt.

⎕PR is a character scalar other than newline character (usually a blank)

User is permitted to modify the prompt, but unmodified positions are replaced by ⎕PR. It is as though, before the user starts to type, the characters displayed in the prompt had been replaced by the ⎕PR character. The result contains the ⎕PR character for all positions on the same line as the input, except where the user moves the cursor back and substitutes others.

⎕PR is a character scalar whose value is newline

The user can see the prompt but cannot modify it. The rule is the same as for any other nonempty ⎕PR. However, since all the characters at the prompt are treated as though they precede the newline prompt, none of them is included in the zone in which the user may write, and hence none is included in the result. The effect is the same as that produced by `⎕arbout ''`.

Any use of ⎕arbout after the ⍞-output but before the ⍞-input discards the entire prompt from the result returned by ⍞-input, regardless of the setting of ⎕PR.

⎕ps Position and Spacing in Display

The system variable ⎕ps is a 4-element integer vector that governs the display and formatting of nested arrays, or of numeric elements in a mixed array.

In the display of any array of rank 2 or greater, the alignment of rows and columns is preserved by padding to match the width of the widest element in its column, and/or padding to match the tallest element in its row.

- The first two elements of ⎕ps determine the *positioning* of an element within the horizontal or vertical band assigned to it when the element itself is smaller than the space reserved for it.
- The second two elements control the amount of additional space inserted vertically or horizontally to separate adjacent rows or columns. For nested arrays only, a negative value in one or both of the last two elements causes a boundary character to be drawn at the vertical or horizontal edges of each box.

The effects of the four positions are as follows:

`0@ ⎕ps` Position, first axis (vertical)

`¯1` At the *top* of the available space
`0` *Centred* in the available space
`1` At the *bottom* of the available space

`1@ ⎕ps` Position, second axis (horizontal)

`¯1` At the *left* of the available space
`0` *Centred* in the available space
`1` At the *right* of the available space

`2@ ⎕ps` **Vertical separation.** The magnitude of the value in this position is the number of additional rows interpolated vertically in the result between the representation of successive rows of the argument.

Vertical box boundaries. When `2@ ⎕ps` is negative and has a magnitude of at least `2`, the *top and bottom edges* of a box are marked by the characters `¯` (at the top of the box) and `_` (at the bottom). Where necessary, these box-edge characters overwrite the interpolated blank rows (which is why there must be at least `2` of them). The result contains extra rows at the top or bottom to accommodate box boundaries, which would otherwise project beyond the edges of the result.

`3@ ⎕ps` **Horizontal separation.** The magnitude of the value in this position is the number of additional columns interpolated horizontally in the result between the representation of successive columns of the argument.

Horizontal box boundaries. When `3@ ⎕ps` is negative and has a magnitude of at least `2`, the *left and right edges* of a box are marked by the character `|` (at each side). Where necessary, these box-edge characters overwrite the interpolated blank columns (which is why there must be at least `2` of them). The result contains extra columns at the left or right edges to accommodate box boundaries, which would otherwise project beyond the edges of the result.

In a clear workspace, the default value of `⎕ps` is `¯1 ¯1 0 1`, which conforms to standard APL display: the position of an array is top left with no spaces between rows and one space between columns.

⎕psel Package Select

Rank: ⋆ ⋆

The right argument of the dyadic system function ⎕psel is a package, and the left argument is a name list (a character vector of names separated by blanks, or a character matrix with one name per row). The result is a package containing the names from ⍺ and their referents from ⍵.

If one or more names in ⍺ do not represent functions in the package ⍵, a *domain error* results.

⎕pval Package Value

Rank: ⋆ ⋆

The right argument of the dyadic system function ⎕pval is a package. The left argument is a character vector containing a single name. (A character scalar is acceptable for a one-character name.) The referent in ⍵ of the name ⍺ must be a variable. The result is the value of the variable in ⍵ named in ⍺.

⎕pw Printing Width

The system variable ⎕pw sets the maximum length of a line displayed by the session manager during default output, or output produced by ⎕ or ⍞.

The value of ⎕pw can be an integer scalar between `30` and `255`, inclusive, or `0`. Attempts to set ⎕pw to other values are rejected by the system with *domain error*. In a clear workspace, the default value of ⎕pw is `80`.

Turning Scissoring Off

The default display of wide arrays is scissored (see "Scissoring Wide Displays" in *Language Guide, Chapter 3*). To turn scissoring off in situations for which a traditional line-by-line display is appropriate, ⎕pw can be set to `0`. This facility is provided to support display of data from external sources, most commonly when using the 8-bit SHARP APL interpreter `apl-8` as a UNIX filter. See the *Handbook*, *Chapter 3* for more information on SHARP APL interpreters.

Session Variable

The behavior of □*pw* in SHARP APL for UNIX differs from that in SHARP APL for OS/390. The UNIX version treats □*pw* as a session variable: its value is preserved across)*load* or)*clear*.

□*qload* Quiet Load a Workspace

Rank: *

See the description of □*load*.

□*rdac* Read Access Matrix of a Tied File

Rank: 1

The monadic system function □*rdac* returns the access matrix of a tied file. This system function is documented in the *File System Manual*, Chapter 2.

□*rdci* Read Component Information

Rank: 1

The monadic system function □*rdci* returns information regarding a particular component of a tied APL file. This system function is documented in the *File System Manual*, Chapter 2.

□*rdfi* Read File Information

Rank: 1

The monadic system function □*rdfi* returns information regarding a tied file. This system function is documented in the *File System Manual*, Chapter 2.

⎕*read* Read Component of a Tied File

Rank: `1`

The monadic system function ⎕*read* lets you read the array stored in a component of a tied file. This system function is documented in the *File System Manual*, Chapter 2.

⎕*rename* Rename a Tied File

Rank: `1 1`

The dyadic system function ⎕*rename* lets you change the name of a tied file. This system function is documented in the *File System Manual*, Chapter 2.

⎕*replace* Replace Variable Stored in a Component

Rank: ∞ 1

The dyadic system function ⎕*replace* lets you replace the contents of a component of a tied file. This system function is documented in the *File System Manual*, Chapter 2.

⎕*resize* Reset File Size Limit

Rank: `0 1`

The dyadic system function ⎕*resize* lets you change the file size limit for a tied file. This system function is documented in the *File System Manual*, Chapter 2.

⎕*rl* Random Link

The value of the system variable ⎕*rl* is the seed value (or random link) used by the pseudorandom number generator. It can be assigned a value in the range described below. If you reset ⎕*rl* to the same value each time, you can reproduce test results. To "randomize" results, set ⎕*rl* to an arbitrary value such as the time of day or the connect time.

The domain for assignment to ⎕*rl* is any integer from `1` to `2147483646` (that is, 2^{31-2}), inclusive. In a clear workspace, the default value is `16807` (which is 7^5).

The value of ⎕*rl* is used in computing the result of the functions *roll* (monadic *?*) and *deal* (dyadic *?*). As each pseudorandom number is generated, the seed (⎕*rl*) used in the computation is changed.

⎕*run* Start N-task

The monadic system function ⎕*run* starts an N-task. The result is a `2`-element integer vector containing the number of the task started (when the request is successful) or a return code indicating the nature of the trouble (when it is unsuccessful).

The right argument is a character vector consisting of five fields separated by blanks. An N-task might be launched by an expression such as the following:

```
⎕run ': 501 Analysis NCrash 40 600'
```

The right argument in the above example is marked by underlines to show the division into fields. The five fields are discussed below.

Owner (Field 1)

Identifies the host machine that the new task will run on and the user number that will own the new task. Ordinarily, these will be the same as the host machine and user that executed ⎕*run*. In that case, a colon suffices (as illustrated). When the host machine is to be different, this field begins with the name of the remote host machine followed by an exclamation point. If the owner's account number is to be different, the name of the remote host and exclamation point, if present, are followed by the account number, a colon, and user's current sign-on password (but not the right parenthesis that accompanies sign-on to a T-task or S-task).

In any case, the field must contain a colon, but the presence of an exclamation point is optional. In fact, the presence of an exclamation point is only permitted if preceded by a host machine name.

Currently in SHARP APL for UNIX, if a remote host name is specified or if the account number specified differs from the running account, then the Initial Workspace (see below) may not specify the Split Workspace method. If this restriction is violated, result code `18 0` will be returned. Specification of a remote host name, even if the name is that of the host machine that executed □*run*, causes the resulting N-task to be defined as a remote N-task, subject to all the restrictions and limitations that pertain to such N-tasks.

Initial workspace (Field 2)

Determines by which of two methods the N-task is started.

Autostart: The name of a workspace, in APL style or UNIX directory style. The N-task starts by loading that workspace. (It is up to the workspace's latent expression to take things from there.) Note that there may be a blank between the library number and the name of the workspace (so that blank is permissible inside this field).
If an APL-style workspace name is specified without a library number, the library number is assumed to be the user number of the owner of the N-task.

Split Workspace: A colon. The N-task starts with a clone of the present active workspace. (It is up to the program now being executed to distinguish the parent workspace from the new N-task.) Currently in SHARP APL for UNIX, the Split Workspace case may not be specified if the Owner field specified a remote host or an account number that differs from the running account. This holds true even if the remote host specified is the same is the host machine that executed □*run*. If this restriction is violated, result code `18 0` will be returned.

Save-name (Field 3)

Name to be assigned to the N-task's active workspace when it is saved at termination of the N-task, in APL style or UNIX directory style. This workspace is presumed to be in a library of the N-task's owner. This field must not be elided.

MaxCPU (Field 4)

The maximum number of CPU units for this task; `0` means no limit. When SHARP APL notices that the task has exceeded its CPU limit, it terminates the task and saves its active workspace under the name indicated.

***Note*:** CPU limits are not honored; however, they are reported in the result of □*runs*.

Max elapsed (Field 5) Maximum elapsed seconds for the task; `0` means no limit. When SHARP APL notices that the task has exceeded its elapsed time limit, it terminates the task and saves its active workspace under the name indicated.

***Note*:** Elapsed time limits are not honored; however, they are reported in the result of □*runs*.

Currently in SHARP APL for UNIX, an N-task can only be started with an owner other than the running account if NSVP is active on the local host machine. Furthermore, an N-task can only be started on a remote host machine if NSVP is active on both the local and remote host machines. An attempt is made to ensure that the execution environment of an N-task matches that of the parent task as closely as possible. There are certain limitations, however. In all cases, all command line parameters of the running APL task will be passed on to the child N-task, except for `-P`, `-Q`, `-R`, `-s`, `-W`.

A non-remote N-task, whose owner is the same as the running account, will inherit all of the UNIX environment variables of the running task. A non-remote N-task, whose owner is different from the running account, will inherit only those of its parent's UNIX environment variables that are explicitly passed to it by □*run* via NSVP. Currently, the list of such variables includes `HOME`, `SAXAP`, `SAXAPDSM`, `SAXAPTAB`, `SAXCNF`, `SAXDIR`, `SAXLOG`, `STARTAPS`, `SYS`, and `TEMPPATH`.

A remote N-task does not inherit any of its parent's UNIX environment variables. Instead, it inherits all the UNIX environment variables of the NSVP Daemon (`NSVPD`) on the remote host; except, `HOME` is set to the home directory of the owner of the N-task. A remote N-task also fails to inherit the value of □*sp* from the parent workspace. It receives instead the default value of □*sp* in a clear workspace. The *Initial workspace* and *Save-name* fields for a remote N-task are interpreted on the remote host machine, not the local host machine.

Finally, the result of □*run* of a remote N-task, if successful, returns the result code `0 0`, which is otherwise an indication of a successfully started split workspace. Currently, a remote N-task cannot be a split workspace, so no ambiguity exists.

Table 5.9 provides an interpretation of the `2`-element integer vector returned by □*run*.

Table 5.9. Result codes for □*run*.

Code	*Meaning*
0 0	New task. This is the value □*run* returns for an N-task started by cloning the parent task's active workspace. It is also the value □*run* returns to the parent of a successfully started remote N-task
0 *n*	*n* has been successfully started. This is the value □*run* returns to the parent task.
1 *n*	*ω* contains an invalid character at position *n*.
2 0	The SHARP APL system's capacity for N-tasks is exhausted.
3 0	The proposed N-task would exceed the owner's task quota.
4 0	The owner has already saved the full quota of saved workspaces.
5 0	Attempt to load the workspace identified by *Initial workspace* produces *ws full*.
6 0	Attempt to load the workspace identified by *Initial workspace* produces *symbol table full*.
10 0	The proposed sign-on number or password is not recognized.
11 0	The proposed sign-on number has been locked out.
12 0	The proposed *Initial workspace* was not found.
13 0	The *Initial workspace* did not have an executable □*lx*.
14 0	The system's capacity for tasks is exhausted.
15 0	The user is not authorized to use □*run*.
16 0	Hardware error.
17 0	Attempted to □*run* onto a constrained account number.
18 0	□*run* nonce error. The specified combination of options is not currently permitted in SHARP APL for UNIX.
19 0	□*run* interrupt.
20 0	*continue* workspace unwritable.
58 0	□*sp* too large.

□runs Report Running Tasks

The niladic system function □runs returns an integer matrix that shows APL tasks running under the same account as the inquirer. The matrix has one row for each task; the task that executed □runs is always in the top row. To be compatible with applications developed for mainframe SHARP APL the matrix consists of eight columns. However, under UNIX, columns 4 and 5 are meaningful only for the first row; these columns contain 0's for all other rows.

Table 5.10. Columns of □runs.

Column	*Meaning*
0	Task number
1	Account number of task's owner
2	Account number of task's initiator
3	Task type (0: T- or S-task. 1: N-task)
4	CPU units since start of the task
5	Elapsed seconds since start of the task
6	Task's CPU limit
7	Task's elapsed time limit

***Note*:** While values for CPU and elapsed limits are maintained, they do not actually restrict an APL task.

□sc Shared Variable State Change

This *interlocked* system variable contains the value 1 when the state of a name shared with another task has changed, or when the number of seconds you have set in □sc has elapsed. This system variable is documented in the *SVP Manual*, *Chapter 3*.

□signal Signal Event

Rank: ⋆ ⋆ ⋆

The monadic system function □signal terminates execution of the user-defined function or ⍎ that contains it, and signals an event to the function that invoked it. This system function is documented in "Chapter 9. Event Handling".

□size Size of Tied File

Rank: 1

The monadic system function □size returns a 4-element integer vector describing the size of a tied file. This system function is documented in the *File System Manual*, Chapter 2.

□sp Session Parameter

The system variable □sp retains its definition for the entire session (T-task, S-task, or N-task) to which it belongs. At the start of a session, □sp is set by the interpreter to its default value (an empty character vector) unless you have supplied a nonempty character vector from UNIX through the startup parameter -s (see the *Handbook*, Chapter 3). In conjunction with □lx (latent expression), an initial setting of the session parameter provides a powerful tool for customization of a SHARP APL application.

At any time during an APL session, you can set □sp to the value of any variable; this value survives in □sp after execution of the system commands)load,)xload, and)clear, as well as that of the system functions □load and □qload. An N-task inherits □sp of the task that □run's it.

The session parameter is often used in conjunction with a package as a means of moving a set of objects from one active workspace to another without having to use the system commands)save and)copy:

```
□sp←□pack 'function variable'
)clear
□pdef □sp
```

As with other variables in the active workspace, the space required to store the value in □sp is subtracted from the working area available (see □wa). When the session ends, the contents in □sp is lost. If you wish to preserve its contents for use in a future session, save the value in □sp under its own name.

⎕*stac* Set Access Matrix of a Tied File

Rank: `2 1`

The dyadic system function ⎕*stac* allows you to modify the access matrix of a tied file. This function is documented in the *File System Manual, Chapter 2*.

⎕*stie* Share-tie

Rank: `1 1`

The dyadic system function ⎕*stie* lets you establish the linkage between the name of an APL component file and the tie number used to refer to the file. The linkage is such that other users can concurrently share the file. This system function is documented in the *File System Manual, Chapter 2*.

⎕*stop* Stop Execution of User-defined Function

Rank: `1 1 1`

The ambivalent system function ⎕*stop* sets, removes, or reports flags on specified lines of an unlocked user-defined function. Execution of the function halts just before any of the flagged lines it reaches, to allow examination and adjustment of the local environment.

The right argument ⍵ is a character vector containing a name whose visible use is to refer to a single user-defined function. Used monadically, ⎕*stop* `'foo'` returns an integer vector containing the numbers of the lines at which the interpreter halts before execution of the defined function *foo*.

Used dyadically, the left argument is an integer vector of the line numbers before which execution should halt. The line numbers in ⍺ need not be in order. A dyadic use resets the list of stops; any integer (including `0`) that is not the number of a line in the function named ⍵ has no effect. A dyadic use also removes any existing stop settings from lines in the function that are not included in the argument: `(⍳0)` ⎕*stop* `'foo'` removes all stop settings from the definition of *foo*.

⎕*stop* is effective only for an unlocked user-defined function.

The result is the *former vector of stops* (that is, stops as they were before the reset), or an empty vector when the definition is locked.

The effect of ⎕*stop* is to halt subsequent executions of the function immediately prior to executing the line(s) specified in the left argument. The system returns to immediate execution mode, preserving the state indicator and all local values and definitions. You can then explore and possibly alter the local environment before branching (→) back into or out of the suspended function. The resulting ability to observe and alter the local environment at those chosen points in execution is a valuable aid in debugging a program. Note, however, that a return to immediate execution is a trappable event, so it is possible to set ⎕*trap* in a way that nullifies the effects of ⎕*stop*.

Stop settings are saved and reloaded with a workspace, but are not copied (by the)*copy* or)*pcopy* system commands) with the function to which they apply.

Redefining a function with ⎕*fx*, ⎕*fd*, or the full-screen editor removes all stop settings from that function. If other lines are inserted or deleted before a line, the setting moves with the line of code, so that the result of monadic ⎕*stop* changes accordingly. Stop settings are not preserved in packages.

⎕*svc* Shared Variable Access Control

Rank: `1 1 1`

The ambivalent system function ⎕*svc* reports or sets the 4-element *access control vector* for a shared variable. This system function is documented in the *SVP Manual*, *Chapter 3* .

⎕*svn* Shared Variable Clone ID

Rank: `1 * *`

The system function ⎕*svn* permits you to set or inspect the identifying number used to distinguish multiple tasks for the same user, and to obtain or remove a sign on index from a remote host. This system function is documented in the *SVP Manual*, *Chapter 3*.

⎕svo Shared Variable Offer

Rank: `1 ¯1 1`

The ambivalent system function ⎕*svo* sets or reports *offers* to share one or more variables with another task. This system function is documented in the *SVP Manual*, *Chapter 3*.

⎕svq Shared Variable Query

Rank: `1`

The monadic system function ⎕*svq* reports the identification of tasks that have made offers to share a variable with your task, and identifies the variables offered by a particular task. This system function is documented in the *SVP Manual*, *Chapter 3*.

⎕svr Shared Variable Retraction

Rank: `1`

The monadic system function ⎕*svr* retracts an offer to share a variable. This system function is documented in the *SVP Manual*, *Chapter 3*.

⎕svs Shared Variable State

Rank: `1`

The monadic system function ⎕*svs* reports the state of one or more shared variables. This system function is documented in the *SVP Manual*, *Chapter 3*.

⎕SYSVER System Version

The niladic system function ⎕*SYSVER* returns a character vector with three fields delimited by two spaces. This experimental system function may not be available in future releases of SHARP APL for UNIX.

The three fields of the result provide the following information:

- version of SHARP APL and the serial number of the copy active on your host, expressed as a single number in the fixed-length form `vv.vv000sssss`, where `vv.vv` is the version number, and `sssss` is a 5-digit serial number for your copy
- date for the version in the fixed-length form `yyyy/mm/dd`
- host identifier, one or more words of varying length

Example:

```
      ⎕SYSVER
4.9.20000000  1998/12/21  SPARC
```

⎕*tie* Exclusive-tie

Rank: `1 1`

The dyadic system function ⎕*tie* lets you establish the linkage between the variable of a file and the tie number used to refer to the file, in such a way that only your task accesses the file. This system function is documented in the *File System Manual*, Chapter 2.

⎕*trace* Trace Execution of a User-defined Function

Rank: `1 1 1`

The ambivalent system function ⎕*trace* permits the user to observe the progress of specified lines in the execution of a user-defined function as a debugging aid, or to determine which lines are flagged to be observed. The arguments and results of ⎕*trace* are the same as those of ⎕*stop*. Both functions are affected by copying, fixing, and editing in the same ways.

During execution of a user-defined function, the interpreter displays the value of the *root expression* in each statement on the traced line. In the display, the value appears after the function's name and the bracketed line number, or after a ⋄. In the case of a branch in the function, a → is displayed before the line number to which execution branched.

The resulting ability to observe the sequence in which the lines are executed and the internal values (not normally displayed) is a valuable aid in debugging a program.

□*trap* Event Trap

By setting the value of the system variable □*trap*, you specify the action to be taken if certain events occur. This system variable is fully documented in "Chapter 9. Event Handling".

□*ts* Time Stamp

The niladic system function □*ts* returns the current date and time of day (represented by the computer's internal clock) as a 7-element integer vector. Table 5.11 shows the information provided in this vector.

Table 5.11. Elements of □ts.

Element	*Meaning*
0	year
1	month
2	day
3	hour
4	minute
5	second
6	millisecond

The first three elements of □*ts* always indicate a date, and the last four elements always indicate a time of less than 24 hours.

⎕*twsid* Termination Workspace Identification

Rank: ⋆

The monadic system function ⎕*twsid* establishes and reports the name of the workspace to be saved in the event of an abnormal termination of the APL task.

If an APL session terminates without manual intervention, as would be the case if the phone connection to a remote terminal is dropped, or your task is bounced, it may be valuable to save a copy of the active workspace automatically. This is especially useful, for example, in debugging disconnected S-tasks. A copy of the workspace saved at the moment of termination allows you to examine ⎕*er* (the event report) and other information that would otherwise be completely lost.

The termination workspace facility in SHARP APL for UNIX follows the design of SHARP APL for OS/390 in all respects. When spontaneous termination occurs, the default action is to save a termination workspace. The default (and traditional) name for the workspace saved by the system is *continue*, which is taken to mean `continue.sw` in your home library.

⎕*twsid* Usage

Use a character right argument to set or reference the name of the termination workspace. Use a numeric right argument to set or reference the action to be taken upon termination. When you change the value of the name or the action, the explicit result of ⎕*twsid* is the former value for the name or action.

Do not specify libraries 1 to 99 as character arguments to ⎕*twsid*, as those are reserved for public libraries.

Termination Workspace Name

To refer to the current setting of the termination workspace name, use ⎕*twsid* with an empty character vector:

```
     ⎕twsid ''
121 continue
```

To set the termination workspace name, use ⎕*twsid* with a valid workspace name presented as a character vector or scalar either in library or directory form:

```
      ⎕twsid '/usr2/cory/tmp/continue'
121 continue
      ⎕twsid ''
/usr2/cory/tmp/continue
```

The actions you control with an integer argument to ⎕*twsid* are listed in Table 5.12.

Table 5.12. Actions for ⎕twsid.

Action	*Value*
Save	`1`
Inquire	`0`
Ignore	`¯1`

The default is `1`: the system automatically saves your active workspace if your session terminates spontaneously. To reference the current action setting, use the value `0` as the argument:

```
      ⎕twsid 0
¯1
```

Use a `1` or a `¯1` to change the action to save or not save your active workspace, respectively:

```
      ⎕twsid 1
¯1
```

Resuming Work

You may want to use the `-W` startup option when you begin your next session to request the system to load your termination workspace automatically at the start of your session:

```
% sax -Wcontinue
```

⎕*ul* User Load

The niladic system function ⎕*ul* returns an integer scalar representing the number of APL tasks running on this host.

⎕*untie* Untie Tied Files

Rank: `0`

The monadic system function ⎕*untie* unties the files whose tie numbers appear in its argument, which is an integer vector or scalar. This system function is documented in the *File System Manual*, *Chapter 2*.

⎕*USERID* UNIX User Name

The system function ⎕*USERID* returns a character vector containing up to the first eight characters of your UNIX user name. This experimental system function may not be available in future releases of SHARP APL for UNIX.

⎕*vi* Verify Input

Rank: `1`

The monadic system function ⎕*vi* can be used in conjunction with ⎕*fi* to provide a validity check on the argument which is a character vector or scalar.

⎕*vi* treats ⍵ as a sequence of fields, in the same fashion as ⎕*fi*. The result is a Boolean vector with a `1` for each field whose characters represent a well-formed number, and a `0` otherwise. See also ⎕*fi* in this chapter.

⎕*wa* Work Area Available

The niladic system function ⎕*wa* returns an integer scalar whose value is the current number of unused bytes of work area available in the active workspace. It is computed by subtracting from the gross size allocated to the workspace the amount currently in use for the definitions of variables and functions, for the symbol table, and for the execution stack and partial results associated with it.

⎕ws Workspace Information

Rank: ∞ ∞

The dyadic system function ⎕ws reports varied information about the current state of the active workspace. The information returned in the result depends on the values of the two arguments, each of which is a numeric scalar. The possible values and the results associated with them are as follows:

`1 ⎕ws ⍵`

Name list. The right argument ⍵ is a numeric scalar identifying the class or classes of names to be included in the result. The effect of `1 ⎕ws ⍵` is essentially the same as that used by ⎕nl, which supersedes it. The right argument identifies the classes to be included in the result by a scheme different from that used by ⎕nl. The argument ⍵ must be an integer scalar whose value is the sum of the numbers representing the desired classes, which are encoded as powers of 2. Values for the argument ⍵ are shown in Table 5.13.

Table 5.13. Right argument of `1 ⎕ws ⍵`.

⍵	*Class*
1	Function
2	Variable
4	Undefined (formerly *group*)
8	Label
16	IF Bind

The result is a character matrix containing one row for each visible name in any of the classes identified in ⍵.

`2 ⎕ws 1`

Workspace ID. The result is a character vector containing the name of the active workspace in the same form as the characters displayed by the system command)wsid (see "Chapter 6. System Commands"). The form of the result varies depending upon whether you are using *directory* or *library* format. See "Chapter 3. Files, Libraries, and Directories" for more information.

`2 ⎕ws 2` **State indicator**. The result is a character matrix containing one row for each element on the state indicator, in the same form as displayed by the system command `)si`. Each row contains the name of a user-defined function followed by a line number in brackets. A function that is suspended is marked by an asterisk following the closing bracket of the line number. The line number is the number of the line currently active (for a function that is not suspended) or the number of the line to be executed next (for a function that is suspended). The rows of the matrix show the most recently invoked functions at the top. Calls to ⍎ or to ⍞-input are shown in the result in the same way as user-defined functions; each has a row on which the symbol ⍎ or ⍞ appears alone.

Table 5.14. Result of `2 ⎕ws 3`**.**

Element	*Meaning*
`0–1`	Unused; value `¯1`
`2`	Bytes used in active workspace
`3–4`	Unused; value `¯1`
`5`	`1` in a *recovered* workspace; `0` otherwise
`6`	Current size of symbol table
`7`	Symbols in use
`8`	Account number of owner (same as `1↑⎕ai`)
`9`	Unused (task's port number)
`10`	Task-ID (assigned by UNIX)
`11`	Output encoding used by session manager:
	`0` N-task
	1 Unused (IBM Correspondence)
	`2` Unused (IBM BCD)
	`3` Unused (ASCII without APL)
	`4` T-task (ASCII with APL)
	`5` Unused (ASCII with APL bit-paired)
	`6` Unused (Telex)
	`7` S-task
	`8` Unused (PC)

`2 ⎕ws 3` **Workspace environment**. The result is a `12`-element integer vector. Table 5.14 describes the meaning of an element at each position in the result. A number of elements are no longer in use and contain `¯1` or `0`, but are preserved to maintain consistent positions for the others. This matrix notes several results which, while defined for the environment of SHARP APL for OS/390, never occur in the UNIX version.

`2 ⎕ws 4` **Workspace information**. The result describes the workspace in a `3`-element integer vector in the same way that `⎕rdci` describes a component of an APL file (see the *File System Manual*, Chapter 2). Table 5.15 describes the meaning of an element at each position in the result.

Table 5.15. Result of `2 ⎕ws 4`**.**

Element	*Meaning*
`0`	Bytes needed to store the contents of the workspace
`1`	Account which saved the workspace
`2`	Timestamp (since 1960-03-01 00:00:00 in 60^{ths} of a second)

`3 ⎕ws ⍵` **Group members**. Since groups are not supported in SHARP APL for UNIX, the result for any argument is a `0`-by-`0` character matrix; except packages, which result in `domain error`.

`4 ⎕ws ⍵` **Storage space**. The number of bytes of workspace storage required to store each of the objects named in `⍵`, which is a name list (a character vector with successive names separated by blanks, or a character matrix with one name per row). The result is the value required if the named object were the only instance of the object (although in fact an object can be stored merely by pointing to another instance of the same value). The value includes space occupied by the header that the interpreter uses to describe the object's type, rank, and shape, but not the space occupied by the name of the entry in the symbol matrix.

`5 ⎕ws ω` **Name class by state indicator level.** The right argument ω is a name list (a character vector with successive names separated by blanks, or a character matrix with one name per row). The result is an integer matrix, having one row for each name in ω, and one column for each level of the state indicator, plus an additional final column representing the global environment (not represented in the state indicator). Values within the result are listed in Table 5.16.

Table 5.16. Result of `5 ⎕ws ω`***.***

Element	*Meaning*
`¯1`	not localized at this level
`0`	undefined use - localized at this level
`1`	function - localized at this level
`2`	variable - localized at this level
`4`	unused (group)
`8`	label - localized at this level
`16`	IF Bind

`6 ⎕ws ω` **Value of variable named in** ω. This facility allows a user-defined function to obtain the value in a variable without causing a reference to a shared variable. You can obtain the value in a name with a shared variable reference when you use the function ⍎ (execute), which was originally introduced into SHARP APL for OS/390 after this system function.

Summary of System Variables and Functions

The seven categories in Table 5.17 show the system variables and functions at a glance. Many other groupings are possible.

Table 5.17. Categories of system variables and functions.

System	Workspace	I/O	Files	Packages	Sharing	Events
□ai	□cr	□	□append	□pack	□sc	□ec
□av	□ct	□	□appendr	□pdef	□svc	□er
□avm	□dl	□arbin	□avail	□pex	□svn	□signal
□bounce	□ex	□arbout	□bind	□pins	□svo	□trap
□load	□fd	□fc	□create	□plock	□svq	
□qload	□fi	□fmt	□drop	□pnames	□svr	
□run	□fx	□fsm	□erase	□pnc	□svs	
□runs	□io	□ht	□fhold	□ppdef		
□sp	□lc	□pp	□hold	□psel		
□SYSVER	□lx	□PR	□lib	□pval		
□ts	□MF	□ps	□names			
□twsid	□NA	□pw	□nums			
□ul	□nc		□paths			
□USERID	□nl		□rdac			
	□PARSE		□rdci			
	□rl		□rdfi			
	□stop		□read			
	□trace		□rename			
	□vi		□replace			
	□wa		□resize			
	□ws		□size			
			□stac			
			□stie			
			□tie			
			□untie			

6
System Commands

A system command is a request to UNIX to perform an action that affects the environment of an active workspace. A system command begins with a right parenthesis followed by the name of the command and its arguments (if any). The names of commands are spelled in the primary alphabet. Where a command takes one or more arguments, the arguments are separated from the command and from each other by blanks. Names of APL objects that appear in system commands are never surrounded by quotes even when related system functions (see "Chapter 5. System Variables and Functions") require quotes around their arguments.

You can execute a system command only when the interpreter is in immediate execution mode (including □-input mode) in response to the interpreter's standard input prompt (newline followed by six blanks). You cannot execute a system command from within the ∇ editor, nor can a system command occur in the definition of a user-defined function. However, a program that controls another APL task by way of a shared name can use the shared name interface to transmit to the other task the characters that form a system command, provided that the other task is itself in immediate-execution mode when it receives the command.

The complete list of SHARP APL system commands is summarized in Table 6.1. Optional portions of a command are shown surrounded in square brackets; the square bracket character itself is never part of the command. A functional description of each command is provided later in this chapter.

Table 6.1. System commands.

Command	*Notes*
)	recall previous line
)*clear* [*size*]	[workspace size in bytes]
)[*p*]*copy* *wsname* [*nm1 nm2 ...*]	*wsname* in path or library format
)*drop* *wsname*	*wsname* in path or library format
)*edit* *name* [*options*]	
)*erase* *nm1 nm2 ...*	
)*fns* [*nm1*]	
)*lib* [*directory*]	*directory* may be path or library format
)*libs* [*all*]	displays path library associations
)[*x*]*load* *wsname*	*wsname* in path or library format
)*off*	
)*reset*	superseded by)*sic*
)*save* [*wsname*]	*wsname* in path or library format
)*sic*	state indicator clear
)*si*[*namelist*]	state indicator [*namelist*]
)*symbols* [*n*]	
)*vars* [*nm1*]	
)*wsid* [*wsname*]	*wsname* in path or library format

Workspace Names

Many of the commands refer to or affect the name of your currently active workspace. You can specify a workspace name by library or by path. Your private workspace library is, by default, the set of workspaces saved in your home directory. If you start your session from your home directory, then the workspaces you have at your disposal without any qualification of the workspace name are files with the suffix `.sw` in that directory.

For example, when you type)*lib* from a SHARP APL session invoked from your home directory:

```
      )lib
utils data
```

you can see the workspaces you saved there. Your home directory is the one associated with your UNIX user ID, the number SHARP APL calls your account number. If your account number is 106, an equivalent statement, which makes the library number explicit, is:

```
      )lib 106
utils data
```

You can request a listing of workspace names from the private home directory of any user on the system if you know their account number:

```
      )lib 103
ststs wsis0 testing
```

You are also free to use UNIX path designations, both relative and absolute, to indicate a directory. Any use of the four special symbols / (slash), . (dot), $ (dollar sign), and ~ (tilde) in forming a directory name show your intention to specify a directory by path and not by library. For example, you can specify the private library of user 103 either by presenting a relative path to that user's home directory:

```
      )lib ../joe
ststs wsis0 testing
```

or by presenting an absolute path:

```
      )lib /usr2/joe
ststs wsis0 testing
```

The same applies to public libraries. For example, the public workspaces distributed with SHARP APL are stored in library 1. The path to the directory associated with this library is revealed through use of the system command)*libs*:

```
      )libs
          1 /usr/sax/rel/lib/wss
```

This association lets you request a listing of the workspaces stored in that directory either using library format:

```
      )lib 1
ap124 del hostap hsp if pc108 socket ssql toolkit tools unix
wssearch wstransfer xfer
```

or using path format:

```
      )lib /usr/sax/rel/lib/wss
ap124 del hostap hsp if pc108 socket ssql toolkit tools unix
wssearch wstransfer xfer
```

Alphabetical List of System Commands

Following is an alphabetically arranged list of all the available system commands. Each entry presents a definition, points to a related definition, or poins to discussions in other chapters.

) Recall Previous Line

Typing) alone on a line causes the interpreter to display the sentence previously typed from the keyboard during immediate execution. You can accept the line as it is, or edit it in the usual way (by inserting or deleting characters). When you press ENTER, the line goes to the interpreter for execution.

The line recalled in response to) is the *APL expression* (not a system command) most recently entered from the keyboard. If you try to recall a line immediately after loading a workspace, the expression ⍎⎕*lx* appears. The OS/390 version of SHARP APL allows you to distinguish between levels of the state indicator in using this command, but the UNIX version does not. You can use the session manager's searching and scrolling features to recall lines in a fashion that was not available on typewriter terminals when this system command was first introduced.

SHARP APL for OS/390 also includes provision to follow) by a number between 0 and ⎕*pw*, which governs position of the cursor when a sentence is recalled. A value greater than 0 invoked a two-pass editor intended to benefit low-speed printing terminals; this editor is not supported in SHARP APL for UNIX. Using the system command)*sic* results in clearing the recall buffer as

well as the state indicator. This command is not available in the interpreter called `apl-8` which, because it has no session manager, has no facility to store lines from your session log.

)*binds* Display Bound Intrinsic Functions

)*binds* [*nm1*]

This command causes SHARP APL to display a list of names whose visible referent is a bound intrinsic function, in alphabetical order, horizontally across the terminal display. For alphabetization, the collating sequence is sorted based on ⎕*av*. When the command is followed by a name, the list shows only those names that occur in the alphabetized list at or after the position of the sample name. The names listed are the same as those returned in the result of `1 ⎕`*ws* `16` (See "Chapter 5. System Variables and Functions" for more information on ⎕*ws*).

)*clear* Clear Active Workspace

)*clear* [*size*]

This command clears the active workspace and restores all system variables, except for ⎕*sp* and ⎕*pw*, to their default values. Table 6.2 presents an overview of the environment in a clear workspace.

The optional argument specifies the size of the new work area in bytes. The minimum you can request is 16384 bytes (2^{14}); the maximum depends on the installation and the resources of the machine. SHARP APL requires you to write the number in decimal form without commas.

Table 6.2. Conditions in a clear workspace.

Meaning	*Name*	*Default Value*
Comparison tolerance	`⎕ct`	`2*¯44`
Environment condition	`⎕ec`	`0`
Event report	`⎕er`	`0 0⍴''`
Format control	`⎕fc`	`.;**_¯`
Horizontal tabs	`⎕ht`	`''`
Index origin	`⎕io`	`1`
Line counter	`⎕lc`	`⍳0`
Latent expression	`⎕lx`	`''`
Printing precision	`⎕pp`	`10`
Prompt replacement	`⎕PR`	`''`
Position and spacing	`⎕ps`	`¯1 ¯1 0 1`
Printing width	`⎕pw` †	`80`
Random link	`⎕rl`	`16807`
State change	`⎕sc`	`0` or `1`
Session parameter	`⎕sp` †	`''`
Trap control	`⎕trap`	`''`
Work-area available	`⎕wa`	As specified.
Shared names		None; previous offers retracted.
Symbol table		Initial size varies with `⎕wa`.
Workspace name		None; reported as `clear ws`.

† ***Session variable, unaffected by*** `)clear` ***or*** `)load`

For example, to get a million bytes, you enter

```
)clear 1000000
```

and not `1,000,000` or `1e6`. If resources are available, the system assigns the space requested and sets an initial size for the symbol table as a function of the size of the work area. Thereafter, the symbol table grows automatically as needed to accommodate names for APL objects defined in the workspace.

When you use)*clear* with no argument, the gross size of the workspace remains as it was before the command. See also □*trap* action *d clear*, (in "Chapter 9. Event Handling") which clears the workspace when certain events are encountered.

)*copy* Copy Objects into Workspace,)*pcopy* Protected Copy

)[*p*]*copy* *wsname* [*nm1 nm2 . . .*]

The name of the workspace you wish to copy in whole or in part is expressed in library or path format (see "Workspace Names" on page 6-2).

The command)*copy* copies global objects from the saved workspace *wsname* into your active workspace. When no names of APL objects follow *wsname*, all global account names (and no system names) are copied. When names of objects do follow *wsname*, only the objects named are copied. The list of names may include those of system variables. After a copy, the list of names of global objects in the active workspace is the union of those previously in the active workspace and those copied from the saved workspace.

The two forms)*copy* and)*pcopy* differ in the way they treat names that occur in both workspaces. Following)*copy*, for global objects having the same name in both workspaces, the version from the saved workspace overwrites the version in the active workspace. When using)*pcopy*, however, global objects already in the active workspace are protected. They remain as they were, and objects from the saved workspace do not overwrite those already present. The system displays a list of the names of objects that were not copied.

The interpreter makes a working copy of a user-defined function's definition when it executes the function, and keeps this instance of the function separate from the reference definition stored in the workspace. When you copy a new definition for a function while a function of the same name is *pendent* (its name is on the execution stack), the reference definition is overwritten, but the interpreter's working copy remains untouched. The effect of copying a new definition for a pendent function does not become apparent until you complete (or abandon) execution of the working copy.

)drop Delete a Workspace

)drop *wsname*

The name of the workspace you wish to drop can be expressed in library or path format (see "Workspace Names" on page 6-2).

The saved workspace named *wsname* is discarded from library storage. The active workspace is unaffected, and so is its name (which may be the same as the name of the workspace thus dropped). The system reports the date and time at which you dropped the saved workspace.

```
      )drop data
1999-03-30 11:47:32
```

)edit Full-Screen Editor

)edit *name*

If you are running the interpreter called `apl`, which has a built-in session manager, the system command)edit activates the full-screen editor on the object named. (See the section "Full-Screen Editor" in "Chapter 7. System Facilities" for descriptions of editor commands.)

)erase Erase Global Objects

)erase *nm1 nm2 . . .*

This command erases the global definitions of the user-defined objects named in the list to the right. The command has no effect on the working copy of a pendent function, but does erase the reference definition of which the working copy is an instance. An erased pendent function's definition completely disappears only when you complete or abandon its execution.

)fns Display Function Names

)fns [nm1]

This command causes SHARP APL to display a list of names whose visible referent is a user-defined function (or function), in alphabetical order, horizontally across the terminal display. For alphabetization, the collating sequence is sorted based on ⎕av. When the command is followed by a name, the list shows only those names that occur in the alphabetized list at or after the position of the sample name. The names listed are the same as those returned in the result of ⎕nl 3 or 1 ⎕ws 1 (see "Chapter 5. System Variables and Functions").

)ifdef Display IF Definition/Redefine IF

)ifdef nm1[.nm2] [defn]

This command manipulates the definitions of either a group of functions or the individual functions in the group.

)ifdef nm1	obtains the current definition of the group of intrinsic functions whose name is *nm1*.
)ifdef nm1.nm2	obtains the current definition of the intrinsic function *nm2* within the group *nm1*.
)ifdef nm1 defn	establishes the definition of the intrinsic function group *nm1*.
)ifdef nm1.nm2 defn	establishes the definition of the intrinsic function *nm2* within the group *nm1*.

In these last 2 forms, *defn* must be a definition of the intrinsic function group or function within the group that is valid according to the rules given for .saxif definition files in the *Intrinsic Functions Manual, Chapter 2*.

This command can be used to establish the definitions of intrinsic functions after starting an APL session, instead of relying on the startup processing associated with an APL session.

Examples:

To define a group,

```
      )ifdef exgrp dynamic /home/joe/ifs/exgrp.si
```

To verify that it is defined,

```
      )ifdef exgrp
dynamic /home/joe/ifs/exgrp.si
```

To define two functions within the group,

```
      )ifdef exgrp.exfn1 Cfunction exfn1(int,char*)
      )ifdef exgrp.exfn2 Cfunction exfn2(int*)
```

To verify that they are both defined,

```
      )ifdef exgrp.exfn1
Cfunction exfn1(int,char*)
      )ifdef exgrp.exfn2
Cfunction exfn2(int*)
```

)iferase Remove IF Definitions

)iferase *nm1*[.*nm2*]

This command removes the definition of either a group of functions or an individual function in the group.

)iferase *nm1*	removes the current definition of the group of intrinsic functions whose name is *nm1*.
)iferase *nm1.nm2*	removes the current definition of the function *nm2* within the group of intrinsic functions whose name is *nm1*.

This command can be used to remove the definitions of intrinsic functions after an APL session has started, when they are no longer required, or possibly in order to redefine them.

Examples:

To remove both functions from a group,

```
)iferase exgrp.exfn1
)iferase exgrp.exfn2
```

To verify that they have been removed,

```
)ifdef exgrp.exfn1
)ifdef exgrp.exfn2
```

To remove the group definition,

```
)iferase exgrp
```

To verify that it has been removed,

```
)ifdef exgrp
```

)iflist List IF Groups and Their Contents

)iflist [*nm1*]

This command displays a list of names of the currently defined intrinsic function groups or the list of names of currently defined functions within an individual defined group, in alphabetical order, horizontally across the terminal display. For alphabetization, the collating sequence is sorted based on □*av*.

)iflist	lists the names of all currently defined intrinsic function groups.
)iflist *nm1*	lists the names of all currently defined intrinsic functions within the group *nm1*.

Examples:

To list all defined groups,

```
      )iflist
Sunix  amap   lf     misc   sapl   socket system unix
```

To list all functions in the `unix` group,

```
      )iflist unix
errno       filesync    path        usefilelocks
```

)lib Display Saved Workspace Names

)lib [*directory*]

This command causes SHARP APL to display a list of the names of saved workspaces, in alphabetical order, horizontally across the terminal display. For alphabetization, the collating sequence is sorted based on □*av*.

Used with no argument, the list shows the names of saved workspaces in your home directory. The optional argument identifies a directory, either your own or that belonging to another user. You can identify the directory by UNIX path:

```
      )lib /usr/fred
```

or by its library number:

```
      )lib 114
```

The display from)*lib* lists each of the files whose name suffix identifies it as a workspace. When it saves a workspace, the system creates a file in a directory belonging to the person who created it, and assigns to it the name that the saver requested.

The full name of the file includes the path leading to the owner's directory. To the name of each workspace the system automatically attaches the suffix *.sw* so it can easily distinguish a file containing an APL workspace from any other type of file. However, SHARP APL neither displays the suffix in lists such as the one returned by)*lib* nor requires you to write the suffix—or even to know what it is—when you refer to a workspace. The suffix is visible only from the UNIX command shell or from AP11.

)libs Display Library Numbers and Directory Paths

)libs [all]

This system command displays a table of library numbers and their associated directory paths. With no arguments, the display contains only those entries which have been explicitly set with the -L startup parameter or explicitly referenced during the session. If used with the optional parameter *all*, the display includes private library numbers for every user enrolled in the system. The table appears similar to the following:

```
)libs all
 1/usr/sax/rel/lib/wss     100 /usr/guest
 109 /home/gsd            108 /home/ayyad
 104 /home/sax            101 /home/users
 103 /home/klh            107 /home/msy
 116 /home/gbe            125 /home/dlf
 111 /home/ema            148 /home/ran
 118 /home/aba            105 /home/bca
 102 /home/mnu            110 /home/radhak
```

)load Load a Workspace
)xload Load a Workspace

)[x]load *wsname* [*size*]

The commands)load and)xload replace the current active workspace by a copy of a saved workspace. The former contents of the active workspace are lost.

The name of the workspace you wish to load can also be expressed in UNIX path format. For example:

```
     )load /usr2/fred/utils
/usr2/fred/utils saved 1999-09-26 10:24:58
```

The optional argument specifies the desired size of the active workspace for contents of the saved workspace to be loaded into, in bytes. The size argument requires the same format as the optional argument to)clear. If resources are not available, or the contents of the saved workspace would not fit in a workspace of the size requested,)load reports *ws too large*.

Without the optional argument,)load attempts to adjust the size of the active workspace to the size of the saved workspace if the two sizes differ.

Following the)load command, the interpreter immediately executes the statement stored as the visible value of the system variable ⎕lx (latent expression). The system functions ⎕load and ⎕qload produce the same effect as the system command)load without an optional size argument; they are described in "Chapter 5. System Variables and Functions".

The)xload variant

The variant command)xload is similar to)load, except that the interpreter does not execute the latent expression. This is intended as a convenience in doing maintenance to a workspace that would otherwise start executing as soon as loaded.

)off Terminate APL Session

This command terminates an APL session and returns control to UNIX. SHARP APL discards the active workspace, retracts any offers to shared names, and unties any tied files. Any auxiliary processors initiated by your task are terminated. If you were running an interpreter with a built-in session manager, it ceases, thus ending interpretation of certain keystrokes as symbols unique to APL.

The form of this command used by SHARP APL for UNIX lacks the variations found in the OS/390 version. Since SHARP APL for UNIX does not use passwords of its own (but leaves that responsibility to UNIX), there is no option to set a password. Similarly, since it is not involved in the handling of terminals (but leaves that to UNIX), there is no variant for holding the communications link for a new sign-on.

See the ⎕trap action 4 off for an alternative method of terminating your session; this is described in "Chapter 9. Event Handling".

)*save* Save Active Workspace

)*save* [*wsname*]

This command causes SHARP APL to store a duplicate of the active workspace in the directory specified by the current workspace ID. It leaves the contents of the active workspace unchanged.

The active workspace resulting from the command)*clear* has no name. You can give your active workspace a name with the command)*wsid*. When you load a duplicate of a saved workspace, the active workspace has the same name as the saved workspace. In either circumstance, the name of the active workspace is only a potential name; that is, the name under which your active workspace is saved if you were to execute the command)*save*.

Used with no name following it, the command)*save* causes the system to save a copy of the active workspace using the ID of the currently active workspace. This new copy updates the saved workspace, replacing any former saved version with a duplicate of the present active workspace. You can save it under the name it currently has, which it may have acquired when you loaded it, or under the name which you later gave it by the)*wsid* command. You can resave a workspace only if it has a name that makes you its owner, or a name that saves it in a UNIX directory to which you have write access.

When the workspace as yet has no name, the system rejects the command)*save* with the words *not saved, this ws is clear ws*. Used with a name following the word)*save*, the system saves a duplicate of the active workspace under the name you state. When there already exists a workspace with the specified name, the save will fail. That is, the workspace *data* cannot be overwritten by typing)*save data* in the active workspace unless the name of the active workspace happens to be *data*. Normally, such a workspace overwrite is performed by first using)*wsid* to change the active workspace name.

When the)*save* command is carried out successfully, the name under which the workspace is saved also becomes the name of the active workspace.

Example:

```
      )save
/usr2/coryc/data saved 1999-03-30 11:41:23
```

The size of the active workspace is stored in the copy created with)*save* and used in subsequent loads of the saved workspace. To alter the size of a saved workspace, you can use)*load* with the size argument and a)*save*. For example:

```
     )load bigws 1000000
bigws saved 1999-07-26 14:14:50
     )save
bigws saved 1999-07-26 14:17:45
```

)*sesm* Session Manager Conditions

)*sesm* [*trans*[*off*|*keyw*[<esc>] | *scissor*[*off*|*on*]]

This command reports or sets session manager conditions. Currently, the arguments available for use with)*sesm* include the "transliteration" setting or the scissoring option.

)*sesm trans*[*off*|*keyw*[<esc>]]

If the session manager is unable to display APL characters,)*sesm trans* provides you with the ability to use keywords in place of APL variables and functions; this is called *transliteration* (available for the `apl-8` interpreter only). This argument turns transliteration on or off, sets the transliteration keyword escape character, and reports the display status regarding transliteration.

For example,

```
     )sesm trans keyw ?
was off
```

turns on transliteration and sets the keyword escape character to *?*. If no character is supplied with)*sesm trans keyw*, the keyword escape character is automatically set to the default character, `#`. The resulting message indicates that transliteration *was* off (it is now turned on). In the following example, the default keyword escape character, `#` ,is used to define keywords in the APL expression ⊢*a*←○2.

```
     #dex a #is #circle 2
6.283185307
```

Due to the frequency of the use of □ in system functions, the interpreter allows a lone keyword escape character (e.g., #) to be used as a shortcut for the keyword #*quad*; i.e., #*av* is equivalent to #*quad av*. For a complete list of available keywords, turn on transliteration and enter □*av* using the transliteration keyword escape character (i.e., #*av*). The following expression checks the status of the current display regarding transliteration:

```
      )sesm trans
trans is keyw ?
```

Transliteration is already turned on, therefore, the resulting message reports the current keyword escape character (*?*). More information on using keywords for APL primitives is provided in the *Handbook, Chapter 4*.

)*sesm scissor*[*off*|*on*]

This option reports or sets the type of display for wide arrays. The following example shows how the result of an APL expression is displayed when scissoring is turned on and when it is turned off:

```
      □pw←30
      )sesm scissor
scissor is on
      2 32⍴'a'
aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaa
aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaa
      aa
      aa

      )sesm scissor off
was on
      2 32⍴'a'
aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaa
      aa
aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaa
      aa
```

More information regarding the folding and scissoring of wide displays appears in the *Language Guide, Chapter 3*.

)sic State Indicator Clear

This command abandons execution of all pendent functions, thereby clearing the state indicator. The referents of all local names are lost. The effect of this command is different from and more stringent than that produced by the naked branch →, which abandons execution only of the function most recently started from immediate execution, together with all the functions that function may have invoked. See also □*trap* action *d exit* in "Chapter 9. Event Handling"

The system command)*reset* still performs its role as an obsolete synonym for)*sic*.

)si State Indicator,
)sinl State Indicator with Name List

This command causes the interpreter to display a list of the functions whose execution is either suspended or pendent, and which are therefore on the execution stack (or state indicator). The list shows each level of execution on a separate line, with the name of the function and the number of the line within it that is currently active. For a function that is suspended, the display shows an asterisk. The line number in that case is the number of the line next to be executed when execution is resumed.

To illustrate, consider the following display:

```
     )si
get[4] *
verify[1]
mean[1] *
analyze[3]
report[2]
```

The function *get* is halted and has not completed execution of line 4. If resumed, it would start at the beginning of line 4. The function *get* was invoked by line 1 of *verify*. The function *verify* is waiting for *get* to complete, at which point *verify* will resume automatically its work on line 1. A function in this state is said to be pendent.

The function *mean* is also halted. Evidently, following that halt, rather than resume its execution, the user invoked *verify* (perhaps in an effort to examine the cause of the interruption in *mean*). The function *mean* was invoked by a

sentence on line 3 of *analyze*, which in turn was invoked by line 2 of *report*. The function *report* must have been invoked in a sentence typed from the keyboard during immediate execution.

The function ⍎ and a request for input via ⍞ appear in the state indicator in the same way as user-defined functions, but without a line number.

A character table having the same appearance as *)si* is returned as the explicit result of the system function *2 ⎕ws 2*. A list of the line numbers that appear in both displays is returned as an integer vector in the system function *⎕lc* (line counter). See "Chapter 5. System Variables and Functions". The command *)si* can also include the trailing letters *nl* (for *n*ame *l*ist) in the form *)sinl*. The display in this case is organized in the same way as for *)si*, but it also includes a list of names local to the function on each line (including names of labels). This command is also available with the older spelling *)siv*, for *s*tate *i*ndicator with *v*ariables).

The same information that *)sinl* returns about the use of names is returned in a different form by the system function *5 ⎕ws ⍵* (see "Chapter 5. System Variables and Functions").

)symbols Set Symbol Table Size

)symbols [*n*]

Used without an argument, this command reports the current size of the symbol table and the number of names in use. The symbol table must contain an entry for every name in the workspace, including the name of every user variable and function. It also contains the name of every name referred to in the definition of any user-defined function that has been executed in the workspace.

Used with an integer item argument, this command causes the interpreter to set aside space in the workspace for a symbol table with the number of entries requested. Although you can request a table smaller than the number of names already in use, it is ignored. Since SHARP APL is able to reset the size of the symbol table automatically, it is not essential to do this. However, it may be desirable to set the size large enough to accommodate all foreseeable names (for example, to avoid some inefficiencies in having to resize the table dynamically), or to make it smaller (where you know few names are needed) in order to free the space for general use.

The number of symbols in the symbol table and the number of symbols actually in use are also available as `6 7@ 2 ⎕ws 3`.

```
     )symbols
is 272; 84 in use
     6 7@ 2 ⎕ws 3
272 84
     )symbols 100
was 272
     )symbols
is 100; 84 in use
     6 7@ 2 ⎕ws 3
100 84
```

)vars Display Variable Names

`)vars [nm1]`

The command `)vars` causes SHARP APL to display a list of those names whose visible referent is a variable (variable). The list appears in alphabetical order, horizontally across the terminal display. For alphabetization, the collating sequence is sorted based on `⎕av`. The names listed are the same as those returned in the result of `⎕nl 2` or `1 ⎕ws 2` (see "Chapter 5. System Variables and Functions").

)wsid Set Workspace Name

`)wsid [name]`

The command `)wsid` either reports or sets the name of the active workspace. Used with no argument, it returns the current name of the currently active workspace; for example:

```
     )wsid
is data
```

Used with an argument, it assigns the name you provide as the current name for the active workspace. The name of the active workspace is relevant only when you subsequently use the system command)*save* to store a copy of the active workspace in your workspace library.

```
      )wsid news
was clear ws
      )save
news saved 1989-3-30 14:43:27
```

7
System Facilities

SHARP APL provides a number of features which, while not strictly speaking primitives of APL, are nevertheless supported as part of the environment:

Session manager. The SHARP APL session manager is available when you run the interpreter called `apl`, which is the default. Other versions of the interpreter which do not contain a session manager, are also available (see the *Handbook, Chapter 3*). The session manager is responsible for the management of an interactive APL task. In particular, the session manager interprets entries from the keyboard (possibly as characters in the APL character set) and displays output. The APL interpreter may place the session in one of several input modes, and the session manager handles input in accordance with them. The session manager maintains the session log (in which past input and output is recorded) so that the screen appears to be a window over entries that extend wider and much taller than the visible screen. In general, you can move back over the log and reuse earlier input or output as part of new input.

The ∇ editor for function definitions. The traditional APL editor facilitates entry and revision of user-defined functions. You invoke it by typing the symbol ∇ followed by the name of the function to be edited.

Control of execution. The session manager includes facilities to start the execution of a user-defined function. If work is halted by an interrupt or error in the course of executing the definition, the system allows you to take diagnostic or corrective action from the keyboard, and to abandon or resume work.

Full-screen editor. The session manager provides a full-screen editor for your APL session, and also makes it available for editing *user-defined functions* and *variables of rank* `1` or `2` character type. A concise summary of editor commands follows the full description.

Session Manager

The session manager handles transmission of characters between the APL interpreter and the input/output device of an interactive task: that is, its screen and keyboard. The session manager treats the display area as if it were a Cartesian plane extending indefinitely upward (off the screen at the top) and rightward (off the screen at the right). The screen is a *window* onto that plane. By default, the window is positioned over the bottom left corner of the plane and moves down as new lines are added at the bottom so that the characters already displayed seem to move upward and disappear off the top of the screen.

In general, when the session manager sends output to the screen, it appends new output at the bottom, adding a new line for each line displayed. The screen window moves down with the new lines, so that the view is usually of the most recent output. The session manager folds each line horizontally to fit within the number of positions specified in ⎕*pw* (printing width). When you work at an `80`-column screen and leave ⎕*pw* at its default value of `80`, all your output fits within the width of the screen. However, if you set ⎕*pw* to a value wider than the screen, the session manager still folds lines at ⎕*pw*. But as far as it is concerned, there are lines that exist off-screen to the right. If you move the window to the right, you'll be able to see them.

A further level of folding may be imposed by the terminal, which is unfamiliar with infinite Cartesian planes and output that cannot be seen. As far as the session manager is concerned, this does not introduce new rows. Thus a logical line from the session manager's point of view may be folded by the terminal so that it occupies more than one line on the screen.

The Session Log

Every line of APL input or output is appended to the session log. The session manager receives all input from the keyboard. Each keystroke causes the session manager to amend the session log, adding new characters, forming overstrikes, scrolling back, and so on. The size of the session log is controlled by one of the startup parameters that may be included in the command to start an APL session, or may be stored in your individual SHARP APL profile (refer to the *Handbook, Chapter 3*). When the log is filled to capacity, lines are deleted from the beginning to make room as new material is added at the end.

Scrolling

The arrow keys ← → ↑ ↓ move the cursor about the screen. They do not generate any character, except that when you use the right-arrow key to move to the right and then type some other key, the session manager treats the intervening gap as if it were filled with blanks.

When you move the cursor so that it *pushes* against a top or bottom edge of the window, the cursor movement *drags* the window with it, so that a new row of characters becomes visible in that direction (and the corresponding row at the other edge of the screen becomes invisible).

Some terminals have keys marked PG UP and PG DN to shift the view by the number of rows on the screen.

Reusing a Line from the Session Log

You may move the cursor upward into the record of earlier entries or displays, and there *edit and reuse* a line that appeared earlier. The procedure is as follows:

- Move the cursor to the logical line you wish to reuse (scroll, press PG UP, or use the search facilities of the editor as necessary to reach it). You can reuse only *one logical line at a time.*
- Insert or delete characters in the line as desired. The changes appear at once, but are not passed to the interpreter until you press ENTER.
- Press ENTER. The session manager restores the logical line that you altered, so that it again appears as it did before you changed it. The session manager copies your edited version of the line to the end of the session log and moves the cursor to the line below it. Then the session manager passes your edited line to the interpreter.

The effect is the same as writing the line yourself at the end of the log—but easier.

Note: Only the *logical line* on which the cursor rests when you press ENTER is sent to the interpreter. A change you make on a line where you do *not* press ENTER is ignored. This warning applies both to changes made during immediate execution and to changes made while using the ∇ editor.

Support of the APL Character Set

The interpreter recognizes a one-byte character set. Of the 256 distinct character encodings, positions `32` through `127` are standard ASCII. A complete table is shown in the *Language Guide, Chapter 3*. More information on keyboards and character support is provided in the *Handbook, Chapter 4*.

Symbols in Displays from APL

Numeric or character displays that make no use of APL's special symbols—for example, the text of a document, the numbers in a report—may use nothing but standard ASCII symbols. In that case, they are readily displayed by standard means. However, the instructions for a calculation or the listing of a program usually contain symbols that are not in the standard ASCII set. To display them on the screen or printer, you have to provide suitable character generators, either in the hardware of the display device or as down-loadable *soft* fonts. Such generators are included with the distributed software for certain devices.

National Use Characters

Like most environments that use one-byte character codes, SHARP APL lacks official encodings for international characters. The only fully correct solution to this problem is to adopt complete support for the ISO/IEC standard IS 10646 (also known as Unicode 1.1) when it becomes available. Currently SHARP APL addresses this issue by reserving some character positions, known as *National Use* characters, whose display is permitted to vary.

National Use characters are fully documented in the *Handbook, Chapter 4*. Any customer that finds a need to customise display or entry of National Use characters should contact the Soliton Associates Technical Support group for assistance.

Insert or Replace while Typing Characters

When you type a character from the keyboard, it is received by the session manager in one of three modes: *insert* mode, *replace* mode, or *APL overstrike* mode. Insert mode places the new character at the cursor position and shifts all characters to the right of the cursor by one position. Replace mode replaces the character currently at the cursor position with the new character you typed. APL overstrike mode is explained later in this section.

The INSERT key toggles between replace mode and insert mode. (If you are using a keyboard that has no INSERT key, or whose INSERT key is not recognized, sending CTRL-T has the same effect.) Each time you press ENTER, the mode reverts to its default.

Entering Symbols from the Keyboard

Because standard keyboards frequently lack some of the symbols used in APL, the front end of the APL system provides means to invoke them from various keyboards. The usual practice is to designate one of the keys a special *APL escape* key, so that, after it has been pressed, pressing another key causes that key to display an APL symbol. In this fashion, the key that otherwise generates *g* or *G* instead generates ∇.

A number of APL characters can also be formed as *overstrikes*, as explained below.

APL Overstrike Mode

With some terminal emulators, when you press the *APL overstrike key* (i.e., ALT-BACKSPACE) the cursor moves one position to the left without erasing the character already there. For the next keystroke, you are in *APL overstrike mode.* The character you type, taken together with the character already there, jointly determine the new character. In that manner, you can overstrike □ with ÷ to form ⍰, and so on.

Once you have entered APL overstrike mode, when your next keystroke:

- forms an APL character, that new character replaces the old.
- is a blank or a movement of the cursor, the character on the screen remains unchanged; the cursor moves in the usual way, and you are no longer in APL overstrike mode.
- is another ALT-BACKSPACE, the character on the screen remains unchanged; the cursor moves in the opposite way, and you are no longer in APL overstrike mode.
- is none of these, it is illegal input; you are still at the same place and still in APL overstrike mode.

Varieties of Backspace

During a session, the standard BACKSPACE key acts as a *destructive* backspace. That is, each time you press it, you delete the character to the left of the cursor. However, when you are already at the beginning of a logical line, BACKSPACE is an illegal entry.

For a legal backspace, the cursor, the character at the cursor, and all characters to the right of it in the same logical line, move leftward together, to take up the position vacated by the character just deleted. The behavior is the same, regardless of whether you are in replace mode or insert mode.

The cursor movement keys and also the APL overstrike (or APL backspace) are *nondestructive*.

Signalling an Interrupt

Interrupts are sent differently, depending on whether the interpreter is executing or is waiting for input.

- When the APL interpreter is *executing* (including executing a sentence that causes it to wait), you can signal an interrupt. From a terminal, that is usually done by pressing the BREAK or ATTN keys. If your terminal does not have a key marked BREAK or ATTN, a control sequence such as CTRL-C may be used. Sending interrupt once is called a *weak* interrupt. In general, it causes the interpreter to halt before starting a new line, but not to stop in mid-line. (This is discussed further in the section "Interrupts" later in this chapter.)
- If the interpreter has not yet halted in response to your weak interrupt, you may send another. That promotes your signal from *weak interrupt* to *strong interrupt*. It is keyed in exactly the same way as a weak interrupt.
- When the interpreter is waiting for input from the keyboard (but before you have pressed ENTER to show that your input is complete), signalling BREAK retracts whatever you have just typed. The effects of signalling interrupt may be deferred until the next pressed character is typed (or the ENTER key pressed). You remain in input mode. If you key the sequence `abc`<CTRL-B>`def` the effect is the same as keying `def`.

- To interrupt the executing line while it is waiting for input in response to ⍞, you signal an input interrupt. This is transmitted by using APL overstrike mode to overstrike the characters `o u t`. This sequence is often available as CTRL-G.

(See the keyboard map provided in the *Handbook, Chapter 4*)

Programming and Immediate Execution

The APL environment provides two main modes of work: *immediate-execution* mode and *definition* mode. There are certain other specialized modes as well; for example, the entry of data. See the discussion of ⎕ and ⍞ in "Chapter 5. System Variables and Functions"

Definition mode. While the system is in definition mode (or *editing mode*), it records each instruction you write as part of a definition, but does not attempt to use the definition until subsequently told to do so.

Immediate-execution mode. While in immediate-execution mode, whenever you enter an instruction (from the keyboard), the system immediately interprets and executes it.

For example, if you write

```
2÷6
```

as soon as you press the ENTER key, the system performs the calculation and displays the answer:

```
0.3333333333
```

Similarly, when you write an instruction such as

```
analyze data
```

the system checks whether the terms `analyze` and `data` have been defined. If it finds that `analyze` has been defined as a function, and that `data` has been given a value as a variable, it executes the definition of `analyze`, applying that function to the variable called `data`. If `analyze` computes a result, the system then formats and displays the result.

Immediate execution is the system's normal mode of operation. When your session first begins, the system is (by default) in immediate-execution mode. The system returns to immediate execution each time it completes work in one of the other modes. It also returns to immediate execution when work is halted by any signal or error.

Defining Variables and Functions

APL programming consists of providing definitions for variables and functions not primitive (that is, not built in) to SHARP APL.

Since APL primitives are denoted by symbols or by *distinguished names* which start with □, any use of an alphabetic name (such as `analyze` or `data` in the example that appeared earlier) is a reference to an object that has been *arbitrarily defined* by some user—perhaps you, perhaps someone else.

You may enter new definitions from the keyboard, or make use of definitions that are already written (perhaps a definition you yourself wrote earlier, or one provided ready-made by others). An *application package* is a set of related definitions that somebody has prepared. Often you can use such a package without knowing how its parts are defined. The package may be written so that, once it starts, it handles all further transactions; some are written so that you don't need to know APL at all, or even to see anything written in APL.

Once a function has been named and defined, you can execute it while the system is in immediate-execution mode by including its name in a sentence you type. In the same fashion, one function can cause execution of another by including in its definition a sentence that mentions the name of the function to be executed. See the *Language Guide, Chapter 1* for more information on primitive functions and variables.

User-Defined Functions

The interpreter executes the lines of a user-defined function in order, from first to last.

Exceptions to the order of execution follow:

→ω ***Branch***. The argument is a vector of non-negative integers. Execution passes immediately to the line in the currently executing user-defined function with a line number identified by the first element in the vector; elements other

than the first are ignored. A label, which is a name separated from the rest of a line by `:` (colon), is often used to make references to line numbers dynamic.

When the vector of line numbers or labels that refer to them is empty, there is no branch, and execution continues to the next line. When the value of the first element in the branch vector is `0` or any integer that is not the number of a line in the definition of the function now being executed, the function's execution is completed and its result passed to the line that invoked it.

The branch arrow does not affect any function other than the one that contains the branch itself: the destination of control is either a line in the present function or a normal exit from the present function.

→ ***Naked branch.*** A branch arrow with no argument causes the function now being executed, and the function or functions that invoked it, to be aborted. The state indicator is cut back to the number of levels it had before the most recent entry in immediate execution. This is less severe a cutback than that produced by the system command `)sic`, which completely clears the state indicator (see "Chapter 6. System Commands").

α `⎕signal` ω

Signal. The right argument of `⎕signal` is a integer vector greater than `0` and less than `1000`. When the ω vector is not empty, execution of the function at the top of the state indicator is aborted, and the event indicated by the value of the first element in ω is signalled in the sentence that invoked the function thus aborted. This makes a user-defined function behave in the same way as a primitive function when it encounters an error. You can use `⎕signal` conditionally by taking advantage of the fact that when the right argument vector is empty, no signal occurs.

The event signalled by `⎕signal` may be an event to which a standard event number has been assigned (see "Chapter 9. Event Handling") or an arbitrary event. If that event number is recognized in a `⎕trap` visible in the environment from which the aborted function was called, the interpreter takes the action indicated in the trap; otherwise it simply halts as it would if no `⎕trap` were in effect.

A user-defined function can signal to itself by using ⍎ (execute) to evaluate an expression containing `⎕signal`: any invocation of ⍎ introduces a new level onto the state indicator; abortion of that level of the state indicator forces a halt in the function that used ⍎.

The left argument of □*signal* is optional; any character vector you supply becomes the message that appears after the event number on the first row of □*er* (event report).

When a Function is Suspended

When the interpreter returns to immediate execution while a function is suspended, you may do any of the following from the keyboard (or the controlling shared name):

- Type any of the controlling options available from within the function's definition (listed above), including →⍵ or → or □*signal* ⍵. The expression →⍵ typed during immediate execution causes the function at the top of the state indicator to resume execution at the start of the line identified in the first element in ⍵. Note that, in immediate execution (as from within the function's definition), →⎕IO does nothing (does not cause a resumption of execution).
- Type a line for immediate execution. This does not remove pendent functions from the state indicator. Note that the new line is executed *within the environment of the suspended functions.* Starting a new execution is not recommended other than as an intermediate step in resolving the causes of the previous suspensions. It is a common error of novices, when work has been suspended, to start a new execution without resolving or abandoning the execution already started.

Interrupts

You can interrupt the interpreter's execution by a signal sent from the keyboard or by a special input to the controlling shared name of an S-task. The specific key or key sequence depends on the terminal and the terminal-support file in use. (See the *Handbook, Chapter 4*.)

The signal sent from the keyboard or through the shared name is called a *break* or *attention* signal. Exactly what happens at the moment the signal occurs depends on the context in the active workspace.

- ***Interrupt current line***. Sending one break signal (sometimes called *soft interrupt*) causes the interpreter to stop execution and return to immediate execution at the end of the *line* it is executing when it receives the signal. The interpreter signals *attention* as event 1002. Sending a second break

signal (sometimes called *hard interrupt*) may be needed to cause an immediate halt in the currently executing primitive function. The interpreter signals `interrupt` as event `1003`.

- ***Interrupt keyboard input***. A special signal (often available as CTRL-G) causes the interpreter to break out of ⍞-input mode and return to immediate execution. The interpreter signals `input interrupt` as event `1004`.
- ***Interrupt shared variable reference***. If the interpreter is waiting for another task to complete a reference to a shared variable, or for a shared name state change to occur, two break signals are required to generate an `sv interrupt` as event `1005`.

The ∇ editor

The ∇ editor (del editor) permits you to enter the definition of a user-defined function or revise an existing definition. The ∇ editor is by no means the only way to edit a definition. The system functions ⎕`cr` and ⎕`fd` return character arrays that represent a definition, and the functions ⎕`fx` and ⎕`fd` convert a text array into a definition. You can use any of a variety of text editors to supply or alter the characters that form a definition. The ∇ editor is provided only for the definitions of functions, and is ill-suited to other kinds of editing.

You invoke the editor by typing, in immediate execution mode, a line whose first nonblank character is ∇, followed by the name of the function to be edited. That puts you into *editing mode.* You remain in editing mode until you type a line whose last nonblank character is ∇ (or ⍢ to lock the function; see below).

You build the header (opening line) differently depending on whether you are opening a definition or reopening an existing definition.

- ***New definition***. The line consists of the character ∇ followed by a paradigm that contains the name of the *function* embedded in an expression that both illustrates its use and provides names for its arguments and its result. The form is

  `z←α function ω`

  where `z` is in the position used for the name of the result, and `α` and `ω` are in the positions reserved for the left and right arguments, respectively. Any well-formed names may be used. The result and either or both of the arguments may be omitted (in which case the function cannot take

arguments or return a result). If a function has only one argument, then the argument must be to the right of the name of the function. Further to the right (beyond the paradigm) is a list of additional names local to the definition, set off from the paradigm and from each other by semicolons.

- ***Existing definition***. The opening line consists of the character ∇ followed by the name of the function and nothing else. It is also permissible (but not required) to include on the opening line the specific editing instruction that would otherwise be typed after the interpreter has accepted the command to open the definition.

While in definition mode, the APL interpreter prompts after each entry by supplying the line number of the next line. Line numbers are at the left edge of each line, enclosed in square brackets. When the editor proposes a line number, you may accept it and, to the right of it, type the definition for that line. Alternatively, you may change the prompt by substituting another line number, or ask the editor to recapitulate all or part of the definition as thus far recorded, by writing one of the following editing instructions in response to one of the editor's prompts:

`[`*n*`]`	Define or redefine line *n*.
`[□]`	Display the entire definition, with lines sorted in ascending order, but *not renumbered.*
`[`*n*`□]`	Display line *n*.
`[□`*n*`]`	Display from line *n* to the end.
`[`*n*`□`*p*`]`	Display line *n* and leave the cursor at position *p*.
`[∆`*n1 n2 . . .*`]`	Delete lines *n1 n2 . . .* from the definition.

At any time while using the editor, you may move the cursor to the display of any of the definition's lines and there insert or delete characters. When you press ENTER while on that line, your corrected version replaces the former text of that line.

If you move the cursor upward to any line displayed earlier (whether or not it was part of the definition) and there press ENTER (first editing the line to remove any leading line number in square brackets, or to modify the line in some other way), the text of that line becomes the text of the next line of your definition. In this way you can capture lines that were entered from the keyboard

while in immediate-execution mode and incorporate them into a function's definition. If the first character of a revised line is not `[`, the line you release goes into the definition as the next available line in the function.

To insert a new line between existing ones, give it an interpolated line number (that is, a fraction between the numbers of the surrounding lines).

When you close the definition and leave the editor (by typing a line whose last non-blank character is ∇ or ⍢), the interpreter sorts the lines in ascending order by line number and then renumbers them so that the line numbers are consecutive integers starting at `1`. The header—the top line containing the paradigm and list of local names—is numbered in the display as line `[0]`.

When you use the ∇ editor to display a function in an expression such as ∇*foo*`[⎕]`∇, the system displays all lines with their line numbers and no opening or closing ∇. This makes editing the header no different from editing any other line. It departs from the convention in SHARP APL for OS/390 of showing the opening ∇ at the beginning of the header line, and the closing ∇ on a line by itself below the display of the body of the definition.

Locking a Function's Definition

Closing a definition with ⍢ indicates that its definition is locked. Henceforth, the ∇ editor refuses to display or edit its definition, and the functions ⎕*cr* and ⎕*fd* return an empty array when asked to represent its definition (see "Chapter 5. System Variables and Functions").

The Full-Screen Editor

SHARP APL provides a full-screen editor for creating and modifying one or more objects (user-defined functions and character variables) in your active workspace. Full-screen editing allows you to:

- create new variables or functions, or modify existing ones
- edit several objects at the same time
- edit an image of your APL session and save lines of text from it into variables or user-defined functions
- use extended cursor movement, scrolling, and paging

- insert and delete characters; copy, move, and delete blocks of text within an edit screen
- move blocks of text between edit screens
- search for specific character sequences, and optionally replace them.

Full-screen editing commands begin with a *command character*, which is the TAB key by default; you can change ***TAB*** to any character of your choice (see the *Handbook, Chapter 3*). Since it is the default, ***TAB*** represents the command character in the discussions that follow.

Activating the Editor

There are three ways to activate the editor:

`)edit` *name*	where *name* is the object you wish to create or modify. This system command is available from immediate execution only.
`⎕EDIT '`*name*`'`	where the right argument *name* is the object you wish to create or modify. `⎕EDIT` is an experimental system function that not be available in future releases of SHARP APL for UNIX.
TAB `e`	where you are automatically prompted for the name of the object you wish to create or modify.

Multiple Screens

You may open more than one edit screen by keying TAB `e` for every additional object you wish to create or modify. The forward (***TAB*** `f`), backward (***TAB*** `b`), and session (***TAB*** `s`) commands allow you to flip between edit screens. You can also transfer text from screen to screen, and from your APL session log. Active edit screens survive system commands such as `)load` and `)clear`.

APL Session versus Editor Screen

There are three main differences in behavior between working in the APL session and the editor screens:

Response to ENTER key.

The APL session treats ENTER as release (read current line and act upon it); the editor treats ENTER as a newline command (move cursor to beginning of next line).

Scrolling memory

The APL session has a fixed amount of scrolling memory (established at sign-on), which forces lines off the top of the scroll as new ones are appended at the bottom; editor lines, however, are not lost: the size of editor memory is controlled by a sign-on parameter (see the *Handbook, Chapter 3*).

Name and numbers

The editor displays the name of the object being edited and provides line numbers along the left side of each screen.

Editor Commands

As mentioned earlier, editor command sequences can only be invoked if they are preceded by the recognized command character; ***TAB*** is the default command character provided by SHARP APL. In some of the following command descriptions, optional parameters are shown in square brackets.

Single Key Commands

Some basic functions are invoked ***without*** use of the ***command character***. The single key commands are listed in Table 7.1.

Table 7.1. Single key commands.

Logical key	*Typical Keystroke*	*Function*
Untype	BackSpace or CTRL-R	Backspace and delete
Delete	Del or CTRL-D	Delete one character
Clear–EOL	Clear line or CTRL-E	Clear to end of line
Undo	CTRL-B	Restore line
Alt	CTRL-A or Esc	APL escape
Overstrike	CTRL-P	Form overstrike
Enter	Return	Release current line
Cursor–Up	Cursor–Up	Cursor up one line
Cursor–Down	Cursor–Down	Cursor down one line
Cursor–Left	Cursor–Left	Cursor left one space
Cursor–Right	Cursor–Right	Cursor right one space
Scroll–Up	Scroll–Up	Scroll up one line
Scroll–Down	Scroll–Down	Scroll down one line
Pg–Up	Pg–Up	Scroll up one page
Pg–Dn	Pg–Dn	Scroll down one page
Ins	Ins or CTRL-T	Insert mode on/off
Interrupt	CTRL-C	Interrupt APL program
O–U–T	CTRL-G	Interrupt ⍞-input
Rfrsh	Rfrsh or CTRL-L	Re-display the screen

Note: All other full-screen editor command sequences must begin with the recognized command character (i.e., TAB).

Screen Management

Note: Optional arguments, *options,* are explained after the list of commands.

`e` *name* [*options*]	***edit.*** Open a screen for the object named.
`w` [*name*] [*options*]	***write.*** Write a copy of the current screen into your active workspace; a new name for the object may be supplied.
`q` [`all`]	***quit.*** Quit current screen (deleting any changes); if `all` is specified, quits all screens.
`z` [*name*] [*options*]	***write & quit.*** Write a copy of the current screen into your active workspace, then quit; a new name for the object may be supplied.
`g` *name*	***get.*** Get named object and copy into the current screen at the current cursor location.
`f` [*name*]	***forward.*** Switch to the next logical screen or to the screen indicated by the object name.
`b`	***back.*** Switch to the previous logical screen.
`s`	***session.*** Move to the APL session.
`r`	***roll.*** Display a list of all active screens

Optional arguments (*options*) must include the dash:

`-fn`	Edit or write the object as an APL user-defined function.
`-char`	Edit or write the object as a character data.
`-vec`	Edit or write the object as a vector.
`-mat`	Edit or write the object as a matrix.
`-ln0`	Edit the object without line numbers.
`-ln1`	Edit the object with line numbers (default).

Cursor Movement and Scrolling

[*n*] ⍘ Move to the left margin or *n* spaces to the left.

[*n*] ⍥ Move to the right end of the line or n spaces to the right.

[*n*] ⍤ Move to the first line of the current screen or *n* lines up.

[*n*] ⍣ Move to the last line of the current screen or *n* lines down.

[*n*] ⍠ *c* Move right to the n^{th} occurrence of character c.

[*n*] ⍗ *c* Move left to the n^{th} occurrence of character c.

[*n*] ⍦ Scroll to the first page of text or up *n* lines from the top of the currently displayed page.

[*n*] ⍢ Scroll down to the last page of text or down *n* lines from the bottom of the currently displayed page.

[*n*] ⍟ Page up to the first page of text or up by *n* full pages.

[*n*] ⍯ Page down to the last page of text or down by *n* full pages.

⍦ *c* Mark the current line with any lowercase letter *c* (⍺−⍳).

⍥ *c* Move to the line marked by the letter *c*.

Copying Lines

Note: The contents of the copy buffer *survives* such actions as ⍟, ⍠, ⍯, or ⍴.

[*n*] [[Save *n* lines downward (including current line) into the copy buffer. Equivalent sequence: [*n*] [⍣.

[*n*] [⍤ Save *n* lines upward (including current line) into the copy buffer.

[*n*] [⍺ Save from current line to line *n*, or all lines to the end of the current screen.

[⍥ *c* Save from current line to the line marked with letter *c* into the copy buffer.

Inserting Lines

`i`	Insert a blank line below the current line.
`I`	Insert a blank line above the current line.
`p`	Put the contents of the copy buffer below the current line.
`P`	Put the contents of the copy buffer above the current line.

Deleting Blocks and Lines

Note: If *n* is omitted, the default value of `1` is used.

[*n*] `d h`	Delete *n* characters (including cursor character) to the left.
[*n*] `d l`	Delete *n* characters (including cursor character) to the right. The default is [*n*] `d` SPACE
[*n*] `d x` c	Delete right to the n^{th} occurrence of character c.
[*n*] `d X` c	Delete left to the n^{th} occurrence of character c.
[*n*] `d d`	Delete *n* lines downward (including current line). Equivalent sequence: [*n*] `d j`.
[*n*] `d k`	Delete *n* lines upward (including current line).
[*n*] `d a`	Delete from current line *upward or downward* to line *n*. If *n* is omitted, all lines downward to the last line of the current screen are deleted.
`d o` c	Delete through line marked with letter *c*.

Splitting and Joining Lines

,	Join next line to the current line to form one line.
.	Split the current line at the cursor position.

Search and Replace

`/` *text*	Search logical screen for next occurrence of *text*.
`/`	Repeat previous search command.
[*n*] `? ?` *string*	Replace *n* occurrences of *text* from the previous search with string. If n is omitted the default is `1`. If *n* is ~ (*tilde*), all occurrences are replaced.
SPACE	Repeat the previous replacement command.
[*n*] `? l` *string*	Delete *n* characters to the right and replace with *string*.
[*n*] `? h` *string*	Delete *n* characters to the left and replace with *string*.
[*n*] `? x` *c string*	Delete right to the n^{th} occurrence of character *c* and replace with *string*.
[*n*] `? X` *c string*	Delete left to the n^{th} occurrence of character *c* and replace with *string*.

Status Line

`t`	Toggle status line on or off.

Errors

An incorrectly formed editor command will produce a warning to the user by flashing the terminal screen or sending a terminal bell character. This may occur when you try to delete more lines than a screen contains or form an illegal overstrike character, or when you attempt various other actions.

Error messages generated by the editor are as follows:

`cannot get undefined object`
: The object named in the ***get*** command does not exist in the active workspace.

`cannot quit al session`
: The ***quit*** command is not valid from the APL session screen.

cannot replace fn with var or var with fn
: The ***write*** command was issued to write an object as a variable (*var*) when it already exists as a function (*fn*) in the active workspace (or vice versa). Use the ***session*** command to move back to the APL session to erase the object, then the ***forward*** command to get back to the original screen to retry the ***write*** command.

defn error
: An error in the structure of the function prevents the function from being defined in the workspace.

domain error
: The object of an editor command is not of an allowable data type (numeric, boxed, and mixed arrays cannot be edited).

help file not found
: In the UNIX version, no help file is provided.

improper name
: An ***edit, get, forward,*** or ***write*** command was issued with an improperly formed variable or function name.

incorrect option to quit command
: The ***quit*** command was used with an option other than *all*.

length error
: The ***edit*** or ***get*** command was issued for a table that has a last dimension greater than `1014` or for a vector that has more than `1014` characters between newline characters.

line too long
: The ***insert*** or ***replace*** command issued would have resulted in a line more than `1014` characters long.

name missing
: An ***edit*** or ***get*** command was issued without specifying the name of the variable or function to be edited.

no previous change
: A ***repeat last change*** command was issued when there is no last change to repeat.

not enough memory available
: The editor has run out of memory and cannot perform the requested operation. You can delete an unneeded logical screen from the edit ring by moving to it and using the ***quit*** command. If this problem occurs in the

delete lines command, it is due to the fact that the deleted lines go into the save buffer. Try deleting fewer lines at once. The memory allocated to the editor can be increased or decreased using the `-M` option when you start your APL session (see the *Handbook, Chapter 3*).

not found
The object of the ***search*** command was not found.

object already exists
The ***write*** command was used to rename an object, but an object with the new name already exists in the active workspace. Use another name, or return to the session with the ***session*** command, erase the object with the conflicting name, and return to perform the rename using the ***forward*** command.

only in immediate execution
During □- or ⍞-input, an editor command was issued that would move to a new logical screen, an action allowed only from immediate execution.

rank error
The ***edit*** or ***get*** command was used to edit an array with more than two dimensions.

si damage: write ignored
The ***write*** command was issued for a function that has been altered in such a way as to affect localization of names or line labels. Use the ***session*** command to move back to the session to clear the suspension (using → (naked branch), *)sic*, or *□signal*). Then use the ***forward*** command to get back to the screen to retry the ***write*** command.

unknown option
An unknown option was specified in a command.

unknown write option
An unidentified option or one invalid for ***write*** was specified in a ***write*** command.

unnamed object, supply name
A ***write*** command was issued for the APL session screen without supplying a name.

8
Public Workspaces

For easy access, commonly used programs may be placed in workspaces that are saved in one or more *public libraries*. In SHARP APL for OS/390, a public library is associated with an account number that falls within the range 1 to 499. In SHARP APL for UNIX, your system administrator can use the startup parameter `-L` to associate up to 15 library numbers with full UNIX paths (see the *Handbook, Chapter 3*). As distributed, SHARP APL for UNIX includes startup parameters to define only ***library*** 1 as public.

Library `1`

The following public workspaces are currently distributed with SHARP APL for UNIX:

`1 ap124`	AP124 standard utility functions.
`1 del`	AP124 based full-screen editor.
`1 hostap`	Utilities to support AP11 commands.
`1 hsp`	SHARP APL for OS/390 HSPRINT facility emulator.
`1 if`	Intrinsic function documentation and system binds.
`1 pc108`	Up-loading and down-loading tools for use with the PC108 terminal emulator for IBM PCs.
`1 socket`	Cover functions for the Socket Intrinsic Function facility.
`1 ssql`	Provides access to the Sharp APL for OS/390 SSQL Auxiliary processor, AP127.

1 *toolkit*	Complimentary set of ISO standard APL utilities (distributed by the Toronto chapter of ACM SIGAPL).
1 *tools*	Various tools for SHARP APL for UNIX.
1 *unix*	Provides access to numerous C library and system calls available in UNIX.
1 *wssearch*	Utilities to search for and replace strings in APL programs and text strings.
1 *wstransfer*	Allows transfer of APL workspaces and objects between different APL systems using the 'Workspace Interchange Convention Version 0'.
1 *xfer*	Converts SHARP APL for UNIX workspaces and files to UNIX text.

9
Event Handling

When an event (an error, interrupt, or signal) occurs, the interpreter identifies it by its ***event number***. It records the event by setting the visible value of the variable □*er* (***event report***).

You may (in advance) set the variable □*trap* to tell the interpreter what it should do when a trappable event occurs. The variable □*trap* contains a set of trap definitions, one for each event or group of events you wish your programs to handle. Each trap definition consists of a list of event numbers, an action code, and a recovery expression. The action code is one of six broad classes of action. The recovery expression is a line of APL to be executed as part of the handling of the event.

Once a nonempty □*trap* is set, nothing happens in the event trapping arena until an event occurs. The event may be any error, interrupt signal, or return to immediate execution. At that point the interpreter looks in the various local and global values of □*trap* for an event number that matches that of the event that has occurred. The interpreter does not confine itself to the visible □*trap* but, if not satisfied in its earlier searches, continues to examine shadowed copies of □*trap* lower on the state indicator.

The system function □*signal* lets a user-defined function signal events which trap definitions may handle appropriately.

The system variable □*ec* provides control over sensitive code by conditioning the environment to respond for user-defined functions as it does for primitive functions.

Events

An *event* is a signal that the interpreter either receives as a result of its own operations or has sent to it from outside. The interpreter records the event by setting the visible value of the variable ⎕*er*.

Errors

The interpreter signals an *error* when it finds that the conditions for executing part of an APL expression are not met. An error may arise from an ill-formed statement, a misspelled name, improper syntax, or an argument outside the function's domain. An error may also arise when there is nothing wrong with the sentence being executed, but there are insufficient resources to carry it out (for example, insufficient memory or file space). As each error occurs, the interpreter categorizes it by attaching to it an ***event number*** from a standard list (see Table 9.1). The first four characters of the ***event report***, which the interpreter writes in ⎕*er*, are the `4 0⍕` of the event number (i.e., four characters, right-justified number, padded with blanks).

Interrupts

An *interrupt* is a signal transmitted to the interpreter from outside the active workspace, usually (but not necessarily) from your keyboard or from the shared name that controls an S-task. An interrupt causes a halt to the interpreter's execution.

In processing an interrupt, the system recognizes two levels of urgency. The first time you send an interrupt signal, the system notes a ***weak interrupt***. It takes that to mean ***halt before starting a new line***. During the interval after you have sent the first interrupt signal (which the system recognizes as a weak interrupt) but before the interpreter has halted work, you may signal a second interrupt. When the system detects a second interrupt, it upgrades the urgency of your request and calls it a ***strong interrupt***. Following a strong interrupt, the interpreter halts as quickly as possible. It is likely that the halt resulting from a strong interrupt occurs in the middle of a line, with some portion executed and some not.

An interrupt is classified by what the interpreter was doing at the time it halted execution. For example, if the interpreter was waiting to access a shared name, it calls the interrupt a ***shared variable interrupt***; if it was processing input from the terminal or controlling shared name, it calls the interrupt an ***input interrupt***; if it

was waiting for a return from the file system, it calls the interrupt a ***file interrupt***, and so on. The term ***interrupt*** (without a qualifying adjective) refers to an interruption during execution of an APL expression that did not involve something external to the workspace.

Signalled Errors

In addition to external events that arise in the ways just described, execution can also be interrupted by a ***signal***, which is produced by the system function ⎕*signal*. Like an error encountered in a primitive function, ⎕*signal* aborts execution of the function that contains it and halts the interpreter with the error number specified in its right argument. The error number may be one of those that the interpreter already recognizes from its standard list. In that case, the interpreter responds just as it would if the error condition had actually occurred. Alternatively, ⎕*signal* may signal an arbitrary number (greater than `0` and less than `1000`) that may mean nothing to the interpreter but may be understood in the context of the application in which it occurs. An event signalled by ⎕*signal* may not be one of the events classified as interrupts.

Unreportable or Untrappable Events

Certain events are outside the scope of the event handling mechanism. In particular, nothing that occurs during execution of a system command can generate a reportable or trappable event; nor can certain actions of the operating system, such as the termination of the APL task by a UNIX `kill` command.

An event encountered during execution of a recovery expression in a ⎕*trap* definition is not trappable. This restriction prevents endless recursion from a defective ⎕*trap*. It also suggests that you should keep your recovery expression short—branch to a line in a user-defined function to perform the recovery you wish, so that a subsequent event may also be trapped in a controlled fashion. An important consideration here is whether or not to redefine (perhaps to empty) ⎕*trap* itself as part of the recovery expression.

Event `2001` (return to immediate execution) is not reportable—it does not set a new value for ⎕*er*, even though it is a trappable event.

An event arising during execution of a system command is never reported in ⎕*er* and cannot be trapped.

Table 9.1. Event numbers.

Errors		*Errors (con'd)*	
0	**Any Error** *	46	*ws not found*
1	*ws full*	47	*ws not readable*
2	*syntax error*	70	*ws too large*
3	*index error*	72	*interface quota exhausted*
4	*rank error*	73	*no shares*
5	*length error*	74	*interface capacity exceeded*
6	*value error*	75	*share table full*
7	*format error*	76	*processor table full*
8	*result error*	77	*identification in use*
10	*association error*	86	*improper library reference*
11	*domain error*	95	*ws not writable*
15	*symbol table full*	96	*remote no shares*
16	*nonce error*	97	*remote share table full*
18	*file tie error*		
19	*file access error*	*Interrupts*	
20	*file index error*	1000	**Any interrupt** **
21	*file full*	1001	*stop*
22	*file name error*	1002	*attention*
24	*file tied*	1003	*interrupt*
25	*file system hardware error*	1004	*input interrupt*
28	*file system not available*	1005	*sv interrupt*
29	*file library not available*	1006	*file interrupt*
30	*file system ties used up*		
31	*file tie quota used up*	*Workspace State*	
33	*file reservation error*	2001	*return to immediate execution* ***
34	*file system no space*		
38	*file component damaged*		

* 0 ***in*** □*trap* ***stands for errors*** 1 ***to*** 999***;*** 0 ***never occurs in*** □*er*.
** 1000 ***in*** □*trap* ***stands for interrupts*** 1001 ***to*** 1999***;*** 1000 ***never occurs in*** □*er*.
*** ***Not reported in*** □*er*.

□*er* Event Report

When any reportable error or interrupt occurs, the interpreter sets the visible value of the system variable □*er* to a character matrix containing a description of the event. A reportable event is any of the events included in Table 9.1, or any event generated by executing □*signal*.

The value of □*er* is, in general, a 3-row character matrix. However, following event `6 value error`, the interpreter may also include a fourth row in special circumstances.

In a workspace in which no space remains to set □*er* in the normal manner, the interpreter sets it as a 1-by-4 character matrix containing only `4 0⍕` of the event number (i.e., four characters, right-justified number, padded with blanks). This short form of □*er* is assured even when the reported event is `ws full`.

The rows of □*er* are as follows:

`0@ □er` The first four characters represent the event number (flush right). Next there is one blank. Then follows the event text as specified in the left argument of □*signal* or from the interpreter's list of standard events and their names. The text may be blank if □*signal* provided no title and the event is one that is not on the interpreter's standard list.

`1@ □er` Text of the line being executed when the event occurred. When the event occurred during execution of the recovery expression of □*trap*, instead of a reference to the function's name and line number, the second row of □*er* contains the characters □*trap* followed by the recovery expression.

`2@ □er` Two carets (which may be superimposed and appear to be a single caret mark) point to the part of the line on which the interpreter was working at the time the event occurred.

`3@ □er` Only when event `6 value error` is trapped by action *i*, a fourth line is included and contains the undefined name.

When the event described in □*er* occurred during execution of a locked function, the system withholds some of the information that it would otherwise include. The second row of □*er* contains the function's name and the line number, followed by the character ⍒. The caret on the third row points to the ⍒ symbol.

When the workspace lacks space to contain a complete □*er*, the system uses a terse form containing only the four characters of the event number as a 1-by-4 matrix.

□*ec* Environment Condition

By setting the system variable □*ec*, you can prevent certain defined functions (or portions of them) from being interrupted. The variable □*ec* may be set only to a numeric scalar, 0 or 1. While the visible value of □*ec* is 1, the effect of an external interrupt is deferred until □*ec* is no longer 1. The effect of an untrapped error while □*ec* is 1 is to abort execution, in the same way as □*signal*. Thus, regardless of the internal errors or external interrupts, the function in which the event occurred cannot be suspended. Under these conditions, a function's execution can be interrupted only by a previously set □*stop*. Note that □*stop* has effect only for a function whose definition is not locked, so that, for a locked definition, even that avenue is not available. Moreover, in a function to which □*ec* is local, or in any situation requiring a value for □*ec* before one has been set, the interpreter supplies the value 1.

This special property of □*ec* makes it impossible to receive an interrupt that occurs after the start of execution of a function that localizes □*ec* but before a value has been set. This makes for consistent handling of the function's behavior, with no possibility of a race between an external interrupt and the setting of □*ec*.

□*trap* Trap Definition

The variable □*trap* contains the definitions of actions to be taken when an event occurs.

When you attempt to set a value in □*trap*, the interpreter first validates your setting before it allows your trap definitions to be assigned. If you submit a properly formed set of trap definitions (with proper action codes properly placed in each definition, etc.), the system acknowledges this by making the contents of □*trap* the same as the variable you have assigned to it. You may set □*trap* to any value; however, if the variable you assign is not valid, the interpreter sets □*trap* to an empty character vector. The value of □*trap* may be either a character matrix or a character vector.

- When it is a matrix, it has one row for each trap definition.

- When it is a vector, the first character serves as a delimiter, and subsequent uses of the delimiter serve to separate the various definitions. The delimiter may *not* be a numeral, a blank, or the character −, which has a special use in the event numbers field of a trap definition.

A trap definition in □*trap* may consist of ***four*** fields, as in the following expression:

```
□trap ← '∇ 11 ⍝ ⍝ →□lc'
```

This example is marked by underlines to show the division into fields: event numbers, qualifier, action, and recovery expression. Use at least one blank to separate the fields. A blank is not required to separate the leading delimiter for each definition in the list form of □*trap* from the definition itself.

A trap definition in □*trap* may also consist of ***three*** fields where the ***qualifier field is omitted***, as in the following expression:

```
□trap ← '∇ 500-504 ⍝ →L3'
```

Descriptions of the fields are explained below.

Event Numbers Field

Event numbers are indicated by numerals. Successive numbers are separated by blanks; for example, `2 4 5`. The order in which events are listed within a trap definition is not significant. A hyphen may be used to indicate all events that fall within a range of numbers. For example, *n1-n2* means, ``All events in the range *n1* to *n2*, inclusive."

- A `0` means any event less than `1000` and has the same effect as `1-999`.
- A `1000` means any event greater than `1000` and less than `2000` and has the same effect as `1001-1999`.

Where no numbers are stated, the trap definition applies to all events for that action.

Qualifier Field

The optional letter *o* indicates that the trap is to be operative only at the function's own level: in the function to which □*trap* is local, but not in any functions invoked within it. The *o* action may be needed because, when functions invoked within the current one encounter an event and fail to trap it, this □*trap* may take effect even when the more local function localizes □*trap*.

Note: Separate the qualifier *o* from the surrounding trap definition by at least one space.

Action Field

The action is indicated by a single letter, as follows:

s **Stop**. Preempts other traps, if any.

n **Next**. Exempts the current trap; leaves more global ones in effect.

e **Execute**. Executes the recovery expression at the level of the state indicator at which the event occurs.

c **Cut and execute**. Cuts back the state indicator to the level of localization of this □*trap*, then executes the recovery expression.

i **Immediate**. Resumes the expression in which the event occurred in mid-line after the recovery expression supplies a definition for the missing name detected as event `6 value error`, the only event trapped with action *i*.

d **Do**. Allows one of four measures to be taken when event `2001`, which indicates the return to immediate execution, occurs.

Recovery Expression Field

The recovery expression is a line of APL in the same form as you might type from the keyboard or assign to □*lx*. In particular, with actions *e* and *c*, if you want execution of the suspended function to be resumed, you must include the branch arrow →. However, with action *i* (used to supply the definition of an undefined name), execution resumes at once when the name is defined, and the recovery

expression should not include a branch arrow. If the recovery expression itself contains an error, that error cannot be trapped, and execution halts. The trap actions and their associated recovery expressions follow.

Action s Stop

This action stops execution when an event occurs. It takes no recovery expression. In effect, it restores the default action of the interpreter, and is often used during debugging. Note that setting `⎕trap←'∘ s'` is stronger than `⎕trap←''` because making the visible `⎕trap` empty does not remove the action of other traps that may be shadowed. For example,

```
⎕trap←'∇ 0 s'
```

sets the environment so that no errors will be trapped, but interrupts are still subject to traps set in any global or shadowed values of `⎕trap`.

Action n Next

The action *n* instructs the interpreter to look no further in the present `⎕trap`, but to move on to the next (more global) `⎕trap`. It has the effect of nullifying trap definitions that follow it in its own `⎕trap`. Like *s*, action *n* takes no recovery expression. For example,

```
⎕trap←'∇ 500-509 n ∇ 0 e logerror'
```

sets the environment so that, for errors in the range `500` through `509`, the search for an appropriate trap is occurs in more global values of `⎕trap`; the second part of the definition handles all other errors.

Action i Immediate

Action *i* is usable only with event `6 value error`. It is designed to supply the value of the undefined name and resume execution in mid-line. Action *i* is the only action that permits resumption of an interrupted line at a point other than the beginning. Action *i* requires a recovery expression, which may not include a branch statement. It should do something to remedy the fact that an undefined name occurs in the sentence being executed (for example, read its definition from a file). For example,

```
⎕trap←'∇ 6 i page 3@ ⎕er'
```

relies on the definition of the function *page* to define the object named in the fourth row of □*er*. The function *page* can take advantage of any definition method appropriate to the application at hand, such as using □*pdef* on a package read from a component file. Such a paging function can also maintain a list of the objects it defines for subsequent erasure, or for monitoring patterns of system usage.

When the recovery expression has been executed, the interpreter tries to resume execution of the sentence it was executing when it found an undefined name. If, at that point, the missing name remains undefined, a true *value error* not subject to □*trap* is signalled.

If the recovery expression is unable to define the name but does not wish to incur the value error, it may explicitly □*signal* an error, which is trappable.

Action e Execute

The interpreter executes the recovery expression in the current environment. To resume work after completing the remedial action, the recovery expression should include a sentence such as →□*lc*. For example,

```
□trap←'∇ 3 e →□lc⊣i←0'
```

causes the recovery expression →□*lc*⊣*i*←0 if an *index error* occurs. The recovery expression sets the variable *i* and branches to the line in the function in which the error occurred.

Action c Cut Back and Execute

The interpreter aborts the execution of any functions invoked by the function that localized □*trap* (thereby cutting back the state indicator so the function to which □*trap* is local is at the top), and then executes the recovery expression.

Action *c* differs from action *e* only when the operative trap is more global than the function being executed. Especially where the recovery expression uses a branch to resume execution, it is essential to use action *c* rather than *e* to avoid the possibility that the branch instruction will be executed in the called function, rather than in the function that set the trap. For example,

```
□trap←'∇ 3 c →L20⊣i←0'
```

causes the recovery expression →*L*20⊣*i*←0 to be executed after the state indicator is cut back to the level of the user-function in which this trap definition is local.

Action d Do

This action invokes one of several options which are not directly attainable by executing an APL sentence (and hence not obtainable from actions *e* or *c*). In addition, action *d* is the only action permitted in response to event 2001 (return to immediate execution), which is not otherwise a trappable event. The recovery expression for action *d* may be empty, or may contain one of the following keywords:

exit *clear* *off*

These actions have the following effects:

d	Action *d* with no keyword has no effect, other than to handle the event and thus prevent more global ⎕*trap*s from being searched.
d exit	Causes execution to exit from the function that contains the ⎕*trap*, and return to the environment in which that function was invoked.
d clear	Clears the workspace. The former contents are lost.
d off	Clears the workspace and terminates the APL session. The action *d off* terminates the APL task without saving a *continue* workspace, whereas ⎕*bounce* saves *continue*, or another termination workspace name, if one was specified by ⎕*twsid*.

For example,

```
⎕trap←'∇ 2001 d clear'
```

would cause the active workspace to be cleared when the task returns to immediate-execution mode.

⎕*signal* Signal Event

The right argument of ⎕*signal* is a an integer scalar or vector of integers greater than `0` and less than `1000`. When the vector is not empty, execution of the function at the top of the state indicator is aborted, and the event indicated by the value of the first element in ⍵ is signalled in the sentence that invoked the function thus aborted. This makes a user-defined function behave in the same way as a primitive function when it encounters an error.

You can use ⎕*signal* conditionally by taking advantage of the fact that when the right argument vector is empty, no signal occurs. In the following example, the function *init* contains the line `'testing' ⎕signal (debug=1)/501`.

```
     start    ⍝ debug is 1
testing
start[1] init
        ^
     ⎕er
501 testing
start[1] init
        ^
```

The left argument of ⎕*signal* is optional. If present, it becomes the error text found in the first row of ⎕*er*. If omitted, the error text in ⎕*er* depends upon whether the signalled event is a standard SHARP APL error, in which case the standard error text appears, or open for use specific to an application, in which case no text appears unless the application places it there.

```
     ⎕signal 18
file tie error
     ⎕signal 18
     ^
```

If the event number that ⎕*signal* issues is recognized in a ⎕*trap* visible in the environment from which the aborted function was called, the interpreter takes the action indicated in the trap; otherwise it simply halts as it would if no ⎕*trap* were in effect.

A user-defined function can signal to itself by using the function ⍎ (execute) to evaluate an expression containing ⎕*signal*: any invocation of ⍎ introduces a new level onto the state indicator, and abortion of that level of the state indicator forces a halt in the function that used ⍎.

```
     ∇function[3⎕]∇
[3] ⍎(0=⎕nc 'n')/'''oops'' ⎕signal 511'

    function
oops
function[3] ⍎(0=⎕nc 'n')/'''oops'' ⎕signal 511'
       ^
```