dyalog *APL*/W

Dyalog *APL*™ for Windows

Language Reference

Dyadic Systems Ltd

Riverside View
Basing Road, Old Basing
Basingstoke
Hampshire, RG24 7AL
United Kingdom

tel: +44 (0)1256 811125
fax: +44 (0)1256 811130
email: support@dyadic.com
http://www.dyadic.com

Overview

Contents

Contents

CHAPTER 4   Primitive Operators

C H A P T E R   1

Introduction

Workspaces

APL expressions are evaluated within a workspace. The workspace may contain objects, namely operators, functions and variables defined by the user. APL expressions may include references to operators, functions and variables provided by APL. These objects do not reside in the workspace, but space is required for the actual process of evaluation to accommodate temporary data. During execution, APL records the state of execution through the STATE INDICATOR which is dynamically maintained until the process is complete. Space is also required to identify objects in the workspace in the SYMBOL TABLE. Maintenance of the symbol table is entirely dynamic. It grows and contracts according to the current workspace contents.

Workspaces may be explicitly saved with an identifying name. The workspace may subsequently be loaded, or objects may be selectively copied from a saved workspace into the current workspace.

Under UNIX, workspace names must be valid file names, but are otherwise unrestricted. See your UNIX documentation for details.

Under Windows, Dyalog APL workspaces are stored in files with the suffix ".DWS". However, they are referred to from within APL by only the first part of the file name which must conform to Windows file naming rules.

Namespaces

A namespace is a name class 9 object in Dyalog APL. Namespaces are analogous to nested workspaces.

'Flat' APL Workspace   **Workspace with Namespaces**

```
DISPLAY


FOO MAT VEC


WsDoc_Init
WsDoc_Tree
WsDoc_Xref


WsDoc_prt_Init
WsDoc_prt_Page
WsDoc_current_line
```

```
FOO MAT VEC

  DISPLAY



  Init
  Tree   Init
  Xref   Page      line
```

They provide the same sort of facility for workspaces as directories do for filing systems. The analogy might prove helpful:

Operation	MS-DOS	Namespace
Create	MKDIR	)*NS* or □*NS*
Change	CD	)*CS*
Relative name	DIR1\DIR2\FILE	*NS1.NS2.OBJ*
Absolute name	\DIR\FILE	*#.NS.OBJ*
Name separator	\	*.*
Top (Root) object	\	*#*
Parent object	..	*##*

Major Benefits of Namespaces

Namespaces provide static (as opposed to dynamic) local names. This means that a defined function can use local variables and functions which persist when it exits and which are available next time it is called.

Just as with the provision of directories in a filing system, namespaces allow us to organise the workspace in a tidy fashion. This helps to promote an object oriented programming style.

APL's traditional name-clash problem is ameliorated in several ways.

- Workspaces can be arranged so that there are many fewer names at each namespace level. This means that when copying objects from saved workspaces there is a much reduced chance of a clash with existing names.

- Utility functions in a saved workspace may be coded as a single namespace and therefore on being copied into the active workspace consume only a single name. This avoids the complexity and expense of a solution which is sometimes used in 'flat' workspaces, where such utilities dynamically fix local functions on each call.

- In flat APL, workspace administration functions such as $WSDOC$ must share names with their subject namespace. This leads to techniques for trying to avoid name clashes such as using obscure name prefixes like $'\Delta\Delta L1'$. This problem is now virtually eliminated because such a utility can operate exclusively in its own namespace.

The programming of GUI objects is considerably simplified.

- An object's callback functions may be localised in the namespace of the object itself.

- Static variables used by callback functions to maintain information between calls may be localised within the object.

This means that the object need use only a single name in its namespace.

Arrays

A Dyalog APL data structure is called an array. An array is a rectangular arrangement of items, each of which may be a single number, a single character, a namespace reference (ref), another array, or the $\square OR$ of an object. An array which is part of another array is also known as a subarray.

An array has two properties; structure and data type. Structure is identified by rank, shape, and depth.

Rank

An array may have 0 or more axes or dimensions. The number of axes of an array is known as its rank. Dyalog APL supports arrays with a maximum of 15 axes.

- An array with 0 axes (rank 0) is called a scalar.

- An array with 1 axis (rank 1) is called a vector.

- An array with 2 axes (rank 2) is called a matrix.

- An array with more than 2 axes is called a multi-dimensional array.

Shape

Each axis of an array may contain zero or more items. The number of items along each axis of an array is called its shape. The shape of an array is itself a vector. Its first item is the length of the first axis, its second item the length of the second axis, and so on. An array whose length along one or more axes is zero, is called an empty array.

Depth

An array whose items are all simple scalars (i.e. single numbers, characters or refs) is called a simple array. If one or more items of an array is not a simple scalar (i.e. is another array, or a $\square OR$), the array is called a nested array. A nested array may contain items which are themselves nested arrays. The degree of nesting of an array is called its depth. A simple scalar has a depth of 0. A simple vector, matrix, or multi-dimensional array has depth 1. An array whose items are all depth 1 subarrays has depth 2; one whose items are all depth 2 subarrays has depth 3, and so forth.

Type

An array whose elements are all numeric, is called a numeric array; its TYPE is numeric. A character array is one in which all items are characters. An array whose items contain both numeric and character elements is of MIXED type.

Numbers

Numbers are entered or displayed in conventional or scaled decimal form. On entry, a decimal point is optional if there is no fractional part. On output, a number with no fractional part (an integer) is displayed without a decimal point.

The scaled form consists of:

a) an integer or decimal number called the mantissa,

b) the letter E or e,

c) an integer called the scale, or exponent.

The scale specifies the power of 10 by which the mantissa is to be multiplied.

Example

```
      12 23.24 23.0 2.145E2
12 23.24 23 214.5
```

Negative numbers are preceded by the high minus (⁻) symbol, not to be confused with the minus (-) function. In scaled form, both the mantissa and the scale may be negative.

Example

```
      ¯22 2.145E¯2 ¯10.25
¯22 0.02145 ¯10.25
```

The empty vector (ι0) may be represented by the numeric constant θ called ZILDE.

Characters

Characters are entered within a pair of APL quotes. The surrounding APL quotes are
not displayed on output. The APL quote character itself must be entered as a pair of
APL quotes.

Examples

```
      'DYALOG APL'
DYALOG APL

      'I DON''T KNOW'
I DON'T KNOW

      '*'
*
```

Enclosed Elements

An array may be enclosed to form a scalar element through any of the following
means:

- by the enclose function (⊂)

- by inclusion in vector notation

- as the result of certain functions when applied to arrays

Examples

```
      (⊂1 2 3),⊂'ABC'
 1 2 3  ABC

      (1 2 3) 'ABC'
 1 2 3  ABC

      ι2 3
 1 1  1 2  1 3
 2 1  2 2  2 3
```

Legal Names

APL objects may be given names. A name may be any sequence of characters, starting with an alphabetic character, selected from the following:

1. *A* to *Z*
2. *a* to *z*
3. <u>*A*</u> to <u>*Z*</u>   (in Dyalog Standard Font)
   or *ÂÃÇÈÊËÌÍÎÏÐÒÓÔÕÙÚÛÝÞ ā ì ð ò ō*   (in Dyalog Alternate Font)
4. _ or ∆ or <u>∆</u>   (underbar, delta, delta underbar characters)
5. National Language Symbols
   (viz. *ÀÄÅÆÉÑÖØÜ ß à á â ä å æ ç è é ê ë í î ï ñ ó ô ö ø ù ú û ü*)
6. *0* to *9*

The first five groups of characters are designated as being alphabetic.

Examples

Legal	Illegal
THIS∆IS∆A∆NAME	*BAD NAME*
X123	*3+21*
SALES	*S!H|PRICE*
pjƀ_1	*1_pjƀ*

Specification of Variables

A variable is a named array. An undefined name or an existing variable may be assigned an array by specification with the left arrow (←).

Examples

```
      A←'CHIPS WITH EVERYTHING'

      A
CHIPS WITH EVERYTHING

      X Y←'ONE' 'TWO'

      X
ONE
      Y
TWO
```

Vector Notation

A series of two or more adjacent expressions results in a vector whose elements are the enclosed arrays resulting from each expression. This is known as VECTOR (or STRAND) NOTATION. Each expression in the series may consist of one of the following:

a)   a single numeric value;

b)   a single character, within a pair of quotes;

c)   more than one character, within a pair of quotes;

d)   the name of a variable;

e)   the evaluated input symbol [];

f)   the quote-quad symbol [];

g)   the name of a niladic, defined function yielding a result;

h)   any other APL expression which yields a result, within parentheses.

Examples

```
      ρA←2 4 10
3


      ρTEXT←'ONE' 'TWO'
2
```

Numbers and characters may be mixed:

```
      ρX←'THE ANSWER IS ' 10
2

      X[1]
 THE ANSWER IS

      X[2] + 32
42
```

Blanks, quotes or parentheses must separate adjacent items in vector notation. Redundant blanks and parentheses are permitted. In this manual, the symbol pair '←→' indicates the phrase 'is equivalent to'.

```
        1   2   ↔  (1)(2) ↔  1   (2)   ↔  (1)   2
2'X'3 ↔  2 'X' 3 ↔  (2) ('X') (3)
        1  (2+2) ↔  (1) ((2+2)) ↔  ((1))  (2+2)
```

Vector notation may be used to define an item in vector notation:

```
      ρX ← 1 (2 3 4) ('THIS' 'AND' 'THAT')
3

      X[2]
 2 3 4

      X[3]
  THIS  AND  THAT
```

Expressions within parentheses are evaluated to produce an item in the vector:

```
      Y ← (2+2) 'IS' 4
      Y
4  IS  4
```

The following identity holds:

```
  A   B   C   ↔  (⊂A), (⊂B),  ⊂C
```

Restructuring Arrays

A class of primitive functions re-structures arrays in some way. Arrays may be input only in scalar or vector form. Structural functions may produce arrays with a higher rank. The Structural functions are reshape (ρ), ravel, laminate and catenate (,), reversal and rotation (⌽), transpose (⍉), mix and take (↑), split and drop (↓), and enclose (⊂). These functions are described in *Chapter 3*.

Examples

```
      2 2ρ1 2 3 4
1 2
3 4

      2 2 4ρ'ABCDEFGHIJKLMNOP'
ABCD
EFGH

IJKL
MNOP
```

```
      ↓2 4ρ'COWSHENS'
COWS  HENS
```

Display of Arrays

Simple scalars and vectors are displayed in a single line beginning at the left margin. A number is separated from the next adjacent element by a single space. The number of significant digits to be printed is determined by the system variable $\square PP$ whose default value is 10. The fractional part of the number will be rounded in the last digit if it cannot be represented within the print precision. Trailing zeros after a decimal point and leading zeros will not be printed. An integer number will display without a decimal point.

Examples

```
      0.1 1.0 1.12
0.1 1 1.12

      'A' 2 'B' 'C'
A 2 BC

      ÷3 2 6
0.3333333333 0.5 0.1666666667
```

If a number cannot be fully represented in $\square PP$ significant digits, or if the number requires more than five leading zeros after the decimal point, the number is represented in scaled form. The mantissa will display up to $\square PP$ significant digits, but trailing zeros will not be displayed.

Examples

```
      □PP←3

      123 1234 12345 0.12345 0.00012345 0.00000012345
123 1.23E3 1.23E4 0.123 0.000123 1.23E¯7
```

Simple matrices are displayed in rectangular form, with one line per matrix row. All elements in a given column are displayed in the same format, but the format and width for each column is determined independently of other columns. A column is treated as numeric if it contains any numeric elements. The width of a numeric column is determined such that the decimal points (if any) are aligned; that the E characters for scaled formats are aligned, with trailing zeros added to the mantissae if necessary, and that integer forms are right-adjusted one place to the left of the decimal point column (if any). Numeric columns are right-justified; a column which contains no numeric elements is left-justified. Numeric columns are separated from their neighbours by a single column of blanks.

Examples

```
      2 4ρ'HANDFIST'
HAND
FIST

      1 2 3 ∘.× 6 2 5
 6  2  5
12  4 10
18  6 15

      2 3ρ2 4 6.1 8 10.24 12
2  4      6.1
8 10.24 12

      2 4ρ4 'A' 'B' 5 ¯0.000000003 'C' 'D' 123.56
 4E0   AB    5
¯3E¯9 CD 123.56
```

In the display of non-simple arrays, each element is displayed within a rectangle such
that the rows and columns of the array are aligned. Simple items within the array are
displayed as above. For non-simple items, this rule is applied recursively, with one
space added on each side of the enclosed element for each level of nesting.

Examples

```
      ι3
1 2 3

      ⊂ι3
 1 2 3

      ⊂⊂ι3
  1 2 3

      ('ONE' 1) ('TWO' 2) ('THREE' 3) ('FOUR' 4)
  ONE  1    TWO  2    THREE  3    FOUR  4

      2 4ρ'ONE' 1 'TWO' 2 'THREE' 3 'FOUR' 4
 ONE    1 TWO    2
THREE   3 FOUR   4
```

Multi-dimensional arrays are displayed in rectangular planes. Planes are separated by
one blank line, and hyper-planes of higher dimensions are separated by increasing
numbers of blank lines. In all other respects, multi-dimensional arrays are displayed in
the same manner as matrices.

Examples

```
      2 3 4ρι24
 1  2  3  4
 5  6  7  8
 9 10 11 12

13 14 15 16
17 18 19 20
21 22 23 24


      3 1 1 3ρ'THEREDFOX'
THE


RED


FOX
```

The power of this form of display is made apparent when formatting reports.

Examples

```
      +AREAS←'West' 'Central' 'East'
West  Central  East

      +PRODUCTS←'Biscuits' 'Cakes' 'Buns' 'Rolls'
Biscuits  Cakes  Buns  Rolls

      SALES←50 5.25 75 250 20.15 900 500
      SALES,←80.98 650 1000 90.03 1200
      +SALES←4 3ρSALES
  50  5.25   75
 250 20.15  900
 500 80.98  650
1000 90.03 1200

      ' ' PRODUCTS ⍪., AREAS SALES
          West  Central  East
Biscuits    50     5.25    75
Cakes      250    20.15   900
Buns       500    80.98   650
Rolls     1000    90.03  1200
```

If the display of an array is wider than the page width, as set by the system variable
⎕PW, it will be folded at or before ⎕PW and the folded portions indented six spaces.

The display of a simple numeric or mixed array may be folded at a width less than `⎕PW` so that individual numbers are not split across a page boundary.

Example

```
      ⎕PW←40

      ?3 20⍴100
54 22  5 68 68 94 39 52 84  4  6 53 68
85 53 10 66 42 71 92 77 27  5 74 33 64
66  8 64 89 28 44 77 48 24 28 36 17 49
        1  39  7 42 69 49 94
       76 100 37 25 99 73 76
       90  91  7 91 51 52 32
```

The Display Function

DISPLAY is a defined function distributed with Dyalog APL which may be used to illustrate the structure of an array. *DISPLAY* is monadic. Its result is a character matrix containing a pictorial representation of its argument. *DISPLAY* is used throughout this manual to illustrate examples. An array is illustrated with a series of boxes bordering each sub-array. Characters embedded in the border indicate rank and type information. The top and left borders contain symbols that indicate its rank. A symbol in the lower border indicates type. The symbols are defined as follows:

→	Vector.
↓	Matrix or higher rank array.
⊖	Empty along last axis.
⌽	Empty along other than last axis.
∊	Nested array.
~	Numeric data.
–	Character data.
+	Mixed character and numeric data.
∇	⎕*OR* object.
#	array of refs.

```
      DISPLAY 'ABC' (1 4⍴1 2 3 4)
.→-----------------.
| .→--.   .→------. |
| |ABC|  ↓1 2 3 4| |
| '---'  '~------' |
'∊-----------------'
```

Expressions

An expression is a sequence of one or more syntactic tokens which may be symbols or constants or names representing arrays (variables) or functions. An expression which produces an array is called an ARRAY EXPRESSION. An expression which produces a function is called a FUNCTION EXPRESSION. Some expressions do not produce a result.

An expression may be enclosed within parentheses.

Evaluation of an expression proceeds from right to left, unless modified by parentheses. If an entire expression results in an array that is not assigned to a name, then that array value is displayed. (Some system functions and defined functions return an array result only if the result is assigned to a name or if the result is the argument of a function or operator.)

Examples

```
      X←2×3-1

      2×3-1
4
      (2×3)-1
5
```

Either blanks or parentheses are required to separate constants, the names of variables, and the names of defined functions which are adjacent. Excessive blanks or sets of parentheses are redundant, but permitted. If *F* is a function, then:

```
      F 2↔ F(2)  ↔  (F)2  ↔  (F) (2)  ↔  F  (2) ↔ F ((2))
```

Blanks or parentheses are not needed to separate primitive functions from names or constants, but they are permitted:

```
      -2  ↔  (-)(2)  ↔  (-) 2
```

Blanks or parentheses are not needed to separate operators from primitive functions, names or constants. They are permitted with the single exception that a dyadic operator must have its right argument available when encountered. The following syntactical forms are accepted:

```
      (+.×)  ↔  (+).×  ↔  +.(×)
```

The use of parentheses in the following examples is not accepted:

```
      +(.)×   or       (+.)×
```

Functions

A function is an operation which is performed on zero, one or two array arguments and may produce an array result. Three forms are permitted:

- NILADIC defined for no arguments

- MONADIC defined for a right but not a left argument

- DYADIC defined for a left and a right argument

The number of arguments is referred to as its VALENCE.

The name of a non-niladic function is AMBIVALENT; that is, it potentially represents both a monadic and a dyadic function, though it might not be defined for both. The usage in an expression is determined by syntactical context. If the usage is not defined an error results.

Functions have long SCOPE on the right; that is, the right argument of the function is the result of the entire expression to its right which must be an array. A dyadic function has short scope on the left; that is, the left argument of the function is the array immediately to its left. Left scope may be extended by enclosing an expression in parentheses whence the result must be an array.

For some functions, the explicit result is suppressed if it would otherwise be displayed on completion of evaluation of the expression. This applies on assignment to a variable name. It applies for certain system functions, and may also apply for defined functions.

Examples

```
      10×5-2×4
¯30
      2×4
8
      5-8
¯3
      10×¯3
¯30
      (10×5)-2×4
42
```

Defined Functions

Functions may be defined with the system function `⎕FX`, or with the function editor. A function consists of a HEADER which identifies the syntax of the function, and a BODY in which one or more APL statements are specified.

The header syntax identifies the function name, its (optional) result and its (optional) arguments. If a function is ambivalent, it is defined with two arguments but with the left argument within braces (`{}`). If an ambivalent function is called monadically, the left argument has no value inside the function. If the explicit result is to be suppressed for display purposes, the result is shown within braces. A function need not produce an explicit result. Refer to *Chapter 2* for further details.

Example

```
      ∇ R←{A} FOO B
[1]     R←⊃'MONADIC' 'DYADIC'[⎕IO+0≠⎕NC'A']
[2]   ∇

       FOO 1
MONADIC

      'X' FOO 'Y'
DYADIC
```

Functions may also be created by using assignment (`←`).

Function Assignment & Display

The result of a function-expression may be given a name. This is known as
FUNCTION ASSIGNMENT (see also *Dynamic Functions*). If the result of a function-
expression is not given a name, its value is displayed. This is termed FUNCTION
DISPLAY.

Examples

```
      PLUS←+
      PLUS
+
      SUM←+/
      SUM
+/
```

Function expressions may include defined functions and operators. These are
displayed as a ∇ followed by their name.

Example

```
      ∇ R←MEAN X      ⍝ Arithmetic mean
[1]    R←(+/X)÷ρX
      ∇

      MEAN
 ∇MEAN

      AVERAGE←MEAN
      AVERAGE
 ∇MEAN

      AVG←MEAN∘,
      AVG
 ∇MEAN ∘,
```

Operators

An operator is an operation on one or two operands which produces a function called a DERIVED FUNCTION. An operand may be a function or an array. Operators are not ambivalent. They require either one or two operands as applicable to the particular operator. However, the derived function may be ambivalent. The derived function need not return a result. Operators have higher precedence than functions. Operators have long scope on the left. That is, the left operand is the longest function or array expression on its left. The left operand may be terminated by:

1. the end of the expression

2. the right-most of two consecutive functions

3. a function with an array to its left

4. an array with a function to its left

5. an array or function to the right of a monadic operator.

A dyadic operator has short scope on the right. That is, the right operand of an operator is the single function or array on its right. Right scope may be extended by enclosing an expression in parentheses.

Examples

```
      ρ¨X←'WILLIAM' 'MARY' 'BELLE'
 7  4  5

      ρ∘ρ¨X
 1  1  1

      (ρ∘ρ)¨X
 1  1  1

      □∘←∘□VR¨'PLUS' 'MINUS'
    ∇ R←A PLUS B
[1]    R←A+B
    ∇
    ∇ R←A MINUS B
[1]    R←A-B
    ∇

    PLUS/1 2 3 4
10
```

Defined Operators

Operators may be defined with the system function $\square FX$, or with the function editor. A defined operator consists of a HEADER which identifies the syntax of the operator, and a BODY in which one or more APL statements are specified.

A defined operator may have one or two operands; and its derived function may have one or two arguments, and may or may not produce a result. The header syntax defines the operator name, its operand(s), the argument(s) to its derived function, and the result (if any) of its derived function. The names of the operator and its operand(s) are separated from the name(s) of the argument(s) to its derived function by parentheses.

Example

```
      ∇ R←A(F AND G)B
[1]     R←(A F B)(A G B)
      ∇
```

The above example shows a dyadic operator called AND with two operands (*F* and *G*). The operator produces a derived function which takes two arguments (*A* and *B*), and produces a result (*R*).

```
      12 +AND÷ 4
16 3
```

Operands passed to an operator may be either functions or arrays.

```
      12 (3 AND 5) 4
12 3 4  12 5 4

      12 (× AND 5) 4
48  12 5 4
```

Namespace Syntax

Names within namespaces may be referenced *explicitly* or *implicitly*. An **explicit** reference requires that you identify the object by its full or relative pathname using a '.' syntax; for example:

```
X.NUMB ← 88
```

sets the variable *NUMB* in namespace *X* to 88.

```
88 UTIL.FOO 99
```

calls dyadic function *FOO* in namespace *UTIL* with left and right arguments of 88 and 99 respectively. The interpreter can distinguish between this use of '.' and its use as the inner product operator, because the leftmost name: *UTIL* is a (class 9) namespace, rather than a (class 3) function.

The general namespace reference syntax is:

```
SPACE . SPACE . (...) EXPR
```

Where *SPACE* is an *expression* which resolves to a namespace reference, and *EXPR* is any APL expression to be resolved in the resulting namespace.

There are two special space names:

> `#` is the top level or 'Root' namespace.
> `##` is the parent or space containing the current namespace.

`⎕SE` is a system namespace which is preserved across workspace load and clear.

Examples

```
WSDOC.PAGE.NO +← 1       ⍝ Increment WSDOC page count

#.⎕NL 2                  ⍝ Variables in root space

UTIL.⎕FX 'Z←DUP A' 'Z←A A'   ⍝ Fix remote function

##.⎕ED'FOO'              ⍝ Edit function in parent space

⎕SE.RECORD ← PERS.RECORD    ⍝ Copy from PERS to ⎕SE

UTIL.(⎕EX ⎕NL 2)         ⍝ Expunge variables in UTIL

(⊃⎕SE #).(⍋⊃↓⎕NL 9).(⎕NL 2)   ⍝ Vars in first ⎕SE
                              ⍝ namespace.
```

```
        UTIL.⍎STRING              ⍝ Execute STRING in UTIL space
```

You may also reference a function or operator in a namespace *implicitly* using the mechanism provided by $\square EXPORT$ and $\square PATH$. If you reference a name that is undefined in the current space, the system searches for it in the list of exported names defined for the namespaces specified by $\square PATH$. See $\square EXPORT$ and $\square PATH$ for further details.

Evaluation

When the interpreter encounters a namespace reference, it :

- Switches to the namespace.
- Evaluates the name.
- Switches back to the original namespace.

If for example, in the following, the current namespace is `#.W`, the interpreter evaluates the line:

```
A ← X.Y.DUP MAT
```

in the following way:

- Evaluate array `MAT` in current namespace `W` to produce argument for function.
- Switch to namespace `X.Y` within `W`,
- Evaluate function `DUP` in namespace `W.X.Y` with argument.
- Switch back to namespace `W`,
- Assign variable `A` in namespace `W`.

Static Vs Dynamic Localisation

The rules for name resolution have been generalised for namespaces.

In flat APL, the interpreter searches the state indicator to resolve names referenced by a defined function or operator. If the name does not appear in the state indicator, then the workspace-global name is assumed.

With namespaces, a defined function or operator is evaluated in its 'home' namespace. When a name is referenced, the interpreter searches only those lines of the state indicator which belong to the home namespace. If the name does not appear in any of these lines, the home namespace-global value is assumed.

For example, if `#.FN1` calls `XX.FN2` calls `#.FN3` calls `XX.FN4`, then:

`FN1:`
> is evaluated in `#`
> can see its own dynamic local names
> can see global names in `#`

`FN2:`
> is evaluated in `XX`
> can see its own dynamic local names
> can see global names in `XX`

`FN3:`
> is evaluated in `#`
> can see its own dynamic local names
> can see dynamic local names in `FN1`
> can see global names in `#`

`FN4:`
> is evaluated in `XX`
> can see its own dynamic local names
> can see dynamic local names in `FN2`
> can see global names in `XX`

Namespace References

A *namespace reference*, or *ref* for short, is a unique data type that is distinct from and in addition to *number* and *character*.

Any expression may result in a *ref*, but the simplest one is the namespace itself:

```
      )NS NS1            ⍝ Make a namespace called NS1
      NS1.A←1            ⍝ and populate it with variables A
      NS1.B←2 3⍴⍳6       ⍝ and B

      NS1               ⍝ expression results in a ref
#.NS1
```

You may assign a ref ; for example:

```
      X←NS1
      X
#.NS1
```

In this case, the display of `X` informs you that `X` refers to the named namespace `#.NS1`.

You may also supply a ref as an argument to a defined or dynamic function:

```
      ∇ FOO ARG
[1]    ARG
      ∇
      FOO NS1
#.NS1
```

The name class of a *ref* is 9.

```
      □NC 'X'
9
```

You may use a ref to a namespace anywhere that you would use the namespace itself. For example:

```
      X.A
1
      X.B
1 2 3
4 5 6
```

Notice that refs are references to namespaces, so that if you make a copy, it is the reference that is copied, not the namespace itself. This is sometimes referred to as a shallow as opposed to a deep copy. It means that if you change a ref, you actually change the namespace that it refers to.

```
      X.A+←1
      X.A
2
      NS1.A
2
```

Similarly, a ref passed to a defined function is call-by-reference, so that modifications to the content or properties of the argument namespace using the passed reference, persist after the function exits. For example:

```
      ∇ FOO nsref
[1]    nsref.B+←nsref.A
      ∇

      FOO NS1
      NS1.B
3 4 5
6 7 8
      FOO X
      NS1.B
```

```
5  6   7
8  9  10
```

Notice that the expression to the right of a dot may be arbitrarily complex and will be executed within the namespace or ref to the left of the dot.

```
      X.(C←A×B)
      X.C
10 12 14
16 18 20
      NS1.C
10 12 14
16 18 20
```

Unnamed Namespaces

The monadic form of `⎕NS` makes a new (and unique) unnamed namespace and returns a ref to it.

One use of unnamed namespaces is to represent hierarchical data structures; for example, a simple employee database:

The first record is represented by `JOHN` which is a ref to an *unnamed* namespace:

```
      JOHN←⎕NS ''
      JOHN
#.[Namespace]

      JOHN.FirstName←'John'
      JOHN.FirstName
John

      JOHN.LastName←'Smith'
      JOHN.Age←50
```

Data variables for the second record, `PAUL`, can be established using strand, or vector, assignment:

```
      PAUL←⎕NS ''
      PAUL.(FirstName LastName Age←'Paul' 'Brown' 44)
```

The function `SHOW` can be used to display the data in each record (the function is split into 2 lines only to fit on the printed page). Notice that its argument is a ref.

```
      ∇ R←SHOW PERSON
[1]    R←PERSON.FirstName,' ',PERSON.LastName
[2]    R, ←' is ',⍕PERSON.Age
      ∇
```

```
      SHOW JOHN
John Smith is 50

      SHOW PAUL
Paul Brown is 44
```

An alternative version of the function illustrates the use of the *:With   :EndWith* control structure to execute an expression, or block of expressions, within a namespace:

```
      ∇ R←SHOW1 PERSON
[1]     :With PERSON
[2]         R←FirstName,' ',LastName,' is ',(⍕Age)
[3]     :EndWith
      ∇

      SHOW1 JOHN
John Smith is 50
```

In this case, as only a single expression is involved, it can be expressed more simply using parentheses.

```
      ∇ R←SHOW2 PERSON
[1]     R←PERSON.(FirstName,' ',LastName,' is ',(⍕Age))
      ∇

      SHOW2 PAUL
Paul Brown is 44
```

Dynamic functions also accept refs as arguments:

```
      SHOW3←{
         ⍵.(FirstName,' ',LastName,' is ',⍕Age)
      }

      SHOW3 JOHN
John Smith is 50
```

Arrays of Namespace References

You may construct arrays of refs using strand notation, catenate (,) and reshape (ρ).

```
      EMP←JOHN PAUL
      ρEMP
2
      EMP
 #.[Namespace]  #.[Namespace]
```

Like any other array, an array of refs has name class 2:

```
      ⎕NC 'EMP'
2
```

Expressions such as indexing and pick return refs that may in turn be used as follows:

```
      EMP[1].FirstName
John
      (2⊃EMP).Age
44
```

The each (¨) operator may be used to apply a function to an array of refs:

```
      SHOW¨EMP
 John Smith is 50  Paul Brown is 44
```

An *array* of namespace references (refs) to the left of a ' . ' is expanded according to
the following rule, where *x* and *y* are refs, and *exp* is an arbitrary expression:

```
      (x y).exp → (x.exp)(y.exp)
```

If *exp* evaluates to a function, the items of its argument array(s) are *distributed* to each
referenced function. In the dyadic case, there is a 3-way distribution among: left
argument, referenced functions and right argument.

Monadic function *f*:   `(x y).f d e → (x.f d)(y.f e)`
Dyadic function *g*:   `a b (x y).g  d e → (a x.g d)(b y.g e)`

An array of refs to the left of an assignment arrow is expanded thus:

```
      (x y).a←c d   →  (x.a←c)(y.a←d)
```

Note that the array of refs can be of any rank. In the limiting case of a simple scalar
array, the *array* construct: `refs.exp` is identical to the *scalar* construct: `ref.exp`.

Note that the expression to the right of the '.' *pervades* a nested array of refs to its left:

$$((u\ v)(x\ y)).exp \rightarrow ((u.exp)(v.exp))((x.exp)(y.exp))$$

Note also that with *successive* expansions (u v).(x y z). ..., the final number of 'leaf' terms is the *product* of the number of refs at each level.

Examples:

```
      JOHN.Children←⎕NS¨'' ''
      ⍴JOHN.Children
2
      JOHN.Children[1].FirstName←'Andy'
      JOHN.Children[1].Age←23

      JOHN.Children[2].FirstName←'Katherine'
      JOHN.Children[2].Age←19

      PAUL.Children←⎕NS¨'' ''
      PAUL.Children[1].(FirstName Age←'Tom' 25)
      PAUL.Children[2].(FirstName Age←'Jamie' 22)

      ⍴EMP
2
      (⊃EMP).Children.(FirstName Age)
  Andy  23   Katherine  19

      DISPLAY (2⊃EMP).Children.(FirstName Age)
.→---------------------------.
| .→---------. .→-----------. |
| | .→--.     | | .→----.    | |
| | |Tom| 25 | | |Jamie| 22 | |
| | '---'     | | '-----'    | |
| '∊---------' '∊-----------' |
'∊---------------------------'

      EMP.Children ⍝ Is an array of refs
  #.[Namespace]  #.[Namespace]   #.[Namespace]  ...

      EMP.Children.(FirstName Age)
   Andy  23   Katherine  19    Tom  25   Jamie  22
```

Distributed Assignment

Assignment pervades nested strands of names to the left of the arrow. The conformability rules are the same as for scalar (pervasive) dyadic primitive functions such as '+'. The mechanism can be viewed as a way of naming the parts of a *structure*.

Examples:

```
      EMP.(FirstName Age)
  JOHN  43    PAUL  44

      EMP.(FirstName Age)←('Jonathan' 21)('Pauline' 22)

      EMP.(FirstName Age)
  Johnathan  21    Pauline  22

      ⍝ Distributed assignment is pervasive
      JOHN.Children.(FirstName Age)
  Andy  23    Katherine  19

      JOHN.Children.(FirstName Age)←('Andrew' 21)('Kate' 19
)

      JOHN.Children.(FirstName Age)
  Andrew  21    Kate  19
```

More Examples:

```
      ((a b)(c d))←(1 2)(3 4)    ⍝ a←1 ◊ b←2 ◊ c←3 ◊ d←4

      ((⎕io ⎕ml)vec)←0 ⎕av      ⍝ ⎕io←0 ◊ ⎕ml←0 ◊ vec←⎕av

      (i (j k))+←1 2            ⍝ i+←1 ◊ j+←2 ◊ k+←2

⍝ Naming of parts:

      ((first last) sex (street city state))←n⊃pvec

⍝ Distributed assignment in :For loop:

      :For (i j)(k l) :In array

⍝ Ref array expansion:

      (x y).(first last)←('John' 'Doe')('Joe' 'Blow')
      (f1 f2).(b1 b2).Caption←⊂'OK' 'Cancel'
```

```
⍝ Structure rearrangement:

    rotate1←{          ⍝ Simple binary tree rotation.
        (a b c)d e←⍵
        a b(c d e)
    }
    rotate3←{          ⍝ Compound binary tree rotation.
        (a b(c d e))f g←⍵
        (a b c)d(e f g)
    }
```

Distributed Functions

Namespace ref array expansion syntax applies to functions too.

```
    JOHN.PLOT←{↑⍵⍴¨'□'}
    JOHN.PLOT ι10
□
□□
□□□
□□□□
□□□□□
□□□□□□
□□□□□□□
□□□□□□□□
□□□□□□□□□
□□□□□□□□□□

    PAUL.PLOT←{(⍵,¨1)⍴¨'□'}
    PAUL.PLOT ι10
 □  □  □  □  □  □  □  □  □  □
    □  □  □  □  □  □  □  □  □
       □  □  □  □  □  □  □  □
          □  □  □  □  □  □  □
             □  □  □  □  □  □
                □  □  □  □  □
                   □  □  □  □
                      □  □  □
                         □  □
                            □

    EMP.PLOT⊂ι10  ⍝ (temporary vector of functions)
□                □  □  □  □  □  □  □  □  □  □
□□               □  □  □  □  □  □  □  □  □
□□□              □  □  □  □  □  □  □  □
□□□□             □  □  □  □  □  □  □
□□□□□            □  □  □  □  □  □
```

□□□□□□   □   □   □   □   □
□□□□□□□   □   □   □   □
□□□□□□□□   □   □   □
□□□□□□□□□   □   □
□□□□□□□□□□   □

```
      (x y).⎕NL 2 3              ⍝ x:vars, y:fns
 varx   funy

      (x y).⎕NL⊂2 3             ⍝ x&y: vars&fns
 funx   funy
 varx   vary

      (x y).(⎕NL¨)⊂2 3          ⍝ x&y: separate vars&fns
  varx   funx     vary   funy

      'v'(x y).⎕NL 2 3          ⍝ x:v-vars, y:v-fns
 varx

      'vf'(x y).⎕NL 2 3         ⍝ x:v-vars, y:f-fns
 varx   funy
                               ⍝ x:v-vars&fns,
      'vf'(x y).⎕NL⊂2 3         ⍝ y:f-vars&fns
 varx   funy

      x.⎕NL 2 3                ⍝ depth 0 ref
 funx
 varx

      (x y).⎕NL⊂2 3            ⍝ depth 1 refs
 funx   funy
 varx   vary

      ((u v)(x y)).⎕NL⊂⊂2 3     ⍝ depth 2 refs
  funu   funv     funx   funy
  varu   varv     varx   vary

      (1 2)3 4(w(x y)z).+1 2(3 4) ⍝ argument distribution.
 2 3   5 5   7 8
```

Operators

A function passed as operand to a primitive or defined operator, carries its namespace context with it. This means that if subsequently, the function operand is applied to an argument, it executes in its home namespace, irrespective of the namespace from which the operator was invoked or defined.

Examples

```
      VAR←99                   ⍝ #.VAR

      )NS X
#.X
      X.VAR←77                 ⍝ X.VAR
      X.⎕FX'Z←FN R' 'Z←R,VAR'

      )NS Y
#.Y
      Y.VAR←88                 ⍝ Y.VAR
      Y.⎕FX'Z←(F OP)R' 'Z←F R'

      X.FN¨⍳3
 1 77   2 77   3 77

      X.FN 'VAR:'
 VAR: 77

      X.FN Y.OP 'VAR:'
 VAR: 77
      ⍎ Y.OP'VAR'
99
```

Summary

Apart from its use as a decimal separator (3.14), '.' is interpreted by looking at the type or *class* of the expression to its left:

Template	Interpretation	Example
`∘.`	Outer product	`2 3 ∘.× 4 5`
`function.`	Inner product	`2 3 +.× 4 5`
`ref.`	Namespace reference	`2 3 x.foo 4 5`
`array.`	Reference array expansion	`(x y).⎕nc⊂'foo'`

MultiThreading

Overview

Dyalog APL supports multithreading - the ability to run more than one APL expression at the same time.

This unique capability allows you to perform background processing, such as printing, database retrieval, database update, calculations, and so forth while at the same time perform other interactive tasks.

Multithreading may be used to improve throughput and system responsiveness.

A *thread* is a strand of execution in the APL workspace.

A thread is created by calling a function *asynchronously*, using the new primitive operator 'spawn': & or by the asynchronous invocation of a callback function.

With a traditional APL *synchronous* function call, execution of the calling environment is paused, *pendent* on the return of the called function. With an *asynchronous* call, both calling environment and called function proceed to execute concurrently.

An asynchronous function call is said to start a new *thread* of execution. Each thread has a unique *thread number*, with which, for example, its presence can be monitored or its execution terminated.

Any thread can spawn any number of sub-threads, subject only to workspace availability. This implies a hierarchy in which a thread is said to be a *child thread* of its *parent thread*. The *base thread* at the root of this hierarchy has thread number 0.

With multithreading, APL's stack or state indicator can be viewed as a branching tree in which the path from the base to each leaf is a thread.

When a parent thread terminates, any of its children which are still running, become the children of (are 'adopted' by) the parent's parent.

Thread numbers are allocated sequentially from 0 to 2147483647. At this point, the sequence 'wraps around' and numbers are allocated from 0 again avoiding any still in use. The sequence is reinitialised when a *)RESET* command is issued, or the active workspace is cleared, or a new workspace is loaded. A workspace may not be saved with threads other than the base thread: 0, running.

Threads introduce new language elements.

- Primitive operator, spawn: *&*.

- System functions: *⎕TID*, *⎕TCNUMS*, *⎕TNUMS*, *⎕TKILL*, *⎕TSYNC*.

- An extension to the GUI Event syntax to allow asynchronous callbacks.

- A control structure: *:Hold*.

- System commands: *)HOLDS*, *)TID*.

- Extended *)SI* and *)SINL* display.

Running CallBack Functions as Threads

A callback function is associated with a particular event via the Event property of the object concerned. A callback function is executed by *⎕DQ* when the event occurs, or by *⎕NQ*.

If you append the character *&* to the name of the callback function in the *Event* specification, the callback function will be executed asynchronously as a thread when the event occurs. If not, it is executed synchronously as before.

For example, the event specification:

```
⎕WS'Event' 'Select' 'DoIt&'
```

tells *⎕DQ* to execute the callback function *DoIt asynchronously as a thread* when a Select event occurs on the object.

Thread Switching

Programming with threads requires care.

The interpreter may switch between running threads at the following points:

- Between any two lines of a defined (or dynamic) function or operator.

- While waiting for a $\square DL$ to complete.

- While waiting for a $\square FHOLD$ to complete.

- While awaiting input from:
  $\square DQ$
  $\square SR$
  $\square ED$
  The session prompt or $\square$: or $\square$.

- While awaiting the completion of an external operation:
  A call on an external (AP) function.
  A call on a $\square NA$ (DLL) function.
  A call on an OLE function.

At any of these points, the interpreter might execute code in other threads. If such threads change the global environment; for example by changing the value of, or expunging a name; then the changes will appear to have happened while the thread in question passes through the switch point. It is the task of the application programmer to organise and contain such behaviour!

You can prevent threads from interacting in critical sections of code by using the :*Hold* control structure.

Name Scope

APL's name scope rules apply whether a function call is synchronous or asynchronous. For example when a defined function is called, names in the calling environment are visible, unless explicitly shadowed in the function header.

Just as with a synchronous call, a function called asynchronously has its own local environment, but can communicate with its parent and 'sibling' functions via local names in the parent.

This point is important. It means that siblings can run in parallel without danger of local name clashes. For example, a GUI application can accommodate multiple concurrent instances of its callback functions.

However, with an asynchronous call, as the calling function continues to execute, both child *and parent functions* may modify values in the calling environment. Both functions see such changes immediately they occur.

If a parent function terminates while any of its children are still running, those children will thenceforward 'see' local names in the environment that called the parent function. In cases where a child function relies on its parent's environment (the setting of a local value of $\Box IO$ for example), this would be undesirable, and the parent function would normally execute a $\Box TSYNC$ in order to wait for its children to complete before itself exiting.

If, on the other hand, after launching an asynchronous child, the parent function calls a *new* function (either synchronously or asynchronously), names in the new function are beyond the purview of the original child. In other words, a function can only ever see its calling stack decrease in size – never increase. This is in order that the parent may call new defined functions without affecting the environment of its asynchronous children.

Debugging

If a thread sustains an untrapped error, its execution is *suspended* in the normal way. At the same time, all other thread activity in the workspace is *paused* until the suspension is cleared or restarted. A consequence of this is that, at any one time, only a single thread can be suspended.

The session is attached to the suspended thread, so that you can:

- Examine and modify local variables.
- Trace and edit functions.
- Clear suspensions.
- Restart execution.

When the last suspension in a thread is cleared or restarted, the session is reattached to the base thread, and any paused threads are resumed.

From the session in a suspended thread, you can spawn a new thread, but its execution is immediately paused until the parent's suspension is removed.

The error message from a thread other than the base is prefixed with its thread number:

```
260:DOMAIN ERROR
Div[2]  rslt←num÷div
      ^
```

State indicator displays: `)SI` and `)SINL` have been extended to show threads' tree-like calling structure.

```
      )SI
·   Calc[1]
&5
·   ·    DivSub[1]
·   &7
·   ·    DivSub[1]
·   &6
·   Div[2]*
&4
Sub[3]
Main[4]
```

Here, `Main` has called `Sub`, which has spawned threads 4 and 5 with functions: `Div` and `Calc`. `Div`, after spawning `DivSub` in each of threads 6 and 7, has been suspended at line `[2]`.

Removing stack frames using 'Quit' from the tracer or '→' from the session affects only the current thread. When the final stack frame in a thread (other than the base thread) is removed, the thread is expunged.

`)RESET` removes all but the base thread.

Using Threads

Put most simply, multithreading allows you to *appear to* run more than one APL function at the same time, just as Windows (or UNIX) *appears to* run more than one application at the same time. In both cases this is something of an illusion, although it does nothing to detract from its usefulness.

Dyalog APL implements an internal timesharing mechanism whereby it shares processing between threads. Although the mechanics are somewhat different, APL multithreading is rather similar to the multitasking provided by Windows 95/98 and NT. If you are running more than one application, Windows switches from one to another, allocating each one a certain *time slice* before switching. At any point in time, only one application is actually running; the others are paused, waiting.

If you execute more than one Dyalog APL thread, only one thread is actually running; the others are paused. Each APL thread has its own State Indicator, or SI stack. When APL switches from one thread to another, it saves the current stack (with all its local variables and function calls), restores the new one, and then continues processing.

Stack Considerations

When you start a thread, it begins with the SI stack of the calling function and sees all of the local variables defined in all the functions down the stack. However, unless the calling function specifically waits for the new thread to terminate (see $\square TSYNC$), the calling functions will (bit by bit, in their turn) continue to execute. The new thread's view of its calling environment may then change. Consider the following example:

Suppose that you had the following functions: RUN[3] calls $INIT$ which in turn calls $GETDATA$ but as 3 separate threads with 3 different arguments:

```
      ∇ RUN;A;B
[1]     A←1
[2]     B←'Hello World'
[3]     INIT
[4]     CALC
[5]     REPORT
      ∇

      ∇ INIT;C;D
[1]     C←D←0
[2]     GETDATA&'Sales'
[3]     GETDATA&'Costs'
[4]     GETDATA&'Expenses'
      ∇
```

When each *GETDATA* thread starts, it immediately *sees* (via □*SI*) that it was called by *INIT* which was in turn called by *RUN*, and it *sees* local variables *A*, *B*, *C* and *D*. However, once *INIT*[4] has been executed, *INIT* terminates, and execution of the root thread continues by calling *CALC*. From then on, each *GETDATA* thread no longer sees *INIT* (it thinks that it was called directly from *RUN*) nor can it see the local variables *C* and *D* that *INIT* had defined. However, it *does* continue to see the locals *A* and *B* defined by *RUN*, until *RUN* itself terminates.

Note that if *CALC* were also to define locals *A* and *B*, the *GETDATA* threads would still see the values defined by *RUN* and not those defined by *CALC*. However, if *CALC* were to modify *A* and *B* (as globals) without localising them, the *GETDATA* threads would see the modified values of these variables, whatever they happened to be at the time.

Globals and the Order of Execution

It is important to recognise that any reference or assignment to a global or semi-global object (including GUI objects) is **inherently dangerous** (i.e. a source of programming error) if more than one thread is running. Worse still, programming errors of this sort may not become apparent during testing because they are dependent upon random timing differences. Consider the following example:

```
      ∇ BUG;SEMI_GLOBAL
[1]    SEMI_GLOBAL←0
[2]    FOO& 1
[3]    GOO& 1
      ∇

      ∇ FOO
[1]    :If SEMI_GLOBAL=0
[2]        DO_SOMETHING SEMI_GLOBAL
[3]    :Else
[4]        DO_SOMETHING_ELSE SEMI_GLOBAL
[5]    :EndIf
      ∇

      ∇ GOO
[1]    SEMI_GLOBAL←1
      ∇
```

In this example, it is formally impossible to predict in which order APL will execute statements in *BUG*, *FOO* or *GOO* from *BUG*[2] onwards. For example, the actual sequence of execution may be:

```
BUG[1] → BUG[2] → FOO[1] → FOO[2] →
          DO_SOMETHING[1]
```

or

```
BUG[1] → BUG[2] → BUG[3] → GOO[1] →
          FOO[1] → FOO[2] → FOO[3] →
          FOO[4] → DO_SOMETHING_ELSE[1]
```

This is because APL may switch from one thread to another between any two lines in a defined function. In practice, because APL gives each thread a significant time-slice, it is likely to execute many lines, maybe even hundreds of lines, in one thread before switching to another. However, you must not rely on this; **thread-switching may occur at any time between lines in a defined function**.

Secondly, consider the possibility that APL switches from the *FOO* thread to the *GOO* thread after *FOO*[1]. If this happens, the value of *SEMI_GLOBAL* passed to *DO_SOMETHING* will be 1 and not 0. Here is another source of error.

In fact, in this case, there are two ways to resolve the problem. To ensure that the value of *SEMI_GLOBAL* remains the same from *FOO*[1] to *FOO*[2], you may use diamonds instead of separate statements, e.g.

```
:If SEMI_GLOBAL=0 ◇ DO_SOMETHING SEMI_GLOBAL
```

Even better, although less efficient, you may use :*Hold* to synchronise access to the variable, for example:

```
      ∇ FOO
[1]    :Hold 'SEMI_GLOBAL'
[2]        :If SEMI_GLOBAL=0
[3]            DO_SOMETHING SEMI_GLOBAL
[4]        :Else
[5]            DO_SOMETHING_ELSE SEMI_GLOBAL
[6]        :EndIf
[7]    :EndHold
      ∇


      ∇ GOO
[1]    :Hold 'SEMI_GLOBAL'
[2]        SEMI_GLOBAL←1
[3]    :EndHold
      ∇
```

Now, although you still cannot be sure which of *FOO* and *GOO* will run first, you can be sure that *SEMI_GLOBAL* will not change (because *GOO* cuts in) within *FOO*.

Note that the string used as the argument to :*Hold* is completely arbitrary, so long as threads competing for the same resource use the same string.

A Caution

These types of problems are inherent in all multithreading programming languages, and not just with Dyalog APL. *If you want to take advantage of the additional power provided by multithreading, it is advisable to think carefully about the potential interaction between different threads.*

Threads & Niladic Functions

In common with other operators, the spawn operator & may accept monadic or dyadic functions as operands, but not niladic functions. This means that, using spawn, you cannot start a thread that consists only of a niladic function

If you wish to invoke a niladic function asynchronously, you have the following choices:

- Turn your niladic function into a monadic function by giving it a dummy argument which it ignores.
- Call your niladic function with a dynamic function to which you give an argument that is implicitly ignored. For example, if the function *NIL* is niladic, you can call it asynchronously using the expression:

    ```
    {NIL}& 0
    ```
- Call your function via a dummy monadic function, e.g.

    ```
        ∇ NIL_M DUMMY
    [1]    NIL
        ∇
        NIL_M& ''
    ```

- Use execute, e.g.

    ```
    ⍎& 'NIL'
    ```

Note that niladic functions *can* be invoked asynchronously as callback functions. For example, the statement:

```
⎕WS'Event' 'Select' 'NIL&'
```

will execute correctly as a thread, even though *NIL* is niladic. This is because callback functions are invoked directly by ⎕*DQ* rather than as an operand to a the spawn operator.

Threads & External Functions

External functions in dynamic link libraries (DLLs) defined using the $\square NA$ interface may be run in separate C threads. Such threads:

- **take advantage of multiple processors** if the operating system permits.
- allow APL to **continue processing in parallel** during the execution of a $\square NA$ function.

When you define an external function using $\square NA$, you may specify that the function be run in a separate C thread by appending an ampersand (&) to the function name, for example:

```
'beep'⎕NA'user32|MessageBeep& i'
⍝ MessageBeep will run in a separate C thread
```

When APL first comes to execute a multi-threaded $\square NA$ function, it starts a new C-thread, executes the function within it, and waits for the result. Other APL threads may then run in parallel.

Note that when the $\square NA$ call finishes and returns its result, its new C-thread is retained to be re-used by any subsequent multithreaded $\square NA$ calls made within the same APL thread. Thus any APL thread that makes any multi-threaded $\square NA$ calls maintains a separate C-thread for their execution. This C-thread is discarded when its APL thread finishes.

Note that there is no point in specifying a $\square NA$ call to be multi-threaded, unless you wish to execute other APL threads at the same time.

In addition, if your $\square NA$ call needs to access an APL GUI object (strictly, a window or other handle) it should normally run within the same C-thread as APL itself, and not in a separate C-thread. This is because Windows associates objects with the C-thread that created them. Although you *can* use a multi-threaded $\square NA$ call to access (say) a Dyalog APL Form via its window handle, the effects may be different than if the $\square NA$ call was not multi-threaded. In general, $\square NA$ calls that access APL (GUI) objects should not be multi-threaded.

If you wish to run the same $\square NA$ call in separate APL threads at the same time, you must ensure that the DLL is *thread-safe*. Functions in DLLs which are not *thread-safe*, must be prevented from running concurrently by using the `:Hold` control structure. Note that all the standard Windows API DLLs **are** *thread safe*.

Notice that you may define two separate functions (with different names), one single-threaded and one multi-threaded, associated with the same function in the DLL. This allows you to call it in either way.

External Variables

An external variable is a variable whose contents (value) reside not in the workspace, but in a file. An external variable is associated with a file by the system function $\square XT$. If at the time of association the file exists, the external variable assumes its value from the contents of the file. If the file does not exist, the external variable is defined but a *VALUE ERROR* occurs if it is referenced before assignment. Assignment of an array to the external variable or to an indexed element of the external variable has the effect of updating the file. The value of the external variable or the value of indexed elements of the external variable is made available in the workspace when the external variable occurs in an expression. No special restrictions are placed on the usage of external variables.

Normally, the files associated with external variables remain permanent in that they survive the APL session or the erasing of the external variable from the workspace. External variables may be accessed concurrently by several users, or by different nodes on a network, provided that the appropriate file access controls are established. Multi-user access to an external variable may be controlled with the system function $\square FHOLD$ between co-operating tasks.

Refer to the sections describing the system functions $\square XT$ and $\square FHOLD$ in *Chapter 5* for further details.

Examples

```
      'ARRAY' □XT 'V'

      V←ι10
      V[2] + 5
7

      □EX'V'

      'ARRAY' □XT 'F'
      F
1 2 3 4 5 6 7 8 9 10
```

Component Files

A component file is a data file maintained by Dyalog APL. It contains a series of APL arrays known as components which are accessed by reference to their relative positions or component number within the file. A set of system functions is provided to perform a range of file operations. (See *Chapter 5*.) These provide facilities to create or delete files, and to read and write components. Facilities are also provided for multi-user access including the capability to determine who may do what, and file locking for concurrent updates. (See *User Guide*.)

Auxiliary Processors

Auxiliary Processors (APs) are non-APL programs which provide Dyalog APL users with additional facilities. They run as separate tasks, and communicate with the Dyalog APL interpreter through pipes (UNIX) or via an area of memory (Windows). Typically, APs are used where speed of execution is critical, such as in screen management software, or for utility libraries. Auxiliary Processors may be written in any compiled language, although 'C' is preferred and is directly supported.

When an Auxiliary Processor is invoked from Dyalog APL, one or more **external functions** are fixed in the active workspace. Each external function behaves as if it were a locked defined function, but is in effect an entry point into the Auxiliary Processor. An external function occupies only a negligible amount of workspace. (See *User Guide*.)

Migration Level

⎕*ML* determines the degree of migration of the Dyalog APL language towards IBM's APL2. Unless otherwise stated, the manual assumes ⎕*ML* has a value of 0.

Key to Notation

The following definitions and conventions apply throughout this manual:

f	A function, or an operator's left argument when a function.
g	A function, or an operator's right argument when a function.
A	An operator's left argument when an array.
B	An operator's right argument when an array.
X	The left argument of a function.
Y	The right argument of a function.
R	The explicit result of a function.
[K]	Axis specification.
[I]	Index specification.
{X}	The left argument of a function is optional.
{R}←	The function may or may not return a result, or the result may be suppressed.

The term 'function' may refer to a primitive function, a system function, a defined (canonical, dynamic or assigned) function or a derived (from an operator) function.

C H A P T E R 2

Defined Functions & Operators

A defined function is a program that takes 0, 1, or 2 arrays as **arguments** and may produce an array as a result. A defined operator is a program that takes 1 or 2 functions or arrays (known as **operands**) and produces a **derived function** as a result. To simplify the text, the term **operation** is used within this chapter to mean function or operator.

Canonical Representation

Operations may be defined with the system function $\Box FX$ (Fix) or by using the editor within definition mode. Applying $\Box CR$ to the character array representing the name of an already established operation will produce its canonical representation. A defined operation is composed of lines. The first line (line 0) is called the operation HEADER. Remaining lines are APL statements, called the BODY.

The operation header consists of the following parts:

1. its model syntactical form,

2. an optional list of local names, each preceded by a semi-colon (;) character,

3. an optional comment, preceded by the symbol ⍝.

Only the model is required. If local names and comments are included, they must appear in the prescribed order.

Model Syntax

The model for the defined operation identifies the name of the operation, its valence, and whether or not an explicit result may be returned. Valence is the number of explicit arguments or operands, either 0, 1 or 2; whence the operation is termed NILADIC, MONADIC or DYADIC respectively. Only a defined function may be niladic. There is no relationship between the valence of a defined operator, and the valence of the derived function which it produces. Defined functions and derived functions produced by defined operators may be ambivalent, i.e. may be executed monadically with one argument, or dyadically with two. An ambivalent operation is identified in its model by enclosing the left argument in braces.

The value of a result-returning function or derived function may be suppressed in execution if not explicitly used or assigned by enclosing the result in its model within braces. Such a suppressed result is termed SHY.

Figures 2(i), 2(ii) and 2(iii) show all possible models for defined functions and operators respectively.

Defined Functions

Result	Niladic	Monadic	Dyadic	Ambivalent
None	*f*	*f Y*	*X f Y*	*{X} f Y*
Explicit	*R←f*	*R←f Y*	*R←X f Y*	*R←{X} f Y*
Suppressed	*{R}←f*	*{R}←f Y*	*{R}←X f Y*	*{R}←{X} f Y*

Figure 2(i) : Models for Defined Functions

Derived Functions produced by Monadic Operator

Result	Monadic	Dyadic	Ambivalent
None	`(A op)Y`	`X(A op)Y`	`{X}(A op)Y`
Explicit	`R←(A op)Y`	`R←X(A op)Y`	`R←{X}(A op)Y`
Suppressed	`{R}←(A op)Y`	`{R}←X(A op)Y`	`{R}←{X}(A op)Y`

Figure 2(ii) : Models for Defined Operators (Monadic)

Derived Functions produced by Dyadic Operator

Result	Monadic	Dyadic	Ambivalent
None	`(A op B)Y`	`X(A op B)Y`	`{X}(A op B)Y`
Explicit	`R←(A op B)Y`	`R←X(A op B)Y`	`R←{X}(A op B)Y`
Suppress	`{R}←(A op B)Y`	`{R}←X(A op B)Y`	`{R}←{X}(A op B)Y`

Figure 2(iii) : Models for Defined Operators (Dyadic)

Statements

A statement is a line of characters understood by APL. It may be composed of:

1. a LABEL (which must be followed by a colon :), or a CONTROL STATEMENT (which is preceded by a colon), or both,

2. an EXPRESSION (see *Chapter 1*),

3. a SEPARATOR (consisting of the diamond character ◇ which must separate adjacent expressions),

4. a COMMENT (which must start with the character ⍝).

Each of the four parts is optional, but if present they must occur in the given order except that successive expressions must be separated by ◇. Any characters occurring to the right of the first comment symbol (⍝) that is not within quotes is a comment.

Comments are not executed by APL. Expressions in a line separated by ◇ are taken in left-to-right order as they occur in the line. For output display purposes, each separated expression is treated as a separate statement.

Examples

```
      5×10
50

      MULT: 5×10
50

      MULT: 5×10 ◇ 2×4
50
8

      MULT: 5×10 ◇ 2×4  ⍝ MULTIPLICATION
50
8
```

Global & Local Names

The following names, if present, are local to the defined operation:

1. the result,

2. the argument(s) and operand(s),

3. additional names in the header line following the model, each name preceded by a semi-colon character,

4. labels,

5. the argument list of the system function $\square SHADOW$ when executed,

6. a name assigned within a Dynamic Function.

All names in a defined operation must be valid APL names. The same name may be repeated in the header line, including the operation name (whence the name is localised). Normally, the operation name is not a local name.

The same name may not be given to both arguments or operands of a dyadic operation. The name of a label may be the same as a name in the header line. More than one label may have the same name. When the operation is executed, local names in the header line after the model are initially undefined; labels are assigned the values of line numbers on which they occur, taken in order from the last line to the first; the result (if any) is initially undefined.

In the case of a defined function, the left argument (if any) takes the value of the array to the left of the function when called; and the right argument (if any) takes the value of the array to the right of the function when called. In the case of a defined operator, the left operand takes the value of the function or array to the left of the operator when called; and the right operand (if any) takes the value of the function or array to the right of the operator when called.

During execution, a local name temporarily excludes from use an object of the same name with an active definition. This is known as LOCALISATION or SHADOWING. A value or meaning given to a local name will persist only for the duration of execution of the defined operation (including any time whilst the operation is halted). A name which is not local to the operation is said to be GLOBAL. A global name could itself be local to a pendent operation. A global name can be made local to a defined operation during execution by use of the system function $\square SHADOW$. An object is said to be VISIBLE if there is a definition associated with its name in the active environment.

Examples

```
      A←1

      ∇F
[1]   A←10
[2]   ∇

      F
      ⍝ <A> NOT LOCALISED IN <F>, GLOBAL VALUE REPLACED
      A
10
      A←1

      )ERASE F

      ∇F;A
[1]   A←10
[2]   ∇
      F
      ⍝ <A> LOCALISED IN <F>, GLOBAL VALUE RETAINED
      A
1
```

Any statement line in the body of a defined operation may begin with a LABEL. A label is followed by a colon (`:`). A label is a constant whose value is the number of the line in the operation defined by system function `⎕FX` or on closing definition mode.

The value of a label is available on entering an operation when executed, and it may be used but not altered in any expression.

Example

```
      ⎕VR'PLUS'
    ∇ R←{A} PLUS B
[1]   →DYADIC ρ¨2=⎕NC'A' ◇ R←B ◇ →END
[2]  DYADIC: R←A+B
[3]  END:
    ∇

      1 ⎕STOP'PLUS'

      2 PLUS 2

PLUS[1]
      DYADIC
2
```

END

3

Idiom Recognition

Idioms are commonly used expressions that are recognised and evaluated internally, providing a significant performance improvement.

For example, the idiom *BV/ιρA* (where *BV* is a Boolean vector and *A* is an array) would previously have been evaluated in 3 steps as follows:

1. Evaluate ρ*A* and store result in temporary variable *temp1* (*temp1* is just an arbitrary name for the purposes of this explanation)
2. Evaluate ι*temp1* and store result in temporary variable *temp2*.
3. Evaluate *BV/temp2*
4. Discard temporary variables

In Version 10, the expression is recognised in its entirety and processed as if it were a single primitive function in a single step. The resultant improvement in performance is between 2 and 4.5.

Idiom recognition is precise; an expression that is almost identical but not exactly identical to an expression given in the Idiom List table will not be recognised.

For example, □*AV*ι will be recognised as an idiom, but (□*AV*)ι will not. Similarly, (,)/ would not be recognized as the Join idiom.

Idiom List

In the following table, arguments to the idiom have types and ranks as follows:

Type	Description		Rank	Description
C	Character		S	Scalar or 1-item vector
B	Boolean		V	Vector
N	Numeric		M	Matrix
P	Nested (pointer)		A	Array (any rank)
A	Any type			

For example: NV: numeric vector, CM: character matrix, PV: nested vector.

Expression	**Description**
$\rho\rho A$	Rank
$BV/\iota NS$	Sequence selection
$BV/\iota\rho A$	Index selection
$NA \supset^{..} \subset A$	Array selection
$A\{\}A$	Sink
$A\{\alpha\}A$	Left (Lev)
$A\{\omega\}A$	Right (Dex)
$A\{\alpha\ \ \omega\}A$	Link
$\{0\}A$	Zero
$\{0\}^{..}A$	Zero Each
$,/PV$	Join
$\supset\phi A$	Upper right item ($\square ml < 2$)
$\uparrow\phi A$	Upper right item ($\square ml \geq 2$)
$\supset\phi,A$	Lower right item ($\square ml < 2$)
$\uparrow\phi,A$	Lower right item ($\square ml \geq 2$)
$0=\rho V$	Zero shape
$0=\rho\rho A$	Zero rank
$0=\equiv A$	Zero depth
$\square AV\iota CA$	Atomic vector index
$CM\{(\downarrow\alpha)\iota\downarrow\omega\}CM$	Matrix Iota
$\downarrow\phi\uparrow PV$	Nested vector transpose ($\square ml < 2$)
$\downarrow\phi\supset PV$	Nested vector transpose ($\square ml \geq 2$)
$\wedge\backslash'\ '=CA$	Mask of leading blanks.
$+/\wedge\backslash'\ '=CA$	Number of leading blanks
$+/\wedge\backslash BA$	Number of leading ones
$\{(\vee\backslash'\ '\neq\omega)/\omega\}CV$	Trim leading blanks
$\{(+/\wedge\backslash'\ '=\omega)\downarrow\omega\}CV$	Trim leading blanks
$\sim\circ'\ '^{..}\downarrow CA$	No-blank split
$\{(+/\vee\backslash'\ '\neq\phi\omega)\uparrow^{..}\downarrow\omega\}CA$	No-trailing-blank split

Notes

Sequence Selection `/ι` and **Index Selection** `/ιρ`, as well as providing an execution time advantage, reduce intermediate workspace usage and consequently, the incidence of memory compactions and the likelihood of a WS FULL.

Array Selection `NV⊃¨⊂A`, is implemented as `A[NV]`, which is significantly faster. The two are equivalent but the former may now be used as a matter of taste with no performance penalty.

Join `,/` is currently special-cased only for vectors of vectors or scalars. Otherwise, the expression is evaluated as a series of concatenations. Recognition of this idiom turns **join** from an "n-squared" algorithm into a linear one. In other words, the improvement factor is proportional to the size of the argument vector.

Upper and **Lower Right Item** now take constant time. Without idiom recognition, the time taken depends linearly on the number of items in the argument.

Zero Depth `0=≡` takes a small constant time. Without idiom recognition, time taken would depend on the size and depth of the argument, which in the case of a deeply nested array, could be significant.

Nested vector transpose `↓⍉↑` is special-cased only for a vector of nested vectors, each of whose items is of the same length.

Matrix Iota `{(↓α)ι↓ω}`. As well as being quicker, the Matrix Iota idiom can accommodate much larger character matrices. It is particularly effective when bound with a left argument using the compose operator:

```
find←cmat∘{(↓α)ι↓ω}      ⍝ find rows in cmat table.
```

In this case, the internal hash table for `cmat` is retained so that it does not need to be generated each time the monadic derived function `find` is applied to a character matrix argument.

Trim leading blanks `{(∨\' '≠ω)/ω}` and `{(+/∧\' '=ω)↓ω}` are two codings of the same idiom. Both use the same C code for evaluation.

No-blank split `~∘' '¨↓` typically takes a character matrix argument and returns a vector of character vectors from which, all blanks have been removed. An example might be the character matrix of names returned by the system function `⎕NL`. In general, this idiom accommodates character arrays of any rank.

No-trailing-blank split `{(+/∨\' '≠⌽ω)↑¨↓ω}` typically takes a character matrix argument and returns a vector of character vectors. Any embedded blanks in each row are preserved but trailing blanks are removed. In general, this idiom accommodates character arrays of any rank.

Search Functions and Hash Tables

Primitive dyadic *search* functions, such as ⍳ (index of) and ∊ (membership) have a *principal* argument in which items of the other *subject* argument are located.

In the case of ⍳, the principal argument is the one on the left and in the case of ∊, it is the one on the right. The following table shows the principal (P) and subject (s) arguments for each of the functions.

`P ⍳ s`	Index of
`s ∊ P`	Membership
`s ∩ P`	Intersection
`s ∪ P`	Union
`s ~ P`	Without
`P {(↓⍺)⍳↓⍵} s`	Matrix Iota (idiom)

The Dyalog APL implementation of these functions already uses a technique known as *hashing* to improve performance over a simple linear search.

Building a *hash table* for the principal argument takes a significant time but is rewarded by a considerably quicker search for each item in the subject. Unfortunately, the hash table is discarded each time the function completes and must be reconstructed for a subsequent call (even if its principal argument is identical to that in the previous one).

For optimal performance of *repeated* search operations, the hash table may be retained between calls, by binding the function with its principal argument using the primitive ∘ (compose) operator. The retained hash table is then used directly whenever this monadic derived function is applied to a subject argument.

Notice that retaining the hash table pays off only on a second or subsequent application of the derived function. This usually occurs in one of two ways: either the derived function is named for later (and repeated) use, as in the first example below or it is applied repeatedly as the operand of a primitive or defined operator, as in the second example.

Example: naming a derived function.

```
words←'red' 'ylo' 'grn' 'brn' 'blu' 'pnk' 'blk'

find←words∘⍳                    ⍝ monadic find function
find'blk' 'blu' 'grn' 'ylo'  ⍝
7 5 3 2
find'grn' 'brn' 'ylo' 'red'  ⍝ fast find
3 4 2 1
```

Example: repeated application by (¨) each operator.

```
∊∘⎕A¨'This' 'And' 'That'
1 0 0 0  1 0 0  1 0 0 0
```

Control Structures

Control structures provide a means to control the flow of execution in your APL programs.

Traditionally, lines of APL code are executed one by one from top to bottom and the only way to alter the flow of execution is using the branch arrow. So how do you handle logical operations of the form "If this, do that; otherwise do the other"?.

In APL this is often not a problem because many logical operations are easily performed using the standard array handling facilities that are absent in other languages. For example, the expression :

```
STATUS←(1+AGE<16)⊃'Adult' 'Minor'
```

sets *STATUS* to '*Adult*' if *AGE* is 16 or more; otherwise sets *STATUS* to '*Minor*'.

Things become trickier if, depending upon some condition, you wish to execute one set of code instead of another, especially when the code fragments cannot conveniently be packaged as functions. Nevertheless, careful use of array logic, defined operators, the execute primitive function and the branch arrow can produce high quality maintainable and comprehensible APL systems.

Control structures provide an additional mechanism for handling logical operations and decisions. Apart from providing greater affinity with more traditional languages, Control structures may enhance comprehension and reduce programming errors, especially when the logic is complex. Control structures are not, however, a replacement for the standard logical array operations that are so much a part of the APL language.

Control Structures are blocks of code in which the execution of APL statements follows certain rules and conditions. Control structures are implemented using a set of *control words* that all start with the colon symbol (:). Control Words are case-insensitive.

There are eight different types of control structures defined by the control words, `:If`, `:While`, `:Repeat`, `:For`, `:Select`, `:With`, `:Trap` and `:Hold` . Each one of these control words may occur only at the beginning of an APL statement and indicates the start of a particular type of control structure.

Within a control structure, certain other control words are used as qualifiers. These are `:Else`, `:ElseIf`, `:AndIf`, `:OrIf`, `:Until`, `:Case` and `:CaseList`.

A third set of control words is used to identify the end of a particular control structure. These are `:EndIf`, `:EndWhile`, `:EndRepeat`, `:EndFor`, `:EndSelect`, `:EndWith`, `:EndTrap` and `:EndHold`. Although formally distinct, these control words may all be abbreviated to `:End`.

Finally, the `:GoTo`, `:Return`, `:Leave` and `:Continue` control words may be used to conditionally alter the flow of execution within a control structure.

Key to Notation

The following notation is used to describe Control Structures within this section:

```
aexp    an expression returning an array,

bexp    an expression returning a single boolean value (0 or 1),

var     loop variable used by :For control structure,

code    0 or more lines of APL code, including other (nested) control structures,

andor   either one or more :AndIf statements, or one or more :OrIf statements.


        |
        .-----------------------.
        |                       |
        |<-------------.        |<-------------.
        |             |         |              |
        :AndIf bexp-----'       :OrIf bexp------'
        |                       |
        |<---------------------'
        |
```

If Statement   *:If bexp*

The simplest *:If* control structure is a single condition of the form:

```
[1]     :If AGE<21
[2]         expr 1
[3]         expr 2
[5]     :EndIf
```

If the test condition (in this case *AGE<21*) is true, the statements between the *:If* and the *:EndIf* will be executed. If the condition is false, none of these statements will be run and execution resumes after the *:EndIf*. Note that the test condition to the right of *:If* must return a single element boolean value 1 (true) or 0 (false).

:If control structures may be considerably more complex. For example, the following code will execute the statements on lines [2-3] if *AGE<21* is 1 (true), **or alternatively**, the statement on line [6] if *AGE<21* is 0 (false).

```
[1]     :If AGE<21
[2]         expr 1
[3]         expr 2
[5]     :Else
[6]         expr 3
[7]     :EndIf
```

Instead of a single condition, it is possible to have multiple conditions using the *:ElseIf* control word. For example,

```
[1]     :If WINEAGE<5
[2]         'Too young to drink'
[5]     :ElseIf WINEAGE<10
[6]         'Just Right'
[7]     :ElseIf WINEAGE<15
[8]         'A bit past its prime'
[9]     :Else
[10]      'Definitely over the hill'
[11]    :EndIf
```

Notice that APL executes the expression(s) associated with the **first** condition that is true or those following the *:Else* if **none** of the conditions are true.

The *:AndIf* and *:OrIf* control words may be used to define a block of conditions and so refine the logic still further. You may qualify an *:If* or an *:ElseIf* with one or more *:AndIf* statements **or** with one or more *:OrIf* statements. You may not however mix *:AndIf* and *:OrIf* in the same conditional block. For example:

```
[1]     :If WINE.NAME≡'Château Lafitte'
[2]     :AndIf WINE.YEAR∈1962 1967 1970
[3]         'The greatest?'
[4]     :ElseIf WINE.NAME≡'Château Latour'
[5]     :Orif WINE.NAME≡'Château Margaux'
[6]     :Orif WINE.PRICE>100
[7]         'Almost as good'
[8]     :Else
[9]         'Everyday stuff'
[10]    :EndIf
```

Please note that in a *:If* control structure, the conditions associated with each of the condition blocks are executed in order until an entire condition block evaluates to true. At that point, the APL statements following this condition block are executed. None of the conditions associated with any other condition block are executed. Furthermore, if an *:AndIf* condition yields 0 (false), it means that the entire block must evaluate to false so the system moves immediately on to the next block without executing the other conditions following the failing *:AndIf*. Likewise, if an *:OrIf* condition yields 1 (true), the entire block is at that point deemed to yield true and none of the following *:OrIf* conditions in the same block are executed.

:If Statement

```
|
:If bexp
|
.-------.
|       |
|      andor
|       |
|<------'
|
code
|
|<---------------------------.
|                            |
.-------.-------.            |
|       |       |            |
|      :Else   :ElseIf bexp  |
|       |       |            |
|       |      .-------.     |
|       |      |       |     |
|       |      |      andor  |
|       |      |       |     |
|       |      |<------'     |
|       |      |            |
|      code   code          |
|       |      |            |
|<------'      `------------'
|
:End[If]
|
```

While Statement   `:While bexp`

The simplest `:While` loop is :

```
[1]    I←0
[2]    :While I<100
[3]        expr1
[4]        expr2
[5]        I←I+1
[6]    :EndWhile
```

Unless `expr1` or `expr2` alter the value of `I`, the above code will execute lines `[3-4]` 100 times. This loop has a single condition; the value of `I`. The purpose of the `:EndWhile` statement is solely to mark the end of the iteration. It acts the same as if it were a branch statement, branching back to the `:While` line.

An alternative way to terminate a `:While` structure is to use a `:Until` statement. This allows you to add a second condition. The following example reads a native file sequentially as 80-byte records until it finds one starting with the string `'Widget'` or reaches the end of the file.

```
[1]    I←0
[2]    :While I<□NSIZE ¯1
[3]        REC←□NREAD ¯1 82 80
[4]        I←I+ρREC
[5]    :Until 'Widget'≡6ρREC
```

Instead of single conditions, the tests at the beginning and end of the loop may be defined by more complex ones using `:AndIf` and `:OrIf`. For example:

```
[1]    :While 100>i
[2]    :AndIf 100>j
[3]        i j←foo i j
[4]    :Until 100<i+j
[5]    :OrIf i<0
[6]    :OrIf j<0
```

In this example, there are complex conditions at both the start and the end of the iteration. Each time around the loop, the system tests that both `i` and `j` are less than or equal to 100. If either test fails, the iteration stops. Then, after `i` and `j` have been recalculated by `foo`, the iteration stops if `i+j` is equal to or greater than 100, or if either `i` or `j` is negative.

:While Statement

```
|
:While bexp
|
.-------.
|       |
|       andor
|       |
|<------'
|
code
|
.--------------.
|              |
:End[While]    :Until bexp
|              |
|              .-------.
|              |       |
|              |       andor
|              |       |
|              |<------'
|              |
|<-------------'
|
```

Repeat Statement   `:Repeat`

The simplest type of `:Repeat` loop is as follows. This example executes lines `[3-5]` 100 times. Notice that as there is no conditional test at the beginning of a `:Repeat` structure, its code statements are executed at least once.

```
[1]    I←0
[2]    :Repeat
[3]        expr1
[4]        expr2
[5]        I←I+1
[6]    :Until I=100
```

You can have multiple conditional tests at the end of the loop by adding `:AndIf` or `:OrIf` expressions. The following example will read data from a native file as 80-character records until it reaches one beginning with the text string `'Widget'` or reaches the end of the file.

```
[1]    :Repeat
[2]        REC←⎕NREAD ¯1 82 80
[3]    :Until 'Widget'≡6⍴REC
[4]    :OrIf 0=⍴REC
```

A `:Repeat` structure may be terminated by an `:EndRepeat` (or `:End`) statement in place of a conditional expression. If so, your code must explicitly jump out of the loop using a `:Leave` statement or by branching. For example:

```
[1]    :Repeat
[2]        REC←⎕NREAD ¯1 82 80
[3]        :If 0=⍴REC
[4]        :OrIf 'Widget'≡6⍴REC
[5]            :Leave
[6]        :EndIf
[7]    :EndRepeat
```

:Repeat Statement

```
|
:Repeat
|
code
|
.--------------.
|              |
:End[Repeat]    :Until bexp
|              |
|              .-------.
|              |       |
|              |       andor
|              |       |
|              |<------'
|              |
|<-------------'
|
```

:Repeat Statement
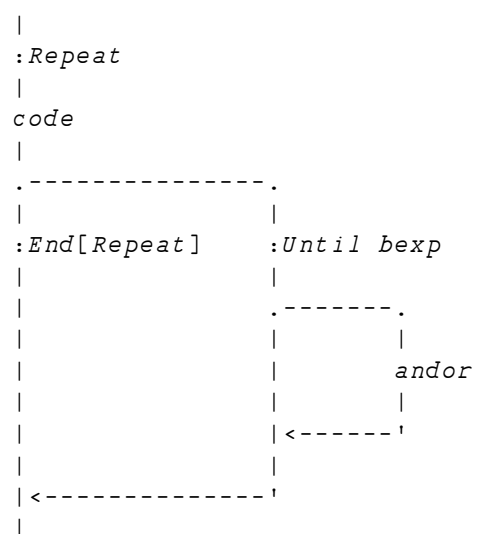
For Statement   `:For var :In[Each] aexp`

Single Control Variable

The `:For` loop is used to execute a block of code for a series of values of a particular control variable. For example, the following would execute lines `[2-3]` successively for values of `I` from `⎕IO` to 10

```
[1]     :For I :In ⍳10
[2]         expr1 I
[3]         expr2 I
[4]     :EndFor
```

The way a `:For` loop operates is as follows. On encountering the `:For`, the expression to the right of `:In` is evaluated and the result stored. This is the *control array*. The *control variable*, named to the right of the `:For`, is then assigned the first value in the control array, and the code between `:For` and `:EndFor` is executed. On encountering the `:EndFor`, the control variable is assigned the next value of the control array and execution of the code is performed again, starting at the first line after the `:For`. This process is repeated for each value in the control array.

Note that if the control array is empty, the code in the `:For` structure is not executed. Note too that the control array may be any rank and shape, but that its elements are assigned to the control variable in ravel order.

The control array may contain any type of data. For example, the following code resizes (and compacts) all your component files

```
[1]     :For FILE :In (↓⎕FLIB '')~¨' '
[2]         FILE ⎕FTIE 1
[3]         ⎕FRESIZE 1
[4]         ⎕FUNTIE 1
[5]     :EndFor
```

You may also nest `:For` loops. For example, the following expression finds the timestamp of the most recently updated component in all your component files.

```
[1]     TS←0
[2]     :For FILE :In (↓⎕FLIB '')~¨' '
[3]         FILE ⎕FTIE 1
[4]         START END←2⍴⎕FSIZE 1
[5]         :For COMP :In (START-1)↓⍳END-1
[6]             TS⌈←¯1↑⎕FREAD FILE COMP
[7]         :EndFor
[8]         ⎕FUNTIE 1
[9]     :EndFor
```

Multiple Control Variables

The :*For* control structure can also take multiple variables. This has the effect of doing a strand assignment each time around the loop.

For example :*For a b c :in* (1 2 3)(4 5 6), sets *a b c*←1 2 3, first time around the loop and *a b c*←4 5 6, the second time.

Another example is :*For i j :In* ιρ*Matrix*, which sets *i* and *j* to each row and column index of *Matrix*.

:InEach Control Word

```
:For var ... :InEach value ...
```

In a :*For* control structure, the keyword :*InEach* is an alternative to :*In*.

For a single control variable, the effect of the keywords is identical but for multiple control variables the values vector is inverted.

The distinction is best illustrated by the following equivalent examples:

```
:For a b c :In (1 2 3)(3 4 5)(5 6 7)(7 8 9)
    ⎕←a b c
:EndFor
```

```
:For a b c :InEach (1 3 5 7)(2 4 6 8)(3 5 7 9)
    ⎕←a b c
:EndFor
```

In each case, the output from the loop is:

```
1 2 3
3 4 5
5 6 7
7 8 9
```

Notice that in the second case, the number of items in the values vector is the same as the number of control variables. A more typical example might be.

```
:For a b c :InEach avec bvec cvec
    ...
:EndFor
```

Here, each time around the loop, control variable *a* is set to the next item of *avec*, *b* to the next item of *bvec* and *c* to the next item of *cvec*.

:For Statement

```
|
:For var :In[Each] aexp
|
code
|
:End[For]
|
```

Select Statement   *:Select aexp*

A *:Select* structure is used to execute alternative blocks of code depending upon the value of an array. For example, the following displays '*I is 1*' if the variable *I* has the value 1, '*I is 2*' if it is 2,or '*I is neither 1 nor 2*' if it has some other value.

```
[1]    :Select I
[2]    :Case 1
[3]        'I is 1'
[4]    :Case 2
[5]        'I is 2'
[6]    :Else
[7]        'I is neither 1 nor 2'
[8]    :EndSelect
```

In this case, the system compares the value of the array expression to the right of the *:Select* statement with each of the expressions to the right of the *:Case* statements and executes the block of code following the one that matches. If none match, it executes the code following the *:Else* (which is optional). Note that comparisons are performed using the ≡ primitive function, so the arrays must match exactly. Note also that not all of the *:Case* expressions are necessarily evaluated because the process stops as soon as a matching expression is found.

Instead of a *:Case* statement, you may also use a *:CaseList* statement. If so, the *enclose of* the array expression to the right of *:Select* is tested for membership of the array expression to the right of the *:CaseList* using the ∊ primitive function.

Example

```
[1]     :Select ?6 6
[2]     :Case 6 6
[3]         'Box Cars'
[4]     :Case 1 1
[5]         'Snake Eyes'
[6]     :CaseList 2ρ¨ι6
[7]         'Pair'
[8]     :CaseList (ι6),¨φι6
[9]         'Seven'
[10]    :Else
[11]        'Unlucky'
[12]    :EndSelect
```

:Select Statement

```
    |
    :Select aexp
    |
    |<---------------------------------------------.
    |                                              |
    .-------.-------.--------------.               |
    |       |       |              |               |
    |     :Else   :Case aexp    :CaseList aexp     |
    |       |       |              |               |
    |       |       |<-------------'               |
    |       |       |                              |
    |     code    code                             |
    |       |       |                              |
    |<------'       `------------------------------'
    |
    :End[Select]
```

With Statement   `:With obj`

`:With` is a control structure that may be used to simplify a series of references to an object or namespace. `:With` changes into the specified namespace for the duration of the control structure, and is terminated by `:End[With]`. For example, you could update several properties of a Grid object `F.G` as follows:

```
:With F.G
    Values←4 3ρ0
    RowTitles←'North' 'South' 'East' 'West'
    ColTitles←'Cakes' 'Buns' 'Biscuits'
:EndWith
```

`:With` is analogous to `⎕CS` in the following senses:

- The namespace argument to `:With` is interpreted relative to the current space.
- Local names in the containing defined function continue to be visible in the new space.
- Global references from within the `:With` control structure are to names in the new space.
- Exiting the defined function from within a `:With` control structure causes the space to revert to the one from which the function was called.

On leaving the `:With` control structure, execution reverts to the original namespace. Notice however that the interpreter does not detect branches (→) out of the control structure.

`:With` control structures can be nested in the normal fashion:

```
[1]   :With 'x'          ⍝ Change to #.x
[2]       :With 'y'      ⍝ Change to #.x.y
[3]           :With ⎕SE  ⍝ Change to ⎕SE
[4]               ...    ⍝ ... in ⎕SE
[5]           :EndWith   ⍝ Back to #.x.y
[6]       :EndWith       ⍝ Back to #.x
[7]   :EndWith           ⍝ Back to #
```

:With Statement

```
|
:With namespace (ref or name)
|
code
|
:End[With]
```

Hold Statement   `:Hold tkns`

Whenever more than one thread tries to access the same piece of data or shared resource at the same time, you need some type of synchronisation to control access to that data. This is provided by `:Hold`.

`:Hold` provides a mechanism to control thread entry into a critical section of code. `tkns` must be a simple character vector or scalar, or a vector of character vectors. `tkns` represents a set of 'tokens', all of which must be acquired before the thread can continue into the control structure. `:Hold` is analogous to the component file system `⎕FHOLD`.

Within the whole active workspace, a token with a particular value may be held only once. If the hold succeeds, the current thread *acquires* the tokens and execution continues with the first phrase in the control structure. On exit from the structure, the tokens are released for use by other threads. If the hold fails, because one or more of the tokens is already in use:

1. If there is no `:Else` clause in the control structure, execution of the thread is blocked until the requested tokens become available.

2. Otherwise, acquisition of the tokens is abandoned and execution resumed immediately at the first phrase in the `:Else` clause.

`tkns` can be either a single token:

```
'a'
'Red'
'#.Util'
''
'Program Files'
```

… Or a number of tokens:

```
'red' 'green' 'blue'
'doe' 'a' 'deer'
,¨'abc'
↓⎕nl 9
```

Pre-processing removes trailing blanks from each token before comparison, so that, for example, the following two statements are equivalent:

```
:Hold 'Red' 'Green'

:Hold ↓2 5⍴'Red  Green'
```

Unlike `⎕FHOLD`, a thread does not release all existing tokens before attempting to acquire new ones. This enables the nesting of holds, which can be useful when multiple threads are concurrently updating parts of a complex data structure.

In the following example, a thread updates a critical structure in a child namespace, and then updates a structure in its parent space. The holds will allow all 'sibling' namespaces to update concurrently, but will constrain updates to the parent structure to be executed one at a time.

```
:Hold ⎕cs''        ⍝ Hold child space
    ...            ⍝ Update child space
   :Hold ##.⎕cs''  ⍝ Hold parent space
       ...         ⍝ Update Parent space
   :EndHold
   ...
:EndHold
```

However, with the nesting of holds comes the possibility of a 'deadlock'. For example, consider the two threads:

Thread 1	**Thread 2**

```
:Hold 'red'            :Hold 'green'
   ...                    ...
   :Hold 'green'          :Hold 'red'
      ...                    ...
   :EndHold               :EndHold
:EndHold               :EndHold
```

In this case if both threads succeed in acquiring their first hold, they will both block waiting for the other to release its token. Fortunately, the interpreter detects such cases and issues an error `(1008) DEADLOCK`. You can avoid deadlock by ensuring that threads always attempt to acquire tokens in the same chronological order, and that threads never attempt to acquire tokens that they already own.

Note that token acquisition for any particular `:Hold` is atomic, that is, either *all* of the tokens or *none* of them are acquired. The following example *cannot* deadlock:

Thread 1	**Thread 2**

```
:Hold 'red'
   ...                 :Hold 'green' 'red'
   :Hold 'green'          ...
      ...              :EndHold
   :EndHold
:EndHold
```

Examples

:Hold could be used for example, during the update of a complex data structure that might take several lines of code. In this case, an appropriate value for the token would be the name of the data structure variable itself, although this is just a programming convention: the interpreter does not associate the token value with the data variable.

```
:Hold'Struct'
    ...                  ⍝ Update Struct
    Struct ← ...
:EndHold
```

The next example guarantees exclusive use of the current namespace:

```
:Hold ⎕CS''          ⍝ Hold current space
    ...
:EndHold
```

The following example shows code that holds two positions in a vector while the contents are exchanged.

```
:Hold ⍕¨to fm
    :If >/vec[fm to]
        vec[fm to]←vec[to fm]
    :End
:End
```

Between obtaining the next available file tie number and using it:

```
:Hold '⎕FNUMS'
    tie←1+⌈/0,⎕FNUMS
    fname ⎕FSTIE tie
:End
```

The above hold is not necessary if the code is combined into a single line:

```
    fname ⎕FSTIE tie←1+⌈/0,⎕FNUMS
or,
    tie←fname ⎕FSTIE 0
```

Note that :*Hold*, like its component file system counterpart □*FHOLD*, is a device to enable *co-operating* threads to synchronise their operation.

:*Hold* does not *prevent* threads from updating the same data structures concurrently, it prevents threads only from :*Hold*-ing the same tokens.

:Hold Statement

```
|
:Hold token(s)
|
code
|
|-------.
|       |
|       :Else
|       |
|       code
|       |
|<------.
|
:End[Hold]
|
```

Trap Statement   :*Trap ecode*

:*Trap* is an error trapping mechanism that can be used in conjunction with, or as an alternative to, the □*TRAP* system variable. It is equivalent to APL2's □*EA*, except that the code to be executed is not restricted to a single expression and is not contained within quotes (and so is slightly more efficient).

Operation

The segment of code immediately following the :*Trap* keyword is executed. On completion of this segment, if no error occurs, control passes to the code following :*End[Trap]*.

If an error does occur, the event code (error number) is noted and:

- If the error occurred within a sub-function, the system cuts the execution stack back to the function containing the :*Trap* keyword. In this respect, :*Trap* behaves like □*TRAP* with a '*C*' qualifier.

- The system searches for a :*Case[List]* representing the event code.

- If there is such a :*Case[List]*, or failing that, an :*Else* keyword, execution continues from this point.

Otherwise, control passes to the code following :*End[Trap]* and no error processing occurs.

Note that the error trapping is in effect **only** during execution of the initial code segment. It is disabled (or surrendered to outer level :*Trap*s or □*TRAP*s) immediately a trapped error occurs. In particular, the error trap is no longer in effect during processing of :*Case[List]*'s argument or in the code following the :*Case[List]* or :*Else* statement. This avoids the situation sometimes encountered with □*TRAP* where an infinite 'trap loop' occurs. If an error which is not specified occurs, it is processed by outer :*Trap*s, □*TRAP*s, or default system processing in the normal fashion.

Examples

```
    ∇ lx
[1]   :Trap 1000        ⍝ Cutback and exit on interrupt
[2]       Main ...
[3]   :EndTrap
    ∇


    ∇ ftie←Fcreate file      ⍝ Create null component file
[1]   ftie←1+⌈/0,⎕fnums      ⍝ next tie number.
[2]   :Trap 22              ⍝ Trap FILE NAME ERROR
[3]       file ⎕fcreate ftie ⍝ Try to create file.
[4]   :Else
[5]       file ⎕ftie ftie    ⍝ Tie the file.
[6]       file ⎕ferase ftie  ⍝ Drop the file.
[7]       file ⎕fcreate ftie ⍝ Create new file.
[8]   :EndTrap
    ∇


    ∇ lx ⍝ Distinguish various cases
[1]   :Trap 0 1000
[2]       Main ...
[3]   :Case 1002
[4]       'Interrupted ...'
[5]   :CaseList 1 10 72 76
[6]       'Not enough resources'
[7]   :CaseList 17+⍳20
[8]       'File System Problem'
[9]   :Else
[10]      'Unexpected Error'
[11]  :EndTrap
    ∇
```

Note that :Traps can be nested:

```
    ∇ ntie←Ntie file              ⍝ Tie native file
[1]   ntie←¯1+⌊/0,⎕nnums          ⍝ Next native tie num
[2]   :Trap 22                    ⍝ Trap FILE NAME ERRO
R
[3]       file ⎕ntie ntie         ⍝ Try to tie file
[4]   :Else
[5]       :Trap 22                ⍝ Trap FILE NAME ERRO
R
[6]           (file,'.txt')⎕ntie ntie ⍝ Try with .txt extn
[7]       :Else
[8]           file ⎕ncreate ntie   ⍝ Create null file.
```

```
[9]          :EndTrap
[10]   :EndTrap
     ∇
```

:Trap Statement

```
|
:Trap <ecode>
|
code
|
|<--------------------------------.
|                                 |
.-------.-------.                 |
|       |       |                 |
|     :Else    :Case[List] <ecode> |
|       |       |                 |
|       |       |                 |
|       |       |                 |
|     code     code               |
|       |       |                 |
|<------'       `-----------------'
|
:End[Trap]
|
```

Where *ecode* is a scalar or vector of $\square TRAP$ event codes (see *Chapter 5*).

Note that within the :*Trap* control structure, :*Case* is used for a single event code and :*CaseList* for a vector of event codes.

GoTo Statement   :*GoTo aexp*

A :*GoTo* statement is a direct alternative to → (branch) and causes execution to jump to the line specified by the first element of *aexp*.

The following are equivalent. See Branch for further details.

```
→Exit
:GoTo Exit

→(N<I←I+1)/End
:GoTo (N<I←I+1)/End

→1+⎕LC
:GoTo 1+⎕LC

→10
:GoTo 10
```

Return Statement   :*Return*

A :*Return* statement causes a function to terminate and has exactly the same effect as →0.

The :*Return* control word takes no argument.

A :*Return* statement may occur anywhere in a function or operator.

Leave Statement   :*Leave*

A :Leave statement is used to explicitly terminate the execution of a block of statements within a :*For*, :*Repeat* or :*While* control structure.

The :*Leave* control word takes no argument.

Continue Statement   :*Continue*

A :*Continue* statement starts the next iteration of the immediately surrounding :*For*, :*Repeat* or :*While* control loop.

When executed within a :*For* loop, the effect is to start the body of the loop with the next value of the iteration variable.

When executed within a :*Repeat* or :*While* loop, if there is a trailing test that test is executed and, if the result is true, the loop is terminated. Otherwise the leading test is executed in the normal fashion.

Locked Functions & Operators

A defined operation may be locked by the system function $\square LOCK$. A locked operation may not be displayed or edited. The system function $\square CR$ returns an empty matrix of shape 0 0 and the system functions $\square NR$ and $\square VR$ return an empty vector for a locked operation.

Stop and trace vectors may be set by the system functions $\square STOP$ and $\square TRACE$ respectively. Existing stop or trace settings are cancelled when an operation is locked.

A locked operation may not be suspended, nor may a locked operation remain pendent when execution is suspended. The state indicator is cut back as described below.

The State Indicator

The state of execution is dynamically recorded in the STATE INDICATOR. The state indicator identifies the chain of execution for operators, functions and the evaluated or character input/output system variables ($\square$ and $\square$). At the top of the state indicator is the most recently activated operation.

Execution may be suspended by an interrupt, induced by the user, the system, or by a signal induced by the system function $\square SIGNAL$ or by a stop control set by the system function $\square STOP$. If the interrupt (or event which caused the interrupt) is not defined as a trappable event by the system variable $\square TRAP$, the state indicator is cut back to the first of either a defined operation or the evaluated input prompt ($\square$) such that there is no locked defined operation in the state indicator. The topmost operation left in the state indicator is said to be SUSPENDED. Other operations in the chain of execution are said to be PENDENT.

The state indicator may be examined when execution is suspended by the system commands $)SI$ and $)SINL$. The line numbers of defined operations in the state indicator are given also by the system variable $\square LC$.

Suspended execution may be resumed by use of the Branch function (see *Chapter 3*). Whilst execution is suspended, it is permitted to enter any APL expression for evaluation, thereby adding to the existing state indicator. Therefore, there may be more than one LEVEL OF SUSPENSION in the state indicator. If the state indicator is cut back when execution is suspended, it is cut back no further than the prior level of suspension (if any).

Examples

```
      ∇ F
[1]    G
      ∇

      ∇ G
[1]    'FUNCTION G'+
      ∇

      ⍕'F'
SYNTAX ERROR
G[1] 'FUNCTION G'+
      ^

      )SI
G[1]*
F[1]
⍕

      □LOCK'G'

      ⍕'F'
SYNTAX ERROR
F[1] G
      ^

      )SI
F[1]*
⍕
G[1]*
F[1]
⍕
```

A suspended or pendent operation may be edited by the system editor or redefined using □FX provided that it is visible and unlocked. However, pendent operations retain their original definition until they complete, or are cleared from the State Indicator. When a new definition is applied, the state indicator is repaired if necessary to reflect changes to the operations, model syntax, local names, or labels.

Dynamic Functions & Operators

A Dynamic Function (operator) is an alternative function definition style suitable for defining small to medium sized functions. It bridges the gap between operator expressions: `rank←ρ∘ρ` and full 'header style' definitions such as:

```
    ∇ rslt←larg func rarg;local...
```

In its simplest form, a dynamic function is an APL expression enclosed in curly braces `{}` possibly including the special characters α and ω to represent the left and right arguments of the function respectively. For example:

```
    {(+/ω)÷ρω} 1 2 3 4      ⍝ Arithmetic Mean (Average)
2.5
    3 {ω*÷α} 64             ⍝ αth root
4
```

Dynamic functions can be named in the normal fashion:

```
    mean←{(+/ω)÷ρω}
    mean¨(2 3)(4 5)
 2.5  4.5
```

Dynamic Functions can be defined and used in any context where an APL function may be found, in particular:

- In immediate execution mode as in the examples above.
- Within a defined function or operator.
- As the operand of an operator such as each (¨).
- Within another dynamic function.

The last point means that it is easy to define nested local functions.

Multi-Line Dynamic Functions

The single expression which provides the result of the Dynamic Function may be preceded by any number of assignment statements. Each such statement introduces a name which is local to the function.

For example in the following, the expressions *sum←* and *num←* create **local** variables *sum* and *num*.

```
mean←{            ⍝ Arithmetic mean
    sum←+/ω       ⍝ Sum of elements
    num←⍴ω        ⍝ Number of elements

    sum÷num       ⍝ Mean
}
```

Note that Dynamic Functions may be commented in the usual way using ⍝.

When the interpreter encounters a local definition, a new local name is created. The name is shadowed dynamically exactly as if the assignment had been preceded by: ⎕*shadow name ◇*.

It is **important** to note the distinction between the two types of statement above. There can be **many** assignment statements, each introducing a new local variable, but only a **single** expression where the result is not assigned. As soon as the interpreter encounters such an expression, it is evaluated and the result returned immediately as the result of the function.

For example, in the following,

```
mean←{            ⍝ Arithmetic mean
    sum←+/ω       ⍝ Sum of elements
    num←⍴ω        ⍝ Number of elements

    sum,num       ⍝ Attempt to show sum,num (wrong)!

    sum÷num       ⍝ ... and return result.
}
```

... as soon as the interpreter encounters the expression *sum,num*, the function terminates with the two element result (*sum,num*) and the following line is not evaluated.

To display arrays to the session from within a Dynamic function, you can use the
explicit display forms ⎕← or ⍞← as in:

```
mean←{              ⍝ Arithmetic mean
    sum←+/ω         ⍝ Sum of elements
    num←ρω          ⍝ Number of elements

    ⎕←sum,num       ⍝ show sum,num.

    sum÷num         ⍝ ... and return result.
}
```

Note that local definitions can be used to specify local nested Dynamic Functions:

```
rms←{                   ⍝ Root Mean Square
    root←{ω*0.5}        ⍝ ∇ Square root
    mean←{(+/ω)÷ρω}     ⍝ ∇ Mean
    square←{ω×ω}        ⍝ ∇ Square

    root mean square ω
}
```

Default Left Argument

The special syntax: α←*expr* is used to give a default value to the left argument if a
Dynamic Function is called monadically. For example:

```
root←{        ⍝ αth root
    α←2       ⍝ default to sqrt
    ω*÷α
}
```

The expression to the right of α← is evaluated *only* if its Dynamic Function is called
with no left argument.

Guards

A Guard is a Boolean-single valued expression followed on the right by a `':'`. For example:

```
0≡≡ω:           ⍝ Right arg simple scalar

α<0:            ⍝ Left arg negative
```

The guard is followed by a single APL expression: the result of the function.

```
ω≥0: ω*0.5     ⍝ Square root if non-negative.
```

A Dynamic function may contain any number of guarded expressions each on a separate line (or collected on the same line separated by diamonds). Guards are evaluated in turn until one of them yields a 1. The corresponding expression to the right of the guard is then evaluated as the result of the function.

If an expression occurs without a guard, it is evaluated immediately as the default result of the function. For example:

```
sign←{
    ω>0: '+ve'      ⍝ Positive
    ω=0: 'zero'     ⍝ zero
         '-ve'      ⍝ Negative (Default)
}
```

Local definitions and guards can be interleaved in any order.

Note again that any code following the first unguarded expression (which terminates the function) could never be executed and would therefore be redundant.

Shy Result

Dynamic Functions are usually 'pure' functions that take arguments and return explicit results. Occasionally, however, the main purpose of the function might be a side-effect such as the display of information in the session, or the updating of a file, and the value of a result, a secondary consideration. In such circumstances, you might want to make the result 'shy', so that it is discarded unless the calling context requires it. This can be achieved by assigning a dummy variable after a (true) guard:

```
log←{                   ⍝ Append ω to file α.
    tie←α ⎕fstie 0      ⍝ tie number for file,
    cno←ω ⎕fappend tie  ⍝ new component number,
    tie←⎕funtie tie     ⍝ untie file,
    1:rslt←cno          ⍝ comp number as shy result.
}
```

Static Name Scope

When an inner (nested) Dynamic Function refers to a name, the interpreter searches for it by looking outwards through enclosing Dynamic Functions, rather than searching back along the execution stack. This regime, which is more appropriate for nested functions, is said to employ **static scope** instead of APL's usual **dynamic scope**. This distinction becomes apparent only if a call is made to a function defined at an outer level. For the more usual inward calls, the two systems are indistinguishable.

For example, in the following function, variable *type* is defined both within *which* itself and within the inner function *fn1*. When *fn1* calls outward to *fn2* and *fn2* refers to *type*, it finds the outer one (with value `'static'`) rather than the one defined in *fn1*:

```
which←{

    type←'static'

    fn1←{
        type←'dynamic'
        fn2 ω
    }

    fn2←{
        type ω
    }

    fn1 ω
}


which'scope'
static  scope
```

Tail Calls

A novel feature of the implementation of Dynamic Functions is the way in which tail calls are optimised.

When a Dynamic Function calls a sub-function, the result of the call may or may not be modified by the calling function before being returned. A call where the result is passed back immediately without modification is termed a tail call.

For example in the following, the first call on function $fact$ is a tail call because the result of $fact$ is the result of the whole expression, whereas the second call isn't because the result is subsequently multiplied by ω.

```
(α×ω)fact ω-1        ⍝ Tail call on fact.

ω×fact ω-1           ⍝ Embedded call on fact.
```

Tail calls occur frequently in Dynamic Functions, and the interpreter optimises them by re-using the current stack frame instead of creating a new one. This gives a significant saving in both time and workspace usage. It is easy to check whether a call is a tail call by tracing it. An embedded call will pop up a new trace window for the called function, whereas a tail call will re-use the current one.

Using tail calls can improve code performance considerably, although at first the technique might appear obscure. A simple way to think of a tail call is as a **branch with arguments**. The tail call, in effect, branches to the first line of the function after installing new values for ω and α.

Iterative algorithms can almost always be coded using tail calls.

In general, when coding a loop, we use the following steps, possibly in a different order depending on whether we want to test at the 'top' or the 'bottom' of the loop.

1.	Initialise loop control variable(s).	⍝	*init*
2.	Test loop control variable.	⍝	*test*
3.	Process body of loop.	⍝	*proc*
4.	Modify loop control variable for next iteration.	⍝	*mod*
5.	Branch to step 2.	⍝	*jump*

For example, in classical APL you might find the following:

```
      ∇ value←limit loop value      ⍝ init
[1]    top:→(⎕CT>value-limit)/0     ⍝ test
[2]     value←Next value           ⍝ proc, mod
[3]     →top                       ⍝ jump
      ∇
```

Control structures help us to package these steps:

```
      ∇ value←limit loop value      ⍝ init
[1]    :While ⎕CT<value-limit       ⍝ test
[2]        value←Next value         ⍝ proc, mod
[3]    :EndWhile                    ⍝ jump
      ∇
```

Using tail calls:

```
loop←{                            ⍝ init
    ⎕CT>α-ω:ω                      ⍝ test
    α ∇ Next ω                     ⍝ proc, mod, jump
}
```

Error-Guards

An **error-guard** is (an expression that evaluates to) a vector of error numbers, followed by the digraph: `::`, followed by an expression, the *body* of the guard, to be evaluated as the result of the function. For example:

```
    11 5 :: ω×0 ⍝ Trap DOMAIN and LENGTH errors.
```

In common with `:Trap` and `⎕TRAP`, error numbers 0 and 1000 are catchalls for synchronous errors and interrupts respectively.

When an error is generated, the system searches statically upwards and outwards for an error-guard that matches the error. If one is found, the execution environment is unwound to its state immediately *prior* to the error-guard's execution and the body of the error-guard is evaluated as the result of the function. This means that, during evaluation of the body, the guard is no longer in effect and so the danger of a hang caused by an infinite 'trap loop', is avoided.

Notice that you can provide 'cascading' error trapping in the following way:

```
0::try_2nd
0::try_1st
   expr
```

In this case, if `expr` generates an error, its immediately preceding: `0::` catches it and evaluates `try_1st` leaving the remaining error-guard in scope. If `try_1st` fails, the environment is unwound once again and `try_2nd` is evaluated, this time with no error-guards in scope.

Examples:

`Open` returns a handle for a component file. If the exclusive tie fails, it attempts a share-tie and if this fails, it creates a new file. Finally, if all else fails, a handle of 0 is returned.

```
open←{                  ⍝ Handle for component file ω.
    0::0                ⍝ Fails:: return 0 handle.
    22::ω ⎕FCREATE 0    ⍝ FILE NAME:: create new one.
    24 25::ω ⎕FSTIE 0   ⍝ FILE TIED:: try share tie.
         ω ⎕FTIE 0      ⍝ Attempt to open file.
}
```

An error in *div* causes it to be called recursively with *improved* arguments.

```
div←{                    ⍝ Tolerant division:: α÷0 → α.
    α←1                  ⍝ default numerator.
    5::↑∇/↓↑α ω          ⍝ LENGTH:: stretch to fit.
    11::α ∇ ω+ω=0        ⍝ DOMAIN:: increase divisor.
    α÷ω                  ⍝ attempt division.
}
```

Notice that some arguments may cause *div* to recur twice:

```
        6 4 2 div 3 2
→       6 4 2 div 3 2 0
→       6 4 2 div 3 2 1
→       2 2 2
```

The final example shows the unwinding of the local environment before the error-guard's body is evaluated. Local name *trap* is set to describe the domain of its following error-guard. When an error occurs, the environment is unwound to expose *trap*'s statically correct value.

```
        add←{
            trap←'domain' ◊ 11::trap
            trap←'length' ◊  5::trap
            α+ω
        }

        2 add 3          ⍝ Addition succeeds
5

        2 add 'three'    ⍝ DOMAIN ERROR generated.
domain

        2 3 add 4 5 6    ⍝ LENGTH ERROR generated.
length
```

Dynamic Operators

The operator equivalent of a dynamic function is distinguished by the presence of either of the compound symbols 'αα' or 'ωω' anywhere in its definition. αα and ωω represent the left and right operand of the operator respectively.

Example

The following monadic `each` operator applies its function operand only to unique elements of its argument. It then distributes the result to match the original argument. This can deliver a performance improvement over the primitive each (¨) operator if the operand function is costly and the argument contains a significant number of duplicate elements. Note however, that if the operand function causes side effects, the operation of dynamic and primitive versions will be different.

```
each←{              ⍝ Fast each:

    shp←⍴ω           ⍝ Shape and ...

    vec←,ω           ⍝ ... ravel of arg.

    nub←∪vec         ⍝ Vector of unique elements.

    res←αα¨nub       ⍝ Result for unique elts.

    idx←nubιvec      ⍝ Indices of arg in nub ...
    shp⍴idx⊃¨⊂res    ⍝ ... distribute result.
}
```

The dyadic `else` operator applies its left (else right) operand to its right argument depending on its left argument.

```
else←{
    α: αα ω      ⍝ True: apply Left operand
       ωω ω      ⍝ Else, ..   Right   ..
}

0 1 ⌈else⌊¨ 2.5    ⍝ Try both false and true.
2 3
```

Recursion

A recursive Dynamic Function can refer to itself using its name explicitly, but because
we allow unnamed functions, we also need a special symbol for implicit self-
reference: '∇'. For example:

```
fact←{              ⍝ Factorial ω.
    ω≤1: 1          ⍝ Small ω, finished,
    ω×∇ ω-1         ⍝ Otherwise recur.
}
```

Implicit self-reference using '∇' has the further advantage that it incurs less
interpretative overhead and is therefore quicker. Tail calls using '∇' are particularly
efficient.

Recursive Dynamic Operators refer to their derived functions, that is the operator
bound with its operand(s) using ∇ or the operator itself using the compound symbol:
∇∇. The first form of self reference is by far the more frequently used.

```
pow←{               ⍝ Function power.
    α=0:ω           ⍝ Apply function operand α times.
    (α-1)∇ αα ω ⍝ αα αα αα ... ω
}
```

The following example shows a rather contrived use of the second form of (operator)
self reference. The *exp* operator composes its function operand with itself on each
recursive call. This gives the effect of an exponential application of the original
operand function:

```
exp←{               ⍝ Exponential fn application.
    α=0:αα ω        ⍝ Apply operand 2*α times.
    (α-1)αα∘αα ∇∇ ω ⍝ (αα∘αα)∘( ... ) ... ω
}

succ←{1+ω}          ⍝ Successor (increment).

10 succ exp 0
1024
```

Example: Pythagorean triples

The following sequence shows an example of combining Dynamic Functions and
Operators in an attempt to find Pythagorean triples: (3 4 5)(5 12 13) ...

```
      sqrt←{ω*0.5}                   ⍝ Square root.

      sqrt 9 16 25
3 4 5

      hyp←{sqrt+/⊃ω*2}               ⍝ Hypoteneuse of triangl
e.

      hyp(3 4)(4 5)(5 12)
5 6.403124237 13

      intg←{ω=⌊ω}                    ⍝ Whole number?

      intg 2.5 3 4.5
0 1 0

      pyth←{intg hyp ω}            ⍝ Pythagorean pair?

      pyth(3 4)(4 9)(5 12)
1 0 1

      pairs←{,⍳ω ω}               ⍝ Pairs of numbers 1..ω.

      pairs 3
 1 1   1 2   1 3   2 1   2 2   2 3   3 1   3 2   3 3

      filter←{(αα ω)/ω}             ⍝ Op: ω filtered by αα.

      pyth filter pairs 12     ⍝ Pythagorean pairs 1..12
 3 4   4 3   5 12   6 8   8 6   9 12   12 5   12 9
```

So far, so good, but we have some duplicates: `(6 8)` is just double `(3 4)`.

```
        rpm←{                          ⍝ Relatively prime?
            ω=0:α=1                     ⍝ C.f. Euclid's gcd.
            ω ∇ ω|α
        }/¨                            ⍝ Note the /¨

        rpm(2 4)(3 4)(6 8)(16 27)
 0 1 0 1

        rpm filter pyth filter pairs 20
 3 4   4 3   5 12   8 15   12 5   15 8
```

We can use an operator to combine the tests:

```
        and←{                          ⍝ Lazy parallel 'And'.
            mask←αα ω                   ⍝ Left predicate selects...
            mask\ωω mask/ω               ⍝ args for right predi
cate.
        }

        pyth and rpm filter pairs 20
 3 4   4 3   5 12   8 15   12 5   15 8
```

Better, but we still have some duplicates: `(3 4)` `(4 3)`.

```
        less←{</⊃ω}
        less(3 4)(4 3)
 1 0

        less and pyth and rpm filter pairs 40
 3 4   5 12   7 24   8 15   9 40   12 35   20 21
```

And finally, as promised, triples:

```
        {ω,hyp ω}¨less and pyth and rpm filter pairs 35
 3 4 5   5 12 13   7 24 25   8 15 17   12 35 37   20 21 29
```

A Larger Example

Function *tokens* uses nested local D-Fns to split an APL expression into its
constituent tokens. Note that all calls on the inner functions: *lex*, *acc*, and the
unnamed D-Fn in each token case, are *tail calls*. In fact, the *only* stack calls are those
on function: *all*, and the unnamed function: {ω∨¯1⌽ω}, within the 'Char literal'
case.

```
tokens←{                              ⍝ Lex of APL src line.
    alph←⎕A,⎕A,'_∆⍙',26↑17↓⎕AV        ⍝ Alphabet for names.
    all←{+/∧\α∊ω}                      ⍝ No. of leading α∊ω.
    acc←{(α,↑/ω)lex⊃↓/ω}              ⍝ Accumulate tokens.
    lex←{
        0=ρω:α ◇ hd←↑ω                ⍝ Next char else done.

        hd=' ':α{                     ⍝ White Space.
            size←ω all' '
            α acc size ω
        }ω

        hd∊alph:α{                    ⍝ Name
            size←ω all alph,⎕D
            α acc size ω
        }ω

        hd∊'⎕:':α{                    ⍝ System Name/Keyword
            size←ω all hd,alph
            α acc size ω
        }ω

        hd='''':α{                    ⍝ Char literal
            size←+/∧\{ω∨¯1⌽ω}≠\hd=ω
            α acc size ω
        }ω

        hd∊⎕D,'¯':α{                  ⍝ Numeric literal
            size←ω all ⎕D,'.¯E'
            α acc size ω
        }ω

        hd='⍝':α acc(ρω)ω             ⍝ Comment
        α acc 1 ω                     ⍝ Single char token.
    }
    (0ρ⊂'')lex,ω
}

    display tokens'xtok←size↑srce ⍝ Next token'
```

```
.→----------------------------------------------------.
| .→---.  .→.  .→---.  .→.  .→---.  .→-.  .→-----------. |
| |xtok|  |←|  |size|  |↑|  |srce|  | |   |⍝ Next token| |
| '----'  '-'  '----'  '-'  '----'  '--'  '------------' |
'∈----------------------------------------------------'
```

Restrictions

Currently **multi-line** Dynamic Functions can't be typed directly into the session. The interpreter attempts to evaluate the first line with its trailing left brace and a *SYNTAX ERROR* results.

Dynamic Functions need not return a result. However even a non-result-returning expression will terminate the function, so you can't, for example, call a non-result-returning function from the middle of a Dynamic Function.

You can trace a Dynamic Function **only** if it is defined on more than one line. Otherwise it is executed atomically in the same way as an execute (⍎) expression. This deliberate restriction is intended to avoid the confusion caused by tracing a line and seeing nothing change on the screen.

Dynamic Functions do not currently support □*cs*.

Supplied Workspaces

You can find more examples of dynamic functions and operators in workspaces in the **samples\dfns** directory.

DFNS.DWS - a selection of utility functions.

MIN.DWS - an example application.

APL Line Editor

The APL Line Editor described herein is included for completeness and for adherence to the ISO APL standard. See *User Guide* for a description of the more powerful full-screen editor, ⎕*ED*.

Using the APL Line Editor, functions and operators are defined by entering Definition Mode. This mode is opened and closed by the del symbol , ∇. Within this mode, all evaluation of input is deferred. The standard APL line editor (described below) is used to create and edit operations within definition mode.

Operations may also be defined using the system function ⎕*FX* (implicit in a ⎕*ED* fix) which acts upon the canonical (character), vector, nested or object representation form of an operation. (See *Chapter 5* and *User Guide* for details.)

Functions may also be created dynamically or by function assignment. (See above and *Chapter 3*.)

The line editor recognises three forms for the opening request.

Creating Defined Operation

The opening ∇ symbol is followed by the header line of a defined operation. Redundant blanks in the request are permitted except within names. If acceptable, the editor prompts for the first statement of the operation body with the line-number 1 enclosed in brackets. On successful completion of editing, the defined operation becomes the active definition in the workspace.

Example

```
      ∇R←FOO
[1]   R←10
[2]   ∇

      FOO
10
```

The given operation name must not have an active referent in the workspace, otherwise the system reports *defn error* and the system editor is not invoked:

```
      )VARS
SALES  X  Y

      ∇R←SALES Y
defn error
```

The header line of the operation must be syntactically correct, otherwise the system reports *defn error* and the system editor is not invoked:

```
      ∇R←A B C D:G
defn error
```

Listing Defined Operation

The ∇ symbol followed by the name of a defined operation and then by a closing ∇, causes the display of the named operation. Omitting the function name causes the suspended operation (ie. the one at the top of the state indicator) to be displayed and opened for editing.

Example

```
      ∇FOO∇
     ∇ R←FOO
[1]    R←10
     ∇

      )SI
FOO[1] *

      ∇
     ∇ R←FOO
[1]    R←10
[2]
```

Editing Active Defined Operation

Definition mode is entered by typing ∇ followed optionally by a name and editing directive.

The ∇ symbol on its own causes the suspended operation (i.e. the one at the top of the state indicator) to be displayed. The editor then prompts for a statement or editing directive with a line-number one greater than the highest line-number in the function. If the state indicator is empty, the system reports *defn error* and definition mode is not entered.

The ∇ symbol followed by the name of an active defined operation causes the display of the named operation. The editor then prompts for input as described above. If the name given is not the name of an active referent in the workspace, the opening request is taken to be the creation of a new operation as described in paragraph 1. If the name refers to a pendent operation, the editor issues the message *warning pendent operation* prior to displaying the operation. If the name refers to a locked operation, the system reports defn error and definition mode is not entered.

The ∇ symbol followed by the name of an active defined operation and an editing directive causes the operation to be opened for editing and the editing directive actioned. If the editing directive is invalid, it is ignored by the editor which then prompts with a line-number one greater than the highest line-number in the operation. If the name refers to a pendent operation, the editor issues the message *warning pendent operation* prior to actioning the editing directive. If the name refers to a locked operation, the system reports *defn error* and definition mode is not entered.

On successful completion of editing, the defined operation becomes the active definition in the workspace which may replace an existing version of the function. Monitors, and stop and trace vectors are removed.

Example

```
      ∇FOO[2]
[2]   R←R*2
[3]   ∇
```

Editing Directives

Editing directives, summarised in Figure 2(iv) are permitted as the first non-blank characters either after the operation name on opening definition mode for an active defined function, or after a line-number prompt.

Syntax	Description
∇	*Closes definition mode*
[□]	*Displays the entire operation*
[□n]	*Displays the operation starting at line n*
[n□]	*Displays only line n*
[∆n]	*Deletes line n*
[n∆m]	*Deletes m lines starting at line n*
[n]	*Prompts for input at line n*
[n]s	*Replaces or inserts a statement at line n*
[n□m]	*Edits line n placing the cursor at character position m*
	Edit control symbols are:
	* / - deletes character above*
	* 1 to 9 - inserts that number of spaces*
	* A to Z - inserts multiples of 5 spaces*
	* .text - inserts the text prior to the character above '.'*
	* ,text - inserts the text as above but continues the edit*

Figure 2(iv) : Editing directives

Line Numbers

Line numbers are associated with lines in the operation. Initially, numbers are assigned as consecutive integers, beginning with `[0]` for the header line. The number associated with an operation line remains the same for the duration of the definition mode unless altered by editing directives. Additional lines may be inserted by decimal numbering. Up to three places of decimal are permitted. On closing definition mode, operation lines are re-numbered as consecutive integers.

The editor always prompts with a line number. The response may be a statement line or an editing directive. A statement line replaces the existing line (if there is one) or becomes an additional line in the operation:

```
      ∇R←A PLUS B
[1]   R←A+B
[2]
```

Position

The editing directive `[n]`, where n is a line number, causes the editor to prompt for input at that line number. A statement or another editing directive may be entered. If a statement is entered, the next line number to be prompted is the previous number incremented by a unit of the display form of the last decimal digit. Trailing zeros are not displayed in the fractional part of a line number:

```
[2]    [0.8]
[0.8] ⍝ MONADIC OR DYADIC +
[0.9] ⍝ A ←→ OPTIONAL ARGUMENT
[1]
```

The editing directive `[n]`s, where *n* is a line number and s is a statement, causes the statement to replace the current contents of line *n*, or to insert line n if there is none:

```
[1] [0] R←{A} PLUS B
[1]
```

Delete

The editing directive `[∆n]`, where *n* is a line number, causes the statement line to be deleted. The form `[n∆m]`, where *n* is a line number and *m* is a positive integer, causes *m* consecutive statement lines starting from line number *n* to be deleted.

Edit

The editing directive [*n*□*m*], where *n* is a line number and *m* is an integer number, causes line number *n* to be displayed and the cursor placed beneath the *m*{th} character on a new line for editing. The response is taken to be edit control symbols selected from:

/	- to delete the character immediately above the symbol.
1 to 9	- to insert from 1 to 9 spaces immediately prior to the character above the digit.
A to Z	- to insert multiples of 5 spaces immediately prior to the character above the letter, where A = 5, B = 10, C = 15 and so forth.
,	- to insert the text after the comma, including explicitly entered trailing spaces, prior to the character above the comma, and then re-display the line for further editing with the text inserted and any preceding deletions or space insertions also effected.
.	- to insert the text after the comma, including explicitly entered trailing spaces, prior to the character above the comma, and then complete the edit of the line with the text inserted and any preceding deletions or space insertions also effected.

Invalid edit symbols are ignored. If there are no valid edit symbols entered, or if there are only deletion or space insertion symbols, the statement line is re-displayed with characters deleted and spaces inserted as specified. The cursor is placed at the first inserted space position or at the end of the line if none. Characters may be added to the line which is then interpreted as seen.

The line number may be edited.

Examples

```
[1]     [1⎕7]
[1]     R←A+B
        ,→(0=⎕NC'A')ρ1←⎕LC ⋄
[1]     →(0=⎕NC'A')ρ1←⎕LC ⋄ R←A+B
                                    .⋄→END
[2]     R←B
[3]     END:
[4]
```

The form `[n⎕0]` causes the line number *n* to be displayed and the cursor to be positioned at the end of the displayed line, omitting the edit phase.

Display

The editing directive `[⎕]` causes the entire operation to be displayed. The form `[⎕n]` causes all lines from line number n to be displayed. The form `[n⎕]` causes only line number *n* to be displayed:

```
[4]     [0⎕]
[0]     R←{A} PLUS B
[0]
[0]     [⎕]
[0]     R←{A} PLUS B
[0.1] ⍝ MONADIC OR DYADIC +
[1]     →(0=⎕NC'A')ρ1+⎕LC ⋄ R←A+B ⋄→END
[2]     R←B
[3]     'END:
[4]
```

Close Definition Mode

The editing directive ∇ causes definition mode to be closed. The new definition of the operation becomes the active version in the workspace. If the name in the operation header (which may or may not be the name used to enter definition mode) refers to a pendent operation, the editor issues the message *warning pendent operation* before exiting. The new definition becomes the active version, but the original one will continue to be referenced until the operation completes or is cleared from the State Indicator.

If the name in the operation header is the name of a visible variable or label, the editor reports *defn error* and remains in definition mode. It is then necessary to edit the header line or quit.

If the header line is changed such that it is syntactically incorrect, the system reports *defn error*, and re-displays the line leaving the cursor beyond the end of the text on the line. Backspace/linefeed editing may be used to alter or cancel the change:

```
[3]    [0⎕]              - display line 0
[0]    R←{A} PLUS B
[0]    R←{A} PLUS B:G;H - put syntax error in line 0
defn error
[0]    R←{A} PLUS B:G;H - line redisplayed
                   ;G;H - backspace/linefeed editing
[1]
```

Local names may be repeated. However, the line editor reports warning messages as follows:

1. If a name is repeated in the header line, the system reports "warning duplicate name" immediately.

2. If a label has the same name as a name in the header line, the system reports "warning label name present in line 0" on closing definition mode.

3. If a label has the same name as another label, the system reports "warning duplicate label" on closing definition mode.

Improper syntax in expressions within statement lines of the function is not detected by the system editor with the following exceptions:

- If the number of opening parentheses in each entire expression does not equal the number of closing parentheses, the system reports "warning unmatched parentheses", but accepts the line.

- If the number of opening brackets in each entire expression does not equal the number of closing brackets, the system reports "warning unmatched brackets", but accepts the line.

These errors are not detected if they occur in a comment or within quotes. Other syntactical errors in statement lines will remain undetected until the operation is executed.

Example

```
[4]    R←(A[;1]=2)≠⍎EXP,'×2
warning unmatched parentheses
warning unmatched brackets
[5]
```

Note that there is an imbalance in the number of quotes. This will result an a *SYNTAX ERROR* when this operation is executed.

Quit Definition Mode

The user may quit definition mode by typing the INTERRUPT character. The active version of the operation (if any) remains unchanged.

C H A P T E R 3

Primitive Functions

Scalar Functions

There is a class of primitive functions termed SCALAR FUNCTIONS. This class is identified in Figure 3(i) below. Scalar functions are **pervasive**, ie. their properties apply at all levels of nesting. Scalar functions have the following properties:

Symbol	Monadic	Dyadic
+	Identity	Plus (Add)
−	Negative	Minus (Subtract)
×	Signum	Times (Multiply)
÷	Reciprocal	Divide
|	Magnitude	Residue
⌊	Floor	Minimum
⌈	Ceiling	Maximum
*	Exponential	Power
⊗	Natural Logarithm	Logarithm
○	Pi Times	Circular
!	Factorial	Binomial
~	Not	$
?	Roll	$
∈	Type (See Enlist)	$
∧		And
∨		Or
⍲		Nand
⍱		Nor
<		Less
≤		Less Or Equal
=		Equal
≥		Greater Or Equal
>		Greater
≠		Not Equal
$ Dyadic form is not scalar		

Figure 3(i) : Scalar primitive functions

Monadic Scalar Functions

a)   The function is applied independently to each simple scalar in its argument.

Example

```
÷2 (1 4)
  0.5  1 0.25
```

b)   The function produces a result with a structure identical to its argument.

c)   When applied to an empty argument, the function produces an empty result. With the exception of + and ∊, the type of this result depends on the function, not on the type of the argument. By definition + and ∊ return a result of the same type as their arguments.

Dyadic Scalar Functions

a)   The function is applied independently to corresponding pairs of simple scalars in its arguments.

Examples

```
      2 3 4 + 1 2 3
3 5 7

      2 (3 4) + 1 (2 3)
3  5 7

      (1 2) 3 + 4 (5 6)
 5 6  8 9
```

b)   A simple scalar will be replicated to conform to the structure of the other argument. If a simple scalar in the structure of an argument corresponds to a non-simple scalar in the other argument, then the function is applied between the simple scalar and the items of the non-simple scalar. Replication of simple scalars is called SCALAR EXTENSION.

Examples

```
      10 × 2 (3 4)
20  30 40

      2 4 = 2 (4 6)
1  1 0
```

c)   A simple unit is treated as a scalar for scalar extension purposes. A UNIT is a single element array of any rank. If both arguments are simple units, the argument with lower rank is extended.

Example

```
      (1 1ρ5) - 1 (2 3)
4  3 2
```

d)   The function produces a result with a structure identical to that of its arguments (after scalar extensions).

e)   If applied between empty arguments, the function produces a composite structure resulting from any scalar extensions, with type appropriate to the particular function. (All scalar dyadic functions return a result of numeric type.)

Examples

```
      1↑''+ι0
0

      1↑(0ρ⊂' ' (0 0))×''
0  0 0
```

Note: The Axis operator applies to all scalar dyadic functions.

Mixed Functions

Mixed rank functions are summarised in Figure 3(ii). For convenience, they are sub-divided into five classes:

Structural	These functions change the structure of the arguments in some way.
Selection	These functions select elements from an argument.
Selector	These functions identify specific elements by a boolean map or by an ordered set of indices.
Miscellaneous	These functions transform arguments in some way, or provide information about the arguments.
Special	These functions have special properties.

In general, the structure of the result of a mixed primitive function is different from that of its arguments.

Scalar extension may apply to some, but not all, dyadic mixed functions.

Mixed primitive functions are not pervasive. The function is applied to elements of the arguments, not necessarily independently.

Examples

```
      'CAT' 'DOG' 'MOUSE'ι⊂'DOG'
2

      3↑ 1 'TWO' 3 'FOUR'
1  TWO  3
```

Class	Symbol	Monadic	Dyadic
Structural	ρ	$	*Reshape*
	,	*Ravel* []	*Catenate* []
			Laminate []
	⍪		*Catenate First*[]
			Laminate []
	⌽	*Reverse* []	*Rotate* []
	⊖	*Reverse First*[]	*Rotate First*[]
	⍉	*Transpose*	*Transpose*
	↑	*Mix/Disclose (First)* []	$
	↓	*Split* []	$
	⊂	*Enclose* []	*Partitioned Enclose* []
	∊	*Enlist (See Type)*	$
Selection	⊃	*Disclose/Mix*	*Pick*
	↑	$	*Take* []
	↓	$	*Drop* []
	/		*Replicate* []
	⌿		*Replicate First*[]
	\		*Expand* []
	⍀		*Expand First* []
	~	$	*Without (Excluding)*
	∩		*Intersection*
	∪	*Unique*	*Union*
Selector	⍳	*Index Generator*	*Index Of*
	∊	$	*Membership*
	⍋	*Grade Up*	*Grade Up*
	⍒	*Grade Down*	*Grade Down*
	?	$	*Deal*
	⍷		*Find*

[ ] Implies axis specification is optional

$ This function is in another class

Figure 3(ii) : Mixed Primitive Functions

Class	Symbol	Monadic	Dyadic
Miscellaneous	ρ	*Shape*	$
	≡	*Depth*	*Match*
	≠		*Not Match*
	⍎	*Execute*	*Execute*
	⍕	*Format*	*Format*
	⊥		*Decode (Base)*
	⊤		*Encode (Representation)*
	⌹	*Matrix Divide*	*Matrix Inverse*
Special	→	*Abort (Niladic)*	
	→	*Branch*	
	←	$	*Assignment*
	[I]←	$	*Assignment (Indexed)*
	(I)←		*Assignment (Selective)*
	[]		*Indexing*
[]	Implies axis specification is optional		
$	This function is in another class		

Figure 3(ii) : Mixed Primitive Functions (Continued)

Conformability

The arguments of a dyadic function are said to be CONFORMABLE if the shape of each argument meets the requirements of the function, possibly after scalar extension.

Fill Elements

Some primitive functions may include fill elements in their result. The fill element for an array is the enclosed type of the disclose of the array ($\subset \epsilon \supset Y$ for array Y). The Type function (ϵ) replaces a numeric value with zero and a character value with $'\ '$.

The Disclose function ($\supset$) returns the first item of an array. If the array is empty, $\supset Y$ is the PROTOTYPE of Y. The prototype is the type of the first element of the original array.

Primitive functions which may return an array including fill elements are Expand (\ or $\backslash$), Replicate (/ or $\neq$), Reshape (ρ) and Take ($\uparrow$).

Examples

```
      ∈ι5
0 0 0 0 0

      ∈⊃(ι3)('ABC')
0 0 0

      ⊂∈⊃(ι3)('ABC')
 0 0 0

      ⊂∈⊃⊂(ι3)('ABC')
  0 0 0

      A←'ABC' (1 2 3)
      A←0ρA
      ⊂∈⊃A

      ' '=⊂∈⊃A
 1 1 1
```

Axis Specification

The axis operator may be applied to all scalar dyadic primitive functions and certain mixed primitive functions. An integer axis identifies a specific axis along which the function is to be applied to one or both of its arguments. If the primitive function is to be applied without an axis specification, a default axis is implied, either the first or last.

Example

```
      1 0 1/[1] 3 2⍴⍳6
1 2
5 6

      1 2 3+[2]2 3⍴10 20 30
11 22 33
11 22 33
```

Sometimes the axis value is fractional, indicating that a new axis or axes are to be created between the axes identified by the lower and upper integer bounds of the value (either of which might not exist).

Example

```
      'NAMES',[0.5]'='
NAMES
=====
```

⎕IO is an implicit argument of an axis specification.

Function Presentation

Scalar and mixed primitive functions are presented in alphabetical order of their descriptive names as shown in Figures 3(i) and 3(ii) respectively. Scalar functions are described in terms of single element arguments. The rules for extension are defined at the beginning of this chapter.

The class of the function is identified in the heading block. The valence of the function is implied by its syntax in the heading block.

Abort: →

This is a special case of the Branch function used in the niladic sense. If it occurs in a statement it must be the only symbol in an expression or the only symbol forming an expression in a text string to be executed by ⍎. It clears the most recently suspended statement and all of its pendent statements from the state indicator.

The Abort function has no explicit result. The function is not in the function domain of operators.

Examples

```
      ∇ F
[1]    'F[1]'
[2]    G
[3]    'F[3]'
      ∇


      ∇ G
[1]    'G[1]'
[2]    →
[3]    'G[3]'
      ∇
       F
F[1]
G[1]


      ⎕VR'VALIDATE'
      ∇ VALIDATE
 [1]   →(12=1↑⎕AI)⍴0 ◇ 'ACCOUNT NOT AUTHORISED' ◇ →
      ∇

      VALIDATE
ACCOUNT NOT AUTHORISED

      1↑⎕AI
52
```

Add: $R \leftarrow X + Y$

Y must be numeric. *X* must be numeric. *R* is the arithmetic sum of *X* and *Y*. *R* is numeric. This function is also known as Plus.

Examples

```
      1 2 + 3 4
4 6

      1 2 + 3,⊂4 5
4   6 7
```

And: $R \leftarrow X \wedge Y$

Y must be boolean. *X* must be boolean. *R* is boolean. *R* assumes the following result for each corresponding value of *X* and *Y*:

X	Y	R
0	0	0
0	1	0
1	0	0
1	1	1

Note that the ASCII caret (^) will also be interpreted as an APL And (∧).

Example

```
      0 1 0 1 ∧ 0 0 1 1
0 0 0 1
```

Assignment: $X \leftarrow Y$

Assignment allocates the result of the expression Y to the ***name*** or ***names*** in X.

If Y is an array expression, X must contain one or more names which are variables, system variables, or are undefined. Following assignment, the name(s) in X become variable(s) with value(s) taken from the result of the expression Y.

If X contains a single name, the variable assumes the value of Y.

The assignment arrow (or specification arrow) is often read as 'Is' or 'Gets'.

Examples

```
      A←2.3
      A
2.3

      A←ι3
      A
1 2 3
```

More than one name may be specified in X by using vector notation. If so, Y must be a vector or a scalar. If Y is a scalar, its value is assigned to all names in X. If Y is a vector, each element of Y is assigned to the corresponding name in X.

Examples

```
      A B←2
      A
2

      B
2

      P ⎕IO Q←'TEXT' 1 (1 2 3)
      P
TEXT
      ⎕IO
1
      Q
1 2 3
```

For compatibility with IBM's APL2, the list of names specified in X may be enclosed in parentheses.

Examples

```
(A B C)←1 2 3
(D E)←'Hello' 'World'
```

Multiple assignments are permitted. The value of *Y* is carried through each assignment:

```
      I←J←K←0

      I,J,K
0 0 0
```

Function Assignment

If *Y* is a function expression, *X* must be a single name which is either undefined, or is the name of an existing function or defined operator. *X* may not be the name of a system function, or a primitive symbol.

Examples

```
      PLUS←+
      PLUS
+

      SUM←+/
      SUM
+/

      MEAN←{(+/ω)÷ρω}
```

Namespace Reference Assignment

If an expression evaluates to a namespace reference, or *ref,* you may assign it to a name. A name assigned to a simple scalar *ref,* has name class 9, whereas one assigned to an *array containing refs* has name class 2.

```
      'f1'⎕WC'Form'
      'ns1' ⎕NS ''

      N←ns1
      ⎕NC'N'           ⍝ name class of a scalar ref
9
      F←f1
      ⎕NC'F'           ⍝ name class of a scalar ref
9
      refs←N F          ⍝ vector of refs.
      ⎕NC'refs'         ⍝ nameclass of vector.
2
```

```
        F2←2⊃refs
        ⎕NC 'F2'
9
```

Assignment (Indexed): $\{R\} \leftarrow X[I] \leftarrow Y$

Indexed Assignment is the Assignment function modified by the Indexing function. The phrase $[I] \leftarrow$ is treated as the function for descriptive purposes.

Y may be any array. X may be the ***name*** of any array. I must be a valid index specification. The shape of Y must conform with the shape (implied) of the indexed structure defined by I. If Y is a scalar or a unit vector it will be extended to conform. A side effect of Indexed Assignment is to change the value of the indexed elements of X.

R is the value of Y. If the result is not explicitly assigned or used it is suppressed.

$\Box IO$ is an implicit argument of Indexed Assignment.

Three forms of indexing are permitted.

Simple Indexed Assignment

For vector X, I is a simple integer array whose items are from the set $\iota \rho R$. Elements of X identified by index positions I are replaced by corresponding elements of Y.

Examples

```
      +A←ι5
1 2 3 4 5

      A[2 3]←10 ◇ A
1 10 10 4 5
```

The last-most element of Y is assigned when an index is repeated in I:

```
      A[2 2]←100 101 ◇ A
1 101 10 4 5
```

For matrix X, I is composed of two simple integer arrays separated by the semicolon character ($;$). The arrays select indices from the rows and columns of X respectively.

Examples

```
      +B←2 3ρ'REDSUN'
RED
SUN

      B[2;2]←'O' ◇ B
RED
SON
```

For higher-order array X, I is a series of simple integer arrays with adjacent arrays
separated by a single semicolon character (;). Each array selects indices from an axis
of X taken in row-major order.

Examples

```
      C
11 12 13
14 15 16

21 22 23
24 25 26

      C[1;1;3]←103 ◇ C
11 12 103
14 15  16

21 22  23
24 25  26
```

An indexing array may be ELIDED. That is, if an indexing array is omitted from the
Kth axis, the indexing vector $\iota(\rho X)[K]$ is implied:

```
      C[;1;2 3]←2 2ρ112 113 122 123 ◇ C
11 112 113
14  15  16

21 122 123
24  25  26

      C[;;]←0 ◇ C
0 0 0
0 0 0

0 0 0
0 0 0
```

Choose Indexed Assignment

The index specification I is a non-simple integer array. Each item identifies a single
element of X by a set of indices with one element per axis of X in row-major order.

Examples

```
      C
11 12 13 14
21 22 23 24

      C[⊂1 1]←101  ◇  C
101 12 13 14
 21 22 23 24

      C[(1 2) (2 3)]←102 203  ◇  C
101 102  13 14
 21  22 203 24

      C[2 2ρ(1 3)(2 4)(2 1)(1 4)]←2 2ρ103 204 201 104◇C
101 102 103 104
201  22 203 204
```

A scalar may be indexed by the enclosed empty vector:

```
      S
10

      S[⊂ι0]←⊂'VECTOR'  ◇  S
 VECTOR

      S[⊂ι0]←5  ◇  S
5
```

Choose Indexed Assignment may be used very effectively in conjunction with Index
Generator (ι) and Structural functions in order to assign into an array:

```
      C
11 12 13 14
21 22 23 24

      ιρC
1 1  1 2  1 3  1 4
2 1  2 2  2 3  2 4

      C[1 1⍉ιρC]←1 2  ◇  C
 1 12 13 14
```

```
21   2 23 24
```

```
      C[2 ¯1↑⍳⍴C]←99 ⋄ C
 1 12 13 99
21  2 23 99
```

Reach Indexed Assignment

The index specification *I* is a non-simple integer array, each of whose items reach down to a nested element of *X*. The items of an item of *I* are simple vectors (or scalars) forming sets of indices that index arrays at successive levels of *X* starting at the top-most level. A set of indices has one element per axis at the respective level of nesting of *X* in row-major order.

Examples

```
      D←(2 3⍴⍳6)(2 2⍴'SMITH' 'JONES' 'SAM' 'BILL')


      D
 1 2 3   SMITH  JONES
 4 5 6   SAM    BILL


      ≡J←⊂2 (1 2)
¯3

      D[J]←⊂'WILLIAMS' ⋄ D
 1 2 3   SMITH  WILLIAMS
 4 5 6   SAM    BILL

      D[(1 (1 1))(2 (2 2) 1)]←10 'W' ⋄ D
10 2 3   SMITH  WILLIAMS
 4 5 6   SAM    WILL


      E
GREEN  YELLOW  RED

      E[⊂2 1]←'M' ⋄ E
GREEN  MELLOW  RED
```

The context of indexing is important. In the last example, the indexing method is determined to be Reach rather than Choose since *E* is a vector, not a matrix as would be required for Choose. Observe that:

```
      ⊂2 1 ↔ ⊂(⊂2),(⊂1)
```

Assignment (Selective): $(EXP\ X)\leftarrow Y$

X is the ***name*** of a variable in the workspace. *EXP* is an expression that **selects** elements of *X*. *Y* is an array expression. The result of the expression *Y* is allocated to the elements of *X* selected by *EXP*.

The following functions may appear in the selection expression. Where appropriate these functions may be used with axis [ ].

↑	Take
↓	Drop
,	Ravel
⌽	Reverse, Rotate
ρ	Reshape
⊃	Disclose, Pick
⍉	Transpose (Monadic and Dyadic)
/	Replicate
\	Expand

Note that Mix and Split (monadic ↑ and ↓) may not be used in the selection expression.

Examples

```
      A←'HELLO'
      ((A∊'AEIOU')/A)←'*'

      A
H*LL*

      Z←3 4ρι12
      (5↑,Z)←0

      Z
0  0  0  0
0  6  7  8
9 10 11 12

      MAT←3 3ρι9
      (1 1⍉MAT)←0

      MAT
0 2 3
4 0 6
7 8 0
```

Binomial: $R \leftarrow X!Y$

Y may be any number other than a negative integer. X may be any number other than a negative integer. R is numeric. An element of R is integer if corresponding elements of X and Y are integers. Binomial is defined in terms of the function Factorial for positive integer arguments:

$$X!Y \leftrightarrow (!Y)\div(!X)\times!Y-X$$

For other arguments, results are derived smoothly from the Beta function:

$$\text{Beta}(X,Y) \leftrightarrow \div Y\times(X-1)!X+Y-1$$

For positive integer arguments, R is the number of selections of X things from Y things.

Example

```
      1 1.2 1.4 1.6 1.8 2!5
5 6.105689248 7.219424686 8.281104786 9.227916704 10
```

Branch: →*Y*

Y may be a scalar or vector which, if not empty, has a simple numeric scalar as its first element. The function has no explicit result. It is used to modify the normal sequence of execution of expressions or to resume execution after a statement has been interrupted. Branch is not in the function domain of operators.

The following distinct usages of the branch function occur:

	Entered in a Statement in a Defined Function	**Entered in Immediate Execution Mode**
→*LINE*	Continue with the specific line	Restart execution at the specific line of the most recently suspended function
→ι 0	Continue with the next expression	No effect

In a defined function, if *Y* is non-empty then the first element in *Y* specifies a statement line in the defined function to be executed next. If the line does not exist, then execution of the function is terminated. For this purpose, line 0 does not exist. (Note that statement line numbers are independent of the index origin ⎕*IO*).

If *Y* is empty, the branch function has no effect. The next expression is executed on the same line, if any, or on the next line if not. If there is no following line, the function is terminated.

The :*GoTo* statement may be used in place of Branch in a defined function.

Example

```
        ∇TEST
  [1]    1
  [2]    →4
  [3]    3
  [4]    4
        ∇


        TEST
 1
 4
```

In general it is better to branch to a LABEL than to a line number. A label occurs in a statement followed by a colon and is assigned the value of the statement line number when the function is defined.

Example

```
        ∇ TEST
[1]    1
[2]    →FOUR
[3]    3
[4]    FOUR:4
        ∇
        TEST
1
4
```

The previous examples illustrate unconditional branching. There are numerous APL idioms which result in conditional branching. Some popular idioms are identified:

Branch Expression	Comment
→*TEST/L1*	*Branches to label L1 if TEST results in 1 but not if TEST results in 0.*
→*TEST*ρ*L1*	*Similar to above.*
→*TEST*↑*L1*	*Similar to above.*
→*L1*ρ¨*TEST*	*Similar to above.*
→*L1*⌈ι*TEST*	*Similar to above but only if* □*IO*↔1
→*L1*×ι*TEST*	*Similar to above but only if* □*IO*↔1
→(*L1,L2,L3*)[*N*]	*Unconditional branch to a selected label.*
→(*T1,T2,T3*)/*L1,L2,L3*	*Branches to the first selected label dependent on tests T1,T2,T3. If all tests result in 0, there is no branch.*
→*N*ϕ*L1,L2,L3*	*Unconditional branch to the first label after rotation.*

A branch expression may occur within a statement including ◇ separators:

```
[5]    →NEXTρ¨TEST  ◇  A←A+1  ◇  →END
[6]  NEXT:
```

In this example, the expressions '*A←A+1*' and '→*END*' are executed only if *TEST* returns the value 1. Otherwise control branches to label *NEXT*.

In immediate execution mode, the branch function permits execution to be continued within the most recently suspended function, if any, in the state indicator. If the state indicator is empty, or if the argument Y is the empty vector, the branch expression has no effect. If a statement line is specified which does not exist, the function is terminated. Otherwise, execution is restarted from the beginning of the specified statement line in the most recently suspended function.

Example

```
        ∇ F
[1]     1
[2]     2
[3]     3
[4]     ∇

        2 □STOP'F'
        F
1

F[2]
        )SI
F[2]*
        →2
2
3
```

The system constant $□LC$ returns a vector of the line numbers of statement lines in the state indicator, starting with that in the most recently suspended function. It is convenient to restart execution in a suspended state by the expression:

```
        →□LC
```

Catenate/Laminate: $R \leftarrow X , [K] Y$

Y may be any array. X may be any array. The axis specification is optional. If specified, K must be a numeric scalar or unit vector which may have a fractional value. If not specified, the last axis is implied.

The form $R \leftarrow X \overline{,} Y$ may be used to imply catenation along the first axis.

Two cases of the function catenate are permitted:

1. With an integer axis specification, or implied axis specification.

2. With a fractional axis specification, also called **laminate**.

Catenation with Integer or Implied Axis Specification

The arrays X and Y are joined along the required axis to form array R. A scalar or unit vector is extended to the shape of the other argument except that the required axis is restricted to a unit dimension. X and Y must have the same shape (after extension) except along the required axis, or one of the arguments may have rank one less than the other, provided that their shapes conform to the prior rule after augmenting the array of lower rank to have a unit dimension along the required axis.

The rank of R is the greater of the ranks of the arguments, but not less than 1.

Examples

```
      'FUR','LONG'
FURLONG

      1,2
1 2

      (2 4ρ'THISWEEK')⊤,'='
THIS
WEEK
====

      S,[1]+⌿S←2 3ρι6
1 2 3
4 5 6
5 7 9
```

If one of X and Y are empty, then the data type of R will be that of the other. If both arguments are empty, then the data type of R will be that of Y.

Lamination with Fractional Axis Specification

The arrays *X* and *Y* are joined along a new axis created before the ⌈*K*th axis. The new axis has a length of 2. *K* must exceed ⎕*IO* (the index origin) minus 1, and *K* must be less than ⎕*IO* plus the greater of the ranks of *X* and *Y*. A scalar or unit vector argument is extended to the shape of the other argument. Otherwise *X* and *Y* must have the same shape.

The rank of *R* is one plus the greater of the ranks of *X* and *Y*.

Examples

```
      'HEADING',[0.5]'-'
HEADING
-------

      'NIGHT',[1.5]'*'
N*
I*
G*
H*
T*

      ⎕IO←0

      'HEADING,[¯0.5]'-'
HEADING
-------
```

Catenate First: $R \leftarrow X \bar{,} [K] Y$

The form $R \leftarrow X \bar{,} Y$ implies catenation along the first axis whereas the form $R \leftarrow X , Y$ implies catenation along the last axis (columns). See Catenate/Laminate above.

Ceiling: $R \leftarrow \lceil Y$

Y must be numeric. R is the least integer greater than or equal to Y.

Example

```
      ⌈¯2.3  0.1  100  3.3
¯2 1 100 4
```

$\Box CT$ is an implied argument of Ceiling.

Circular: $R \leftarrow X \circ Y$

Y must be numeric. X must be an integer in the range $^{-}7 \leq X \leq 7$. R is numeric.

X determines which of a family of trigonometric functions to apply to Y, from the following table:

Range	Domain	$(-X) \circ Y$	X	$X \circ Y$	Range
$0 \leq R \leq 1$	$(\|Y) \leq 1$	$(1-Y*2)*.5$	0	$(1-Y*2)*.5$	$0 \leq R \leq 1$
$(-\circ.5) < R \leq \circ.5$	$(\|Y) \leq 1$	$Arcsin\ Y$	1	$Sine\ Y$	$(\|R) \leq 1$
$0 \leq R \leq \circ 1$	$(\|Y) \leq 1$	$Arccos\ Y$	2	$Cosine\ Y$	$(\|R) \leq 1$
$(\_R) \leq \circ.5$		$Arctan\ Y$	3	$Tangent\ Y$	
$R \geq 0$	$(\|Y) \geq 1$	$(^{-}1+Y*2)*.5$	4	$(1+Y*2)*.5$	
		$Arcsinh\ Y$	5	$Sinh\ Y$	
$R \geq 0$	$Y \geq 1$	$Arccosh\ Y$	6	$Cosh\ Y$	$R > 1$
	$(\|Y) < 1$	$Arctanh\ Y$	7	$Tanh\ Y$	$(\|R) < 1$

Examples

```
      0 ¯1 ○ 1
0 1.570796327

      1○(PI←○1)÷2 3 4
1 0.8660254038  0.7071067812

      2○PI÷3
0.5
```

Deal: $R \leftarrow X?Y$

Y must be a simple scalar or unit vector containing a non-negative integer. *X* must be a simple scalar or unit vector containing a non-negative integer and $X \leq Y$.

R is an integer unit vector obtained by making *X* random selections from ιY without repetition.

Examples

```
      13?52
7 40 24 28 12 3 36 49 20 44 2 35 1

      13?52
20 4 22 36 31 49 45 28 5 35 37 48 40
```

$\square IO$ and *$\square RL$* are implicit arguments of Deal. A side effect of Deal is to change the value of *$\square RL$*.

Decode: $R \leftarrow X \perp Y$

Y must be a simple numeric array. *X* must be a simple numeric array. *R* is the numeric array which results from the evaluation of *Y* in the number system with radix *X*.

X and *Y* are conformable if the length of the last axis of *X* is the same as the length of the first axis of *Y*. A scalar or unit vector is extended to a vector of the required length. If the last axis of *X* or the first axis of *Y* has a length of 1, the array is extended along that axis to conform with the other argument.

The shape of *R* is the catenation of the shape of *X* less the last dimension with the shape of *Y* less the first dimension. That is:

$$\rho R \leftrightarrow (\bar{1} \downarrow \rho X), 1 \downarrow \rho Y$$

For vector arguments, each element of *X* defines the ratio between the units for corresponding pairs of elements in *Y*. The first element of *X* has no effect on the result.

This function is also known as Base Value.

Examples

```
      60 60⊥3 13
193

      0 60⊥3 13
193

      60⊥3 13
193

      2⊥1 0 1 0
10
```

For higher order array arguments, each of the vectors along the last axis of `X` is taken as the radix vector for each of the vectors along the first axis of `Y`.

Examples

```
      M
0 0 0 0 1 1 1 1
0 0 1 1 0 0 1 1
0 1 0 1 0 1 0 1

      A
1 1 1
2 2 2
3 3 3
4 4 4

      A⊥M
0 1 1 2  1  2  2  3
0 1 2 3  4  5  6  7
0 1 3 4  9 10 12 13
0 1 4 5 16 17 20 21
```

Scalar extension may be applied:

```
      2⊥M
0 1 2 3 4 5 6 7
```

Extension along a unit axis may be applied:

```
      +A←2 1ρ2 10
 2
10
```

```
      A⊥M
0 1  2  3   4   5   6   7
0 1 10 11 100 101 110 111
```

Depth: (⎕ML)   R←≡Y

Y may be any array. *R* is a simple integer scalar which indicates the number of levels of nesting for the simple scalar item in the array with the greatest number of levels of nesting. A simple scalar has a depth of 0. A simple array has a depth of 1. An array containing non-simple elements has a depth whose magnitude is greater than 1. The array has uniform depth if all simple scalars in the array have the same depth.

If ⎕ML < 2, a negative value of *R* indicates non-uniform depth.

Examples

```
      ≡1
0
      ≡'A'
0
      ≡'ABC'?
1
      ≡1 'A'
1

      ⎕ML←0

      ≡A←(1 2)(3 (4 5))
¯3
      ≡¨A
1 ¯2
      ≡····A
 0 0  0 1

      ⎕ML←2

      ≡A
3
      ≡¨A
1 2
      ≡····A
 0 0  0 1
```

Disclose: (*□ML*)   *R←⊃Y* or *R←↑Y*

The symbol chosen to represent Disclose depends on the current Migration Level.

If *□ML* < 2, Disclose is represented by the symbol: ⊃.
If *□ML* ≥ 2, Disclose is represented by the symbol: ↑.

Y may be any array. *R* is an array. If *Y* is non-empty, *R* is the value of the first item of *Y* taken in ravel order. If *Y* is empty, *R* is the prototype of *Y*.

Disclose is the inverse of enclose. The identity *R←→⊃⊂R* holds for all *R*. Disclose is also referred to as First.

Examples

```
      ⊃1
1

      ⊃2 4 6
2

      ⊃'MONDAY' 'TUESDAY'
MONDAY

      ⊃(1 (2 3))(4 (5 6))
1  2 3

      ⊃⍳0
0

      ' '=⊃''
1

      ⊃1↓⊂1,⊂2 3
0  0 0
```

Divide: $R \leftarrow X \div Y$

Y must be a numeric array. *X* must be a numeric array. *R* is the numeric array resulting from *X* divided by *Y*. System variable $\Box DIV$ is an implicit argument of Divide.

If $\Box DIV = 0$ and $Y = 0$ then if $X = 0$, the result of $X \div Y$ is 1; if $X \neq 0$ then $X \div Y$ is a *DOMAIN ERROR*.

If $\Box DIV = 1$ and $Y = 0$, the result of $X \div Y$ is 0 for all values of *X*.

Examples

```
      2 0 5÷4 0 2
0.5 1 2.5

      ⎕DIV←1
      2 0 5÷4 0 0
0.5 0 0
```

Drop: $R \leftarrow X \downarrow Y$

Y may be any array. *X* must be a simple scalar or vector of integers. If *X* is a scalar, it is treated as a one-element vector. If *Y* is a scalar, it is treated as an array whose shape is `(ρX)ρ1`. After any scalar extensions, `ρX` must equal `ρρY` (the shape of *X* must equal the rank of *Y*).

R is an array with the same rank as *Y* but with elements removed from the vectors along each of the axes of *Y*. For the *I*th axis:

1.   if *X[I]* is positive, all but the first *X[I]* elements of the vectors result.

2.   if *X[I]* is negative, all but the last *X[I]* elements of the vectors result.

If the magnitude of *X[I]* exceeds the length of the *I*th axis, the result is an empty array with zero length along that axis.

Examples

```
      4↓'OVERBOARD'
BOARD

      ¯5↓'OVERBOARD'
OVER

      ρ10↓'OVERBOARD'
0

      M
ONE
FAT
FLY

      0 ¯2↓M
O
F
F

      ¯2 ¯1↓M
ON
```

Drop with Axes: $R \leftarrow X \downarrow [K] Y$

Y may be any non scalar array. X must be a simple integer scalar or vector. K is a vector of zero or more axes of Y.

R is an array of the elements of Y with the first or last $X[i]$ elements removed. Elements are removed from the beginning or end of Y according to the sign of $X[i]$.

The rank of R is the same as the rank of Y:

$\rho\rho R \leftrightarrow \rho\rho Y$

The size of each axis of R is determined by the corresponding element of X:

$$(\rho R)[,K] \leftrightarrow 0\lceil(\rho Y)[,K]-|,X$$

Examples

```
[]←M←2 3 4ρι24
 1  2  3  4
 5  6  7  8
 9 10 11 12

13 14 15 16
17 18 19 20
21 22 23 24

      1↓[2]M
 5  6  7  8
 9 10 11 12

17 18 19 20
21 22 23 24

      2↓[3]M
 3  4
 7  8
11 12

15 16
19 20
23 24

      2 1↓[3 2]M
 7  8
11 12
```

```
      19 20
      23 24
```

Enclose: $R \leftarrow \subset Y$

Y may be any array. R is a scalar array whose item is the array Y. If Y is a simple scalar, R is the simple scalar unchanged. Otherwise, R has a depth whose magnitude is one greater than the magnitude of the depth of Y.

Examples

```
      ⊂1
1

      ⊂'A'
A

      ⊂1 2 3
 1 2 3

      ⊂1,⊂'CAT'
 1  CAT

      ⊂2 4ρι8
 1 2 3 4
 5 6 7 8

      ⊂ι0

      ⊂⊂ι0

      ⊂⊂10
10
```

Enclose with Axes: $R \leftarrow \subset [K] Y$

`Y` may be any array. `K` is a vector of zero or more axes of `Y`. `R` is an array of the elements of `Y` enclosed along the axes `K`. The shape of `R` is the shape of `Y` with the `K` axes removed:

$$\rho R \quad \leftrightarrow \quad (\rho Y)[(\iota \rho \rho R) \sim K]$$

The shape of each element of `R` is the shape of the `K`'th axes of `Y`:

$$\rho \supset R \quad \leftrightarrow \quad (\rho Y)[,K]$$

Examples

```
      DISPLAY A←2 3 4ρ'DUCKSWANBIRDWORMCAKESEED'
..→---.
↓↓DUCK|
||SWAN|
||BIRD|
||    |
||WORM|
||CAKE|
||SEED|
''----'
      DISPLAY ⊂[3]A
.→--------------------.
↓ .→---. .→---. .→---. |
| |DUCK| |SWAN| |BIRD| |
| '----' '----' '----' |
| .→---. .→---. .→---. |
| |WORM| |CAKE| |SEED| |
| '----' '----' '----' |
'∊--------------------'
      DISPLAY⊂[2 3]A
.→--------------.
| .→---. .→---. |
| ↓DUCK| ↓WORM| |
| |SWAN| |CAKE| |
| |BIRD| |SEED| |
| '----' '----' |
'∊--------------'
      DISPLAY ⊂[1 3]A
.→--------------------.
| .→---. .→---. .→---. |
| ↓DUCK| ↓SWAN| ↓BIRD| |
| |WORM| |CAKE| |SEED| |
```

```
|  '----'  '----'  '----'  |
'∈----------------------'
```

Encode: $R \leftarrow X \top Y$

Y must be a simple numeric array. X must be a simple numeric array. R is the numeric array which results from the representation of Y in the number system defined by X.

The shape of R is $(\rho X), \rho Y$ (the catenation of the shapes of X and Y).

If X is a vector or a scalar, the result for each element of Y is the value of the element expressed in the number system defined by radix X. If Y is greater than can be expressed in the number system, the result is equal to the representation of the residue $(\times/X)|Y$. If the first element of X is 0, the value will be fully represented.

This function is also known as Representation.

Examples

```
      10⊤5 15 125
5 5 5

      0 10⊤5 15 125
0 1 12
5 5  5
```

If X is a higher order array, each of the vectors along the first axis of X is used as the radix vector for each element of Y.

Examples

```
         A
2  0   0
2  0   0
2  0   0
2  0   0
2  8   0
2  8   0
2  8  16
2  8  16

         A⊤75
0  0   0
1  0   0
0  0   0
0  0   0
1  0   0
0  1   0
1  1   4
1  3  11
```

The example shows binary, octal and hexadecimal representations of the decimal number 75.

Examples

```
       0 1⊤1.25 10.5
 1      10
 0.25   0.5

       4 13⊤13?52
 3  1  0   2  3  2  0  1   3  1  2  3  1
12  2  4  12  1  7  6  3  10  1  0  3  8
```

Enlist: ($\Box ML \geq 1$) $R \leftarrow \in Y$

Migration level must be such that $\Box ML \geq 1$ (otherwise see function Type).

Y may be any array, R is a simple vector created from all the elements of Y in ravel order.

Examples

```
      ⎕ML←1            ⍝  Migration level 1

      MAT←2 2ρ'MISS' 'IS' 'SIP' 'PI' ◇ MAT
MISS   IS
SIP    PI


      ∈MAT
MISSISSIPPI

      M←1 (2 2ρ2 3 4 5) (6(7 8))
      M
1   2 3   6   7 8
    4 5
      ∈M
1 2 3 4 5 6 7 8
```

Equal: $R \leftarrow X = Y$

Y may be any array. X may be any array. R is boolean. $\Box CT$ is an implicit argument of Equal.

If X and Y are character, then R is 1 if they are the same character. If X is character and Y is numeric, or vice-versa, then R is 0.

If X and Y are numeric, then R is 1 if X and Y are within comparison tolerance of each other. That is, X is considered equal to Y if $(|X-Y)$ is not greater than $\Box CT \times (|X) \lceil |Y$.

Examples

```
      3=3.1 3 ¯2 ¯3
0 1 0 0

      'CAT'='FAT'
0 1 1

      'CAT'=1 2 3
0 0 0

      'CAT'='C' 2 3
1 0 0

      ⎕CT←1E¯10
      1=1.000000000001
1

      1=1.0000001
0
```

Excluding: $R \leftarrow X \sim Y$

X must be a scalar or vector. *R* is a vector of the elements of *X* excluding those elements which occur in *Y* taken in the order in which they occur in *X*.

Elements of *X* and *Y* are considered the same if *X*≡*Y* returns 1 for those elements.

⎕*CT* is an implicit argument of Excluding.

This function is also known as Without.

Examples

```
      'HELLO'~'GOODBYE'
HLL
      'MONDAY' 'TUESDAY' 'WEDNESDAY'~'TUESDAY' 'FRIDAY'
 MONDAY  WEDNESDAY

      5 10 15~ι10
15
```

For performance information, see *Search Functions and Hash Tables* in Chapter 2.

Execute (Monadic): $R \leftarrow \underline{\pm} Y$

Y must be a simple character scalar or vector. If Y is an empty vector, it is treated as an empty character vector. Y is taken to be an APL statement to be executed. R is the result of the last-executed expression. If the expression has no value, then $\underline{\pm} Y$ has no value. If Y is an empty vector or a vector containing only blanks, then $\underline{\pm} Y$ has no value.

If Y contains a branch expression which evaluates to a non-empty result, R does not yield a result. Instead, the branch is effected in the environment from which the Execute was invoked.

Examples

```
      ±'2+2'
4
      4=±'2+2'
1
      A
1 2 3
4 5 6
      ±'A'
1 2 3
4 5 6
      ±'A←2|¯1↑□TS  ◇  →0ρ∴A  ◇  A'
0
      A
0
```

Execute (Dyadic): $R \leftarrow X \underline{\pm} Y$

Y must be a simple character scalar or vector. If Y is an empty vector, it is treated as an empty character vector. X must be a simple character scalar or vector representing the name of a namespace. Y is then taken to be an APL statement to be executed in namespace X. R is the result of the last-executed expression. If the expression has no value, then $X \underline{\pm} Y$ has no value.

Example

```
      '□SE'  ±  '□NL 9'
```

Expand: $R \leftarrow X \backslash [K] Y$

Y may be any array. X is a simple integer scalar or vector. The axis specification is optional. If present, K must be a simple integer scalar or unit vector. The value of K must be an axis of Y. If absent, the last axis of Y is implied. The form $R \leftarrow X \backslash Y$ implies the first axis. If Y is a scalar, it is treated as a one-element vector.

The number of positive elements in X must be the length of Kth (or implied) axis of Y.

R is composed from the sub-arrays along the Kth axis of Y. If $X[I]$ (an element of X) is the Jth positive element in X, then the Jth sub-array along the Kth axis of Y is replicated $X[I]$ times. If $X[I]$ is negative, then a sub-array of fill elements of Y ($\subset \in \supset Y$) is replicated $|X[I]$ times and inserted in relative order along the Kth axis of the result. If $X[I]$ is zero, it is treated as the value $^-1$. The shape of R is the shape of Y except that the length of the Kth axis is $+/1\lceil|X$.

Examples

```
      0\ι0
0

      1 ¯2 3 ¯4 5\'A'
A  AAA    AAAAA

      M
1 2 3
4 5 6

      1 ¯2 2 0 1\M
1 0 0 2 2 0 3
4 0 0 5 5 0 6

      1 0 1\M
1 2 3
0 0 0
4 5 6

      1 0 1\[1]M
1 2 3
0 0 0
4 5 6

      1 ¯2 1\(1 2)(3 4 5)
 1 2  0 0  0 0  3 4 5
```

Expand First: $R \leftarrow X \mathbin{\backslash\mkern-10mu-} Y$

The form $R \leftarrow X \mathbin{\backslash\mkern-10mu-} Y$ implies expansion along the first axis whereas the form $R \leftarrow X \backslash Y$ implies expansion along the last axis (columns). See Expand above.

Exponential: $R \leftarrow \star Y$

Y must be numeric. R is numeric and is the Yth power of e, the base of natural logarithms.

Example

```
      *1 0
2.718281828 1
```

Factorial: $R \leftarrow !Y$

Y must be numeric excluding negative integers. R is numeric. R is the product of the first Y integers for positive integer values of Y. For non-integral values of Y, $!Y$ is equivalent to the gamma function of $Y+1$.

Examples

```
      !1 2 3 4 5
1 2 6 24 120

      !¯1.5 0 1.5 3.3
¯3.544907702 1 1.329340388 8.85534336
```

Find: $R \leftarrow X \underline{\in} Y$

X and *Y* may be any arrays. *R* is a simple boolean array the same shape as *Y* which identifies occurrences of *X* within *Y*

If the rank of *X* is smaller than the rank of *Y*, *X* is treated as if it were the same rank with leading axes of size 1. For example a vector is treated as a 1-row matrix.

If the rank of *X* is larger than the rank of *Y*, no occurrences of *X* are found in *Y*.

□*CT* and □*IO* are implicit arguments to Find.

Examples

```
      'AN'∈'BANANA'
0 1 0 1 0 0

      'ANA'∈'BANANA'
0 1 0 1 0 0

      'BIRDS' 'NEST'∈'BIRDS' 'NEST' 'SOUP'
1 0 0

    MAT
IS YOU IS
OR IS YOU
IS'NT

      'IS'∈MAT
1 0 0 0 0 0 0 0 1 0
0 0 0 1 0 0 0 0 0 0
1 0 0 0 0 0 0 0 0 0

      'IS YOU'∈MAT
1 0 0 0 0 0 0 0 0 0
0 0 0 1 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0
```

First: ($\Box ML$) $R \leftarrow \supset Y$ or $R \leftarrow \uparrow Y$

See function Disclose.

Floor: $R \leftarrow \lfloor Y$

Y must be numeric. R is the largest integer value less than or equal to Y.

Examples

```
      ⌊¯2.3 0.1 100 3.3
¯3 0 100 3

      ⌊0.5 + 0.4 0.5 0.6
0 1 1
```

$\Box CT$ is an implicit argument of Floor.

Format (Monadic): $R \leftarrow \overline{\Phi} Y$

Y may be any array. R is a simple character array which will display identically to the display produced by Y. The result is independent of $\Box PW$. If Y is a simple character array, then R is Y.

Example

```
      +B←⍕A←2 6ρ'HELLO PEOPLE'
HELLO
PEOPLE

      B ≡ A
1
```

If Y is a simple numeric scalar, then R is a vector containing the formatted number without any spaces. A floating point number is formatted according to the system variable $\Box PP$. $\Box PP$ is ignored when formatting integers.

Examples

```
      ⎕PP←5

      ρC←⍕⍳0
0

      ρC←⍕10
2

      C
10

      ρC←⍕12.34
5

      C
12.34

      ⍕123456789
123456789

      ⍕123.456789
123.46
```

Scaled notation is used if the magnitude of the non-integer number is too large to represent with ⎕PP significant digits or if the number requires more than five leading zeroes after the decimal point.

Examples

```
      ⍕123456.7
1.2346E5


      ⍕0.0000001234
1.234E¯7
```

If Y is a simple numeric vector, then R is a character vector in which each element of Y is independently formatted with a single separating space between formatted elements.

Example

```
      ⍴C←⍕¯123456 1 22.5 ¯0.000000667 5.00001
27

      C
¯1.2346E5 1 22.5 ¯6.67E¯7 5
```

If Y is a simple numeric array rank higher than one, R is a character array with the same shape as Y except that the last dimension of Y is determined by the length of the formatted data. The format width is determined independently for each column of Y, such that:

a)   the decimal points for floating point or scaled formats are aligned.

b)   the E characters for scaled formats are aligned, with trailing zeros added to the mantissae if necessary.

c)   integer formats are aligned to the left of the decimal point column, if any, or right-adjusted in the field otherwise.

d)   each formatted column is separated from its neighbours by a single blank column.

e)   the exponent values in scaled formats are left-adjusted to remove any blanks.

Examples

```
      C←22 ¯0.000000123 2.34 ¯212 123456 6.00002 0

      ρC←⍕2 2 3ρC
2 2 29

      C
  22    ¯1.2300E¯7  2.3400E0
¯212     1.2346E5   6.0000E0

   0     2.2000E1  ¯1.2300E¯7
   2.34 ¯2.1200E2   1.2346E5
```

If Y is non-simple, and all items of Y at any depth are scalars or vectors, then R is a vector.

Examples

```
      B←⍕A←'ABC' 100 (1 2 (3 4 5)) 10

      ρA
4
      ≡A
¯3

      ρB
26
      ≡B
1

      A
 ABC  100  1 2  3 4 5   10

      B
 ABC  100  1 2  3 4 5   10
```

By replacing spaces with ∧, it is clearer to see how the result of ⍕ is formed:

∧ABC∧∧100∧∧1∧2∧∧3∧4∧5∧∧∧10

If Y is non-simple, and all items of Y at any depth are not scalars, then R is a matrix.

Example

```
      D←⍕C←1 'AB' (2 2⍴1+⍳4) (2 2 3⍴'CDEFGHIJKLMN')

      C
1   AB   2 3   CDE
         4 5   FGH

               IJK
               LMN

      ⍴C
4

      ≡C
¯2

      D
1   AB   2 3   CDE
         4 5   FGH

               IJK
               LMN

      ⍴D
5 16

      ≡D
1
```

By replacing spaces with ∧, it is clearer to see how the result of ⍕ is formed:

```
1∧∧AB∧∧2∧3∧∧CDE∧
∧∧∧∧∧∧∧4∧5∧∧FGH∧
∧∧∧∧∧∧∧∧∧∧∧∧∧∧∧∧
∧∧∧∧∧∧∧∧∧∧∧∧IJK∧
∧∧∧∧∧∧∧∧∧∧∧∧LMN∧
```

⎕PP is an implicit argument of Monadic Format.

Format (Dyadic): $R \leftarrow X \, \overline{\Phi} \, Y$

Y must be a simple numeric array. X must be a simple integer scalar or vector. R is a character array displaying the array Y according to the specification X. R has rank $1 \lceil \rho \rho Y$ and $\bar{1} \downarrow \rho R$ is $\bar{1} \downarrow \rho Y$.

Conformability requires that if X has more than two elements, then ρX must be $2 \times \bar{1} \uparrow \rho Y$. If X contains one element, it is extended to $(2 \times \bar{1} \uparrow \rho Y) \rho 0, X$. If X contains 2 elements, it is extended to $(2 \times \bar{1} \uparrow \rho Y) \rho X$.

X specifies two numbers (possibly after extension) for each column in Y. For this purpose, scalar Y is treated as a one-element vector. Each pair of numbers in X identifies a format width (W) and a format precision (P).

If P is 0, the column is to be formatted as integers.

Examples

```
      5 0 ⍕ 2 3⍴⍳6
 1     2     3
 4     5     6

      4 0⍕1.1 2 ¯4 2.547
 1     2    ¯4     3
```

If P is positive, the format is floating point with P significant digits to be displayed after the decimal point.

Example

```
      4 1⍕1.1 2 ¯4 2.547
 1.1 2.0¯4.0 2.5
```

If P is negative, scaled format is used with $| P$ digits in the mantissa.

Example

```
      7 ¯3⍕5 15 155 1555
5.00E0 1.50E1 1.55E2 1.56E3
```

If W is 0 or absent, then the width of the corresponding columns of R are determined by the maximum width required by any element in the corresponding columns of Y, plus one separating space.

Example

```
      3⍕2 3⍴10 15.2346 ¯17.1 2 3 4
  10.000   15.235  ¯17.100
   2.000    3.000    4.000
```

If a formatted element exceeds its specified field width when $W>0$, the field width for that element is filled with asterisks.

Example

```
      3 0 6 2 ⍕ 3 2⍴10.1 15 1001 22.357 101 1110.1
 10 15.00
*** 22.36
101******
```

If the format precision exceeds the internal precision, low order digits are replaced by the symbol '_'.

Example

```
      26⍕2*100
 12676506002282_____._____
_

      ⍴26⍕2*100
59

      0 20⍕÷3
 0.33333333333333_____

      0 ¯20⍕÷3
 3.333333333333333____E¯1
```

The shape of R is the same as the shape of Y except that the last dimension of Y is the sum of the field widths specified in X or deduced by the function. If Y is a scalar, the shape of R is the field width.

```
      ⍴5 2 ⍕ 2 3 4⍴⍳24
2 3 20
```

Grade Down (Monadic): $R \leftarrow \psi Y$

Y must be a simple character or simple numeric array of rank greater than 0. R is an integer vector being the permutation of $\iota 1 \uparrow \rho Y$ that places the sub-arrays of Y along the first axis in descending order. The indices of any set of identical sub-arrays in Y occur in R in ascending order.

If Y is a numeric array of rank greater than 1, the elements in each of the sub-arrays along the first axis are compared in ravel order with greatest weight being given to the first element and least weight being given to the last element.

Example

```
          M
2  5  3  2
3  4  1  1
2  5  4  5
2  5  3  2
2  5  3  4

          ψM
2  3  5  1  4

          M[ψM;]
3  4  1  1
2  5  4  5
2  5  3  4
2  5  3  2
2  5  3  2
```

If Y is a character array, the implied collating sequence is the ordering of characters in $\square AV$ (see *Chapter 8*).

Example

```
        A
SUCH
SURE
SIRE
SUMP
SAFE


        ⍒A
5 2 4 1 3


        A[⍒A;]
SAFE
SURE
SUMP
SUCH
SIRE
```

⎕*IO* is an implicit argument of Grade Down.

Grade Down (Dyadic): $R \leftarrow X \psi Y$

Y must be a simple character array of rank greater than 0. X must be a simple character array of rank 1 or greater. R is a simple integer vector of shape $1 \uparrow \rho Y$ containing the permutation of $\iota 1 \uparrow \rho Y$ that places the sub-arrays of Y along the first axis in descending order according to the collation sequence X. The indices of any set of identical sub-arrays in Y occur in R in ascending order.

If X is a vector, the following identity holds:

$$X \psi Y \leftrightarrow \psi X \iota Y$$

A left argument of rank greater than 1 allows successive resolution of duplicate orderings in the following way.

Starting with the last axis:

- The characters in the right argument are located along the current axis of the left argument. The position of the first occurrence gives the ordering value of the character.

- If a character occurs more than once in the left argument its lowest position along the current axis is used.

- If a character of the right argument does not occur in the left argument, the ordering value is one more than the maximum index of the current axis - as with dyadic iota.

The process is repeated using each axis in turn, from the last to the first, resolving duplicates until either no duplicates result or all axes have been exhausted.

For example, if index origin is 1:

Left argument:	**Right argument:**
abc	*ab*
ABA	*ac*
	Aa
	Ac

Along last axis:

Character:	Value:	Ordering:
ab	1 2	3
ac	1 3	=1 *A* ⇐ duplicate ordering with
Aa	1 1	4
Ac	1 3	=1 ⇐ respect to last axis.

Duplicates exist, so resolve these with respect to the first axis:

Character:	Value:	Ordering:
ac	1 1	2
Ac	2 1	1

So the final row ordering is:

ab	3
ac	2
Aa	4
Ac	1

That is, the order of rows is 4 2 1 3 which corresponds to a descending row sort of:

Ac	1
ac	2
ab	3
Aa	4

Examples

```
      ρS1
2 27

      S1
ABCDEFGHIJKLMNOPQRSTUVWXYZ
A̲B̲C̲D̲E̲F̲G̲H̲I̲J̲K̲L̲M̲N̲O̲P̲Q̲R̲S̲T̲U̲V̲W̲X̲Y̲Z̲

      S2
ABCDEFGHIJKLMNOPQRSTUVWXYZA̲B̲C̲D̲E̲F̲G̲H̲I̲J̲K̲L̲M̲N̲O̲P̲Q̲R̲S̲T̲U̲V̲W̲X̲Y̲Z̲

      S3
AA̲BB̲CC̲DD̲EE̲FF̲GG̲HH̲II̲JJ̲KK̲LL̲MM̲NN̲OO̲PP̲QQ̲RR̲SS̲TT̲UU̲VV̲WW̲XX̲YY̲ZZ̲

      S4
ABCDEFGHIJKLMNOPQRSTUVWXYZA̲B̲C̲D̲E̲F̲G̲H̲I̲J̲K̲L̲M̲N̲O̲P̲Q̲R̲S̲T̲U̲V̲W̲X̲Y̲Z̲
A̲B̲C̲D̲E̲F̲G̲H̲I̲J̲K̲L̲M̲N̲O̲P̲Q̲R̲S̲T̲U̲V̲W̲X̲Y̲Z̲ABCDEFGHIJKLMNOPQRSTUVWXYZ
```

The following results are tabulated for comparison:

X	$X[S1{\Downarrow}X;]$	$X[S2{\Downarrow}X;]$	$X[S3{\Downarrow}X;]$	$X[S4{\Downarrow}X;]$
FIRS̲T̲	TAPE	R̲A̲T	TAPE	TAPE
TAP	TAP	F̲I̲RST	TAP	TAP
RATE	RATE	TAPE	R̲A̲T	RATE
FI̲RST	R̲A̲T	TAP	RATE	R̲A̲T
FIRST	RAT	RATE	RAT	RAT
R̲A̲T	MAT	RAT	MAT	MAT
F̲I̲RST	F̲I̲RST	MAT	F̲I̲RST	FIRST
TAPE	FI̲RST	FI̲RST	FI̲RST	FI̲RST
MAT	FIRS̲T̲	FIRS̲T̲	FIRS̲T̲	FIRST
RAT	FIRST	FIRST	FIRST	F̲I̲RST

`⎕IO` is an implicit argument of Grade Down.

Grade Up (Monadic): $R \leftarrow \spadesuit Y$

Y must be a simple character or simple numeric array of rank greater than 0. *R* is an integer vector being the permutation of $\iota 1 \uparrow \rho Y$ that places the sub-arrays along the first axis in ascending order.

If *Y* is a numeric array of rank greater than 1, the elements in each of the sub-arrays along the first axis are compared in ravel order with greatest weight being given to the first element and least weight being given to the last element.

Examples

```
      ▲22.5 1 15 3 ¯4
5 2 4 3 1


      M
2 3 5
1 4 7

2 3 5
1 2 6

2 3 4
5 2 4


      ▲M
3 2 1
```

If *Y* is a character array, the implied collating sequence is the ordering of characters in *⎕AV* (see *Chapter 8*).

```
      A
A̲A̲1∇φ


      ▲A
3 1 2 4 5
```

⎕IO is an implicit argument of Grade Up.

Grade Up (Dyadic): $R \leftarrow X \triangle Y$

Y must be a simple character array of rank greater than 0. X must be a simple character array of rank 1 or greater. R is a simple integer vector being the permutation of $\iota 1 \uparrow \rho Y$ that places the sub-arrays of Y along the first axis in ascending order according to the collation sequence X.

If X is a vector, the following identity holds:

$$X \triangle Y \quad \leftrightarrow \quad \triangle X \iota Y$$

If X is a higher order array, each axis of X represents a grading attribute in decreasing order of importance. If a character is repeated in X, it is treated as though it were located at the position in the array determined by the lowest index in each axis for all occurrences of the character. The character has the same weighting as the character located at the derived position in X.

Refer to dyadic Grade Down for an illustrative example of these principles.

Greater: $R \leftarrow X > Y$

Y must be numeric. X must be numeric. R is boolean. R is 1 if X is greater than Y and $X = Y$ is 0. Otherwise R is 0.

$\square CT$ is an implicit argument of Greater.

Examples

```
      1 2 3 4 5 > 2
0 0 1 1 1

      □CT←1E¯10

      1 1.00000000001 1.000000001 > 1
0 0 1
```

Greater Or Equal: $R \leftarrow X \geq Y$

Y must be numeric. X must be numeric. R is boolean. R is 1 if X is greater than Y or $X=Y$. Otherwise R is 0.

$\Box CT$ is an implicit argument of Greater Or Equal.

Examples

```
      1 2 3 4 5 ≥ 3
0 0 1 1 1

      ⎕CT←1E¯10

      1≥1
1

      1≥1.00000000001
1

      1≥1.00000001
0
```

Identity: $R \leftarrow +Y$

Y is any array. R is the same array unchanged. Identity may be used to generate a printed result after an assignment or from a function with an otherwise suppressed result.

Examples

```
      +A←⍳5
1 2 3 4 5

      +⎕EX'A'
1
```

Index Generator: $R \leftarrow \iota Y$

Y must be a simple scalar or vector array of non-negative numbers. *R* is a numeric array composed of the set of all possible coordinates of an array of shape *Y*. The shape of *R* is *Y* and each element of *R* occurs in its self-indexing position in *R*. In particular, the following identity holds:

        ιY ↔ (ιY)[ιY]

□IO is an implicit argument of Index Generator. This function is also known as Interval.

Examples

```
      □IO
1
      ρι0
0
      ι5
1 2 3 4 5

      ι2 3
 1 1   1 2   1 3
 2 1   2 2   2 3

      +A←2 4ρ'MAINEXIT'
MAIN
EXIT

      A[ιρA]
MAIN
EXIT

      □IO←0

      ι5
0 1 2 3 4

      ι2 3
 0 0   0 1   0 2
 1 0   1 1   1 2

      A[ιρA]
MAIN
EXIT
```

Index Of: $R \leftarrow X \iota Y$

Y may be any array. X may be any vector. R is a simple integer array with the same shape as Y identifying where elements of Y are first found in X. If an element of Y cannot be found in X, then the corresponding element of R will be $\Box IO + \rho X$.

Elements of X and Y are considered the same if $X \equiv Y$ returns 1 for those elements.

$\Box IO$ and $\Box CT$ are implicit arguments of Index Of.

Examples

```
      ⎕IO←1

      2 4 3 1 4⍳1 2 3 4 5
4 1 3 2 6

      'CAT' 'DOG' 'MOUSE'⍳'DOG' 'BIRD'
2 4
```

For performance information, see *Search Functions and Hash Tables* in Chapter 2.

Indexing: $R \leftarrow X [Y]$

X may be any array. Y must be a valid index specification. R is an array composed of elements indexed from X and the shape of X is determined by the index specification.

Bracket Indexing does not follow the normal syntax of a dyadic function.

$\Box IO$ is an implicit argument of Indexing.

Three forms of indexing are permitted. The form used is determined by context.

Simple Indexing

For vector X, Y is a simple integer array composed of items from the set $\iota \rho X$.

R consists of elements selected according to index positions in Y. R has the same shape as Y.

Examples

```
      A←10 20 30 40 50

      A[2 3ρ1 1 1 2 2 2]
10 10 10
20 20 20

      A[3]
30

      'ONE' 'TWO' 'THREE'[2]
 TWO
```

For matrix *X*, *Y* is composed of two simple integer arrays separated by the semicolon character (;). The arrays select indices from the rows and columns of *X* respectively.

Examples

```
      +M←2 4ρ10×ι8
10 20 30 40
50 60 70 80

      M[2;3]
70
```

For higher order array *X*, *Y* is composed of a simple integer array for each axis of *X* with adjacent arrays separated by a single semicolon character (;). The arrays select indices from the respective axes of *X*, taken in row-major order.

Examples

```
      +A←2 3 4ρ10×ι24
 10   20   30   40
 50   60   70   80
 90  100  110  120

130  140  150  160
170  180  190  200
210  220  230  240

      A[1;1;1]
10

A[2;3 2;4 1]
240 210
200 170
```

If an indexing array is omitted for the Kth axis, the index vector $\iota(\rho X)[K]$ is assumed for that axis.

Examples

```
      A[;2;]
 50  60  70  80
170 180 190 200

      M
10 20 30 40
50 60 70 80

      M[;]
10 20 30 40
50 60 70 80

      M[1;]
10 20 30 40

      M[;1]
10 50
```

Choose Indexing

The index specification Y is a non-simple array. Each item identifies a single element of X by a set of indices with one element per axis of X in row-major order.

Examples

```
      M
10 20 30 40
50 60 70 80

      M[⊂1 2]
20

      M[2 2ρ⊂2 4]
80 80
80 80

      M[(2 1)(1 2)]
50 20
```

A scalar may be indexed by the enclosed empty vector:

```
      S←'Z'

      S[3ρ⊂ι0]
```

zzz

Simple and Choose indexing are indistinguishable for vector *X*:

```
      V←10 20 30 40

      V[⊂2]
20
      ⊂2
2
      V[2]
20
```

Reach Indexing

The index specification *Y* is a non-simple integer array, each of whose items reach down to a nested element of *X*. The items of an item of *Y* are simple vectors (or scalars) forming sets of indices that index arrays at successive levels of *X* starting at the top-most level. A set of indices has one element per axis at the respective level of nesting of *X* in row-major order.

Examples

```
      G←('ABC' 1)('DEF' 2)('GHI' 3)('JKL' 4)

      G←2 3ρG,('MNO' 5)('PQR' 6)

      G
 ABC  1    DEF  2    GHI  3
 JKL  4    MNO  5    PQR  6

      G[((1 2)1)((2 3)2)]
 DEF  6

      G[2 2ρ⊂(2 2)2]
5 5
5 5
      G[⊂⊂1 1]
  ABC  1

      G[⊂1 1]
  ABC  1

      V←,G

      V[⊂⊂1]
  ABC  1

      V[⊂1]
```

```
ABC   1

      V[1]
ABC   1
```

Intersection: $R \leftarrow X \cap Y$

Y must be a scalar or vector. *X* must be a scalar or vector. A scalar *X* or *Y* is treated as a one-element vector. *R* is a vector composed of items occurring in both *X* and *Y* in the order of occurrence in *X*. If an item is repeated in *X* and also occurs in *Y*, the item is also repeated in *R*.

Items in *X* and *Y* are considered the same if *X* ≡ *Y* returns 1 for those items.

$\square CT$ is an implicit argument of Intersection.

Examples

```
      'ABRA'∩'CAR'
ARA

      1 'PLUS' 2 ∩ ⍳5
1 2
```

For performance information, see *Search Functions and Hash Tables* in Chapter 2.

Less: $R \leftarrow X < Y$

Y may be any numeric array. *X* may be any numeric array. *R* is boolean. *R* is 1 if *X* is less than *Y* and *X* = *Y* is 0. Otherwise *R* is 0.

$\square CT$ is an implicit argument of Less.

Examples

```
      (2 4) (6 8 10) < 6
 1 1  0 0 0

      ⎕CT←1E¯10

      1 1.00000000001 1.000000001 < 1
0 0 1
```

Less Or Equal: $R \leftarrow X \le Y$

Y may be any numeric array. X may be any numeric array. R is boolean. R is 1 if X is less than Y or $X = Y$. Otherwise R is 0.

$\square CT$ is an implicit argument of Less Or Equal.

Examples

```
      2 4 6 8 10 ≤ 6
1 1 1 0 0

      ⎕CT←1E¯10

      1  1.00000000001 1.00000001 ≤ 1
1 1 0
```

Logarithm: $R \leftarrow X \circledast Y$

Y must be a positive numeric array. X must be a positive numeric array. X cannot be 1 unless Y is also 1. R is the base X logarithm of Y.

Examples

```
      10 ⊛ 100 2
2 0.3010299957

      1 ⊛ 1
1

      2 ⊛ 1
0
```

Magnitude: $R \leftarrow | Y$

Y may be any numeric array. R is numeric composed of the absolute (unsigned) values of Y.

Example

```
      |2 ¯3.4 0 ¯2.7
2 3.4 0 2.7
```

Note that the ASCII pipe (¦) is also interpreted as Magnitude (|).

Match: $R \leftarrow X \equiv Y$

Y may be any array. X may be any array. R is a simple boolean scalar. If X is identical to Y, then R is 1. Otherwise R is 0.

Non-empty arrays are identical if they have the same structure and the same values in all corresponding locations. Empty arrays are identical if they have the same shape and the same prototype (disclosed nested structure).

$\square CT$ is an implicit argument of Match.

Examples

```
      θ≡ι0
1
      ''≡ι0
0
      A
THIS
WORD

      A≡2 4ρ'THISWORD'
1
      A≡ι10
0
      +B←A A
 THIS  THIS
 WORD  WORD

      A≡⊃B
1

      (0ρA)≡0ρB
0

      ' '=⊃0ρB
1 1 1 1
1 1 1 1

      ' '=⊃0ρA
1
```

Matrix Divide: $R \leftarrow X \boxdot Y$

Y must be a simple numeric array of rank 2 or less. *X* must be a simple numeric array of rank 2 or less. *Y* must be non-singular. A scalar argument is treated as a matrix with one-element. If *Y* is a vector, it is treated as a single column matrix. If *X* is a vector, it is treated as a single column matrix. The number of rows in *X* and *Y* must be the same. *Y* must have at least the same number of rows as columns.

R is the result of matrix division of *X* by *Y*. That is, the matrix product $Y+.\times R$ is *X*. *R* is determined such that $(X-Y+.\times R)*2$ is minimised.

The shape of *R* is $(1\downarrow\rho Y),1\downarrow\rho X$.

Examples

```
      ⎕PP←5

      B
3 1 4
1 5 9
2 6 5

      35 89 79 ⌹ B
2.1444 8.2111 5.0889

      A
35 36
89 88
79 75

      A ⌹ B
2.1444 2.1889
8.2111 7.1222
5.0889 5.5778
```

If there are more rows than columns in the right argument, the least squares solution results. In the following example, the constants a and b which provide the best fit for the set of equations represented by P = a + bQ are determined:

```
      Q
1 1
1 2
1 3
1 4
1 5
1 6

      P
12.03 8.78 6.01 3.75 ¯0.31 ¯2.79

      P⌹Q
14.941 ¯2.9609
```

Matrix Inverse: $R \leftarrow \boxminus Y$

Y must be a simple array of rank 2 or less. Y must be non-singular. If Y is a scalar, it is treated as a one-element matrix. If Y is a vector, it is treated as a single-column matrix. Y must have at least the same number of rows as columns.

R is the inverse of Y if Y is a square matrix, or the left inverse of Y if Y is not a square matrix. That is, $R+.\times Y$ is an identity matrix.

The shape of R is $\phi \rho Y$.

Examples
```
      M
4 ¯1
2  1

      +A←⌹M
 0.1666666667 0.1666666667
¯0.3333333333 0.6666666667
```

Within calculation accuracy, $A+.\times M$ is the identity matrix.

```
      A+.×M
1 0
0 1
```

Maximum: $R \leftarrow X \lceil Y$

Y may be any numeric array. *X* may be any numeric array. *R* is numeric. *R* is the
larger of the numbers *X* and *Y*.

Example

```
      ¯2.01 0.1 15.3 ⌈ ¯3.2 ¯1.1 22.7
¯2.01 0.1 22.7
```

Membership: $R \leftarrow X \in Y$

Y may be any array. *X* may be any array. *R* is boolean. An element of *R* is 1 if the
corresponding element of *X* can be found in *Y*.

An element of *X* is considered identical to an element in *Y* if *X* ≡ *Y* returns 1 for those
elements.

⎕CT is an implicit argument of Membership.

Examples

```
      'THIS NOUN' ∊ 'THAT WORD'
1 1 0 0 1 0 1 0 0

      'CAT' 'DOG' 'MOUSE' ∊ 'CAT' 'FOX' 'DOG' 'LLAMA'
1 1 0
```

For performance information, see *Search Functions and Hash Tables* in Chapter 2.

Minimum: $R \leftarrow X \lfloor Y$

Y may be any numeric array. *X* may be any numeric array. *R* is numeric. *R* is the
smaller of *X* and *Y*.

Example

```
      ¯2.1 0.1 15.3 ⌊ ¯3.2 1 22
¯3.2 0.1 15.3
```

Minus: $R \leftarrow X - Y$

See function Subtract.

Mix:	(*☐ML*)	*R*←↑[*K*]*Y* or *R*←⊃[*K*]*Y*

The symbol chosen to represent Mix depends on the current Migration Level.

If *☐ML*<2, Mix is represented by the symbol: ↑.
If *☐ML*≥2, Mix is represented by the symbol: ⊃.

Y may be any array. All of the items of *Y* must be scalars and/or arrays of the same rank. It is not necessary that nonscalar items have the same shape.

K is an optional axis specification. If present it must be a scalar or unit vector. The value of *K* must be a fractional number indicating the two axes of *Y* between which new axes are to be inserted. If absent, new ones are added at the beginning.

R is an array composed from the items of a *Y* assembled into a higher order array with one less level of nesting. If items of *Y* have different shapes, each is padded with the corresponding prototype to a shape that represents the greatest length along each axis of all items in *Y*. The shape of *R* is the shape of *Y* with the shape of a typical (extended) item of *Y* inserted between the ⌊*K*th and the ⌈*K*th axes of *Y*.

Examples

```
      ↑(1)(1 2)(1 2 3)
1 0 0
1 2 0
1 2 3


      ↑[0.5](1) (1 2) (1 2 3)
1 1 1
0 2 2
0 0 3


      A←('andy' 19)('geoff' 37)('pauline' 21)


      ↑A
 andy    19
 geoff   37
 pauline 21


      ↑[0.5]A
 andy   geoff   pauline
   19      37        21
```

Multiply: $R \leftarrow X \times Y$

Y may be any numeric array. X may be any numeric array. R is the arithmetic product of X and Y.

This function is also known as Times.

Example

```
      3 2 1 0 × 2 4 9 6
6 8 9 0
```

Nand: $R \leftarrow X \barwedge Y$

Y must be a boolean array. X must be a boolean array. R is boolean. The value of R is the truth value of the proposition "not both X and Y", and is determined as follows:

X	Y	R
0	0	1
0	1	1
1	0	1
1	1	0

Example

```
      (0 1)(1 0) ⍲ (0 0)(1 1)
 1 1  0 1
```

Natural Logarithm: $R \leftarrow \circledast Y$

Y must be a positive numeric array. R is numeric. R is the natural (or Napierian) logarithm of Y whose base is the mathematical constant *e=2.71828....*

Example

```
      ⍟1 2
0 0.6931471806
```

Negative: $R \leftarrow - Y$

Y may be any numeric array. R is numeric. R is the negative value of Y.

Example

```
      -4  2  0 ¯3 ¯5
¯4 ¯2  0  3  5
```

Nor: $R \leftarrow X \barwedge Y$

Y must be a boolean array. X must be a boolean array. R is boolean. The value of R is the truth value of the proposition "neither X nor Y", and is determined as follows:

X	Y	R
0	0	1
0	1	0
1	0	0
1	1	0

Example

```
      0  0  1  1 ⍱ 0  1  0  1
1  0  0  0
```

Not: $R \leftarrow \sim Y$

Y must be a boolean array. R is boolean. The value of R is 0 if Y is 1, and R is 1 if Y is 0.

Example

```
      ~0  1
1  0
```

Not Equal: $R \leftarrow X \neq Y$

Y may be any array. X may be any array. R is boolean. R is 0 if $X = Y$. Otherwise R is 1.

For boolean X and Y, the value of R is the "exclusive or" result, determined as follows:

```
X    Y  │   R
────────┼──────
0    0  │   0
0    1  │   1
1    0  │   1
1    1  │   0
```

$\square CT$ is an implicit argument of Not Equal.

Examples

```
      1 2 3 ≠ 1.1 2 3
1 0 0

      □CT←1E¯10

      1≠1 1.00000000001 1.0000001
0 0 1

      1 2 3 ≠'CAT'
1 1 1
```

Not Match: $R \leftarrow X \neq Y$

Y may be any array. X may be any array. R is a simple boolean scalar. If X is identical to Y, then R is 0. Otherwise R is 1.

Non-empty arrays are identical if they have the same structure and the same values in all corresponding locations. Empty arrays are identical if they have the same shape and the same prototype (disclosed nested structure).

$\square CT$ is an implicit argument of Not Match.

Examples

```
      θ≠ι0
0
      ''≠ι0
1

      +A←⊂(ι3) 'ABC'
  1 2 3   ABC

      A≠(ι3)'ABC'
1
      A≠⊂(ι3) 'ABC'
0
      θ≠0ρA
1
      (1↑0ρA)≠⊂(0 0 0) '   '
1
```

Or: $R \leftarrow X \vee Y$

Y must be a boolean array. X must be a boolean array. R is boolean. The value of R is determined as follows:

X	Y	R
0	0	0
0	1	1
1	0	1
1	1	1

Example

```
      0 0 1 1 ∨ 0 1 0 1
0 1 1 1
```

Partition: ($\Box ML \geq 3$) $R \leftarrow X \subset [K] Y$

Y may be any non scalar array.
X must be a simple scalar or vector of non-negative integers.

The axis specification is optional. If present, it must be a simple integer scalar or one element array representing an axis of Y. If absent, the last axis is implied.

R is an array of the elements of Y partitioned according to X.

A new partition is started in the result whenever the corresponding element in X is greater than the previous one. Items in Y corresponding to 0s in X are not included in the result.

Examples

```
      ⎕ML←3

      DISPLAY 1 1 1 2 2 3 3 3⊂'NOWISTHE'
.→----------------.
| .→--. .→-. .→--. |
| |NOW| |IS| |THE| |
| '---' '--' '---' |
'∈----------------'

      DISPLAY 1 1 1 0 0 3 3 3⊂'NOWISTHE'
.→------------.
| .→--. .→--. |
| |NOW| |THE| |
| '---' '---' |
'∈------------'

      TEXT←'   NOW    IS     THE     TIME    '
      DISPLAY(' '≠TEXT)⊂TEXT
.→----------------------.
| .→--. .→-. .→--. .→---. |
| |NOW| |IS| |THE| |TIME| |
| '---' '--' '---' '----' |
'∈----------------------'

      DISPLAY CMAT←⎕FMT(' ',ROWS),COLS,NMAT
.→----------------------.
↓          Jan  Feb  Mar |
| Cakes       0  100  150 |
| Biscuits    0    0  350 |
| Buns        0 1000  500 |
```

'- -'

'- -'

```
        DISPLAY (∨≠' '≠CMAT)⊂CMAT   ⍝ Split at blank cols.
.→-----------------------------.
↓ .→-------.  .→--.  .→----.  .→--. |
| |        |  |Jan|  | Feb|  |Mar| |
| '--------'  '---'  '----'  '---' |
| .→-------.  .→--.  .→----.  .→--. |
| |Cakes   |  |  0|  | 100|  |150| |
| '--------'  '---'  '----'  '---' |
| .→-------.  .→--.  .→----.  .→--. |
| |Biscuits|  |  0|  |   0|  |350| |
| '--------'  '---'  '----'  '---' |
| .→-------.  .→--.  .→----.  .→--. |
| |Buns    |  |  0|  |1000|  |500| |
| '--------'  '---'  '----'  '---' |
'∊-----------------------------'


        DISPLAY N←4 4⍴⍳16
.→---------.
↓ 1  2  3  4|
| 5  6  7  8|
| 9 10 11 12|
|13 14 15 16|
'~---------'


        DISPLAY 1 1 0 1⊂N
.→------------.
↓ .→--.    .→. |
| |1 2|    |4| |
| '~--'    '~' |
| .→--.    .→. |
| |5 6|    |8| |
| '~--'    '~' |
| .→---.   .→-. |
| |9 10|   |12| |
| '~---'   '~-' |
| .→----.  .→-. |
| |13 14|  |16| |
| '~-----' '~-' |
'∊------------'


        DISPLAY 1 1 0 1⊂[1]N
.→----------------------.
↓ .→--.  .→--.  .→--.  .→--. |
| |1 5|  |2 6|  |3 7|  |4 8| |
| '~--'  '~--'  '~--'  '~--' |
```

```
|  .→-.   .→-.   .→-.   .→-.   |
|  |13|   |14|   |15|   |16|   |
|  '~-'   '~-'   '~-'   '~-'   |
'∊----------------------'
```

Partitioned Enclose: ($\Box ML < 3$) $R \leftarrow X \subset [K] Y$

Y may be any array. *X* must be a simple boolean scalar or vector.

The axis specification is optional. If present, it must be a simple integer scalar or one-element vector. The value of *K* must be an axis of *Y*. If absent, the last axis of *Y* is implied.

X must have the same length as the *K*th axis of *Y*. However, if *X* is a scalar or one-element vector, it will be extended to the length of the *K*th axis of *Y*.

R is a vector of items selected from *Y*. The sub-arrays identified along the *K*th axis of *Y* at positions corresponding to each 1 in *X* up to the position before the next 1 in *X* (or the last element of *X*) become the successive items of *Y*. The length of *R* is $+/X$ (after possible extension).

Examples

```
      0 1 0 0 1 1 0 0 0 ⊂⍳9
 2 3 4  5  6 7 8 9

      1 0 1 ⊂[1] 3 4⍴⍳12
1 2 3 4  9 10 11 12
5 6 7 8

      1 0 0 1 ⊂[2]3 4⍴⍳12
1  2  3  4
5  6  7  8
9 10 11  12
```

Pi Times: $R \leftarrow \circ Y$

Y may be any numeric array. *R* is numeric. The value of *R* is the product of the mathematical constant $\pi = 3.14159...$ (Pi), and *Y*.

Example

```
      ○0.5 1 2
1.570796327 3.141592654 6.283185307
```

Pick: $R \leftarrow X \supset Y$

Y may be any array. X is a scalar or vector of indices of Y, viz. $\iota \rho Y$. R is an item selected from the structure of Y according to X.

Elements of X select from successively deeper levels in the structure of Y. The items of X are simple integer scalars or vectors which identify a set of indices, one per axis at the particular level of nesting of Y in row-major order. Simple scalar items in Y may be picked by empty vector items in X to any arbitrary depth.

$\square IO$ is an implicit argument of Pick.

Examples

```
      G←('ABC' 1)('DEF' 2)('GHI' 3)('JKL' 4)

      G←2 3ρG,('MNO' 5)('PQR' 6)

      G
  ABC  1   DEF  2   GHI  3
  JKL  4   MNO  5   PQR  6

      ((⊂2 1),1)⊃G
JKL

      (⊂2 1)⊃G
 JKL  4

      ((2 1)1 2)⊃G
K

      (5ρ⊂⍳0)⊃10
10
```

Plus: $R \leftarrow X + Y$

See function Add.

Power: $R \leftarrow X \star Y$

Y must be a numeric array. X must be a numeric array. R is numeric. The value of R is X raised to the power of Y.

If Y is zero, R is defined to be 1.

If X is zero, Y must be non-negative.

If X is negative, and Y can be approximated as a rational number of the form $P \div Q$ where P and Q are relatively prime integers, then:

a)   if Q is even, $X \star Y$ gives a *DOMAIN ERROR*

b)   if Q is odd and P is even, then $X \star Y \leftrightarrow (|X) \star Y$

c)   if Q and P are both odd, then $X \star Y \leftrightarrow -(|X) \star Y$

If X is negative, and Y cannot be approximated as a rational number, then $X \star Y \leftrightarrow -(|X) \star Y$.

Examples

```
      2*2 ¯2
4 0.25

      9 64*0.5
3 8

      ¯27*2 3,(1 2÷3),1.2
729 ¯19683 3 9  52.19591521
```

Ravel: $R \leftarrow , Y$

Y may be any array. R is a vector of the elements of Y taken in row-major order.

Examples

```
      M
1 2 3
4 5 6

      ,M
1 2 3 4 5 6

      A
ABC
DEF
GHI
JKL

      ,A
ABCDEFGHIJKL

      ρ,10
1
```

Ravel with Axes: $R \leftarrow , [K] Y$

Y may be any array.

K is either:

- A simple fractional scalar adjacent to an axis of Y, or
- A simple integer scalar or vector of axes of Y, or
- An empty vector.

Ravel with axis can be used with selective specification.

R depends on the case of K above.

If K is a fraction, the result R is an array of the same shape as Y, but with a new axis of length 1 inserted at the K'th position.

```
ρρR  ↔  1+ρρY
ρR   ↔  (1,ρY)[⍋K,ιρρY]
```

Examples

```
      ,[0.5]'ABC'
ABC
      ρ,[0.5]'ABC'
1 3
      ,[1.5]'ABC'
A
B
C
      ρ,[1.5]'ABC'
3 1


      MAT←3 4ρι12
      ρ,[0.5]MAT
1 3 4
      ρ,[1.5]MAT
3 1 4
      ρ,[2.5]MAT
3 4 1
```

If `K` is an integer scalar or vector of axes of `Y`, then:

`K` must contain contiguous axes of `Y` in ascending order.
`R` contains the elements of `Y` raveled along the indicated axes.

Note that if `K` is a scalar or single element vector, `R ←→ Y`.

```
      ρρR ←→ 1+(ρρY)-ρ,K
```

Examples

```
        M
  1  2  3  4
  5  6  7  8
  9 10 11 12

13 14 15 16
17 18 19 20
21 22 23 24
        ρM
2 3 4
```

```
        ,[1 2]M
  1   2   3   4
  5   6   7   8
  9  10  11  12
 13  14  15  16
 17  18  19  20
 21  22  23  24
        ρ,[1 2]M
 6 4


        ,[2 3]M
  1   2   3   4   5   6   7   8   9  10  11  12
 13  14  15  16  17  18  19  20  21  22  23  24


        ρ,[2 3]M
 2 12
```

If K is an empty vector a new last axis of length 1 is created.

```
        ρR  ↔  (ρY),1
```

Examples

```
        Q1←'January' 'February' 'March'
        DISPLAY Q1
.→----------------------------.
| .→------. .→--------. .→----. |
| |January| |February| |March| |
| '-------' '--------' '-----' |
'∈----------------------------'
        DISPLAY ,[ι0]Q1
.→-----------.
↓ .→------.   |
| |January|   |
| '-------'   |
| .→--------. |
| |February| |
| '--------' |
| .→----.    |
| |March|    |
| '-----'    |
'∈-----------'
```

Reciprocal: $R \leftarrow \div Y$

Y must be a numeric array. *R* is numeric. *R* is the reciprocal of *Y*; that is $1 \div Y$. If
$\Box DIV = 0$, $\div 0$ results in a *DOMAIN ERROR*. If $\Box DIV = 1$, $\div 0$ returns 0.

$\Box DIV$ is an implicit argument of Reciprocal.

Examples

```
      ÷4 2 5
0.25 0.5 0.2

      □DIV ← 1

      ÷0 0.5
0 2
```

Replicate: $R \leftarrow X / [K] Y$

Y may be any array. *X* is a simple integer vector or scalar.

The axis specification is optional. If present, *K* must be a simple integer scalar or unit
vector. The value of *K* must be an axis of *Y*. If absent, the last axis of *Y* is implied.
The form $R \leftarrow X \neq Y$ implies the first axis of *Y*. If *Y* is a scalar, it is treated as a one-
element vector.

The length of *X* must be the length of the *K*th (or implied) axis of *Y*. However, if *X* is a
scalar or one-element vector, it will be extended to the length of the *K*th axis.

R is composed from sub-arrays along the *K*th axis of *Y*. If *X[I]* (an element of *X*) is
positive, then the corresponding sub-array is replicated *X[I]* times. If *X[I]* is zero,
then the corresponding sub-array of *Y* is excluded. If *X[I]* is negative, then the fill
element of *Y* ($\subset \in \supset Y$) is replicated $|X[I]$ times. Each of the (replicated) sub-arrays
and fill items are joined along the *K*th axis in the order of occurrence. The shape of *R*
is the shape of *Y* except that the length of the (implied) *K*th axis is $+ / |X$ (after possible
extension).

This function is sometimes called Compress when *X* is boolean.

Examples

```
      1 0 1 0 1/ι5
1 3 5

      1 ¯2 3 ¯4 5/ι5
1 0 0 3 3 3 0 0 0 0 5 5 5 5 5

      M
1 2 3
4 5 6

      2 0 1/M
1 1 3
4 4 6

      0 1≠M
4 5 6

      0 1/[1]M
4 5 6
```

Replicate First: $R \leftarrow X \neq [K] Y$

The form $R \leftarrow X \neq Y$ implies replication along the first axis of Y. See Replicate above.

Reshape: $R \leftarrow X \rho Y$

Y may be any array. X must be a simple scalar or vector of non-negative integers. R is an array of shape X whose elements are taken from Y in row-major sequence and repeated cyclically if required. If Y is empty, R is composed of fill elements of Y ($\subset \epsilon \supset Y$). If X contains at least one zero, then R is empty. If X is an empty vector, then R is scalar.

Examples

```
      2 3ρι8
1 2 3
4 5 6

      2 3ρι4
1 2 3
4 1 2
```

```
      2 3⍴⍳0
0 0 0
0 0 0
```

Residue: $R \leftarrow X \mid Y$

Y may be any numeric array. X may be any numeric array.

For positive arguments, R is the remainder when Y is divided by X. If $X = 0$, R is Y. For other argument values, R is $Y - N \times X$ where N is some integer such that R lies between 0 and X, but is not equal to X.

$\Box CT$ is an implicit argument of Residue.

Examples

```
      3 3 ¯3 ¯3|¯5 5 ¯4 4
1 2 ¯1 ¯2

      0.5|3.12 ¯1 ¯0.6
0.12 0 0.4

      ¯1 0 1|¯5.25 0 2.41
¯0.25 0 0.41
```

Note that the ASCII pipe (¦) may also be interpreted as Residue (|).

Reverse: $R \leftarrow \phi[K]Y$

Y may be any array. The axis specification is optional. If present, K must be an integer scalar or one-element vector. The value of K must be an axis of Y. If absent, the last axis is implied. The form $R \leftarrow \ominus Y$ implies the first axis.

R is the array Y rotated about the Kth or implied axis.

Examples

```
      φ1 2 3 4 5
5 4 3 2 1

      M
1 2 3
4 5 6
      φM
3 2 1
6 5 4
```

```
        ⊖M
4  5  6
1  2  3
        ⌽[1]M
4  5  6
1  2  3
```

Reverse First: $R \leftarrow \ominus[K]Y$

The form $R \leftarrow \ominus Y$ implies reversal along the first axis. See Reverse above.

Roll: $R \leftarrow ?Y$

Y may be any positive integer array. R is an integer, pseudo-randomly selected from the integers ιY with each integer in this population having an equal chance of being selected.

$\Box IO$ and $\Box RL$ are implicit arguments of Roll. A side effect of Roll is to change the value of $\Box RL$.

Examples

```
    ?9  9  9
2  7  5
```

Rotate: $R \leftarrow X\phi[K]Y$

Y may be any array. X must be a simple integer array. The axis specification is optional. If present, K must be a simple integer scalar or one-element vector. The value of K must be an axis of Y. If absent, the last axis of Y is implied. The form $R \leftarrow X\ominus Y$ implies the first axis.

If Y is a scalar, it is treated as a one-element vector. X must have the same shape as the rank of Y excluding the Kth dimension. If X is a scalar or one-element vector, it will be extended to conform. If Y is a vector, then X may be a scalar or a one-element vector.

R is an array with the same shape as Y, with the elements of each of the vectors along the Kth axis of Y rotated by the value of the corresponding element of X. If the value is positive, the rotation is in the sense of right to left. If the value is negative, the rotation is in the sense of left to right.

Examples

```
      3 ⌽ 1 2 3 4 5 6 7
4 5 6 7 1 2 3
      ¯2 ⌽ 1 2 3 4 5
4 5 1 2 3

      M
 1  2  3  4
 5  6  7  8

 9 10 11 12
13 14 15 16

      I
0 1 ¯1 0
0 3 2 1
      I⌽[2]M
 1  6  7  4
 5  2  3  8

 9 14 11 16
13 10 15 12

      J
2 ¯3
3 ¯2
      J⌽M
 3  4  1  2
 6  7  8  5

12  9 10 11
15 16 13 14
```

Rotate First: $R \leftarrow X \ominus [K] Y$

The form $R \leftarrow X \ominus Y$ implies rotation along the first axis. See Rotate above.

Shape: $R \leftarrow \rho Y$

Y may be any array. R is a non-negative integer vector whose elements are the dimensions of Y. If Y is a scalar, then R is an empty vector. The rank of Y is given by $\rho\rho Y$.

Examples

```
      ρ10

      ρ'CAT'
3

      ρ3 4ρι12
3 4

      +G←(2 3ρι6)('CAT' 'MOUSE' 'FLEA')
 1 2 3   CAT  MOUSE  FLEA
 4 5 6

      ρG
2

      ρρG
1

      ρ¨G
 2 3   3

      ρ····G
     3   5   4
```

Signum: $R \leftarrow \times Y$

Y may be any numeric array. R is an integer array whose value indicates whether the value of Y is negative (¯1), zero (0) or positive (1).

Example

```
      ×¯15.3 0 101
¯1 0 1
```

Split: $R \leftarrow \downarrow [K] Y$

Y may be any array. The axis specification is optional. If present, *K* must be a simple integer scalar or one-element vector. The value of *K* must be an axis of *Y*. If absent, the last axis is implied.

The items of *R* are the sub-arrays of *Y* along the *K*th axis. *R* is a scalar if *Y* is a scalar. Otherwise *R* is an array whose rank is $^{-}1 + \rho\rho Y$ and whose shape is $(K \neq \iota\rho Y) / \rho Y$.

Examples

```
      ↓3 4ρ'MINDTHATSTEP'
 MIND  THAT  STEP

      ↓2 5ρι10
 1 2 3 4 5  6 7 8 9 10

      ↓[1]2 5ρι10
 1 6  2 7  3 8  4 9  5 10
```

Subtract: $R \leftarrow X - Y$

Y may be any numeric array. *X* may be any numeric array. *R* is numeric. The value of *R* is the difference between *X* and *Y*.

This function is also known as Minus.

Example

```
      3 ¯2 4 0 - 2 1 ¯2 4
 1 ¯3 6 ¯4
```

Take: $R \leftarrow X \uparrow Y$

Y may be any array. X must be a simple integer scalar or vector.

If Y is a scalar, it is treated as a one-element array of shape $(\rho,X)\rho1$. The length of X must be the same as the rank of Y. However, if X is a scalar or a one-element vector, it will be extended to conform.

R is an array of the same rank as Y (after possible extension), and of shape $|X$. If $X[I]$ (an element of X) is positive, then $X[I]$ sub-arrays are taken from the beginning of the Ith axis of Y. If $X[I]$ is negative, then $X[I]$ sub-arrays are taken from the end of the Ith axis of Y.

If more elements are taken than exist on axis I, then the extra positions in R are filled with the fill element of Y $(\subset\in\supset Y)$.

Examples

```
      5↑'ABCDEF'
ABCDE

      5↑1 2 3
1 2 3 0 0

      ¯5↑1 2 3
0 0 1 2 3

      5↑(ι3) (ι4) (ι5)
 1 2 3  1 2 3 4  1 2 3 4 5  0 0 0  0 0 0

      M
1 2 3 4
5 6 7 8

      2 3↑M
1 2 3
5 6 7

      ¯1 ¯2↑M
7 8
```

Take with Axes: $R \leftarrow X \uparrow [K] Y$

Y may be any non scalar array. X must be a simple integer scalar or vector. K is a vector of zero or more axes of Y.

R is an array of the first or last elements of Y taken along the axes K depending on whether the corresponding element of X is positive or negative respectively.

The rank of R is the same as the rank of Y:

```
ρρR  ↔  ρρY
```

The size of each axis of R is determined by the corresponding element of X:

```
(ρR)[,K]  ↔  |,X
```

Examples

```
        ⎕←M←2 3 4ρι24
 1  2  3  4
 5  6  7  8
 9 10 11 12

13 14 15 16
17 18 19 20
21 22 23 24

        2↑[2]M
 1  2  3  4
 5  6  7  8

13 14 15 16
17 18 19 20

        2↑[3]M
 1  2
 5  6
 9 10

13 14
17 18
21 22
```

```
      2 ¯2↑[3 2]M
 5  6
 9 10

17 18
21 22
```

Times: $R \leftarrow X \times Y$

See function Multiply.

Transpose (Monadic): $R \leftarrow \lozenge Y$

Y may be any array. R is an array of shape $\phi \rho Y$, similar to Y with the order of the axes reversed.

Examples

```
        M
1 2 3
4 5 6

        ⍉M
1 4
2 5
3 6
```

Transpose (Dyadic): $R \leftarrow X \lozenge Y$

Y may be any array. X must be a simple scalar or vector whose elements are included in the set $\iota \rho \rho Y$. Integer values in X may be repeated but all integers in the set $\iota \lceil / X$ must be included. Also the number of elements in X must not exceed the rank of Y.

R is an array formed by the transposition of the axes of Y as specified by X. The Ith element of X gives the new position for the Ith axis of Y. If X repositions two or more axes of Y to the same axis, the elements used to fill this axis are those whose indices on the relevant axes of Y are equal.

$\square IO$ is an implicit argument of Dyadic Transpose.

Examples

```
      A
 1  2  3  4
 5  6  7  8
 9 10 11 12

13 14 15 16
17 18 19 20
21 22 23 24

      2 1 3⍉A
 1  2  3  4
13 14 15 16

 5  6  7  8
17 18 19 20

 9 10 11 12
21 22 23 24

      1 1 1⍉A
1 18

      1 1 2⍉A
 1  2  3  4
17 18 19 20
```

Type: ($\Box ML < 1$) $R \leftarrow \epsilon Y$

Migration level must be such that $\Box ML < 1$ (otherwise ϵ means Enlist).

Y may be any array. R is an array with the same shape and structure as Y in which a numeric value is replaced by 0 and a character value is replaced by ' '.

Examples

```
      ∊(2 3⍴⍳6)(1 4⍴'TEXT')
 0 0 0
 0 0 0

      ' '=∊'X'
1
```

Union: $R \leftarrow X \cup Y$

Y must be a vector. *X* must be a vector. If either argument is a scalar, it is treated as a one-element vector. *R* is a vector of the elements of *X* catenated with the elements of *Y* which are not found in *X*.

Items in *X* and *Y* are considered the same if *X* ≡ *Y* returns 1 for those items.

⎕*CT* is an implicit argument of Union.

Examples

```
      'WASH' ∪ 'SHOUT'
WASHOUT

      'ONE' 'TWO' ∪ 'TWO' 'THREE'
 ONE   TWO   THREE
```

For performance information, see *Search Functions and Hash Tables* in Chapter 2.

Unique: $R \leftarrow \cup Y$

Y must be a vector. *R* is a vector of the elements of *Y* omitting non-unique elements after the first.

⎕*CT* is an implicit argument of Unique.

Examples

```
      ∪ 'CAT' 'DOG' 'CAT' 'MOUSE' 'DOG' 'FOX'
 CAT   DOG   MOUSE   FOX

      ∪ 22 10 22 22 21 10 5 10
22 10 21 5
```


Without: $R \leftarrow X \sim Y$

See function Excluding.

Zilde: $R \leftarrow \theta$

The empty vector ($\iota 0$) may be represented by the numeric constant θ called ZILDE.

C H A P T E R   4

Primitive Operators

Operator Syntax

Operators take one or two operands. An operator with one operand is monadic. The operand of a monadic operator is to the left of the operator. An operator with two operands is dyadic. Both operands are required for a dyadic operator.

Operators have long scope to the left. That is, the left operand is the longest function or array expression to its left (see *Chapter 1*). A dyadic operator has short scope on the right. Right scope may be extended by the use of parentheses.

An operand may be an array, a primitive function, a system function, a defined function or a derived function. An array may be the result of an array expression.

An operator with its operand(s) forms a DERIVED FUNCTION. The derived function may be monadic or dyadic and it may or may not return an explicit result.

Examples

```
      +/ι5
15
      (*∘2)ι3
1 4 9

      PLUS ← + ◇ TIMES ← ×
      1 PLUS.TIMES 2
2

      ⎕NL 2
A
X
      ⎕EX¨↓⎕NL 2
      ⎕NL 2
```

Axis Specification

Some operators may include an axis specification. Axis is itself an operator. However the effect of axis is described for each operator where its specification is permitted. `⎕IO` is an implicit argument of the function derived from the Axis operator.

The description for each operator follows in alphabetical sequence. The valence of the derived function is specifically identified to the right of the heading block.

Class of Operator	Name	Producing Monadic derived function	Producing Dyadic derived function

Monadic	*Assignment* *Assignment* *Assignment* *Commute* *Each* *Reduction* *Scan* *Spawn*	`f¨Y` `f/Y  [ ]` `f⌿Y  [ ]` `f\Y  [ ]` `f⍀Y  [ ]` `f&Y`	`Xf←Y` `X[I]f←Y` `(EXP X)f←Y` `Xf⍨Y` `Xf¨Y`
			
			 `Xf&Y`

Dyadic	*Axis* *Composition*
	
	 Inner Product *Outer Product*
		`f[B]Y` `f∘gY` `A∘gY` `(f∘B)Y`	`Xf[B]Y` `Xf∘gY`
			
			 `Xf.gY`
			
			 `X∘.gY`

`[ ]` Indicates optional axis specification			

Figure 4(i) : Primitive Operators

Operator Presentation

Monadic and Dyadic primitive operators are presented in alphabetical order of their descriptive names as shown in Figure 4(i).

The valence of the operator and the derived function are implied by the syntax in the heading block.

Assignment (Modified):   `{R}←X f←Y`

f may be any dyadic function which returns an explicit result. Y may be any array whose items are appropriate to function f. X must be the **name** of an existing array whose items are appropriate to function f.

R is the "pass-through" value, that is, the value of Y. If the result of the derived function is not assigned or used, there is no explicit result.

The effect of the derived function is to reset the value of the array named by X to the result of $X f Y$.

Examples

```
      A
1  2  3  4  5

      A+←10

      A
11  12  13  14  15

      []←A×←2
2
      A
22  24  26  28  30

      vec←¯4+9?9  ◇  vec
3  5  1  ¯1  ¯2  4  0  ¯3  2
      vec/̈←vec>0  ◇vec
3  5  1  4  2
```

Assignment (Indexed Modified): $\{R\}\leftarrow X[I]f\leftarrow Y$

f may be any dyadic function which returns an explicit result. Y may be any array whose items are appropriate to function f. X must be the **name** of an existing array. I must be a valid index specification. The items of the indexed portion of X must be appropriate to function f.

R is the "pass-through" value, that is, the value of Y. If the result of the derived function is not assigned or used, there is no explicit result.

The effect of the derived function is to reset the indexed elements of X, that is $X[I]$, to the result of $X[I]fY$. This result must have the same shape as $X[I]$.

Examples

```
      A
1 2 3 4 5

      +A[2 4]+←1
1

      A
1 3 3 5 5

      A[3]÷←2

      A
1 3 1.5 5 5
```

If an index is repeated, function f will be applied to the successive values of the indexed elements of X, taking the index occurrences in left-to-right order.

Example

```
      B←5ρ0

      B[2 4 1 2 1 4 2 4 1 3]+←1

      B
3 3 1 3 0
```

Assignment (Selective Modified): $\{R\}\leftarrow(EXP\ X)f\leftarrow Y$

f may be any dyadic function which returns an explicit result. Y may be any array whose items are appropriate to function f. X must be the ***name*** of an existing array. EXP is an expression that **selects** elements of X. (See Selective Assignment in *Chapter 3* for a list of allowed selection functions.) The selected elements of X must be appropriate to function f.

R is the "pass-through" value, that is, the value of Y. If the result of the derived function is not assigned or used, there is no explicit result.

The effect of the derived function is to reset the selected elements of X to the result of $X[I]fY$ where $X[I]$ defines the elements of X selected by EXP.

Example

```
      A
12 36 23 78 30

      ((A>30)/A) ×← 100
      A
12 3600 23 7800 30
```

Axis (with Monadic Operand): $R \leftarrow f[B]Y$

f must be a monadic primitive mixed function taken from those shown in Figure 4(ii) below, or a function derived from the operators Reduction (/) or Scan (\). B must be a numeric scalar or vector. Y may be any array whose items are appropriate to function f. Axis does not follow the normal syntax of an operator.

Function	Name	Range of B
⌽ or ⊖	*Reverse*	$B \in \iota \rho \rho Y$
↑	*Mix*	$(0 \neq 1|B) \wedge (B > \square IO-1) \wedge (B < \square IO + \rho \rho Y)$
↓	*Split*	$B \in \iota \rho \rho Y$
,	*Ravel*	*fraction, or zero or more axes of Y*
⊂	*Enclose*	$(B \equiv \iota 0) \vee (\wedge / B \in \iota \rho \rho Y)$

Figure 4(ii) : Primitive monadic mixed functions with optional axis.

In most cases, B is required to be an integer which identifies a specific axis of Y. An exception occurs when f is the Mix function (↑) in which case B is a fractional value whose lower and upper integer bounds select an adjacent pair of axes of Y or an extreme axis of Y. For Ravel (,) and Enclose (⊂) , B can be a **vector** of two or more axes.

$\square IO$ is an implicit argument of the derived function which determines the meaning of B.

Examples

```
      ⌽[1]2 3ρι6
4 5 6
1 2 3

      ↑[.1]'ONE' 'TWO'
OT
NW
EO
```

Axis (with Dyadic Operand): $R \leftarrow X f [B] Y$

f must be a dyadic primitive scalar function, or a dyadic primitive mixed function taken from Figure 4(iii) below. B must be a numeric scalar or vector. X and Y may be any arrays whose items are appropriate to function f. Axis does not follow the normal syntax of an operator.

Function	Name	Range of B
/ or ⌿	*Replicate*	$B \in \iota \rho \rho Y$
\ or ⍀	*Expand*	$B \in \iota \rho \rho Y$
⊂	*Partitioned Enclose*	$B \in \iota \rho \rho Y$
⌽ or ⊖	*Rotate*	$B \in \iota \rho \rho Y$
, or ⍪	*Catenate/ Laminate*	$(0 \neq 1 | B) \wedge (B > \Box IO - 1) \wedge (B < \Box IO + (\rho \rho X) \lceil \rho \rho Y)$
↑	*Take*	*zero or more axes of Y*
↓	*Drop*	*zero or more axes of Y*

Figure 4(iii) : Primitive dyadic mixed functions with optional axis.

In most cases, B must be an integer value identifying the axis of X and Y along which function f is to be applied. Exceptionally, B must be a fractional value for the Laminate function (,) whose upper and lower integer bounds identify a pair of axes or an extreme axis of X and Y. For Take (↑) and Drop (↓), B can be a **vector** of two or more axes.

$\Box IO$ is an implicit argument of the derived function which determines the meaning of B.

Examples

```
      1 4 5 =[1] 3 2ρι6
1 0
0 1
1 0

      2 ¯2 1/[2]2 3ρ'ABCDEF'
AA  C
DD  F

      'ABC',[1.1]'='
A=
B=
C=

      'ABC',[0.1]'='
ABC
===

      ⎕IO←0

      'ABC',[¯0.5]'='
ABC
===
```

Axis with Scalar Dyadic Functions

The axis operator `[X]` can take a scalar dyadic function as operand. This has the effect of 'stretching' a lower rank array to fit a higher rank one. The arguments must be conformable along the specified axis (or axes) with elements of the lower rank array being replicated along the other axes.

For example, if `H` is the higher rank array; `L` the lower rank one; `X` is an axis specification; and f a scalar dyadic function, then the expression: `Hf[X]L` is conformable if: `(ρL)↔(ρH)[X]`. Each element of L is replicated along the remaining `(ρH)~X` axes of `H`.

Examples

```
      mat
10 20 30
40 50 60

      mat+[1]1 2       ⍝ add along first axis
11 21 31
42 52 62

      mat+[2]1 2 3     ⍝ add along last axis
11 22 33
41 52 63

      cube
 100  200  300
 400  500  600

 700  800  900
1000 1100 1200

      cube+[1]1 2
 101  201  301
 401  501  601

 702  802  902
1002 1102 1202

      cube+[3]1 2 3
 101  202  303
 401  502  603

 701  802  903
1001 1102 1203

      cube+[2 3]mat
 110  220  330
 440  550  660

 710  820  930
1040 1150 1260

      cube+[1 3]mat
 110  220  330
 410  520  630
```

```
    740   850   960
   1040  1150  1260
```

Commute: $\{R\}\leftarrow Xf\ddot{\sim}Y$

f may be any dyadic function. *X* and *Y* may be any arrays whose items are appropriate to function *f*.

The derived function is equivalent to *Yf X*. The derived function need not return a result.

Example

```
      N
3 2 5 4 6 1 3

      N/⍨2|N
3 5 1 3
```

The following statements are equivalent:

```
F/⍨←I
F←F/⍨I
F←I/F
```

Commute often eliminates the need for parentheses. The pair ρ⍨ may be read as "if" in the following context:

```
→LABELρ⍨A=1
```

Composition (Form I): $\{R\} \leftarrow f \circ g Y$

f may be any monadic function. *g* may be any monadic function which returns a result. *Y* may be any array whose items are appropriate to function *g*. The items of *gY* must be appropriate to function *f*.

The derived function is equivalent to *fgY*. The derived function need not return a result.

Composition allows functions to be "glued" together to build up more complex functions.

Examples

```
      RANK ← ρ∘ρ
      RANK ¨ 'JOANNE' (2 3ρι6)
 1   2

      +/∘ι¨2 4 6
3 10 21


      □VR'SUM'
    ∇ R←SUM X
[1]    R←+/X
    ∇

      SUM∘ι¨2 4 6
3 10 21
```

Composition (Form II): $\{R\}\leftarrow A\circ gY$

g may be any dyadic function. *A* may be any array whose items are appropriate to function g. *Y* may be any array whose items are appropriate to function *g*.

The derived function is equivalent to *AgY*. The derived function need not return a result.

Examples

```
      2 2ρ ¨ 'AB'
 AA   BB
 AA   BB


      SINE ← 1∘○


      SINE 10 20 30
¯0.5440211109 0.9129452507 ¯0.9880316241
```

The following example uses Composition Forms I and II to list functions in the workspace:

```
      ⎕NL 3
ADD
PLUS


      ⎕∘←∘⎕VR¨↓⎕NL 3
      ∇ ADD X
[1]    →LABρ¨0≠⎕NC'SUM' ◇ SUM←0
[2]   LAB:SUM←SUM++/X
      ∇
      ∇ R←A PLUS B
[1]    R←A+B
      ∇
```

Composition (Form III): $\{R\}\leftarrow(f\circ B)Y$

f may be any dyadic function. *B* may be any array whose items are appropriate to function *f*. *Y* may be any array whose items are appropriate to function *f*.

The derived function is equivalent to *Y f B*. The derived function need not return a result.

Examples

```
      (*∘0.5)4 16 25
2 4 5

      SQRT ← *∘.5

      SQRT 4 16 25
2 4 5
```

The parentheses are required in order to distinguish between the operand *B* and the argument *Y*.

Composition (Form IV): $\{R\}\leftarrow Xf\circ gY$

f may be any dyadic function. *g* may be any monadic function which returns a result. *Y* may be any array whose items are appropriate to function *g*. Also *gY* must return a result whose items are appropriate as the right argument of function *f*. *X* may be any array whose items are appropriate to function *f*.

The derived function is equivalent to *X f g Y*. The derived function need not return a result.

Examples

```
      +∘÷/40ρ1        ⍝ Golden Ratio! (Bob Smith)
1.618033989

      0,∘ι¨ι5
0 1  0 1 2  0 1 2 3  0 1 2 3 4  0 1 2 3 4 5
```

Each (with Monadic Operand): $\{R\}\leftarrow f\ddot{\ }Y$

f may be any monadic function. Y may be any array, each of whose items are separately appropriate to function f.

The derived function applies function f separately to each item of Y. The derived function need not return a result. If a result is returned, R has the same shape as Y, and its elements are the items produced by the application of function f to the corresponding items of Y.

If Y is empty, the derived function is applied once to the prototype of Y, and the shape of R is the shape of Y.

Examples

```
      G←('TOM' (ι3))('DICK' (ι4))('HARRY' (ι5))
      ρG
3
      ρ¨G
 2  2  2

      ρ¨¨G
  3  3    4  4    5  5

      +⎕FX¨('FOO1' 'A←1')('FOO2' 'A←2')
FOO1 FOO2
```

Each (with Dyadic Operand): $\{R\}\leftarrow X f\ddot{} Y$

f may be any dyadic function. X and Y may be any arrays whose corresponding items (after scalar extension) are appropriate to function f when applied separately.

The derived function is applied separately to each pair of corresponding elements of X and Y. If X or Y is a scalar or single-element array, it will be extended to conform with the other argument. The derived function need not produce an explicit result. If a result is returned, R has the same shape as Y (after possible scalar extension) whose elements are the items produced by the application of the derived function to the corresponding items of X and Y.

If X or Y is empty and scalar conformable, the derived function is applied once to the prototypes of X and Y, and the shape of R is determined by the rules for scalar conformability.

Examples

```
      +G←(1 (2 3))(4 (5 6))(8 9)10
1   2 3   4   5 6   9 8   10

      1⌽¨G
 2 3   1   5 6   4   9 8   10

      1⌽¨¨G
1   3 2   4   6 5   8 9   10

      1⌽¨¨¨G
1   2 3   4   5 6   8 9   10

      1 2 3 4↑¨G
1   4   5 6   8 9 0   10 0 0 0

      'ABC',¨'XYZ'
AX   BY   CZ
```

Inner Product: $R \leftarrow Xf.gY$

f must be a dyadic function. g may be any dyadic function which returns a result. The last axis of X must have the same length as the first axis of Y.

The result of the derived function has shape $(^-1\downarrow\rho X),1\downarrow\rho Y$. Each item of R is the result of $f/xg\ddot{\ }y$ where x and y are typical vectors taken from all the combinations of vectors along the last axis of X and the first axis of Y respectively.

Function f (and the derived function) need not return a result in the exceptional case when $2=^-1\uparrow\rho X$. In all other cases, function f must return a result.

If the result of $xg\ddot{\ }y$ is empty, for any x and y, a $DOMAIN\ ERROR$ will be reported unless function f is a primitive scalar dyadic function with an identity element shown in Figure 4(iv).

Examples

```
      1 2 3+.×10 12 14
76

      1 2 3 PLUS.TIMES 10 12 14
76

      +/1 2 3×10 12 14
76

      NAMES
HENRY
WILLIAM
JAMES
SEBASTIAN

      NAMES∧.='WILLIAM  '
0 1 0 0
```

Outer Product: $\{R\}\leftarrow X\circ.gY$

g may be any dyadic function. The left operand of the operator is the symbol $\circ.$ X and Y may be any arrays whose elements are appropriate to the function g.

Function g is applied to all combinations of the elements of X and Y. If function g returns a result, the shape of R is $(\rho X),\rho Y$. Each element of R is the item returned by function g when applied to the particular combination of elements of X and Y.

Examples

```
      1 2 3∘.×10 20 30 40
10 20 30  40
20 40 60  80
30 60 90 120

      1 2 3∘.ρ'AB'
 A    B
 AA   BB
 AAA  BBB

      1 2∘.,1 2 3
 1 1  1 2  1 3
 2 1  2 2  2 3

      (ι3)∘.=ι3
1 0 0
0 1 0
0 0 1
```

Reduction: $R \leftarrow f / [K] Y$

f must be a dyadic function. Y may be any array whose items in the sub-arrays along the Kth axis are appropriate to function f.

The axis specification is optional. If present, K must identify an axis of Y. If absent, the last axis of Y is implied. The form $R \leftarrow f \neq Y$ implies the first axis of Y.

R is an array formed by applying function f between items of the vectors along the Kth (or implied) axis of Y.

Function	APL	Identity
Plus (Add)	+	0
	−	0
Minus (Subtract)	×	1
	÷	1
Times (Multiply)			0
	⌊	*M*
Divide	⌈	*-M*
Residue	*	1
Minimum	!	1
Maximum	∧	1
Power	∨	0
Binomial	<	0
And	≤	1
Or	=	1
Less	>	0
Less Or Equal	≥	1
Equal	≠	0
Greater	⊤	0
Greater Or Equal	∪	~~0~~
Not Equal	/≠	1
Encode	\⍀	1
Union	φ⊖	0
Replicate		
Expand		
Rotate		

M is the largest number which is representable on the machine.

Figure 4(iv) : Identity Elements

For a typical vector Y, the result is:

$$⊂(1⊃Y)f(2⊃Y)f......f(n⊃Y)$$

The shape of R is the shape of Y excluding the Kth axis. If Y is a scalar then R is a scalar. If the length of the Kth axis is 1, then R is the same as Y. If the length of the Kth axis is 0, then $DOMAIN\ ERROR$ is reported unless function f occurs in Figure 4(iv), in which case its identity element is returned in each element of the result.

Examples

```
      ∨/0 0 1 0 0 1 0
1

      M
1 2 3
4 5 6

      +/M
6 15

      +⌿M
5 7 9

      +/[1]M
5 7 9

      +/(1 2 3)(4 5 6)(7 8 9)
 12 15 18

      ,/'ONE' 'NESS'
 ONENESS

      +/⍳0
0

      ,/''
DOMAIN ERROR
      ,/''
      ∧
```

Reduce First: $R←f⌿Y$

The form $R←f⌿Y$ implies reduction along the first axis of Y. See Reduce above.

Reduce N-Wise: $R \leftarrow Xf/[K]Y$

f must be a dyadic function. *X* must be a simple scalar or one-item integer array. *Y* may be any array whose sub-arrays along the *K*th axis are appropriate to function *f*.

The axis specification is optional. If present, *K* must identify an axis of *Y*. If absent, the last axis of *Y* is implied. The form $R \leftarrow Xf \neq Y$ implies the first axis of *Y*.

R is an array formed by applying function *f* between items of sub-vectors of length *X* taken from vectors along the *K*th (or implied) axis of *Y*.

X can be thought of as the width of a 'window' which moves along vectors drawn from the *K*th axis of *Y*.

If *X* is zero, the result is a $(\rho Y)+(\rho\rho Y)=\iota\rho\rho Y$ array of identity elements for the function *f*. See Figure 4(iv).

If *X* is negative, each sub-vector is reversed before being reduced.

Examples

```
      ι4
1 2 3 4

      3+/ι4          ⍝ (1+2+3) (2+3+4)
6 9
      2+/ι4          ⍝ (1+2) (2+3) (3+4)
3 5 7
      1+/ι4          ⍝ (1) (2) (3) (4)
1 2 3 4

      0+/ι4          ⍝ Identity element for +
0 0 0 0 0
      0×/ι4          ⍝ Identity element for ×
1 1 1 1 1

      2,/ι4          ⍝ (1,2) (2,3) (3,4)
 1 2  2 3  3 4
      ¯2,/ι4         ⍝ (2,1) (3,2) (4,3)
 2 1  3 2  4 3
```

Scan: $R \leftarrow f \setminus [K] Y$

f may be any dyadic function that returns a result. *Y* may be any array whose items in the sub-arrays along the *K*th axis are appropriate to the function *f*.

The axis specification is optional. If present, *K* must identify an axis of *Y*. If absent, the last axis of *Y* is implied. The form $R \leftarrow f \diagdown Y$ implies the first axis of *Y*.

R is an array formed by successive reductions along the *K*th axis of *Y*. If *V* is a typical vector taken from the *K*th axis of *Y*, then the *I*th element of the result is determined as $f / I \uparrow V$.

The shape of *R* is the same as the shape of *Y*. If *Y* is an empty array, then *R* is the same empty array.

Examples

```
      ∨\0 0 1 0 0 1 0
0 0 1 1 1 1 1

      ∧\1 1 1 0 1 1 1
1 1 1 0 0 0 0

      +\1 2 3 4 5
1 3 6 10 15

      +\(1 2 3)(4 5 6)(7 8 9)
 1 2 3  5 7 9  12 15 18
```

```
        M
1  2  3
4  5  6

        +\M
1  3   6
4  9  15

        +⍀M
1  2  3
5  7  9

        +\[1]M
1  2  3
5  7  9

      ,\'ABC'
A  AB   ABC

        T←'ONE(TWO) BOOK(S)'

        ≠\T∊'()'
0 0 0 1 1 1 1 0 0 0 0 0 0 1 1 0

        ((T∊'()')⍱≠\T∊'()')/T
ONE BOOK
```

Scan First: $R←f⍀Y$

The form $R←f⍀Y$ implies scan along the first axis of Y. See Scan above.

Spawn: $\{R\}\leftarrow\{X\}f\&Y$

& is a monadic operator with an ambivalent derived function. & spawns a new thread in which f is applied to its argument Y (monadic case) or between its arguments X and Y (dyadic case). The shy result of this application is the number of the newly created thread.

When function f terminates, its result (if any), the **thread result**, is returned. If the thread number is the subject of an active $\square TSYNC$, the thread result appears as the result of $\square TSYNC$. If no $\square TSYNC$ is in effect, the thread result is displayed in the session in the normal fashion.

Note that & can be used in conjunction with the **each** operator ¨ to launch many threads in parallel.

Examples

```
      ÷&4           ⍝ Reciprocal in background
0.25

      □←÷&4          ⍝ Show thread number
1
0.25

      FOO&88         ⍝ Spawn monadic function.

      2 FOO&3        ⍝ dyadic

      {NIL}&0        ⍝ niladic

      ⍎&'NIL'        ⍝ ..

      X.GOO&99       ⍝ thread in remote space.

      ⍎&'□dl 2'      ⍝ Execute async expression.

      'NS'⍎&'FOO' ⍝ .. remote .. .. ..

      PRT&¨↓□nl 9 ⍝ PRT spaces in parallel.
```

C H A P T E R 5

System Functions & Variables

System Functions, Variables, Constants and Namespaces provide information and services within the APL environment. Their case-insensitive names begin with ⎕.

⎕	⎕	⎕*A*	⎕*A*	⎕*AI*
⎕*AN*	⎕*ARBIN*	⎕*ARBOUT*	⎕*AT*	⎕*AV*
⎕*CLEAR*	⎕*CMD*	⎕*CR*	⎕*CS*	⎕*CT*
⎕*CY*	⎕*D*	⎕*DIV*	⎕*DL*	⎕*DM*
⎕*DQ*	⎕*DR*	⎕*ED*	⎕*EM*	⎕*EN*
⎕*EX*	⎕*EXCEPTION*	⎕*EXPORT*	⎕*FAPPEND*	⎕*FAVAIL*
⎕*FCREATE*	⎕*FDROP*	⎕*FERASE*	⎕*FHOLD*	⎕*FLIB*
⎕*FMT*	⎕*FNAMES*	⎕*FNUMS*	⎕*FRDAC*	⎕*FRDCI*
⎕*FREAD*	⎕*FRENAME*	⎕*FREPLACE*	⎕*FRESIZE*	⎕*FSIZE*
⎕*FSTAC*	⎕*FSTIE*	⎕*FTIE*	⎕*FUNTIE*	⎕*FX*
⎕*IO*	⎕*KL*	⎕*LC*	⎕*LOAD*	⎕*LOCK*
⎕*LX*	⎕*MAP*	⎕*ML*	⎕*MONITOR*	⎕*NA*
⎕*NAPPEND*	⎕*NC*	⎕*NCREATE*	⎕*NERASE*	⎕*NL*
⎕*NLOCK*	⎕*NNAMES*	⎕*NNUMS*	⎕*NQ*	⎕*NR*
⎕*NREAD*	⎕*NRENAME*	⎕*NREPLACE*	⎕*NRESIZE*	⎕*NS*
⎕*NSI*	⎕*NSIZE*	⎕*NTIE*	⎕*NULL*	⎕*NUNTIE*
⎕*NXLATE*	⎕*OFF*	⎕*OR*	⎕*PATH*	⎕*PFKEY*
⎕*PP*	⎕*PW*	⎕*REFS*	⎕*RL*	⎕*RTL*
⎕*SAVE*	⎕*SD*	⎕*SE*	⎕*SH*	⎕*SHADOW*
⎕*SI*	⎕*SIGNAL*	⎕*SIZE*	⎕*SM*	⎕*SR*
⎕*STACK*	⎕*STATE*	⎕*STOP*	⎕*SVC*	⎕*SVO*
⎕*SVQ*	⎕*SVR*	⎕*SVS*	⎕*TC*	⎕*TCNUMS*
⎕*TID*	⎕*TKILL*	⎕*TNUMS*	⎕*TRACE*	⎕*TRAP*
⎕*TS*	⎕*TSYNC*	⎕*USING*	⎕*VFI*	⎕*VR*

⎕WA	⎕WC	⎕WG	⎕WN	⎕WS
⎕WSID	⎕WX	⎕XSI	⎕XT	

System Variables

System variables retain information used by the system in some way, usually as implicit arguments to functions.

The characteristics of an array assigned to a system variable must be appropriate, otherwise an error will be reported immediately.

Example

```
      ⎕IO←3
DOMAIN ERROR
      ⎕IO←3
      ^
```

System variables may be localised by inclusion in the header line of a defined function or in the argument list of the system function `⎕SHADOW`. When a system variable is localised, it retains its previous value until it is assigned a new one. This feature is known as "pass-through localisation". The exception to this rule is `⎕TRAP`.

A system variable can never be undefined. Default values are assigned to all system variables in a clear workspace.

Unless otherwise stated, system variables are associated with namespaces.

System variables are summarised in the following Figure.

Name	Description	Scope
⎕	Character Input/Output	Session
⎕	Evaluated Input/Output	Session
⎕CT	Comparison Tolerance	Namespace
⎕DIV	Division Method	Namespace
⎕IO	Index Origin	Namespace
⎕LX	Latent Expression	Workspace
⎕ML	Migration Level	Namespace
⎕PATH	Search Path	Session
⎕PP	Print Precision	Namespace
⎕PW	Print Width	Session
⎕RL	Random Link	Namespace
⎕RTL	Response Time Limit	Namespace
⎕SM	Screen Map	Workspace
⎕TRAP	Event Trap	Workspace
⎕USING	Microsoft .Net Search Path	Namespace
⎕WSID	Workspace Identification	Workspace
⎕WX	Window Expose	Namespace

In other words, ⎕, ⍞, ⎕*SE*, ⎕*PATH* and ⎕*PW* relate to the session. ⎕*LX*, ⎕*SM*, ⎕*TRAP* and ⎕*WSID* relate to the active workspace. All the other system variables relate to the current namespace.

Session	Workspace	Namespace
⍞	⎕*LX*	⎕*CT*
⎕	⎕*SM*	⎕*DIV*
⎕*PATH*	⎕*TRAP*	⎕*IO*
⎕*PW*	⎕*WSID*	⎕*ML*
⎕*SE*		⎕*PP*
		⎕*RL*
		⎕*RTL*
		⎕*USING*
		⎕*WX*

System Namespaces

⎕*SE* is currently the only system namespace.

System Constants

System constants, which can be regarded as niladic system functions, return information from the system. They have distinguished names, beginning with the quad symbol, ⎕. A system constant may **not** be assigned a value. System constants may not be localised or erased.

System constants are summarised in the following Figure:

Name	Description
⎕*A̲*	Underscored Alphabetic upper case characters
⎕*A*	Alphabetic upper case characters
⎕*AI*	Account Information
⎕*AN*	Account Name
⎕*AV*	Atomic Vector
⎕*D*	Digits
⎕*EN*	Event Number
⎕*EXCEPTION*	Reports the most recent Microsoft .net Exception
⎕*LC*	Line Count
⎕*NULL*	Null Item
⎕*SD*	Screen (or window) Dimensions
⎕*TC*	Terminal Control (backspace, linefeed, newline)
⎕*TS*	Time Stamp
⎕*WA*	Workspace Available

System Functions

System functions provide various services related to both the APL and the external environment. System functions have distinguished names beginning with the □ symbol. They are implicitly available in a clear workspace.

The following Figure identifies system functions divided into relevant categories. Each function is described in alphabetical order in this chapter

System Commands

These functions closely emulate system commands (see *Chapter 6*)

Name	Description
□*CLEAR*	Clear workspace (WS)
□*CY*	Copy objects into active WS
□*EX*	Expunge objects
□*LOAD*	Load a saved WS
□*NL*	Name List
□*OFF*	End the session
□*SAVE*	Save the active WS

External Environment

These functions provide access to the the external environment, such as file systems, Operating System facilities, and input/output devices.

Name	Description
□*ARBIN*	Arbitrary Input
□*ARBOUT*	Arbitrary Output
□*CMD*	Execute a (DOS) Command or another program
□*CMD*	Start a (DOS) AP
□*MAP*	Map a file
□*NA*	Declare a DLL function
□*SH*	Execute a (UNIX) command or another program
□*SH*	Start a (UNIX) AP

Defined Functions and Operators

These functions provide services related to defined functions and operators.

Name	Description
`□AT`	Object Attributes
`□CR`	Canonical Representation
`□CS`	Change Space
`□ED`	Edit one or more objects
`□EXPORT`	Export objects
`□FX`	Fix definition
`□LOCK`	Lock a function
`□MONITOR`	Monitor set
`□MONITOR`	Monitor query
`□NR`	Nested Representation
`□NS`	Create Namespace
`□OR`	Object Representation
`□PATH`	Search Path
`□REFS`	Local References
`□SHADOW`	Shadow names
`□STOP`	Set Stop vector
`□STOP`	Query Stop vector
`□TRACE`	Set Trace vector
`□TRACE`	Query Trace vector
`□VR`	Vector Representation

Error Trapping

These functions are associated with event trapping and the system variable `□TRAP`.

Name	Description
`□EM`	Event Messages
`□SIGNAL`	Signal event

Shared Variables

These functions provide the means to communicate between APL tasks and with other applications.

Name	Description
`⎕SVC`	Set access Control
`⎕SVC`	Query access Control
`⎕SVO`	Shared Variable Offer
`⎕SVO`	Query degree of coupling
`⎕SVQ`	Shared Variable Query
`⎕SVR`	Retract offer
`⎕SVS`	Query Shared Variable State

Graphical User Interface

These functions provide access to GUI components.

Name	Description
`⎕DQ`	Await and process events
`⎕NQ`	Place an event on the Queue
`⎕WC`	Create GUI object
`⎕WG`	Get GUI object properties
`⎕WN`	Query GUI object Names
`⎕WS`	Set GUI object properties
`⎕WX`	Expose GUI property names

External Variables

These functions are associated with using external variables.

Name	Description
`⎕XT`	Associate External variable
`⎕XT`	Query External variable
`⎕FHOLD`	External variable Hold

MultiThreading

These functions are associated with using the Spawn operator (&).

Name	Description
□*TID*	Current Thread Identity
□*TCNUMS*	Thread Child Numbers
□*TKILL*	Kill Threads
□*TNUMS*	Thread Numbers
□*TSYNC*	Wait for Threads to Terminate

Component Files

The functions provide the means to store and retrieve data on APL Component Files. See *User Guide* for further details.

Name	Description
□*FAPPEND*	Append a component to File
□*FAVAIL*	File system Availability
□*FCREATE*	Create a File
□*FDROP*	Drop a block of components
□*FERASE*	Erase a File
□*FHOLD*	File Hold
□*FLIB*	List File Library
□*FNAMES*	Names of tied Files
□*FNUMS*	Tie Numbers of tied Files
□*FRDAC*	Read File Access matrix
□*FRDCI*	Read Component Information
□*FREAD*	Read a component from File
□*FRENAME*	Rename a File
□*FREPLACE*	Replace a component on File
□*FRESIZE*	File Resize
□*FSIZE*	File Size
□*FSTAC*	Set File Access matrix
□*FSTIE*	Share-Tie a File
□*FTIE*	Tie a File exclusively
□*FUNTIE*	Untie Files

Native Files

The functions provide the means to store and retrieve data on native files.

Name	Description
`⎕NAPPEND`	Append to File
`⎕NCREATE`	Create a File
`⎕NERASE`	Erase a File
`⎕NLOCK`	Lock a region of a file
`⎕NNAMES`	Names of tied Files
`⎕NNUMS`	Tie Numbers of tied Files
`⎕NREAD`	Read from File
`⎕NRENAME`	Rename a File
`⎕NREPLACE`	Replace data on File
`⎕NRESIZE`	File Resize
`⎕NSIZE`	File Size
`⎕NTIE`	Tie a File exclusively
`⎕NUNTIE`	Untie Files
`⎕NXLATE`	Specify Translation Table

Miscellaneous

These functions provide various miscellaneous services.

Name	Description
`⎕DL`	Delay execution
`⎕DM`	Diagnostic Message
`⎕FMT`	Resolve display
`⎕FMT`	Format array
`⎕KL`	Key Labels
`⎕NC`	Name Classification
`⎕NSI`	Namespace Indicator
`⎕PFKEY`	Programmable Function Keys
`⎕SI`	State Indicator
`⎕SIZE`	Size of objects
`⎕SR`	Screen Read
`⎕STACK`	Report Stack
`⎕STATE`	Return State of an object
`⎕VFI`	Verify and Fix numerics
`⎕XSI`	Extended State Indicator

Programming Reference A-Z

There follows an alphabetical list of system functions and variables which are available in **any** Dyalog APL workspace. Apart from quote-quad (⍞) below, their names all begin with the quad symbol (⎕).

Character Input/Output: ⍞

⍞ is a variable which communicates between the user's terminal and APL. Its behaviour depends on whether it is being assigned or referenced.

When ⍞ is assigned with a vector or a scalar, the array is displayed without the normal ending new-line character. Successive assignments of vectors or scalars to ⍞ without any intervening input or output cause the arrays to be displayed on the same output line.

Example

```
      ⍞←'2+2'  ◇  ⍞←'='  ◇  ⍞←4
2+2=4
```

Output through ⍞ is independent of the print width in ⎕*PW*. The way in which lines exceeding the print width of the terminal is treated is dependent on the characteristics of the terminal. Numeric output is formatted in the same manner as direct output (see "Display of Arrays" in *Chapter 1*).

When ⍞ is assigned with a higher-order array, the output is displayed in the same manner as for direct output except that the print width ⎕*PW* is ignored.

When ⍞ is referenced, terminal input is expected without any specific prompt, and the response is returned as a character vector.

If the ⍞ request was preceded by one or more assignments to ⍞ without any intervening input or output, the last (or only) line of the output characters are returned as part of the response.

Example

```
      mat←↑φ⍞⍞⍞⍞⍞
```

Examples

```
      ⎕←'OPTION : ' ◇ R←⎕
OPTION : INPUT


      R
OPTION : INPUT


      ⍴R
14
```

The output of simple arrays of rank greater than 1 through ⎕ includes a new-line character at the end of each line. Input through ⎕ includes the preceding output through ⎕ since the last new-line character. The result from ⎕, including the prior output, is limited to 256 characters.

A soft interrupt causes an *INPUT INTERRUPT* error if entered while ⎕ is awaiting input, and execution is then suspended (unless the interrupt is trapped):

```
      R←⎕
(Interrupt)
INPUT INTERRUPT
```

A time limit is imposed on input through ⎕ if *⎕RTL* is set to a non-zero value:

```
      ⎕RTL←5 ◇ ⎕←'PASSWORD ? ' ◇ R←⎕
PASSWORD ?
TIMEOUT
      ⎕RTL←5 ◇ ⎕←'PASSWORD : ' ◇ R←⎕
                                    ^
```

The *TIMEOUT* interrupt is a trappable event.

Evaluated Input/Output: ⎕

⎕ is a variable which communicates between the users terminal and APL. Its
behaviour depends on whether it is being assigned or referenced.

When ⎕ is assigned an array, the array is displayed at the terminal in exactly the same
form as is direct output (see "Display of Arrays" in *Chapter 1*).

Example

```
      ⎕←2+ι5
3 4 5 6 7


      ⎕←2 4ρ'WINEMART'
WINE
MART
```

When ⎕ is referenced, a prompt (⎕:) is displayed at the terminal, and input is
requested. The response is evaluated and an array is returned if the result is valid. If
an error occurs in the evaluation, the error is reported as normal (unless trapped by a
⎕*TRAP* definition) and the prompt (⎕:) is again displayed for input. An EOF interrupt
reports *INPUT INTERRUPT* and the prompt (⎕:) is again displayed for input. A soft
interrupt is ignored and a hard interrupt reports *INTERRUPT* and the prompt (⎕:) is
redisplayed for input.

Examples

```
      10×⎕+2
⎕:
      ι3
30 40 50


      2+⎕
⎕:
      X
VALUE ERROR
      X
     ^
⎕:
      2+ι3
5 6 7
```

A system command may be entered. The system command is effected and the prompt
is displayed again (unless the system command changes the environment):

```
      ρ3,⎕
⎕:
      )WSID
WS/MYWORK
⎕:
      )SI
⎕
⎕:
      )CLEAR
CLEAR WS
```

If the response to a ⎕: prompt is an abort statement (→), the execution will be aborted:

```
      1 2 3 = ⎕
⎕:
      →
```

A trap definition on interrupt events set for the system variable ⎕TRAP in the range
1000-1006 has no effect whilst awaiting input in response to a ⎕: prompt.

Example

```
      ⎕TRAP←(11 'C' '''ERROR''')(1000 'C' '''STOP''')

      2+⎕
⎕:
      (Interrupt Signal)
INTERRUPT
⎕:
      'C'+2
ERROR
```

A time limit set in system variable ⎕RTL has no effect whilst awaiting input in
response to a ⎕: prompt.

Underscored Alphabetic Characters: *R←⎕A̲*

This is a simple character vector, composed of the letters of the alphabet with underscores (or the extra National Language characters in the Dyalog Alt Font).

Example

```
      ⎕A̲
A̲B̲C̲D̲E̲F̲G̲H̲I̲J̲K̲L̲M̲N̲O̲P̲Q̲R̲S̲T̲U̲V̲W̲X̲Y̲Z̲
```

(or *ÁÂÃÇÈÊËÌÍÎÏÐÒÓÔÕÙÚÛÝÞāìðòō*)

Alphabetic Characters: *R←⎕A*

This is a simple character vector, composed of the letters of the alphabet.

Example

```
      ⎕A
ABCDEFGHIJKLMNOPQRSTUVWXYZ
```

Account Information: *R←⎕AI*

This is a simple integer vector, whose four elements are:

> *⎕AI[1]* - user identification.
>> Under DOS & Windows, this is the `aplnid` (network ID from configuration dialog box).
>> Under UNIX, this is the UID of the account.

> *⎕AI[2]* - compute time for the APL session in milliseconds.

> *⎕AI[3]* - connect time for the APL session in milliseconds.

> *⎕AI[4]* - keying time for the APL session in milliseconds.

Elements beyond 4 are not defined but reserved.

Example

```
      ⎕AI
52 7396 2924216 2814831
```

Account Name: $R \leftarrow \Box AN$

In Windows 95 and Windows NT this is a simple character vector containing the user name. Under UNIX this is the login name for the account. The length of the vector is the number of characters in the login name.

Under DOS and Windows 3.1 $\Box AN$ returns an empty character vector.

Example

```
        □AN
PROJECT1

        ρ□AN
8
```

Arbitrary Input: $R \leftarrow \{X\}\Box ARBIN\ Y$

This transmits the prompt Y to an output device specified by X prior to reading from an input device specified by X.

Under Windows, the use of $\Box ARBIN$ to the screen or in conjunction with RS232 ports is not supported.

Y may be a scalar, a simple vector, or a vector of simple scalars or vectors. The items of the simple arrays of Y must each be a character, or a number in the range 0 to 255. Numbers are sent to the output device without translation. Characters undergo the standard $\Box AV$ to ASCII translation. If Y is an empty vector, no codes are sent to the output device.

X may take several forms:

terminate (input output)	□ARBIN	prompt
terminate input	□ARBIN	prompt
terminate	□ARBIN	prompt
	□ARBIN	prompt

Each of these elements is discussed separately.

<terminate>

This defines how the read should be terminated.

If it is omitted, the read terminates on receipt of a Newline character.

If supplied, it must be a simple numeric scalar or vector.

- If it is a numeric scalar, it defines the number of characters to be read.

- If it is a numeric vector, it defines a set of terminating characters.

- If it is the null vector, the read terminates on Newline.

<input>

This defines the input device.

If this is omitted, input is taken from standard input (usually the keyboard).

If supplied, it must be a simple numeric scalar or a simple text vector.

- If it is a numeric scalar, it must correspond to a DOS handle or UNIX stream number (see *User Guide*).

- If it is a text vector, it must correspond to a valid device or file name.

You must have permission to read from the chosen device.

<output>

This defines the output device.

If this is omitted, output is sent to standard output (usually the screen).

If supplied, it must be a simple numeric scalar or a simple text vector.

- If it is a numeric scalar, it must correspond to a DOS handle or UNIX stream number (see *User Guide*).

- If it is a text vector, it must correspond to a valid device or file name.

You must have permission to write to the chosen device.

The result R is a simple numeric vector. Each item of R is the numeric representation of an 8-bit code in the range 0 to 255 received from the input device. The meaning of the code is dependent on the characteristics of the input device. If a set of delimiters was defined by <terminate>, the last code returned will belong to that set.

⎕RTL (Response Time Limit) is an implicit argument of ⎕ARBIN. This allows a time limit to be imposed on input. If the time limit is reached, ⎕ARBIN returns with the codes read up to that point.

Examples

Write *HELLO* on the screen, and read a line of input
from the keyboard:

```
R ← ⎕ARBIN 'HELLO'
```

Beep three times, send *ARE YOU AWAKE?* to the screen and wait for a 1 character
answer from the keyboard:

```
R ← 1 ⎕ARBIN 7 7 7 'ARE YOU AWAKE (Y/N)'
```

Read a line from *MYFILE*:

```
R ← '' 'MYFILE' ⎕ARBIN ''
```

Read *MYFILE* until a SPACE (code 32):

```
R ← (,32) 'MYFILE' ⎕ARBIN ''
```

Read *MYFILE* until a SPACE (code 32) or a TAB (code 9):

```
R ← (32 9) 'MYFILE' ⎕ARBIN ''
```

Write *HELLO* on /dev/tty1 (a UNIX terminal screen), then read a line from /dev/tty1 (a
UNIX terminal keyboard):

```
R ← '' ('/dev/tty1' '/dev/tty1') ⎕ARBIN 'HELLO'
```

Write *TITLE* to LPT1 (a DOS printer device), then read from COM1 (a DOS serial
port) up to the first NEWLINE character:

```
R ← (,13) ('COM1' 'LPT1') ⎕ARBIN 'TITLE'
```

Read 100 characters from COM1; timeout after 10 secs:

```
⎕RTL←10
R←100 'COM1' ⎕ARBIN ''
⍎(100≠⍴R)/'''Read timed out'''
```

Read until DELIM from COM1; timeout after 10 secs:

```
⎕RTL←10
R←(,DELIM) 'COM1' ⎕ARBIN ''
⍎(DELIM≠¯1↑R)/'''Read timed out'''
```

Arbitrary Output: *{X}⎕ARBOUT Y*

This transmits *Y* to an output device specified by *X*.

Under Windows, the use of ⎕*ARBOUT* to the screen or to RS232 ports is not supported.

Y may be a scalar, a simple vector, or a vector of simple scalars or vectors. The items of the simple arrays of *Y* must each be a character or a number in the range 0 to 255. Numbers are sent to the output device without translation. Characters undergo the standard ⎕*AV* to ASCII translation. If *Y* is an empty vector, no codes are sent to the output device.

X defines the output device. If *X* is omitted, output is sent to standard output (usually the screen). If *X* is supplied, it must be a simple numeric scalar or a simple text vector.

If it is a numeric scalar, it must correspond to a DOS handle or UNIX stream number (see *User Guide*).

If it is a text vector, it must correspond to a valid device or file name.

You must have permission to write to the chosen device.

Examples

Write ASCII digits '123' to UNIX stream 9:

```
    9 ⎕ARBOUT 49 50 51
```

Write ASCII characters '*ABC*' to *MYFILE*:

```
    'MYFILE' ⎕ARBOUT 'ABC'
```

Beep 3 times:

```
    ⎕ARBOUT 7 7 7
```

Prompt for input:

```
    ⎕← 'Prompt: '◇⎕arbout 12◇ans←⎕
```

Attributes: $R \leftarrow \{X\} \square AT \ Y$

Y can be a simple character scalar, vector or matrix, or a vector of character vectors representing the names of 0 or more defined functions or operators. Used dyadically, this function closely emulates the APL2 implementation. Used monadically, it returns information that is more appropriate for Dyalog APL.

Monadic Use

If X is omitted, R is a 4 column matrix with the same number of rows as functions in Y containing the following attribute information:

$R[;1]$ Each item is a 3-element integer vector representing the function header syntax:

            $Item[1]$ result:
                 0 - No result
                 1 - Explicit result
               ‾1 - Shy result

            $Item[2]$ Function valence:
                 0 - Niladic function
                 1 - Monadic function
                 2 - Dyadic function
               ‾2 - Ambivalent function

            $Item[3]$ Operator valence
                 0 - Not an operator
                 1 - Monadic operator
                 2 - Dyadic operator

The following values correspond to the syntax shown alongside:

```
 0   0   0      ∇ FOO
 1   0   0      ∇ Z←FOO
¯1   0   0      ∇ {Z}←FOO
 0  ¯2   0      ∇ {A} FOO B
¯1   1   2      ∇ {Z}←(F OP G)B
```

R[;2]	Each item is the (□*TS* form) timestamp of the time the function was last fixed.

R[;3]	Each item is an integer reporting the current □*LOCK* state of the function:

                   0 – Not locked
                   1 – Cannot display function
                   2 – Cannot suspend function
                   3 – Cannot display or suspend.

R[;4]	Each item is a character vector - the network ID of the user who last fixed (edited) the function.

Example

```
     ∇ {z}←{l}(fn myop)r
[1]    ...

     ∇ z←foo
[1]    ...

     ∇ z←{larg}util rarg
[1]    ...

     □LOCK'foo'

     util2←util

     DISPLAY □AT 'myop' 'foo' 'util' 'util2'
.→------------------------------------------.
↓ .→------. .→-----------------.   .→---.    |
| |¯1 ¯2 1| |1996 8 2 2 13 56 0|   0 |john|   |
| '~------' '~-----------------'     '----'   |
| .→----.   .→------------.           .⊖.     |
| |1 0 0|   |0 0 0 0 0 0 0|         3 | |     |
| '~----'   '~------------'           '-'     |
| .→-----.  .→-----------------.    .→---.    |
| |1 ¯2 0|  |1996 3 1 14 12 10 0|  0 |pete|   |
| '~-----'  '~-----------------'     '----'   |
| .→-----.  .→------------------.   .→-----.  |
| |1 ¯2 0|  |1998 8 26 16 16 42 0| 0 |graeme| |
| '~-----'  '~------------------'    '------'  |
'∈------------------------------------------'
```

Dyadic Use

The dyadic form of □*AT* emulates APL2. It returns the same rank and shape result containing information that matches the APL2 implementation as closely as possible.

If *Y* specifies as single name, the result *R* is a vector. If *Y* specifies more than one name, *R* is a matrix with one row per name in *Y*. The number of elements (columns) and their meaning depends upon the value of *X* which may be 1, 2, 3 or 4.

If *X* is 1, *R* specifies *valences* and contains 3 elements (or columns) whose meaning is as follows:

1	Explicit result	1 if the object has an explicit result or is a variable; 0 otherwise
2	Function valence	0 if the object is a niladic function or not a function
		1 if the object is a monadic function
		2 if the object is an ambivalent function
3	Operator valence	0 if the object is not an operator
		1 if the object is a monadic operator
		2 if the object is a dyadic operator

If *X* is 2, *R* specifies *fix times* (the time the object was last updated) for functions and operators named in *Y*. The time is reported as 7 integer elements (or columns) whose meaning is as follows. The fix time reported for names in *Y* which are not defined functions or operators is 0.

1	Year
2	Month
3	Day
4	Hour
5	Minute
6	Second
7	Milliseconds (this is always reported as 0)

If *X* is 3, *R* specifies *execution properties* and contains 4 elements (or columns) whose meaning is as follows:

1	Displayable	0 if the object is displayable
		1 if the object is not displayable
2	Suspendable	0 if execution will suspend in the object
		1 if execution will not suspend in the object
3	Weak Interrupt behaviour	0 if the object responds to interrupt
		1 if the object ignores interrupt
4		(always 0)

Note that the execution properties for primitive and system functions are 0 1 1 0.

If *X* is 4, *R* specifies *object size* and contains 2 elements (or columns) which both report the □*SIZE* of the object.

Atomic Vector: $R←⎕AV$

This is a simple character vector of all 256 characters in the Dyalog APL character set (see *Chapter 8* and *User Guide*).

Examples

```
      ⎕AV[48+⍳10]
0123456789
```

```
      5 52⍴12↓⎕av
%'αω_abcdefghijklmnopqrstuvwxyz⎕⎕¯.θ0123456789⎕⊢¥$£¢
∆ABCDEFGHIJKLMNOPQRSTUVWXYZ⎕⎕ý·≡∆ABCDEFGHIJKLMNOPQRS
TUVWXYZ{∊}⍙⎕¨ÀÄÅÆ¨ÉÑÖØÜßàáâäåæçèéêëíîïñ[/≠\⍀<≤=≥>≠∨∧
-+÷×?∊⍴~↑↓⍳○*⌈⌊∇∘(⊂⊃∩∪⊥⊤|;,⍲⍱⍴⍦⍧φ⊖⌹!⍕⍔⍒⍖≡≢óôöø"#⎕&'
⌶⍝⌷⊣⊢⊦⊧⊤|@ùúû^ü`¦⍨:∊¿¡◇←→ₐ)]≡≢ɪ⎕⎕⍣%'αω_abcdefghijk
```

Clear Workspace: $⎕CLEAR$

A clear workspace is activated, having the name $CLEAR\ WS$. The active workspace is lost. All system variables assume their default values. The maximum size of workspace is available.

The contents of the session namespace $⎕SE$ are not affected.

Example

```
      ⎕CLEAR
      ⎕WSID
CLEAR WS
```

Execute (DOS) Command: $R \leftarrow \Box CMD \ Y$

$\Box CMD$ executes a DOS or UNIX shell command or starts another Windows application program. $\Box CMD$ is a synonym of $\Box SH$. Either system function may be used in either environment (Windows or UNIX) with exactly the same effect. $\Box CMD$ is probably more natural for the Windows user. This section describes the behaviour of $\Box CMD$ and $\Box SH$ under Windows. See $\Box SH$ for a discussion of the behaviour of these system functions under UNIX.

The system commands $)CMD$ and $)SH$ provide similar facilities but may only be executed from the APL Session.

Executing a DOS Command

If Y is a simple character vector, $\Box CMD$ invokes the DOS command processor (normally COMMAND.COM) and passes Y to it for execution. R is a vector of character vectors containing the result of the command. Each element in R corresponds to a line of output produced by the command.

The implementation of this form of $\Box CMD$ differs between Dyalog APL Version 7 and Version 8. Specifically, Version 8 takes advantage of Windows 95 and Windows NT facilities that are not available in Windows 3.x. The two different implementations are described below.

Version 7 Implementation

Under Windows 3.x, the execution of a DOS command must be done using a .PIF file. $\Box CMD$ first creates a temporary .PIF file called QCMD1234.PIF in the root directory of your default drive (usually C:). This file is initially a copy of WDYALOG\QUADCMD.PIF which is then dynamically changed to run *your* DOS shell together with the particular command that you have asked to be run; i.e. the right argument you supplied to $\Box CMD$. The DOS shell to be run is obtained from your COMSPEC environment variable. If this is not defined, it defaults to C:\COMMAND.COM.

The temporary .PIF file is then executed in a new (but **invisible**) window. The result of the command (if any) is captured by redirecting output to a temporary file and then reading and deleting the file on completion of the command. The result is then displayed in the APL session. The temporary file is created in the directory specified by your TMP environment variable, or in the root directory of the current drive if TMP is undefined. The file is called POPEN.TMP prefixed by your Task ID and the value of the **aplnid** entry in APL.INI.

Example

```
    Z←⎕CMD'DIR'
    ρZ
8
    ↑Z
```

```
 Volume in drive C has no label
 Directory of C:\DYALOG

.              <DIR>      5-07-89   3.02p
..             <DIR>      5-07-89   3.02p
SALES    DWS    110092 5-07-89   3.29p
EXPENSES DWS    154207 5-07-89   3.29p
```

If the command specified in *Y* already contains the redirection symbol (>) the capture of output through a temporary file is avoided and the result *R* is empty. If the command specified by *Y* issues prompts and expects user input, it is **ESSENTIAL** to explicitly redirect output to the console. If this is done, APL detects the presence of a ">" in the command line, runs COMMAND.COM in a **visible** window, and does not direct output to the temporary file. If you fail to do this your system will appear to hang because there is no mechanism for you to receive or respond to the prompt.

Example

```
    ⎕CMD 'DATE>CON'
```
(DOS window appears)
Current date is Wed 19-07-1995
Enter new date (dd-mm-yy): 20-07-95
(DOS window disappears)

Errors

cannot create temp file
APL was unable to create a temporary .PIF file or could not find a shell to run.

could not execute command
APL has successfully started the .PIF file but has timed out waiting for the DOS window to be created. The time-out is defined by the `qcmd_timeout=` parameter in APL.INI (default 5000 msecs).

not enough memory
not enough memory or corrupt shell
invalid shell

Version 8 Implementation

Under Version 8, the right argument of ⎕*CMD* is simply passed to the appropriate command processor for execution and its output is received using an *unnamed pipe*.

By default, ⎕*CMD* will execute the string `('command.com /c',Y)` under Windows 95 and the string `('cmd.exe /c',Y)` under Windows NT; where *Y* is the argument given to ⎕*CMD*. However, the implementation permits the use of alternative command processors as follows.

Before execution, the argument is prefixed and postfixed with strings defined by the APL parameters CMD_PREFIX and CMD_POSTFIX. The former specifies the name of your command processor and any parameters that it requires. The latter specifies a string which may be required. If CMD_PREFIX is not defined, it defaults to the name defined by the environment variable COMSPEC followed by " /c". If COMSPEC is not defined, it defaults to COMMAND.COM or CMD.EXE as appropriate. If CMD_POSTFIX is not defined, it defaults to an empty vector.

⎕*CMD* treats certain characters as having special meaning as follows:

#	marks the start of a trailing comment,
;	divides the command into sub-commands,
>	if found within the last sub-command, causes ⎕*CMD* to use a visible window.

If you simply wish to open a DOS command window, you may (in Version 8) execute the command as a Windows Program (see below). For example:

```
⎕CMD 'command.com' ''
```

Executing a Windows Program

If Y is a 2-element vector of character vectors, $\square CMD$ starts the executable program named by $Y[1]$ with the initial window parameter specified by $Y[2]$. The shy result is an integer scalar containing the window handle allocated by the window manager.

$Y[1]$ must specify the name or complete pathname of an executable program. If the name alone is specified, Windows will search the following directories :

1. the current directory,
2. the Windows directory,
3. the Windows system directory,
4. the directories specified by the PATH variable,
5. the list of directories mapped in a network.

Note that $Y[1]$ may contain the complete command line, including any suitable parameters for starting the program. If Windows fails to find the executable program, $\square CMD$ will fail and report *FILE ERROR* 2.

$Y[2]$ specifies the window parameter and may be one of the following. If not, a *DOMAIN ERROR* is reported.

'*Normal*'	Application is started in a normal window, which is given the input focus.
''	Same as *Normal*
'*Unfocused*'	Application is started in a normal window, which is NOT given the input focus.
'*Hidden*'	Application is run in an invisible window.
'*Minimized*'	Application is started as an icon which is NOT given the input focus.
'*Minimised*'	Same as *Minimized*.
'*Maximized*'	Application is started maximized (full screen) and is given the input focus.
'*Maximised*'	Same as *Maximized*.

An application started by `⎕CMD` may ONLY be terminated by itself or by the user. There is no way to close it from APL. Furthermore, if the window parameter is HIDDEN, the user is unaware of the application (unless it makes itself visible) and has no means to close it.

Examples

```
      ⎕←⎕CMD 'C:\EXCEL\EXCEL' 'Normal'
2390
      ⎕CMD 'winword /mMyMacro' 'Minimized'
```

Start (DOS) Auxiliary Processor: *X* `⎕CMD` *Y*

Used dyadically, `⎕CMD` starts an Auxiliary Processor. The effect, as far as the APL workspace is concerned, is identical under both DOS and UNIX, although the method of implementation differs. `⎕CMD` is a synonym of `⎕SH`. Either function may be used in either environment (DOS or UNIX) with exactly the same effect. `⎕CMD` is probably more natural for the DOS user. This section describes the behaviour of `⎕CMD` and `⎕SH` under DOS. See `⎕SH` for a discussion of the behaviour of these system functions under UNIX.

X must be a simple character vector containing the name (or pathname) of a Dyalog APL Auxiliary Processor (AP). See *User Guide* for details of how to write an AP.

Y may be a simple character scalar or vector, or a vector of character vectors. Under DOS the contents of *Y* are ignored.

`⎕CMD` loads the Auxiliary Processor into memory. If no other APs are currently running, `⎕CMD` also allocates an area of memory for communication between APL and its APs.

The effect of starting an AP is that one or more **external functions** are defined in the workspace. These appear as locked functions and may be used in exactly the same way as regular defined functions.

When an external function is used in an expression, the argument(s) (if any) are passed to the AP for processing via the communications area described above. APL halts whilst the AP is processing, and waits for a result. Under DOS, unlike under UNIX, it is not possible for external functions to run in parallel with APL.

Canonical Representation: $R←\Box CR\ \ Y$

Y must be a simple character scalar or vector which represents the name of a defined function or operator.

If Y is a name of a defined function or operator, R is a simple character matrix. The first row of R is the function or operator header. Subsequent rows are lines of the function or operator. R contains no unnecessary blanks, except for leading indentation of control structures, trailing blanks that pad each row, and the blanks in comments. If Y is the name of a variable, a locked function or operator, an external function, or is undefined, R is an empty matrix whose shape is 0 0.

Example

```
      ∇R←MEAN X      ⍝ Arithmetic mean
[1]   R←(+/X)÷ρX
[2]   ∇
      +F←⎕CR'MEAN'
R←MEAN X      ⍝ Arithmetic mean
R←(+/X)÷ρX

      ρF
2 30
```

The definition of $\Box CR$ has been extended to names assigned to functions by specification (←), and to local names of functions used as operands to defined operators.

If Y is a name assigned to a primitive function, R is a one-element vector containing the corresponding function symbol. If Y is a name assigned to a system function, R is a one element nested array containing the name of the system function.

Examples

```
      PLUS←+
      +F←⎕CR'PLUS'
+
      ρF
1
      C←⎕CR
      C'C'
 ⎕CR
      ρC'C'
1
```

```
        ∇R←CONDITION (FN1 ELSE FN2) X
[1]     →CONDITION/L1
[2]     R←FN2 X ◊ →0
[3]     L1:R←FN1 X
[4]     ∇

        2 □STOP 'ELSE'
        (X≥0) ⌊ ELSE ⌈ X←¯2.5

ELSE[2]
        X
¯2.5
        □CR'FN2'
⌈
        →□LC
¯2
```

If `Y` is a name assigned to a derived function, `R` is a vector whose elements represent the arrays, functions, and operators from which `Y` was constructed. Constituent functions are represented by their own □CRs, so in this respect the definition of □CR is recursive. Primitive operators are treated like primitive functions, and are represented by their corresponding symbols. Arrays are represented by themselves.

Example

```
        BOX←2 2∘ρ
        +F←□CR'BOX'
 2 2 ∘ρ
        ρF
3
        DISPLAY F
.→----------.
| .→--.      |
| |2 2| ∘ ρ |
| '~--' - - |
'∈----------'
```

If `Y` is a name assigned to a defined function, `R` is the □CR of the defined function. In particular, the name that appears in the function header is the name of the original defined function, not the assigned name `Y`.

Example

```
        AVERAGE←MEAN
        □CR'AVERAGE'
R←MEAN X     ⍝ Arithmetic mean
R←(+/X)÷ρX
```

Change Space: $\{R\}\leftarrow\{X\}\square CS\ Y$

Y must be namespace reference (ref) or a simple character scalar or vector identifying the name of a namespace.

If specified, X is a simple character scalar, vector, matrix or a nested vector of character vectors identifying zero or more workspace objects to be *exported* into the namespace Y.

The identifiers in X and Y may be simple names or compound names separated by '.' and including the names of the special namespaces '$\square SE$', '#', '##' and ''.

The result R is the full name (starting "#.") of the space in which the function or operator was executing prior to the $\square CS$.

$\square CS$ changes the space in which the current function or operator is running to the namespace Y and returns the original space, in which the function was previously running, as a shy result. **After the** $\square CS$, references to *global* names (with the exception of those specified in X) are taken to be references to *global* names in Y. References to local names (i.e. those local to the current function or operator) are unaffected.

When the function or operator terminates, the calling function resumes execution in its original space.

The names listed in X are temporarily *exported* to the namespace Y. If objects with the same name exist in Y, these objects are effectively *shadowed* and are inaccessible.

Note that calling $\square CS$ with an empty argument Y obtains the namespace in which a function is currently executing.

Example

This simple example illustrates how $\square CS$ may be used to avoid typing long pathnames when building a tree of GUI objects. Note that the objects NEW and $OPEN$ are created as children of the $FILE$ menu as a result of using $\square CS$ to change into the $F.MB.FILE$ namespace.

```
      ∇ MAKE_FORM;F;OLD
[1]     'F'⎕WC'Form'
[2]     'F.MB'⎕WC'MenuBar'
[3]     'F.MB.FILE'⎕WC'Menu' '&File'
[4]
[5]     OLD←⎕CS'F.MB.FILE'
[6]     'NEW'⎕WC'MenuItem' '&New'
[7]     'OPEN'⎕WC'MenuItem' '&Open'
[8]     ⎕CS OLD
[9]
[10]    'F.MB.EDIT'⎕WC'Menu' '&Edit'
[11]
[12]    OLD←⎕CS'F.MB.EDIT'
[13]    'UNDO'⎕WC'MenuItem' '&Undo'
[14]    'REDO'⎕WC'MenuItem' '&Redo'
[15]    ⎕CS OLD
[16]    ...
      ∇
```

Example

This example illustrates how ⎕CS can be used to handle multiple *instances* of an object

Suppose a form F1 contains buttons B1 and B2. Each button maintains a count of the number of times it has been pressed, and the form maintains a count of the total number of button presses. The single callback function PRESS and its subfunction FMT can reside in the form itself

```
      )CS F1
#.F1
      ⍝ Note that both instances reference
      ⍝ the same callback function
      'B1'⎕WS'Event' 'Select' 'PRESS'
      'B2'⎕WS'Event' 'Select' 'PRESS'

      ⍝ Initialise total and instance counts.
      TOTAL ← B1.COUNT ← B2.COUNT ← 0


    ∇ PRESS MSG
[1]   'FMT TOTAL'⎕CS⊃MSG ⍝     Switch to instance space.
[2]   (TOTAL COUNT)+←1      ⍝   Incr total & instance count
.
[3]   ⎕WS'Caption'(COUNT FMT TOTAL) ⍝ Set instance caption
.
    ∇
```

```
     ∇ CAPT←INST FMT TOTL        ⍝ Format button caption.
[1]    CAPT←(⍕INST),'/',⍕TOTL   ⍝ E.g. 40/100.
```

Example

This example uses ⎕CS to explore a namespace tree and display the structure. Note that it must export its own name (tree) each time it changes space, because the name tree is global.

```
      ∇ tabs tree space;subs      ⍝ Display namespace tree
[1]    tabs,space
[2]    'tree'⎕CS space
[3]    →(⍴subs←↓⎕NL 9)↓0
[4]    (tabs,'.  ')∘tree¨subs
    ∇

      )ns x.y
#.x.y
      )ns z
#.z
      ''tree '#'
#
.   x
.   .   y
.   z
```

Comparison Tolerance: ⎕CT

The value of ⎕CT determines the precision with which two numbers are judged to be equal. Two numbers, X and Y, are judged to be equal if:

$$(|X-Y) ≤ ⎕CT × (|X) ⌈ |Y$$

where ≤ is applied without tolerance.

⎕CT may be assigned any value in the range from 0 to 16*¯8. A value of 0 ensures exact comparison. The value in a clear workspace is 1E¯14.

⎕CT is an implicit argument of the monadic primitive functions Ceiling (⌈), Floor (⌊) and Unique (∪), and of the dyadic functions Equal (=), Excluding (~), Find (⍷), Greater (>), Greater or Equal (≥), Index of (⍳), Intersection (∩), Less (<), Less or Equal (≤), Match (≡), Membership (∊), Not Match (≢), Not Equal (≠), Residue (|) and Union (∪), as well as ⎕FMT O-format.

Examples

```
      ⎕CT←1E¯10
      1.00000000001 1.0000001 = 1
```

```
     1  0
```

Copy Workspace: $\{X\}\square CY\ Y$

Y must be a simple character scalar or vector identifying a saved workspace. *X* is optional. If present, it must be a simple character scalar, vector or matrix. A scalar or vector is treated as a single row matrix. Each (implied) row of *X* is interpreted as an APL name.

Each (implied) row of *X* is taken to be the name of an active object in the workspace identified by *Y*. If *X* is omitted, the names of all defined active objects in that workspace are implied (defined functions and operators, variables, labels and namespaces).

Each object named in *X* (or implied) is copied from the workspace identified by *Y* to become the active object referenced by that name in the active workspace if the object can be copied. A copied label is re-defined to be a variable of numeric type. If the name of the copied object has an active referent in the active workspace, the name is disassociated from its value and the copied object becomes the active referent to that name. In particular, a function in the state indicator which is disassociated may be executed whilst it remains in the state indicator, but it ceases to exist for other purposes, such as editing.

You may copy an object from a namespace by specifying its full pathname. The object will be copied to the current namespace in the active workspace, losing its original parent and gaining a new one in the process. You may only copy a GUI object into a namespace that is a suitable parent for that object. For example, you could only copy a Group object from a saved workspace if the current namespace in the active workspace is itself a Form, SubForm or Group.

A *DOMAIN ERROR* is reported in any of the following cases:

1. *Y* is ill-formed, or is not the name of a workspace with access authorised for the active user account.

2. Any name in *X* is ill-formed.

3. An object named in *X* does not exist as an active object in workspace named in *Y*.

4. An object being copied has the same name as an active label.

A *WS FULL* is reported if the active workspace becomes full during the copying process.

Example

```
      ⎕VR'FOO'
    ∇ R←FOO
[1]   R←10
    ∇
    'FOO' ⎕CY 'BACKUP'
     ⎕VR'FOO'
    ∇ R←FOO X
[1]   R←10×X
    ∇
```

System variables are copied if explicitly included in the left argument, but not if the left argument is omitted.

Example

```
      ⎕LX

      (2 3ρ'⎕LX  X')⎕CY'WS/CRASH'
      ⎕LX
→RESTART
```

A copied object may have the same name as an object being executed. If so, the name is disassociated from the existing object, but the existing object remains defined in the workspace until its execution is completed.

Example

```
      )SI
FOO[1]*

      ⎕VR'FOO'
    ∇ R←FOO
[1]   R←10
    ∇

    'FOO'⎕CY'WS/MYWORK'

      FOO
1 2 3
      )SI
FOO[1]*

      →⎕LC
```

```
      1 0
```

Digits: $R←⎕D$

This is a simple character vector of the digits from 0 to 9.

Example

```
      ⎕D
0123456789
```

Division Method: $⎕DIV$

The value of $⎕DIV$ determines how division by zero is to be treated. If $⎕DIV$ = 0, division by 0 produces a *DOMAIN ERROR* except that the special case of 0÷0 returns 1.

If $⎕DIV$ = 1, division by 0 returns 0.

$⎕DIV$ may be assigned the value 0 or 1. The value in a clear workspace is 0.

$⎕DIV$ is an implicit argument of the monadic function Reciprocal (÷) and the dyadic function Divide (÷).

Examples

```
      ⎕DIV←0

      1 0 2 ÷ 2 0 1
0.5 1 2

      ÷0 1
DOMAIN ERROR
      ÷0 1
      ^

      ⎕DIV←1

      ÷0 2
0 0.5

      1 0 2 ÷ 0 0 4
0 0 0.5
```

Delay: $\{R\}\leftarrow\square DL\ Y$

Y must be a simple non-negative numeric scalar or one element vector. A pause of approximately Y seconds is caused.

The shy result R is an integer scalar value indicating the length of the pause in seconds.

The pause may be interrupted by a strong interrupt.

Diagnostic Message: $R\leftarrow\square DM$

This niladic function returns the last reported APL error as a three-element vector, giving error message, line in error and position of caret pointer.

Example

```
      2÷0
DOMAIN ERROR
      2÷0
     ^

      □DM
 DOMAIN ERROR          2÷0          ^
```

Dequeue Events: $\{R\}\leftarrow\Box DQ\ Y$

$\Box DQ$ awaits and processes events. Y specifies the GUI objects(s) for which events are to be processed. Objects are identified by their names, as character scalars/vectors, or by namespace references. These may be objects of type Root, Form, Locator, Filebox, MsgBox, PropertySheet, TCPSocket, Timer, Clipboard and pop-up Menu. Sub-objects (children) of those named in Y are also included. However, any objects which exist, but are not named in Y, are effectively disabled (do not respond to the user).

If Y is '.', all objects currently owned and subsequently created by the current thread are included in the $\Box DQ$. Note that because the Root object is owned by thread 0, events on Root are reported only to thread 0.

If Y is empty it specifies the object associated with the current namespace and is only valid if the current space is one of the objects listed above.

Otherwise, Y contains the name(s) of or reference(s) to the objects for which events are to be processed. Effectively, this is the list of objects with which the user may interact. A *DOMAIN ERROR* is reported if an element of Y refers to anything other than an existing "top-level" object.

Associated with every object is a set of events. For every event there is defined an "action" which specifies how that event is to be processed by $\Box DQ$. The "action" may be a number with the value 0, 1 or ¯1, or a character vector containing the name of a "callback function", or a character vector containing the name of a callback function coupled with an arbitrary array. Actions can be defined in a number of ways, but the following examples will illustrate the different cases.

```
OBJ ⎕WS 'Event' 'Select' 0

OBJ ⎕WS 'Event' 'Select' 1

OBJ ⎕WS 'Event' 'Select' 'FOO'

OBJ ⎕WS 'Event' 'Select' 'FOO' 10

OBJ ⎕WS 'Event' 'Select' 'FOO&'
```

These are treated as follows :

Action = 0 (the default)

> ⎕DQ performs "standard" processing appropriate to the object and type of event. For example, the standard processing for a KeyPress event in an Edit object is to action the keypress, i.e. to echo the character on the screen.

Action = ¯1

> This disables the event. The "standard" processing appropriate to the object and type of event is **not** performed, or in some cases is reversed. For example, if the "action code" for a KeyPress event (22) is set to ¯1, ⎕DQ simply ignores all keystrokes for the object in question.

Action = 1

> ⎕DQ terminates and returns information pertaining to the event (the **event message** in *R* as a nested vector whose first two elements are the name of the object (that generated the event) and the event code. *R* may contain additional elements depending upon the type of event that occurred.

Action = *fn {larg}*

> *fn* is a character vector containing the name of a *callback* function. This function is automatically invoked by ⎕DQ whenever the event occurs, and **prior** to the standard processing for the event. The callback is supplied the **event message** (see above) as its right argument, and, if specified, the array *larg* as its left argument. If the callback function fails to return a result, or returns the scalar value 1, ⎕DQ then performs the standard processing appropriate to the object and type of event. If the callback function returns a scalar 0, the standard processing is not performed or in some cases is reversed.

> If the callback function returns its event message with some of the parameters changed, these changes are incorporated into the standard processing. An example would be the processing of a keystroke message where the callback function substitutes upper case for lower case characters. The exact nature of this processing is described in the reference section on each event type.

Action = ⍎*expr*

> If *Action* is set to a character vector whose first element is the execute symbol (⍎) the remaining string will be executed automatically whenever the event occurs. The default processing for the event is performed first and may not be changed or inhibited in any way.

Action = *fn& {larg}*

fn is a character vector containing the name of a *callback* function. The function is executed in a new thread. The default processing for the event is performed first and may not be changed or inhibited in any way.

□DQ terminates in one of four instances. Note that its result is shy.

 Firstly, *□DQ* terminates when an event occurs whose "action code" is 1. In this case, its result is a nested vector containing the **event message** associated with the event. The structure of an event message varies according to the event type (see *Object Reference*). However, an event message has at least two elements of which the first is a character vector containing the name of the object, and the second is a numeric code specifying the event type.

□DQ also terminates if all of the objects named in *Y* have been deleted. In this case, the result is an empty character vector. Objects are deleted either using *□EX*, or on exit from a defined function or operator if the names are localised in the header, or on closing a form using the system menu.

Thirdly, *□DQ* terminates if the object named in its right argument is a special *modal* object, such as a *MsgBox*, *FileBox* or *Locator*, and the user has finished interacting with the object (e.g. by pressing an "OK" button). The return value of *□DQ* in this case depends on the action code of the event.

Finally, *□DQ* terminates with a *VALUE ERROR* if it attempts to execute a callback function that is undefined.

Data Representation (Monadic): R←□DR Y

Monadic *□DR* returns the type of its argument *Y*. The result *R* is an integer scalar containing one of the following values:

Value	Data Type
11	1 bit Boolean
82	8 bits Character
83	**8 bits Integer**
163	**16 bits Integer**
323	32 bits Integer
326	32 bits Pointer
645	64 bits Floating

Note that types **83** and **163** are exclusive to Dyalog APL.

Data Representation (Dyadic):   *R←X ⎕DR Y*

Dyadic *⎕DR* converts the data type of its argument *Y* according to the type specification *X*. See monadic *⎕DR* on the previous page for a list of data types.

Case 1: *X* is a single integer value. The bits in the right argument are interpreted as elements of an array of type *X*. The shape of the resulting new array will typically be changed along the last axis. For example, a character array seen as Boolean will have 8 times as many elements along the last axis.

Case 2: *X* is a 2-element integer value. The bits in the right argument are interpreted as type *X[1]*. The system then attempts to convert the elements of the resulting array to type *X[2]* without loss of precision. The result *R* is a two element nested array comprised of:

[1]   The converted elements or a fill element (0 or blank) where the conversion failed

[2]   A Boolean array of the same shape indicating which elements were successfully converted.

Case 3: *X* is a 3-element integer value and *X[2 3]* is 163 82. The bits in the right argument are interpreted as elements of an array of type *X[1]*. The system then converts them to the character representation of the corresponding 16 bit integers. This case is provided primarily for compatibility with APL*PLUS. For new applications, the use of the [conv] field with *⎕NAPPEND* and *⎕NREPLACE* is recommended.

Conversion to and from character (data type 82) uses the translate vector given by *⎕NXLATE* 0. By default this is the mapping defined by the current output translate table (usually WIN.DOT).

Edit Object: $\{R\} \leftarrow \{X\} \square ED\ Y$

$\square ED$ invokes the Editor. Y is a simple character vector, a simple character matrix, or a vector of character vectors, containing the name(s) of objects to be edited. The optional left argument X is a character scalar or character vector with as many elements as there are names in Y. Each element of X specifies the type of the corresponding (new) object named in Y, where :

'∇'	-	function or operator
'→'	-	simple character vector
'∈'	-	vector of character vectors
'−'	-	simple character matrix

If an object named in Y already exists, the corresponding type specification in X is ignored.

If $\square ED$ is called from the Session, it opens Edit windows for the object(s) named in Y and returns a null result. The cursor is positioned in the first of the Edit windows opened by $\square ED$, but may be moved to the Session or to any other window which is currently open. The effect is almost identical to using $)ED$.

If $\square ED$ is called from a defined function or operator, its behaviour is different. On asynchronous terminals, and in Dyalog APL for DOS/386, the Edit windows are automatically displayed in "full-screen" mode (ZOOMED). In all implementations, the user is restricted to those windows named in Y. The user may not skip to the Session even though the Session may be visible

$\square ED$ terminates and returns a result ONLY when the user explicitly closes all the windows for the named objects. In this case the result contains the names of any objects which have been changed, and has the same structure as Y.

Event Message: $R \leftarrow \Box EM \ Y$

Y must be a simple non-negative integer scalar or vector of event codes. If *Y* is a scalar, *R* is a simple character vector containing the associated event message. If *Y* is a vector, *R* is a vector of character vectors containing the corresponding event messages.

If *Y* refers to an undefined error code "*n*", the event message returned is "*ERROR NUMBER n*".

Example

```
      ⎕EM 11
DOMAIN ERROR
```

Event Number: $R \leftarrow \Box EN$

This simple integer scalar reports the identification number for the most recent event which occurred, caused by an APL action or by an interrupt or by the $\Box SIGNAL$ system function. Its value in a clear workspace is `0`.

Exception: $R \leftarrow \Box EXCEPTION$

This is a system object that identifies the most recent *Exception* thrown by a Microsoft .Net object.

$\Box EXCEPTION$ derives from the Microsoft .Net class System.Exception. Among its properties are the following, all of which are strings:

Source	The name of the .Net namespace in which the exception was generated
StackTrace	The calling stack
Message	The error message

```
      ⎕USING←'System'
      DT←DateTime.New 100000 0 0
EXCEPTION
      DT←DateTime.New 100000 0 0

      ⎕EN
90
```

```
      □EXCEPTION.Message
Specified argument was out of the range of valid values.

Parameter name: Year, Month, and Day parameters describe
an unrepresentable DateTime.

      □EXCEPTION.Source
mscorlib

      □EXCEPTION.StackTrace
   at System.DateTime.DateToTicks(Int32 year, Int32 month,
                                  Int32 day)

   at System.DateTime..ctor(Int32 year, Int32 month,
                            Int32 day)
```

Expunge Object: {R}←□EX Y

Y must be a simple character scalar, vector or matrix. A scalar or vector is treated as a
single row matrix. Each row of *Y* is interpreted as an APL name. *R* is a simple logical
vector with one element per row of *Y*.

Each name in *Y* is disassociated from its value if the active referent for the name is a
defined function, operator, variable or namespace.

The value of an element of *R* is 1 if the corresponding name in *Y* is now available for
use. This does not necessarily mean that the existing value was erased for that name.
A value of 0 is returned for an ill-formed name or for a distinguished name in *Y*. The
result is suppressed if not used or assigned.

Examples

```
      □EX'VAR'
      +□EX¨'FOO' '□IO' 'X' '123'
 1  0  1  0
```

If a named object is being executed the existing value will continue to be used until its
execution is completed. However, the name becomes available immediately for other
use.

Examples

```
      )SI
FOO[1]*

      ⎕VR'FOO'
    ∇ R←FOO
[1]   R←10
    ∇
    +⎕EX'FOO'
1
    )SI
FOO[1]*


    ∇FOO[⎕]
defn error

    FOO←1 2 3
    →⎕LC
10
    FOO
1 2 3
```

If a named object is an external variable, the external array is disassociated from the name:

```
      ⎕XT'F'
FILES/COSTS
      ⎕EX'F' ◇ ⎕XT'F'
```

If the named object is a GUI object, the object and all its children are deleted and removed from the screen. The expression `⎕EX'.'` deletes all objects owned by the current thread **except** for the Root object itself. In addition, if this expression is executed by thread 0, it resets all the properties of `'.'` to their default values. Furthermore, any unprocessed events in the event queue are discarded.

If the named object is a shared variable, the variable is retracted.

If the named object is the last remaining external function of an auxiliary process, the AP is terminated.

If the named object is the last reference into a dynamic link library, the DLL is freed.

Export Object: $\{R\} \leftarrow \{X\} \square EXPORT \ Y$

$\square EXPORT$ is used to set or query the export type of a defined function (or operator) referenced by the $\square PATH$ mechanism.

Y is a character matrix or vector-of-vectors representing the names of functions and operators whose export type is to be set or queried.

X is an integer scalar or vector (one per name in the namelist) indicating the export type. X can currently be one of the values:

> 0 - not exported.
> 1 - exported (default).

A scalar or 1-element-vector type is replicated to conform with a multi-name list.

The result R is a vector that reports the export type of the functions and operators named in Y. When used dyadically to set export type, the result is shy.

When the path mechanism locates a referenced function (or operator) in the list of namespaces in the $\square PATH$ system variable, it examines the function's export type:

0   This instance of the function is ignored and the search is resumed at the next namespace in the $\square PATH$ list. Type-0 is typically used for functions residing in a utility namespace which are not themselves utilities, for example the private sub-function of a utility function.

1   This instance of the function is executed in the namespace in which is was found and the search terminated. The effect is exactly as if the function had been referenced by its full path name.

Warning: The left domain of $\square EXPORT$ may be extended in future to include extra types 2, 3, ... (for example, to change the behaviour of the function). This means that, while $\square EXPORT$ returns a boolean result in the first version, this may not be the case in the future. If you need a boolean result, use $0 \neq$ or an equivalent.

```
      (0≠⎕EXPORT ⎕nl 3 4)/⎕nl 3 4   ⍝ list of exported
                                    ⍝ functions and operators.
```

File Append Component: $\quad\quad\quad \{R\}\leftarrow X\ \square FAPPEND\ Y$

Access code 8

Y must be a simple integer scalar or a 1 or 2 element vector containing the file tie number followed by an optional passnumber. If the passnumber is omitted it is assumed to be zero. X may be any array including, for example, the $\square OR$ of a namespace.

The shy result R is the number of the component to which X is written, and is 1 greater than the previously highest component number in the file, or 1 if the file is new.

Examples

```
      (1000?1000) ⎕FAPPEND 1

      ⎕←(2 3ρι6) 'Geoff' (⎕OR'FOO') ⎕FAPPEND 1
12

      ⎕←A B C ⎕FAPPEND¨1
13 14 15

Dump←{
    tie←α ⎕FCREATE 0              ⍝ create file.
    (⎕FUNTIE tie){}ω ⎕FAPPEND tie ⍝ append and untie.
}
```

File System Available: $\quad\quad\quad\quad\quad\quad R\leftarrow\square FAVAIL$

This niladic function returns the boolean value 1 unless the component file system is unavailable for some reason, in which case it returns 0. If $\square FAVAIL$ does return 0, most of the component file system functions will generate the error message:

FILE SYSTEM NOT AVAILABLE

See *User Guide* for further details.

File Create: $\{R\} \leftarrow X \ \square FCREATE \ Y$

Y must be 0 or a simple integer scalar or 1 element vector containing an available file tie number to be associated with the file for further file operations. It must not be the tie number associated with another tied file. X must be either

a)   a simple character scalar or vector which specifies the name of the file to be created. See *User Guide* for file naming conventions under UNIX and DOS.

or

b)   a vector of length 1 or 2 whose items are:

            - a simple character scalar or vector as above.

             - an integer scalar specifying the file size limit in bytes.
                This element is ignored, but is included for compatibility with other APLs.

The newly created file is tied for exclusive use.

The shy result of $\square FCREATE$ is the tie number of the new file.

Automatic Tie Number Allocation

A tie number of 0 as argument to a create or tie operation, allocates, and returns as an explicit result, the first (closest to zero) available tie number. This allows you to simplify code. For example:

from:

```
tie←1+⌈/0,⎕FNUMS     ⍝ With next available number,
file ⎕FCREATE tie    ⍝ ... create file.
```

to:

```
tie←file ⎕FCREATE 0  ⍝ Create with first available..
```

Examples

```
      ⎕FSIZE 'SALES' ⎕FCREATE 0
1 1 76 4294967295

      '../budget/SALES'.85' ⎕FCREATE 2    ⍝ UNIX

      '..\BUDGET\SALES'     ⎕FCREATE 2    ⍝ DOS

      'COSTS' 200000 ⎕FCREATE 4
```

File Drop Component: $\{R\}←⎕FDROP\ Y$

Access code 32

Y must be a simple integer vector of length 2 or 3 whose elements are:

[1] a file tie number

[2] a number specifying the position and number of components to be dropped. A positive value indicates that components are to be removed from the beginning of the file; a negative value indicates that components are to be removed from the end of the file

[3] an optional passnumber which if omitted is assumed to be zero

The shy result of a ⎕FDROP is a vector of the numbers of the dropped components. This is analogous to ⎕FAPPEND in that the result is potentially useful for updating some sort of dictionary:

```
      cnos,←vec ⎕FAPPEND¨tie ⍝ Append index to dictionary

      cnos~←⎕FDROP tie,-⍴vec ⍝ Remove index from dictionar
y
```

Note that the result vector, though potentially large, is generated only on request.

Examples

```
      ⎕FSIZE 1
1 21 5436 4294967295

      ⎕FDROP 1 3 ◇ ⎕FSIZE 1
4 21 5436 4294967295
```

```
      ⎕FDROP 1 ¯2 ◇ ⎕FSIZE 1
4 19 5436 4294967295
```

File Erase: {R}←X □FERASE Y

Access code 4

Y must be a simple integer scalar or 1 or 2 element vector containing the file tie number followed by an optional passnumber. If the passnumber is omitted it is assumed to be zero. *X* must be a character scalar or vector containing the name of the file associated with the tie number *Y*. This name must be identical with the name used to tie the file, and the file must be exclusively tied. The file named in *X* is erased and untied. See *User Guide* for file naming conventions under UNIX and DOS.

The shy result of □*FERASE* is the tie number of the erased file.

Examples

```
      'SALES'□FERASE 'SALES' □FTIE 0

      './temp' □FCREATE 1
      'temp' □FERASE 1
FILE NAME ERROR
      'temp'□FERASE 1
      ^

      './TEMP' □FERASE 1
FILE NAME ERROR
      './TEMP'□FERASE 1
      ^

      './temp' □FERASE 1
```

File Hold: {R}←□FHOLD Y

Access code 2048

This function holds component file(s) and/or external variable(s).

If applied to component files, then *Y* is an integer scalar, vector, or one-row matrix of file tie numbers, or a two-row matrix whose first row contains file tie numbers and whose second row contains passnumbers.

If applied to external variables, then *Y* is a non-simple scalar or vector of character vectors, each of which is the name of an external variable. (NOT the file names associated with those variables).

If applied to component files **and** external variables, *Y* is a vector whose elements are either integer scalars representing tie numbers, or character vectors containing names of external variables.

The effect is as follows :

1. The user's preceding holds (if any) are released.

2. Execution is suspended until the designated files are free of holds by any other task.

3. When all the designated files are free, execution proceeds. Until the hold is released, other tasks using *⎕FHOLD* on any of the designated files will wait.

If *Y* is empty, the user's preceding hold (if any) is released, and execution continues.

A hold is released by any of the following :

1. Another *⎕FHOLD*

2. Untying or retying all the designated files. If some but not all are untied or retied, they become free for another task but the hold persists for those that remain tied.

3. Termination of APL.

4. Any untrapped error or interrupt.

5. A return to immediate execution.

Note that a hold is not released by a request for input through ⎕ or ⍞.

Note also that point 5 above implies that *⎕FHOLD* is generally useful only when called from a defined function, as holds set in immediate execution (desk calculator) mode are released immediately.

The shy result of *⎕FHOLD* is a vector of tie numbers of the files held.

Examples :

```
      ⎕FHOLD 1

      ⎕FHOLD θ

      ⎕FHOLD ⊂'XTVAR'

      ⎕FHOLD 1 2,[0.5]0 16385

      ⎕FHOLD 1 'XTVAR'
```

Component File Library: $R \leftarrow \Box FLIB \ Y$

Y must be a simple character scalar or vector which specifies the name of the directory whose APL component files are to be listed. If Y is empty, the current working directory is assumed.

The result R is a character matrix containing the names of the component files in the directory with one row per file. The number of columns is given by the longest file name. Each file name is prefixed by Y followed by a directory delimiter character. The ordering of the rows is not defined.

If there are no APL component files accessible to the user in the directory in question, the result is an empty character matrix with 0 rows and 0 columns.

Note that if a file is exclusively tied (as opposed to share tied) then it is not reported by $\Box FLIB$.

Examples

```
      ⎕FLIB ''
SALESFILE
COSTS

      ⎕FLIB '.'
./SALESFILE
./COSTS

      ⎕FLIB '../budget'
../budget/SALES.85
../budget/COSTS.85
```

Format (Monadic): $R \leftarrow \square FMT \ Y$

Y may be any array. *R* is a simple character matrix which appears the same as the default display of *Y*. If *Y* contains control characters from $\square TC$, they will be resolved.

Examples

```
      A←⎕FMT '∩' ,⎕TC[1],'∘'

      ρA
1 1
      A
∩

      A←⎕VR 'FOO'

      A
    ∇ R←FOO
[1]    R←10
    ∇

      ρA
31
      B←⎕FMT A

      B
    ∇ R←FOO
[1]    R←10
    ∇

      ρB
3 12
```

Format (Dyadic): $R \leftarrow X \quad \square FMT \quad Y$

Y must be a simple array of rank not exceeding two, or a non-simple scalar or vector whose items are simple arrays of rank not exceeding two. The simple arrays in Y must be homogeneous, either character or numeric. X must be a simple character vector. R is a simple character matrix.

X is a format specification that defines how columns of the simple arrays in Y are to appear. A simple scalar in Y is treated as a one-element matrix. A simple vector in Y is treated as a one-column matrix. Each column of the simple arrays in Y is formatted in left-to-right order according to the format specification in X taken in left-to-right order and used cyclically if necessary.

R has the same number of columns as the longest column (or implied column) in Y, and the number of rows is determined from the format specification.

The **format specification** consists of a series of control phrases, with adjacent phrases separated by a single comma, selected from the following:

`rAw`	Alphanumeric format
`rEw.s`	Scaled format
`rqFw.d`	Decimal format
`rqG⎕pattern⎕`	Pattern
`rqIw`	Integer format
`Tn`	Absolute tabulation
`Xn`	Relative tabulation
`⎕t⎕`	Text insertion.

(Alternative surrounding pairs for Pattern or Text insertion are
 `< >`, `⊂ ⊃`, `⎕ ⎕` or `¨ ¨`.)

where:

`r`	is an optional repetition factor indicating that the format phrase is to be applied to r columns of Y.
`q`	is an optional usage of qualifiers or affixtures from those described below.
`w`	is an integer value specifying the total field width per column of Y, including any affixtures.

s	is an integer value specifying the number of significant digits in Scaled format; *s* must be less than *w*-1.
d	is an integer value specifying the number of places of decimal in Decimal format; *d* must be less than *w*.
n	is an integer value specifying a tab position relative to the notional left margin (for *T*-format) or relative to the last formatted position (for *X*-format) at which to begin the next format.
t	is any arbitrary text excluding the surrounding character pair. Double quotes imply a single quote in the result.
pattern	see following section **G format**

Qualifiers *q* are as follows:

B	leaves the field blank if the result would otherwise be zero.
C	inserts commas between triads of digits starting from the rightmost digit of the integer part of the result.
Km	scales numeric values by 1*Em* where m is an integer; negation may be indicated by ¯ or - preceding the number.
L	left justifies the result in the field width.
Ov⎕t⎕	replaces specific numeric value *v* with the text *t*.
S⎕p⎕	substitutes standard characters. p is a string of pairs of symbols enclosed between any of the Text Insertion delimiters. The first of each pair is the standard symbol and the second is the symbol to be substituted. Standard symbols are:

*	overflow fill character
.	decimal point
,	triad separator for *C* qualifier
0	fill character for *Z* qualifier
_	loss of precision character

Z	fills unused leading positions in the result with zeros (and commas if C is also specified).
9	digit selector

Affixtures are as follows:

$M⍞t⍞$	prefixes negative results with the text t instead of the negative sign.
$N⍞t⍞$	post-fixes negative results with the text t.
$P⍞t⍞$	prefixes positive or zero results with the text t.
$Q⍞t⍞$	post-fixes positive or zero results with the text t.
$R⍞t⍞$	presets the field with the text t which is repeated as necessary to fill the field. The text will be replaced in parts of the field filled by the result, including the effects of other qualifiers and affixtures except the B qualifier.

The surrounding affixture delimiters may be replaced by the alternative pairs described for Text Insertion.

Examples

A vector is treated as a column:

```
      'I5' ⎕FMT 10 20 30
   10
   20
   30
```

The format specification is used cyclically to format the columns of the right argument:

```
      'I3,F5.2' ⎕FMT 2 4⍴⍳8
 1 2.00  3 4.00
 5 6.00  7 8.00
```

The columns of the separate arrays in the items of a non-simple right argument are formatted in order. Rows in a formatted column beyond the length of the column are left blank:

```
      '2I4,F7.1' ⎕FMT (ι4)(2 2ρ 0.1×ι4)
   1   0    0.2
   1   0    0.4
   3
   4
```

Characters are right justified within the specified field width, unless the *L* qualifier is specified:

```
      'A2' ⎕FMT 1 6ρ'SPACED'
 S P A C E D
```

If the result is too wide to fit within the specified width, the field is filled with asterisks:

```
      'F5.2' ⎕FMT 0.1×5 1000 ¯100
 0.50
*****
*****
```

Relative tabulation (*X*-format) identifies the starting position for the next format phrase relative to the finishing position for the previous format, or the notional left margin if none. Negative values are permitted providing that the starting position is not brought back beyond the left margin. Blanks are inserted in the result, if necessary:

```
      'I2,X3,3A1' ⎕FMT (ι3)(2 3ρ'TOPCAT')
   1   TOP
   2   CAT
   3
```

Absolute tabulation (*T*-format) specifies the starting position for the next format relative to the notional left margin. If position 0 is specified, the next format starts at the next free position as viewed so far. Blanks are inserted into the result as required. Over-written columns in the result contain the most recently formatted array columns taken in left-to-right order:

```
      X←'6I1,T5,A1,T1,3A1,T7,F5.1'

      X ⎕FMT (1 6ρι6)('*')(1 3ρ'ABC')(22.2)
ABC4*6 22.2
```

If the number of specified significant digits exceeds the internal precision, low order digits are replaced by the symbol _:

```
      'F20.1' □FMT 1E18÷3
3333333333333333__._
```

The Text Insertion format phrase inserts the given text repeatedly in all rows of the result:

```
      MEN←3 5ρ'FRED BILL JAMES'

      WOMEN←2 5ρ'MARY JUNE '

      '5A1,<|>' □FMT MEN WOMEN
FRED |MARY |
BILL |JUNE |
JAMES|     |
```

The last example also illustrates that a Text Insertion phrase is used even though the data is exhausted. The following example illustrates some effects of the various qualifiers:

```
      X←'F5.1,BF6.1,X1,ZF5.1,X1,LF5.1,K3CS<.,,.>F10.1'

      X □FMT ⍉5 3ρ¯1.5 0 25
 ¯1.5   ¯1.5 ¯01.5 ¯1.5    ¯1.500,0
  0.0        000.0  0.0          0,0
 25.0   25.0 025.0 25.0    25.000,0
```

Affixtures allow text to be included within a field. The field width is not extended by the inclusion of affixtures. *N* and *Q* affixtures shift the result to the left by the number of characters in the text specification. Affixtures may be used to enclose negative results in parentheses in accordance with common accounting practice:

```
      'M<(>N<)>Q< >F9.2' □FMT 150.3 ¯50.25 0 1114.9
  150.30
  (50.25)
    0.00
 1114.90
```

One or more format phrases may be surrounded by parentheses and preceded by an optional repetition factor. The format phrases within parentheses will be re-used the given number of times before the next format phrase is used. A Text Insertion phrase will not be re-used if the last data format phrase is preceded by a closing parenthesis:

```
      'I2,2(</>,ZI2)' □FMT 1 3ρφ100|3↑□TS
```

20/07/89

G Format

Only the *B, K, S* and *O* qualifiers are valid with the *G* option

�*pattern*⌷ is an arbitrary string of characters, excluding the delimiter characters. Characters '9' and 'Z' (unless altered with the *S* qualifier) are special and are known as **digit selectors.**

The result of a *G* format will have length equal to the length of the pattern.

The data is rounded to the nearest integer (after possible scaling). Each digit of the rounded data replaces one digit selector in the result. If there are fewer data digits than digit selectors, the data digits are padded with leading zeros. If there are more data digits than digit selectors, the result will be filled with asterisks.

A '9' digit selector causes a data digit to be copied to the result.

A 'Z' digit selector causes a non-zero data digit to be copied to the result. A zero data digit is copied if and only if digits appear on either side of it. Otherwise a blank appears. Similarly text between digit selectors appears only if digits appear on either side of the text. Text appearing before the first digit selector or after the last will always appear in the result.

Examples

```
      'G⊂99/99/99⊃'⎕FMT 0 100 100 ⊥8 7 89
08/07/89

      'G⊂ZZ/ZZ/ZZ⊃'⎕FMT 80789 + 0 1
 8/07/89
 8/07/9

      'G⊂Andy ZZ Pauline ZZ⊃' ⎕FMT 2721.499 2699.5
Andy 27 Pauline 21
Andy 27

      ρ⎕←'K2G⊂DM Z.ZZZ.ZZ9,99⊃' ⎕FMT 1234567.89 1234.56
DM 1.234.567,89
DM    1.234,56
2 15
```

An error will be reported if:

1. Numeric data is matched against an *A* control phrase.
2. Character data is matched against other than an *A* control phrase.
3. The format specification is ill-formed.
4. For an F control phrase, *d > w - 2*

5. For an E control phrase, $s>w-2$

O Format Qualifier

The O format qualifier replaces a specific numeric value with a text string and may be used in conjunction with the E, F, I and G format phrases.

An O-qualifier consists of the letter "O" followed by the optional numeric value which is to be substituted (if omitted, the default is 0) and then the text string within pairs of symbols such as "< >". For example:

O - qualifier	Description
O<nil>	Replaces the value 0 with the text "nil"
O42<N/A>	Replaces the value 42 with the text "N/A"
O0.001<1/1000>	Replaces the value 0.001 with the text "1/1000"

The replacement text is inserted into the field in place of the numeric value. The text is normally right-aligned in the field, but will be left-aligned if the L qualifier is also specified.

It is permitted to specify more than one O-qualifier within a single phrase.

The O-qualifier is ⎕*CT* sensitive.

Examples

```
      'O<NIL>F7.2'⎕FMT 12.3 0 42.5
  12.30
    NIL
  42.50

      'O<NIL>LF7.2'⎕FMT 12.3 0 42.5

12.30
NIL
42.50

      'O<NIL>O42<N/A>I6'⎕FMT 12 0 42 13
    12
   NIL
   N/A
    13

     'O99<replace>F20.2'⎕fmt 99 100 101
            replace
            100.00
            101.00
```

File Names: $R \leftarrow \square FNAMES$

The result is a character matrix containing the names of all tied files, with one file name per row. The number of columns is that required by the longest file name.

A file name is returned precisely as it was specified when the file was tied. If no files are tied, the result is a character matrix with 0 rows and 0 columns. The rows of the result are in the order in which the files were tied.

Examples

```
      '/usr/pete/SALESFILE' ⎕FSTIE 16

      '../budget/COSTFILE' ⎕FSTIE 2

      'PROFIT' ⎕FCREATE 5

      ⎕FNAMES
/usr/pete/SALESFILE
../budget/COSTFILE
PROFIT

      ρ⎕FNAMES
3 19
      ⎕FNUMS,⎕FNAMES
16 /usr/pete/SALESFILE
 2 ../budget/COSTFILE
 5 PROFIT
```

File Numbers: $R \leftarrow \Box FNUMS$

The result is an integer vector of the **tie numbers** of all tied files. If no files are tied, the result is empty. The elements of the result are in the order in which the files were tied.

Examples

```
      '/usr/pete/SALESFILE' ⎕FSTIE 16

      '../budget/COSTFILE' ⎕FSTIE 2

      'PROFIT' ⎕FCREATE 5

      ⎕FNUMS
16 2 5

      ⎕FNUMS,⎕FNAMES
16 /usr/pete/SALESFILE
 2 ../budget/COSTFILE
 5 PROFIT

      ⎕FUNTIE ⎕FNUMS
      ρ⎕FNUMS
0
```

File Read Access: $R \leftarrow \Box FRDAC \ Y$

Access code 4096

Y must be a simple integer scalar or 1 or 2 element vector containing the file tie number followed by an optional passnumber. If the passnumber is omitted it is assumed to be zero. The result is the access matrix for the designated file.

See "File Access Control" in *User Guide* for further details.

Examples

```
      ⎕FRDAC 1
28 2105 16385
 0 2073 16385
31   ¯1     0
```

File Read Component Information: *R←⎕FRDCI Y*

Access code 512

Y must be a simple integer vector of length 2 or 3 containing the file tie number, component number and an optional passnumber. If the passnumber is omitted it is assumed to be zero.

The result is a 3 element numeric vector containing the following information:

a)   the size of the component in bytes (i.e. how much workspace it would occupy).

b)   the user number of the user who last updated the component.

c)   the time of the last update in 60ths of a second since 1st January 1970.

Example

```
      ⎕FRDCI 1 13
2200 207 3.702094494E10
```

File Read Component: *R←⎕FREAD Y*

Access code 1

Y must be a simple 2 or 3 element integer vector containing the file tie number, the component number, and an optional passnumber. If the passnumber is omitted it is assumed to be zero.

The result is the value of the array stored on the tied file at the given component number.

Examples

```
      ρSALES←⎕FREAD 1 241
3 2 12

GetFile←{⎕io←0                ⍝ Extract contents.
   tie←ω ⎕fstie 0             ⍝ new tie number.
   fm to←2↑⎕fsize tie         ⍝ first and next component.
   cnos←fm+ιto-fm             ⍝ vector of component nos.
   cvec←{⎕fread tie ω}¨cnos   ⍝ vector of components.
   cvec{α}⎕funtie tie         ⍝ ... untie and return.
```

```
}
```

File Rename: $\{R\}\leftarrow X \ \Box FRENAME \ Y$

Access code 128

Y must be a simple 1 or 2 element integer vector containing a file tie number and an optional passnumber. If the passnumber is omitted it is assumed to be zero.

X must be a simple character scalar or vector containing the new name of the file. This name must be in accordance with the operating system's conventions, and may be specified with a relative or absolute pathname.

The file being renamed must be tied exclusively.

The shy result of $\Box FRENAME$ is the tie number of the file.

Examples

```
      'SALES' ⎕FTIE 1
      'PROFIT' ⎕FTIE 2

      ⎕FNAMES
SALES
PROFIT

      'SALES.85' ⎕FRENAME 1
      '../profits/PROFIT.85' ⎕FRENAME 2

      ⎕FNAMES
SALES.85
../profits/PROFITS.85

Rename←{
    fm to←ω
    ⎕FUNTIE to ⎕FRENAME fm ⎕FTIE 0
}
```

File Replace Component: $\{R\} \leftarrow X$ ⎕$FREPLACE$ Y

Access code 16

Y must be a simple 2 or 3 element integer vector containing the file tie number, the component number, and an optional passnumber. If the passnumber is omitted it is assumed to be zero. The component number specified must lie within the file's component number limits.

X is any array (including, for example, the ⎕OR of a namespace), and overwrites the value of the specified component. The component information (see ⎕$FRDCI$) is also updated.

The shy result of ⎕$FREPLACE$ is the file index (component number of replaced record).

Example

```
      SALES←⎕FREAD 1 241

      (SALES×1.1) ⎕FREPLACE 1 241
```

Define a function to replace (index, value) pairs in a component file JMS.DCF:

```
Frep←{
    tie←α ⎕FTIE 0
    _←{ω ⎕FREPLACE tie α}/¨ω
    ⎕FUNTIE tie
}

      'jms'Frep(0 'abc')(29 'xxx')(7 'yyy')
```

File Resize: $\{R\}\leftarrow\{X\}\square FRESIZE\ Y$

Access code 1024

Y must be a simple integer scalar or 1 or 2 element vector containing the file tie number followed by an optional passnumber. If the passnumber is omitted it is assumed to be zero.

X is included for compatibility with other APL systems.

A side effect of $\square FRESIZE$ is to cause the file to be compacted. Any interrupt entered at the keyboard during the compaction is ignored.

The shy result of $\square FRESIZE$ is the tie number of the file.

Example

```
      □FSIZE 1
1 21 65271 4294967295

      200000 □FRESIZE 1 ◇ □FSIZE 1
1 21 41456 4294967295

      θ □FRESIZE 1      ⍝ Force file compaction.
```

File Size: $R\leftarrow\square FSIZE\ Y$

Y must be a simple integer scalar or 1 or 2 element vector containing the file tie number followed by an optional passnumber. If the passnumber is omitted it is assumed to be zero. The result is a 4 element numeric vector containing the following:

Element	Description
1	the number of first component,
2	1 + the number of the last component (i.e. the result of the next $\square FAPPEND$),
3	the current size of the file in bytes,
4	the file size limit in bytes.

Example

```
      □FSIZE 1
1 21 65271 4294967295
```

File Set Access: $\{R\} \leftarrow X \;\; \Box FSTAC \;\; Y$

Access code 8192

Y must be a simple integer scalar or 1 or 2 element vector containing the file tie number followed by an optional passnumber. If the passnumber is omitted it is assumed to be zero.

X must be a valid access matrix, i.e. a 3 column integer matrix with any number of rows.

See "File Access Control" in *User Guide* for further details.

The shy result of $\Box FSTAC$ is the tie number of the file.

Examples

```
SALES □FCREATE 1
(3 3ρ28 2105 16385 0 2073 16385 31 ¯1 0) □FSTAC 1
((□FRDAC 1),21 2105 16385) □FSTAC 1

(1 3ρ0 ¯1 0)□FSTAC 2
```

File Share Tie: $\{R\} \leftarrow X \ \square FSTIE \ Y$

Y must be 0 or a simple 1 or 2 element integer vector containing an available file tie number to be associated with the file for further file operations, and an optional passnumber. If the passnumber is omitted it is assumed to be zero. The tie number must not already be associated with a tied file.

X must be a simple character scalar or vector which specifies the name of the file to be tied. The file must be named in accordance with the operating system's conventions, and may be specified with a relative or absolute pathname.

The file must exist and be accessible by the user. If it is already tied by another task, it must not be tied exclusively.

The shy result of $\square FSTIE$ is the tie number of the file.

Automatic Tie Number Allocation

A tie number of 0 as argument to a create, share tie or exclusive tie operation, allocates the first (closest to zero) available tie number and returns it as an explicit result. This allows you to simplify code. For example:

from:

```
tie←1+⌈/0,⎕FNUMS  ⍝ With next available number,
file ⎕FSTIE tie   ⍝ ... share tie file.
```

to:

```
tie←file ⎕FSTIE 0 ⍝ Tie with first available number.
```

Example

```
'SALES' ⎕FSTIE 1

'../budget/COSTS' ⎕FSTIE 2
```

Exclusive File Tie: $\{R\} \leftarrow X$ $\square FTIE$ Y

Access code 2

Y must be 0 or a simple 1 or 2 element integer vector containing an available file tie number to be associated with the file for further file operations, and an optional passnumber. If the passnumber is omitted it is assumed to be zero. The tie number must not already be associated with a share tied or exclusively tied file.

X must be a simple character scalar or vector which specifies the name of the file to be exclusively tied. The file must be named in accordance with the operating system's conventions, and may be a relative or absolute pathname.

The file must exist and be accessible by the user. It may not already be tied by another user.

Automatic Tie Number Allocation

A tie number of 0 as argument to a create, share tie or exclusive tie operation, allocates the first (closest to zero) available tie number, and returns it as an explicit result,. This allows you to simplify code. For example:

from:

```
tie←1+⌈/0,⎕FNUMS ⍝ With next available number,
file ⎕FTIE tie   ⍝ ... tie file.
```

to:

```
tie←file ⎕FTIE 0 ⍝ Tie with first available number.
```

The shy result of $\square FTIE$ is the tie number of the file.

Examples

```
'SALES' ⎕FTIE 1

'../budget/COSTS' ⎕FTIE  2

'../budget/expenses' ⎕FTIE 0
3
```

File Untie: {R}←⎕FUNTIE Y

Y must be a simple integer scalar or vector (including Zilde). Files whose tie numbers occur in *Y* are untied. Other elements of *Y* have no effect.

If *Y* is empty, no files are untied, but all internal file buffers are written to disk. Under UNIX this is achieved with `fsync`; under DOS with the Commit File function (Int 21h Function 68h); on Win32 systems it is achieved with `FlushFileBuffers()`. This special facility allows the programmer to add extra security (at the expense of performance) for application data files.

The shy result of ⎕*FUNTIE* is a vector of tie numbers of the files **actually untied**.

Example

```
⎕FUNTIE ⎕FNUMS ⍝ Unties all tied files

⎕FUNTIE ⍬      ⍝ Flushes all buffers to disk
```

Fix Definition: {R}←⎕FX Y

Y is the representation form of a function or operator which may be:

1. its canonical representation form similar to that produced by ⎕*CR* except that redundant blanks are permitted other than within names and constants.

2. its nested representation form similar to that produced by ⎕*NR* except that redundant blanks are permitted other than within names and constants.

3. its object representation form produced by ⎕*OR*.

4. its vector representation form similar to that produced by ⎕*VR* except that additional blanks are permitted other than within names and constants.

R is either a simple character vector or an integer scalar. A result is not returned if a result is not explicitly used or assigned.

A side effect of ⎕*FX* is to create (fix) a function or operator in the workspace or current namespace from the definition given by *Y*. ⎕*IO* is an implicit argument of ⎕*FX*.

If the function or operator is successfully fixed, *R* is its name. Otherwise *R* is the row of the canonical representation form in which the first error preventing its definition is detected. An integer *R* is dependent on ⎕*IO*.

Functions and operators which are pendent, that is, in the State Indicator without a suspension mark (*), retain their original definition until they complete, or are cleared from the State Indicator. All other occurrences of the function or operator assume the new definition. The function or operator will fail to fix if it has the same name as an existing variable, or a visible label.

Index Origin: `⎕IO`

`⎕IO` determines the index of the first element of a non-empty vector.

`⎕IO` may be assigned the value 0 or 1. The value in a clear workspace is 1.

`⎕IO` is an implicit argument of any function derived from the Axis operator (`[K]`), of the monadic functions Fix (`⎕FX`), Grade Down (`▼`), Grade Up (`▲`), Index Generator (`ι`), Roll (`?`), and of the dyadic functions Deal (`?`), Find (`∈`), Grade Down (`▼`), Grade Up (`▲`), Index Of (`ι`), Indexed Assignment, Indexing, Pick (`⊃`) and Transpose (`⍉`).

Examples

```
      ⎕IO←1
      ι5
1 2 3 4 5

      ⎕IO←0
      ι5
0 1 2 3 4

      +/[0]2 3ρι6
3 5 7

      'ABC',[¯.5]'='
ABC
===
```

Key Label: `R←⎕KL Y`

`Y` is a simple character vector or a vector of character vectors containing codes defined for the Input Translate Table associated with your keyboard.

`R` is a simple character vector or a vector of character vectors containing the labels associated with the codes. If `Y` specifies codes that are not defined, the corresponding elements of `R` are empty.

`⎕KL` provides the information required to build device-independent help messages into applications, particularly full-screen applications using `⎕SM` and `⎕SR`.

Examples :

```
      ⎕KL 'RC'
Right
```

```
      ⎕KL 'ER' 'EP' 'QT' 'F1' 'F13'
 Enter  Esc  Shift-Esc  F1  Shift-F1
```

Line Count: $R \leftarrow \Box LC$

This is a simple vector of line numbers drawn from the state indicator (See *Chapter 2*). The most recently activated line is shown first. If a value corresponds to a defined function in the state indicator, it represents the current line number where the function is either suspended or pendent.

The value of $\Box LC$ changes immediately upon completion of the most recently activated line, or upon completion of execution within ⍎ or ⎕. If a $\Box STOP$ control is set, $\Box LC$ identifies the line on which the stop control is effected. In the case where a stop control is set on line 0 of a defined function, the first entry in $\Box LC$ is 0 when the control is effected.

The value of $\Box LC$ in a clear workspace is the null vector.

Examples

```
      )SI
TASK1[5]*
⍎
BEGIN[3]

      ⎕LC
5 3

      →⎕LC

      ⎕LC

      ρ⎕LC
0
```

Load Workspace: `⎕LOAD Y`

Y must be a simple character scalar or vector containing the identification of a saved workspace.

If *Y* is ill-formed or does not identify a saved workspace or the user account does not have access permission to the workspace, a *DOMAIN ERROR* is reported.

Otherwise, the active workspace is replaced by the workspace identified in *Y*. The active workspace is lost. If the loaded workspace was saved by the `)SAVE` system command, the latent expression (`⎕LX`) is immediately executed, unless APL was invoked with the -x option. If the loaded workspace was saved by the `⎕SAVE` system function, execution resumes from the point of exit from the `⎕SAVE` function, with the result of the `⎕SAVE` function being 0.

The workspace identification and time-stamp when saved is not displayed.

If the workspace contains any GUI objects whose *Visible* property is 1, these objects will be displayed. If the workspace contains a non-empty `⎕SM` but does not contain an SM GUI object, the form defined by `⎕SM` will be displayed in a window on the screen.

Lock Definition: `{X}⎕LOCK Y`

Y must be a simple character scalar, or vector which is taken to be the name of a defined function or operator in the active workspace.

The active referent to the name in the workspace is locked. Stop or trace vectors, formerly set by the `⎕STOP` and `⎕TRACE` functions, are cancelled.

The optional left argument *X* specifies to what extent the function code is hidden. *X* may be 1, 2 or 3 (the default) with the following meaning:

1. The object may not be displayed and you may not obtain its character form using `⎕CR`, `⎕VR` or `⎕NR`.

2. Execution cannot be suspended with the locked function or operator in the state indicator. On suspension of execution the state indicator is cut back to the statement containing the call to the locked function or operator.

3. Both 1 and 2 apply. You can neither display the locked object nor suspend execution within it.

Locks are additive, so that

```
1 ⎕LOCK'FOO'
2 ⎕LOCK'FOO'
```

is the same as:

```
3 ⎕LOCK'FOO'
```

A *DOMAIN ERROR* is reported if:

1. *Y* is ill-formed.

2. The name in *Y* is not the name of a visible defined function or operator which is not locked.

Examples

```
      ⎕VR'FOO'
    ∇ R←FOO
[1]    R←10
    ∇

      ⎕LOCK'FOO'
      ⎕VR'FOO'


      ⎕LOCK'FOO'
DOMAIN ERROR
      ⎕LOCK'FOO'
      ^
```

Latent Expression: $\Box LX$

This may be a character vector or scalar representing an APL expression. The expression is executed automatically when the workspace is loaded. If APL is invoked using the -x flag, this execution is suppressed.

The value of $\Box LX$ in a clear workspace is `''`.

Example

```
      ⎕LX←'''GOOD MORNING PETE'''

      )SAVE GREETING
GREETING saved Tue Sep 8 10:49:29 1998

      )LOAD GREETING
./GREETING saved Tue Sep 8 10:49:29 1998
GOOD MORNING PETE
```

Map File: $R \leftarrow \{X\}\Box MAP\ Y$

This function associates a mapped file with an APL array in the workspace.

Two types of mapped files are supported; *APL* and *raw*. An *APL* mapped file contains the binary representation of a Dyalog APL array, including its header. A file of this type must be created using the supplied utility function $\Delta MPUT$. When you map an APL file, the rank, shape and data type of the array is obtained from the information on the file.

A *raw* mapped file is an arbitrary collection of bytes. When you map a raw file, you must specify the characteristics of the APL array to be associated with this data. In particular, the data type and its shape.

The type of mapping is determined by the presence (raw) or absence (APL) of the left argument to $\Box MAP$.

The right argument Y specifies the name of the file to be mapped and, optionally, the access type and a start byte in the file. Y may be a simple character vector, or a 2 or 3-element nested vector containing:
1. file name (character scalar/vector)
2. access code (character scalar/vector) : one of : 'R' , 'W', 'r' or 'w'
3. start byte offset (integer scalar/vector). This must be a multiple of 4 (default is 0)

If X is specified, it defines the type and shape to be associated with *raw* data on file. X must be an integer scalar or vector. The first item of X specifies the data type and must be one of: 11, 82, 83, 163, 323 or 645.

Following items determine the shape of the mapped array. A value of ¯1 on any (but normally the first) axis in the shape is replaced by the system to mean: read as many complete records from the file as possible. Only one axis may be specified in this way.

If no left argument is given, file is assumed to contain a simple APL array, complete with header information (type, rank, shape, etc).

Mapped files may be updated by changing the associated array using indexed assignment: $var[a] \leftarrow b$.

Note that a *raw* mapped file may be updated *only* if its *file offset* is 0.

Examples

Map raw file as a read-only *vector* of doubles:

```
        vec←645 ¯1 ⎕MAP'c:\myfile'
```

Map raw file as a 20-column read-write *matrix* of 1-byte integers:

```
        mat←83 ¯1 20 ⎕MAP'c:\myfile' 'W'
```

Replace some items in mapped file:

```
        mat[2 3;4 5]←2 2⍴⍳4
```

Map bytes 100-180 in raw file as a 5×2 read-only matrix of doubles:

```
        dat←645 5 2 ⎕MAP'c:\myfile' 'R' 100
```

Put simple 4-byte integer array on disk ready for mapping:

```
        (83 323 ⎕DR 2 3 4⍴⍳24)∆MPUT'c:\myvar'
```

Then, map a read-write variable:

```
        var←⎕MAP'c:\myvar' 'w'
```

Note that a mapped array need not be *named*. In the following example, a 'raw' file is mapped, summed and released, all in a single expression:

```
      +/163 ¯1 ⎕MAP'c:\shorts.dat'
42
```

Migration Level: ⎕ML

⎕ML determines the degree of migration of the Dyalog APL language towards IBM's APL2. Setting this variable to other than its default value of 0 changes the interpretation of certain symbols and language constructs.

⎕ML←0		Native Dyalog (Default)
⎕ML←1	Z←∊R	Monadic '∊' is interpreted as 'enlist' rather than 'type'.
⎕ML←2	Z←↑R	Monadic '↑' is interpreted as 'first' rather than 'mix'.
	Z←⊃R	Monadic '⊃' is interpreted as 'mix' rather than 'first'.
	Z←≡R	Monadic '≡' returns the absolute value of the depth of its argument, rather than a negative value if the depths of its subarrays are unequal
⎕ML←3	R←X⊂[K]Y	Dyadic '⊂' follows the APL2 (rather than the original Dyalog APL) convention
	⎕TC	The order of the elements of ⎕TC is the same as in APL2

Subsequent versions of Dyalog APL will provide further migration levels.

Examples

```
      X←2(3 4)

      □ML←0
      ∊X
0   0 0
      ↑X
2 0
3 4
      ⊃X
2
      ≡X
¯2

      □ML←1
      ∊X
2 3 4
      ↑X
2 0
3 4
      ⊃X
2
      ≡X
¯2

      □ML←2
      ∊X
2 3 4
      ↑X
2
      ⊃X
2 0
3 4
      ≡X
2
```

Set Monitor: $\{R\} \leftarrow X$ `⎕MONITOR Y`

Y must be a simple character scalar or vector which is taken to be the name of a visible defined function or operator. *X* must be a simple non-negative integer scalar or vector. *R* is a simple integer vector of non-negative elements.

X identifies the numbers of lines in the function or operator named by *Y* on which a monitor is to be placed. Numbers outside the range of line numbers in the function or operator (other than 0) are ignored. The number 0 indicates that a monitor is to be placed on the function or operator as a whole. The value of *X* is independent of `⎕IO`.

R is a vector of numbers on which a monitor has been placed in ascending order. The result is suppressed unless it is explicitly used or assigned.

The effect of `⎕MONITOR` is to accumulate timing statistics for the lines for which the monitor has been set. See Monitor Query for details.

Examples

```
      +(0,ι10) ⎕MONITOR 'FOO'
0 1 2 3 4 5
```

Existing monitors are cancelled before new ones are set:

```
      +1 ⎕MONITOR 'FOO'
1
```

All monitors may be cancelled by supplying an empty vector:

```
      θ ⎕MONITOR 'FOO'
```

Monitors may be set on a locked function or operator, but no information will be reported. Monitors are saved with the workspace.

Query Monitor: $R ← \square MONITOR \ Y$

Y must be a simple character scalar or vector which is taken to be the name of a visible defined function or operator. R is a simple non-negative integer matrix of 5 columns with one row for each line in the function or operator Y which has the monitor set, giving:

Column 1 :   Line number
Column 2 :   Number of times the line was executed
Column 3 :   CPU time in milliseconds
Column 4 :   Elapsed time in milliseconds
Column 5 :   Reserved

The value of 0 in column one indicates that the monitor is set on the function or operator as a whole.

Example

```
      ∇ FOO
[1]    A←?25 25ρ100
[2]    B←⌹A
[3]    C←⌹B
[4]    R1←⌊0.5+A+.×B
[5]    R2←A=C
      ∇

      (0,ι5) □MONITOR 'FOO' ⍝ Set monitor

      FOO                   ⍝ Run function

      □MONITOR 'FOO'        ⍝ Monitor query
0 1 1418 1000 0
1 1   83    0 0
2 1  400    0 0
3 1  397    0 0
4 1  467 1000 0
5 1  100    0 0
```

Name Association: $\qquad \{R\}\leftarrow\{X\}\Box NA\ \ Y$

$\Box NA$ provides access from APL to compiled functions within a **Dynamic Link Library** (**DLL**). A DLL is a collection of functions typically written in C (or C++) each of which may take arguments and return a result.

Instructional examples using $\Box NA$ can be found in supplied workspace: *QUADNA.DWS*.

The DLL may be part of the standard operating system software, purchased from a third party supplier, or one that you have written yourself.

The right argument *Y* is a character vector that identifies the name and syntax of the function to be associated. The left argument *X* is a character vector that contains the name to be associated with the external function. If the $\Box NA$ is successful, a function (name class 3) is established in the active workspace with name *X*. If *X* is omitted, the name of the external function itself is used for the association.

The shy result *R* is a character vector containing the name of the external function that was fixed.

For example, math.dll might be a library of mathematical functions containing a function divide. To associate the APL name *div* with this external function:

```
'div' □NA 'F8 math|divide I4 I4'
```

where *F8* and *I4*, specify the types of the result and arguments expected by divide. The association has the effect of establishing a new function: *div* in the workspace, which when called, passes its arguments to divide and returns the result.

```
      )fns
div
      div 10 4
2.5
```

Type Declaration

In a compiled language such as C, the types of arguments and results of functions must be declared explicitly. Typically, these types will be published with the documentation that accompanies the DLL. For example, function divide might be declared:

```
double divide(long int, long int);
```

which means that it expects two long (4-byte) integer arguments and returns a double (8-byte) floating point result. Notice the correspondence between the C declaration and the right argument of ⎕NA:

```
C:            double      divide (long int, long int);

APL: 'div' ⎕NA 'F8   math|divide      I4        I4 '
```

It is imperative that care be taken when coding type declarations. A DLL *cannot* check types of data passed from APL. A wrong type declaration will lead to erroneous results or may even cause the workspace to become corrupted and crash.

The full syntax for the right argument of ⎕NA is:

[result] library|function [arg1] [arg2] ...

Note that functions associated with DLLs are never dyadic. All arguments are passed as items of a (possibly nested) vector on the right of the function.

Locating the DLL

The DLL may be specified using a full pathname, file extension, and function type.

Pathname: If the full pathname is omitted, APL looks for the DLL in 'standard' Windows directories. Specifically, APL uses the LoadLibrary system function to locate the library, the exact workings of which can be found in the appropriate Windows documentation for the Software Development Kit.

Alternatively, a full pathname may be supplied in the usual way:

```
⎕NA'... c:\mydir\mydll|foo ...'
```

Errors: If the specified DLL (or a dependent DLL) fails to load it will generate:

FILE ERROR 1 No such file or directory

If the DLL loads successfully, but the specified library function is not accessible, it will generate:

VALUE ERROR

File Extension: If the file extension is omitted, **.dll** is assumed. Note that some DLLs are in fact **.exe** files, and in this case the extension must be specified explicitly:

> *⎕NA'... mydll.exe|foo ...'*

Function Type: On an Intel machine, code can run in 16 or 32 bit mode. Further, two distinct conventions, namely 'C' and 'Pascal' are in use for passing of arguments and receipt of results. If the type of the function you are calling differs from the default, for your version of Dyalog APL (see below) you must specify the function type explicitly immediately following the DLL name. Combinations are

.C32	32 bit, C calling convention	Default for APL/W version 8.
.C16	16 bit, C calling convention	
.P32	32 bit, Pascal calling convention	
.P16	16 bit, Pascal calling convention	Default for APL/W version 7.

Example

> *⎕NA'... mydll.exe.P16|foo ...' ⍝ 16 bit Pascal*

Call by Ordinal Number

A DLL may associate an *ordinal number* with any of its functions. This number may then be used to call the function as an alternative to calling it by name. Using *⎕NA* to call by ordinal number uses the same syntax but with the function name replaced with its ordinal number. For example:

> *⎕NA'... mydll|57 ...'*

Multi-Threading

Appending the '*&*' character to the function name causes the external function to be run in its own C thread. For example:

> *⎕NA'... mydll|foo& ...'*

This means that other APL threads can run concurrently with the one that is calling the ⎕*NA* function.

Data Type Coding Scheme

The type coding scheme introduced above is of the form:

[direction] [special] type [width] [array]

The options are summarised in the following table and their functions detailed below.

Description	Symbol	Meaning

Direction	`<`
	`>`
	`=`
		Pointer to array *input* to DLL function.
		Pointer to array *output* from DLL function.
		Pointer to input/output array.

Special	`0` `#`	Null-terminated string. Byte-counted string.

Type	`I`
	`U`
	`C`
	`T`
	`F`
	`A`
		int
		unsigned int
		char (untranslated)
		char (translated to/from ANSI)
		float
		APL array

Width	`1`
	`2`
	`4`
	`8`
		1-byte
		2-byte
		4-byte
		8-byte

Array	`[n]` `[]`	Array of length *n* elements Array, length determined at call-time.
Structure	`{...}`	Structure.

Direction

C functions accept data arguments either by *value* or by *address*. This distinction is indicated by the presence of a '`*`' character in the argument declaration:

```
int  num1;          // value of num1 passed.
int *num2;          // Address of num2 passed.
```

An argument (or result) of an external function of type pointer, must be matched in the *⎕NA* call by a declaration starting with one of the characters: `<`, `>`, or `=`.

In C, when an address is passed, the corresponding value can be used as either an *input* or an *output* variable. An output variable means that the C function overwrites values at the supplied address. Because APL is a call-by-value language, and doesn't have pointer types, we accommodate this mechanism by distinguishing output variables, and having them returned explicitly as part of the result of the call.

This means that where the C function indicates a *pointer type*, we must code this as starting with one of the characters: `<`, `>` or `=` .

`<` indicates that the address of the argument will be used by C as an input variable and values at the address will *not* be over-written.

`>` indicates that C will use the address as an output variable. In this case, APL must allocate an output array over which C can write values. After the call, this array will be included in the nested result of the call to the external function.

`=` indicates that C will use the address for both input and output. In this case, APL duplicates the argument array into an output buffer whose address is passed to the external function. As in the case of an output only array, the newly modified copy will be included in the nested result of the call to the external function.

Examples

`<I2`	Pointer to 2-byte integer - *input* to external function
`>C`	Pointer to character *output* from external function.
`=T`	Pointer to translated character *input* to and *output* from function.
`=A`	Pointer to APL array *modified* by function.

Special

In C it is common to represent character strings as *null-terminated* or *byte counted* arrays. These special data types are indicated by inserting the symbol 0 (null-terminated) or # (byte counted) between the direction indicator (<, >, =) and the type (*T* or *C*) specification. For example, a pointer to a null-terminated input character string is coded as `<0T[]`, and an output one coded as `>0T[]`.

Note that while appending the array specifier '`[]`' is formally correct, because the presence of the special qualifier (0 or #) *implies* an array, the '`[]`' may be omitted: `<0T`, `>0T`, `=#C`, etc.

Note also that the 0 and # specifiers may be used with data of all types and widths. For example, `<0I2` may be useful for dealing with Unicode.

Type

The data type of the argument is represented by one of the symbols:

I   **Integer**. The value is interpreted as a 2's complement signed integer.

U   **Unsigned integer**. The value is interpreted as an unsigned integer.

C   **Untranslated character**. The value is interpreted as a character and passed untranslated. This means that □*AV positions* map onto corresponding ANSI *positions*. For example, with □*IO*=0:

   □*AV*[35] = '*s*',   maps to ANSI[35] = '#'.

T   **Translated character**. The value is interpreted as a character, and translated using standard Dyalog □*AV* to ANSI translation. This means that □*AV characters* map onto corresponding ANSI *characters*. For example, with □*IO*=0:

   □*AV*[35] = '*s*',   maps to ANSI[115] = 's'.

F   **Floating point number**. The value is interpreted as an IEEE floating point number.

A   **APL array**. A pointer to the whole array (including header information) is passed. This type is used to communicate with DLL functions which have been written specifically to work with Dyalog APL. See the *User Guide* section on Writing Auxiliary Processors. Note that type A is always passed as a pointer, so is of the form `<A`, `=A` or `>A`.

   Note that the type declaration may be given in **lower case**: *i*, *u*, *c*, *t*, *f*, *a*.

Width

The type specifier may be followed by the width of the value in bytes. For example:

*I*4   4-byte signed integer.
*U*2   2-byte unsigned integer.
*F*8   8-byte floating point number.
*F*4   4-byte floating point number.

Type	Possible values for Width	Default value for Width
I	1, 2, 4.	2 for 16-bit DLLs 4 for 32-bit DLLs
U	1, 2, 4.	2 for 16-bit DLLs 4 for 32-bit DLLs
C	1	1
T	1	1
F	4, 8.	8
A	Not applicable.	

Examples

 *I*2   16-bit integer
<*I*4   Pointer to input 4-byte integer
 U   Default width unsigned integer.
=*F*4   Pointer to input/output 4-byte floating point number.

Arrays

Arrays are specified by following the basic data type with `[n]` or `[]`, where *n* indicates the number of elements in the array. In the C declaration, the number of elements in an array may be specified explicitly at compile time, or determined dynamically at runtime. In the latter case, the size of the array is often passed along with the array, in a separate argument. In this case, *n*, the number of elements is omitted from the specification. Note that C deals only in scalars and rank 1 (vector) arrays.

```
int vec[10];            // explicit vector length.
unsigned size, list[];  // undetermined length.
```

would be coded as:

`I[10]`   vector of 10 ints.
`U U[]`   unsigned integer followed by an array of unsigned integers.

Confusion sometimes arises over a difference in the declaration syntax between C and `⎕NA`. In C, an argument declaration may be given to receive a pointer to either a single scalar item, or to the first element of an array. This is because in C, the address of an array is deemed to be the address of its first element.

```
void foo (char *string);

char ch = 'a', ptr = "abc";

foo(&ch);            // call with address of scalar.
foo(ptr);            // call with address of array.
```

However, from APL's point of view, these two cases are distinct and if the function is to be called with the address of (pointer to) a *scalar*, it must be declared: `'<T'`. Otherwise, to be called with the address of an *array*, it must be declared: `'<T[]'`. Note that it is perfectly acceptable in such circumstances to define more than one name association to the same DLL function specifying different argument types:

```
'FooScalar'⎕NA'mydll|foo <T'    ◇ FooScalar'a'
'FooVector'⎕NA'mydll|foo <T[]'  ◇ FooVector'abc'
```

Structures

Arbitrary data structures, which are akin to nested arrays, are specified using the symbols `{}`. For example, the code `{F8 I2}` indicates a structure comprised of an 8-byte *float* followed by a 2-byte *int*. Furthermore, the code `<{F8 I2}[3]` means an input pointer to an array of 3 such structures.

For example, this structure might be defined in C thus:

```
typedef struct
{
      double  f;
      short   i;
} mystruct;
```

A function defined to receive a count followed by an *input* pointer to an array of such structures:

```
void foo(unsigned count, mystruct *str);
```

An appropriate `⎕NA` declaration would be:

```
⎕NA'mydll.foo U <{F8 I2}[]'
```

A call on the function with two arguments - a count followed by a vector of structures:

```
foo 4,⊂(1.4 3)(5.9 1)(6.5 2)(0 0)
```

Notice that for the above call, APL converts the two Boolean (`0 0`) elements to an 8-byte float and a 2-byte int, respectively.

Specifying Pointers Explicitly

⎕NA syntax enables APL to pass arguments to DLL functions by *value* or *address* as appropriate. For example if a function requires an integer followed by a *pointer* to an integer:

```
void fun(int valu, int *addr);
```

You might declare and call it:

> *⎕NA'mydll|fun I <I' ◇ fun* 42 42

The interpreter passes the *value* of the first argument and the *address* of the second one.

Two common cases occur where it is necessary to pass a pointer explicitly. The first is if the DLL function requires a *null pointer*, and the second is where you want to pass on a pointer which itself is a result from a DLL function.

In both cases, the pointer argument should be coded as *I4*. This causes APL to pass the pointer unchanged, *by value*, to the DLL function.

In the previous example, to pass a null pointer, (or one returned from another DLL function), you must code a separate *⎕NA* definition.

> *'fun_null'⎕NA'mydll|fun I I4' ◇ fun_null* 42 0

Now APL passes the *value* of the second argument (in this case 0 - the null pointer), rather than its address.

Using a Function

A DLL function may or may not return a result, and may take zero or more arguments. This syntax is reflected in the coding of the right argument of $\square NA$. Notice that the corresponding associated APL function is niladic or monadic (never dyadic), and that it *always* returns a vector result - a null one if there is no output from the function. See Result Vector section below. Examples of the various combinations are:

DLL function Non-result-returning:

```
⎕NA     'mydll|fn1'          ⍝ Niladic.
⎕NA     'mydll|fn2 <0T'      ⍝ Monadic - 1-element argument
.
⎕NA     'mydll|fn3 =0T <0T'  ⍝ Monadic - 2-element argument
.
```

DLL function Result-returning:

```
⎕NA 'I4 mydll|fn4'           ⍝ Niladic.
⎕NA 'I4 mydll|fn5 F8'        ⍝ Monadic - 1-element argument
.
⎕NA 'I4 mydll|fn6 >I4[] <0T'⍝ Monadic - 2-element argument
.
```

When the external function is called, the number of elements in the argument must match the number defined in the $\square NA$ definition. Using the example functions defined above:

```
        fn1                  ⍝ Niladic function.
        fn2 ⊂'Single String' ⍝ 1-element argument.
        fn3 'This' 'That'    ⍝ 2-element argument.
```

Note in the second example, that you must enclose the argument string to produce a single item (nested) array in order to match the declaration. Dyalog converts the type of a numeric argument if necessary, so for example in fn5 defined above, a Boolean value would be converted to double floating point (F8) prior to being passed to the DLL function.

Pointer Arguments

When passing pointer arguments there are three cases to consider.

< **Input pointer:** In this case you must supply the data array itself as argument to the function. A pointer to its first element is then passed to the DLL function.

```
      fn2 ⊂'hello'
```

> **Output pointer:** Here, you must supply the **number of elements** that the output will need in order for APL to allocate memory to accommodate the resulting array.

```
      fn6 10 'world'  ⍝ 1st arg needs space for 10 ints.
```

Note that if you were to reserve fewer elements than the DLL function actually used, the DLL function would write beyond the end of the reserved array and may cause the interpreter to crash with a System Error.

= **Input/Output:** As with the input-only case, a pointer to the first element of the argument is passed to the DLL function. The DLL function then overwrites some or all of the elements of the array, and the new value is passed back as part of the result of the call. As with the output pointer case, if the input array were too short, so that the DLL wrote beyond the end of the array, the interpreter would almost certainly crash.

```
      fn3 '.....' 'hello'
```

Result Vector

In APL, a function cannot overwrite its arguments. This means that any output from a DLL function must be returned as part of the explicit result, and this includes output via 'output' or 'input/output' pointer arguments.

The general form of the result from calling a DLL function is a nested vector. The first item of the result is the defined explicit result of the external function, and subsequent items are implicit results from output, or input/output pointer arguments.

The length of the result vector is therefore: 1 (if the function was declared to return an explicit result) + the number of output or input/output arguments.

⎕NA **Declaration**	**Result**	**Output Arguments**	**Result Length**
mydll|fn1	0		0
mydll|fn2 <0T	0	0	0
mydll|fn3 =0T <0T	0	1 0	1
I4 mydll|fn4	1		1
I4 mydll|fn5 F8	1	0	1
I4 mydll|fn6 >I4[] <0T	1	1 0	2

As a convenience, if the result would otherwise be a 1-item vector, it is disclosed. Using the third example above:

        ρfn3 '.....' 'abc'
5

fn3 has no explicit result; its first argument is input/output pointer; and its second argument is input pointer. Therefore as the length of the result would be 1, it has been disclosed.

ANSI versus Unicode Characters

There is movement within Microsoft to support 16-bit Unicode (or "wide") characters. To this end, some library functions that take character arguments, or return character results have been replaced with two new functions: one Wide and one ANSI. For example, a function such as MessageBox, will have been replaced with MessageBoxA and MessageBoxW. The 'A' stands for 'ANSI' (1-byte) characters, and the 'W' for 'wide' (2-byte) characters. Dyalog APL currently supports only the 1-byte version. Unfortunately, Windows documentation for many of these DLLs omits the suffix, so that, for example, when you try to associate `MessageBox`, you get a *VALUE ERROR*. In this case, try repeating the ⎕*NA* call with an 'A' appended to the function name.

Type Definitions (typedefs)

The C language encourages the assignment of defined names to primitive and complex data types using its `#define` and `typedef` mechanisms. Using such abstractions enables the C programmer to write code that will be portable across many operating systems and hardware platforms.

Windows software uses many such names and Microsoft documentation will normally refer to the type of function arguments using defined names such as `HANDLE` or `LPSTR` rather than their equivalent C primitive types: `int` or `char*`.

It is beyond the scope of this manual to list *all* the Microsoft definitions and their C primitive equivalents, and indeed, DLLs from sources other than Microsoft may well employ their own distinct naming conventions.

In general, you should consult the documentation that accompanies the DLL in order to convert typedefs to primitive C types and thence to ⎕*NA* declarations. The documentation may well refer you to the 'include' files which are part of the Software Development Kit, and in which the types are defined.

The following table of some commonly encountered Windows typedefs and their ⎕*NA* equivalents might prove useful.

Windows typedef	*□NA* equivalent
HWND	*I*
HANDLE	*I*
GLOBALHANDLE	*I*
LOCALHANDLE	*I*
DWORD	*U*4
WORD	*U*2
BYTE	*U*1
LPSTR	=0*T*[ ] (note 1)
LPCSTR	<0*T*[ ] (note 2)
WPARAM	*U*
LPARAM	*U*4
LRESULT	*I*4
BOOL	*I*
UINT	*U*
ULONG	*U*4
ATOM	*U*2
HDC	*I*
HBITMAP	*I*
HBRUSH	*I*
HFONT	*I*
HICON	*I*
HMENU	*I*
HPALETTE	*I*
HMETAFILE	*I*
HMODULE	*I*
HINSTANCE	*I*
COLORREF	{*U*1[4]}
POINT	{*I I*}
POINTS	{*I*2 *I*2}
RECT	{*I I I I*}
CHAR	*T* or *C*

Notes

1. LPSTR is a pointer to a null-terminated string. The definition does not indicate whether this is input or output, so the safest coding would be =0*T*[ ] (providing the vector you supply for input is long enough to accommodate the result). You may be able to improve simplicity or performance if the documentation indicates that the pointer is 'input only' (<0*T*[ ]) or 'output only' (>0*T*[ ]). See **Direction** above.

2. `LPCSTR` is a pointer to a *constant* null-terminated string and therefore coding `<0T[]` is safe.

Dyalog32.dll & Dyalog16.dll

Included with Dyalog APL is a utility DLL which is called dyalog32.dll in version 8
and dyalog16.dll in version 7.

The DLL contains two functions: MEMCPY and STRNCPY.

MEMCPY

MEMCPY is an extremely versatile function used for moving arbitrary data between
memory buffers.
Its C definition is:

```
void MEMCPY(              // copy memory
     void *to,            // target address
     void *fm,            // source address
     unsigned size        // number of bytes to copy
);
```

MEMCPY copies `size` bytes starting from source address `fm`, to destination address
`to`. If the source and destination areas overlap, the result is undefined.

MEMCPY's versatility stems from being able to associate to it using many different type
declarations.

Example

Suppose a global buffer (at address: *addr*) contains (*numb*) double floating point
numbers. To copy these to an APL array, we could define the association:

> *'doubles' ⎕NA 'dyalog32|MEMCPY >F8[] I4 U4'*
> *doubles numb addr (numb×8)*

Notice that:

As the first argument to *doubles* is an output argument, we must supply the number
of elements to reserve for the output data.

MEMCPY is defined to take the number of *bytes* to copy, so we must multiply the
number of elements by the element size in bytes.

Example

Suppose that a database application requires that we construct a record in global memory prior to writing it to file. The record structure might look like this:

```
typedef struct {
      int empno;          // employee number.
      float salary;       // salary.
      char name[20];      // name.
} person;
```

Then, having previously allocated memory (*addr*) to receive the record, we can define:

```
        'prec' ⎕NA 'dyalog32|MEMCPY I4 <{I4 F4 T[20]} U4'
        prec addr(99 12345.60 'Charlie Brown        ')(4+4+20
)
```

STRNCPY

STRNCPY is used to copy null-terminated strings between memory buffers.
Its C definition is:

```
void STRNCPY(            // copy null-terminated string
      char *to,          // target address
      char *fm,          // source address
      unsigned size      // MAX number of chars to copy
);
```

STRNCPY copies a maximum of `size` characters from the null-terminated source string at address `fm`, to the destination address `to`. If the source and destination strings overlap, the result is undefined.

If the source string is shorter than `size`, null characters are appended to the destination string.

If the source string (including its terminating null) is longer than `size`, only `size` characters are copied and the resulting destination string is not null-terminated

Example

Suppose that a database application returns a pointer (*addr*) to a structure that contains two pointers to (max 20-char) null-terminated strings.

```
typedef struct {     // null-terminated strings:
      char *first;  // first name (max 19 chars + 1 null).
      char *last;   // last name. (max 19 chars + 1 null).
} name;
```

To copy the names *from* the structure:

```
      'get'⎕NA'dyalog32|STRNCPY >0T[] I4 U4'
      get 20 addr 20
Charlie
      get 20 (addr+4) 20
Brown
```

To copy data *from* the workspace *into* an already allocated (*new*) structure:

```
      'put'⎕NA'dyalog32|STRNCPY I4 <0T[] U4'
      put new 'Bo' 20
      put (new+4) 'Peep' 20
```

Notice in this example that you must ensure that names no longer than 19 characters are passed to *put*. More than 19 characters would not leave STRNCPY enough space to include the trailing null, which would probably cause the application to fail.

Examples

The following examples all use functions from the Microsoft Windows 'USER'. DLL.

The file name is different for different versions of Dyalog APL:

Version 7:   USER16.EXE
Version 8:   USER32.DLL

On all three systems, the DLL is should be found in a 'standard' place, and so you shouldn't normally need to give the full path name of the library. However if trying these examples results in the error message 'FILE ERROR 1 No such file or directory', you must locate the DLL and supply the full path name (and possibly extension).

The first example will show all 3 versions of the $\square NA$ call, but to avoid confusion, subsequent ones will assume you are using Windows 95/98. If you are using one of the other operating systems, you must make the necessary adjustments to the call.

Note that in the following examples, DLL functions ending in 'A' are specific to Windows 95/98/NT. If you are using Windows 3.x, you should remove the 'A' from the function name.

Example 1

The Windows function "GetCaretBlinkTime" retrieves the caret blink rate. It takes no arguments and returns an unsigned *int* and is declared as follows:

```
UINT GetCaretBlinkTime(void);
```

The following statements would provide access to this routine through an APL function of the same name. Assuming you are using Windows 95/98.

```
      ⎕NA 'U User32|GetCaretBlinkTime'
      GetCaretBlinkTime
530
```

Using windows 3.x:

```
      ⎕NA 'U User16.exe|GetCaretBlinkTime'
```

Using Windows NT, note that as the file name is User32.DLL, there is no need to specify the file extension.:

```
      ⎕NA 'U User32|GetCaretBlinkTime'
```

The following statement (in Windows 95/98) would achieve the same thing, but using an APL function called *BLINK*.

```
      'BLINK' ⎕NA 'U User32|GetCaretBlinkTime'
      BLINK
530
```

Example 2

The Windows function "SetCaretBlinkTime" sets the caret blink rate. It takes a single unsigned *int* argument, does not return a result and is declared as follows:

```
void SetCaretBlinkTime(UINT);
```

The following statements would provide access to this routine through an APL function of the same name :

```
      ⎕NA 'User32|SetCaretBlinkTime U'
      SetCaretBlinkTime 1000
```

Example 3

The Windows function "MessageBoxA" displays a standard dialog box on the screen and awaits a response from the user. It takes 4 arguments. The first is the window handle for the window that owns the message box. This is declared as an unsigned *int*. The second and third arguments are both pointers to null-terminated strings containing the message to be displayed in the Message Box and the caption to be used in the window title bar. The 4th argument is an unsigned *int* that specifies the Message Box type. The result is an *int* which indicates which of the buttons in the message box the user has pressed. The function is declared as follows:

```
int MessageBoxA(HWND, LPCSTR, LPCSTR, UINT);
```

The following statements provide access to this routine through an APL function of the same name. Note that the 2nd and 3rd arguments are both coded as input pointers to **translated** null-terminated character arrays. This forces the message and caption to go through the standard Windows translation, so that the correct ASCII characters are shown on the screen.

> ⎕*NA 'I User*32|*MessageBoxA U <0T <0T U'*

The following statement displays a Message Box with a stop sign icon together with 2 push buttons labelled OK and Cancel (this is specified by the value 19).

> *MessageBoxA* 0 *'Message' 'Title'* 19

Note that a simpler, portable (and safer) method for displaying a Message Box is to use Dyalog APL's primitive *MsgBox* object.

Example 4

The Windows function "`FindWindowA`" obtains the window handle of a window which has a given character string in its title bar. The function takes two arguments. The first is a pointer to a null-terminated character string that specifies the window's class name. However, if you are not interested in the class name, this argument should be a NULL pointer. The second is a pointer to a character string that specifies the title that identifies the window in question. This is an example of a case described above where two instances of the function must be defined to cater for the two different types of argument. However, in practice this function is most often used without specifying the class name. The function is declared as follows:

```
HWND FindWindowA(LPCSTR, LPCSTR);
```

The following statement associates the APL function *FW* with the second variant of the FindWindow call, where the class name is specified as a NULL pointer. To indicate that APL is to pass the *value* of the NULL pointer, rather than its address, we need to code this argument as *I4*.

```
        'FW' ⎕NA 'U User32|FindWindowA I4 <0T'
```

To obtain the handle of the window entitled "CLEAR WS - Dyalog APL/W":

```
        ⎕←HNDL←FW 0 'CLEAR WS - Dyalog APL/W'
59245156
```

Example 5

The Windows function "`GetWindowTextA`" retrieves the caption displayed in a window's title bar. It takes 3 arguments. The first is an unsigned *int* containing the window handle. The second is a pointer to a buffer to receive the caption as a null-terminated character string. This is an example of an output array. The third argument is an *int* which specifies the maximum number of characters to be copied into the output buffer. The function returns an *int* containing the actual number of characters copied into the buffer and is declared as follows:

```
int GetWindowTextA(HWND, LPSTR, int);
```

The following associates the "GetWindowText" DLL function with an APL function of the same name. Note that the second argument is coded as "`>0T`" indicating that it is a pointer to a character output array.

```
      ⎕NA 'I User32|GetWindowTextA U >0T I'
```

Now change the Session caption using `)WSID`:

```
      )WSID MYWS
was CLEAR WS
```

Then retrieve the new caption (max length 255) using window handle `HNDL` from the previous example:

```
      DISPLAY GetWindowTextA HNDL 255 255
.→------------------------.
|     .→-----------------. |
| 19 |MYWS - Dyalog APL/W| |
|     '-----------------' |
'∈------------------------'
```

There are three points to note. Firstly, the number 255 is supplied as the second argument. This instructs APL to allocate a buffer large enough for a 255-element character vector into which the DLL routine will write. Secondly, the result of the APL function is a nested vector of 2 elements. The first element is the result of the DLL function. The second element is the output character array. Finally, notice that although we reserved space for 255 elements, the result reflects the length of the actual text (19).

An alternative way of coding and using this function is to treat the second argument as an input/output array.

e.g.

```
      ⎕NA 'I User32|GetWindowTextA U =0T I'

      DISPLAY GetWindowTextA HNDL (255ρ' ') 255
.→------------------------.
|     .→-----------------. |
| 19 |MYWS - Dyalog APL/W| |
|     '-----------------' |
'∈------------------------'
```

In this case, the second argument is coded as `=0T`, so when the function is called an array of the appropriate size must be supplied. This method uses more space in the

workspace, although for small arrays (as in this case) the real impact of doing so is negligible.

Example 6

The function "GetCharWidthA" returns the width of each character in a given range
Its first argument is a device context (handle). Its second and third arguments specify
font positions (start and end). The third argument is the resulting integer vector that
contains the character widths (this is an example of an output array). The function
returns a boolean value to indicate success or failure. The function is defined as
follows. Note that this function is provided in the library: GDI32.DLL.

```
BOOL GetCharWidthA(HDC, UINT, UINT, int FAR*);
```

The following statements provide access to this routine through an APL function of the
same name :

```
      ⎕NA 'U4 GDI32|GetCharWidthA I U U >I[]'

     'P'⎕WC'Printer'

      DISPLAY GetCharWidthA ('P' ⎕WG 'Handle') 65 67 3
.→-------------.
|    .→-------. |
| 1 |50 50 50| |
|   '~-------' |
'∊-------------'
```

Note: 'P'⎕WG'Handle' returns a 32-bit handle which, if the *top bit* is set, will
appear in APL as a negative integer. Attempting to supply such a negative number as
an argument to a DLL function when the argument is declared *unsigned* will result in a
DOMAIN ERROR. Window handles should therefore be declared as *I* rather than *U*.

Example 7

The following example from the supplied workspace: *QUADNA.DWS* illustrates several
techniques which are important in advanced □*NA* programming. Function
DllVersion returns the major and minor version number for a given DLL.

In advanced DLL programming, it is often necessary to administer memory outside
APL's workspace. In general, the procedure for such use is:

1.   Allocate global memory.
2.   Lock the memory.
3.   Copy any DLL input information from workspace into memory.
4.   Call the DLL function.
5.   Copy any DLL output information from memory to workspace.
6.   Unlock the memory.
7.   Free the memory.

Notice that steps 1 and 7, and steps 2 and 6 complement each other. That is, if you
allocate global system memory, you must free it after you have finished using it. If you
continue to use global memory without freeing it, your system will gradually run out of
resources. Similarly, if you lock memory (which you must do before using it), then you
should unlock it before freeing it. Although on some versions of Windows, freeing the
memory will include unlocking it, in the interests of good style, maintaining the
symmetry is probably a good thing.

```
      ∇ version←DllVersion file;Alloc;Free;Lock;Unlock;Size
                ;Info;Value;Copy;size;hndl;addr;buff;ok
[1]
[2]    'Alloc'⎕NA'U4 kernel32|GlobalAlloc U4 U4'
[3]     'Free'⎕NA'U4 kernel32|GlobalFree U4'
[4]     'Lock'⎕NA'U4 kernel32|GlobalLock U4'
[5]   'Unlock'⎕NA'U4 kernel32|GlobalUnlock U4'
[6]
[7]    'Size'⎕NA'U4 version|GetFileVersionInfoSizeA <0T >U4'
[8]    'Info'⎕NA'U4 version|GetFileVersionInfoA<0T U4 U4 U4'
[9]   'Value'⎕NA'U4 version|VerQueryValueA U4 <0T >U4 >U4'
[10]
[11]  'Copy'⎕NA'dyalog32|MEMCPY >U4[] U4 U4'
[12]
[13]   :If ×size←⊃Size file 0              ⍝ Size of info
[14]   :AndIf ×hndl←Alloc 0 size           ⍝ Alloc memory
[15]     :If ×addr←Lock hndl               ⍝ Lock memory
[16]       :If ×Info file 0 size addr      ⍝ Version info
[17]         ok buff size←Value addr'\' 0 0 ⍝ Version value
[18]         :If ok
[19]           buff←Copy(size÷4)buff size  ⍝ Copy info
[20]           version←(2/2*16)⊤⊃2↓buff    ⍝ Split version
[21]         :EndIf
[22]       :EndIf
[23]       ok←Unlock hndl                  ⍝ Unlock memory
[24]     :EndIf
[25]     ok←Free hndl                      ⍝ Free memory
[26]   :EndIf
      ∇
```

Lines [2-11] associate APL function names with the DLL functions that will be used.

Lines [2-5] associate functions to administer global memory.
Lines [7-9] associate functions to extract version information from a DLL.
Line[11] associates *Copy* with *MEMCPY* function from **dyalog32.dll**.

Lines [13-26] call the DLL functions.

Line [13] requests the size of buffer required to receive version information for the DLL. A size of 0 will be returned if the DLL does not contain version information.

Notice that care is taken to balance memory allocation and release:

On line [14], the :If clause is taken only if the global memory allocation is successful, in which case (and only then) a corresponding Free is called on line [25].

Unlock on line[23] is called if and only if the call to Lock on line [15] succeeds.

A result is returned from the function *only* if all the calls are successful Otherwise, the calling environment will sustain a *VALUE ERROR*.

More Examples

```
⎕NA'U4 ADVAPI32 |RegCloseKey        U4'
⎕NA'I  ADVAPI32 |RegCreateKeyExA    U <0T I <0T I I I >U >U'
⎕NA'U  ADVAPI32 |RegEnumValueA      U U >0T =U U >U >0T =U'
⎕NA'I  ADVAPI32 |RegOpenKeyA        U <0T >U'
⎕NA'U4 ADVAPI32 |RegOpenKeyExA      U4 <0T U4 U4 =U4'
⎕NA'I  ADVAPI32 |RegQueryValueExA   U <0T U4 >U4 >0T =I4'
⎕NA'I4 ADVAPI32 |RegSetValueExA     U <0T U4 U4 <0T U4'
⎕NA'I  DOS_U32  |Copy               <0T <0T'
⎕NA'U  DOS_U32  |Dir                <0T U >0T'
⎕NA'U  DOS_U32  |DirMore            U >0T'
⎕NA'   DOS_U32  |DirClose'
⎕NA'I  DOS_U32  |Rename             <0T <0T'
⎕NA'I  DOS_U32  |Erase              <0T'
⎕NA'   dyalog32 |STRNCPY            >0T I4 U4'
⎕NA'   dyalog32 |MEMCPY             >{U1[4]}[16] I4 U4'
⎕NA'I  gdi32    |AddFontResourceA   <0T'
⎕NA'   GDI32    |BitBlt             U U U U U U U U U'
⎕NA'U4 gdi32    |GetPixel           U4 U4 U4'
⎕NA'U4 GDI32    |GetStockObject     U4'
⎕NA'I  gdi32    |RemoveFontResourceA <0T'
⎕NA'U4 gdi32    |SetPixel           U4 U4 U4 U4'
⎕NA'U4 GLU32    |gluPerspective     F8 F8 F8 F8'
⎕NA'I  kernel32 |_lclose            U'
⎕NA'I  kernel32 |_lcreat            <0T I'
⎕NA'U4 kernel32 |_llseek            I U4 I'
⎕NA'I  kernel32 |_lopen             <0T I'
⎕NA'U  kernel32 |_lread             U >U1[] U'
⎕NA'U  kernel32 |_lwrite            U <U1[] U'
⎕NA'I  kernel32 |CopyFileA          <0T <0T I'
⎕NA'I4 kernel32 |GetEnvironmentStrings'
⎕NA'I4 kernel32 |GetLastError'
⎕NA'I  kernel32 |GetPrivateProfileIntA <0T <0T I <0T'
⎕NA'I  kernel32 |GetProfileStringA <0T <0T <0T >0T I'
⎕NA'U2 kernel32 |GetSystemDirectoryA >0T U2'
⎕NA'I2 KERNEL32 |GetTempPathA       U4 >0T'
⎕NA'U  KERNEL32 |GetWindowsDirectoryA >0T U'
⎕NA'U  kernel32 |GlobalAlloc        U U4'
⎕NA'U  kernel32 |GlobalFree         U'
⎕NA'   Kernel32 |GlobalMemoryStatus ={U4 U4 U4 U4 U4 U4 U4 U4}'
⎕NA'U  KERNEL32 |WritePrivateProfileStringA <0T <0T <0T <0T'
⎕NA'U4 OpenGL32 |glClearColor       F4 F4 F4 F4'
⎕NA'U4 OpenGL32 |glClearDepth       F4'
⎕NA'U4 OpenGL32 |glEnable           U4'
⎕NA'U4 OpenGL32 |glMatrixMode       U4'
⎕NA'U4 USER32   |ClientToScreen     U =U4[2]'
⎕NA'I  user32   |FindWindowA        I4 <0T'
⎕NA'   user32   |ShowWindow         I I'
⎕NA'I2 USER32   |GetAsyncKeyState   I'
⎕NA'U4 user32   |GetDC              U4'
⎕NA'I4 User32   |GetDialogBaseUnits'
⎕NA'I  user32   |GetFocus'
⎕NA'U4 user32   |GetSysColor        I'
⎕NA'I  user32   |GetSystemMetrics   I'
⎕NA'   USER32   |InvalidateRgn      I4 U4 U4'
⎕NA'I  user32   |MessageBoxA        I <0T <0T I'
```

```
⎕NA'U4 user32    |ReleaseDC        U4 U4'
⎕NA'U4 USER32    |SendMessageA     I4 U4 U4 <I[]'
⎕NA'I2 user32    |SetFocus         I'
⎕NA'I  user32    |WinHelpA         I <0T I I4'
⎕NA'U  WINMM     |sndPlaySoundA    <0T U'
```

Native File Append: $\{R\}\leftarrow X \;\; \Box NAPPEND \;\; Y$

This function appends the ravel of its left argument X to the end of the designated native file. X must be a simple homogeneous APL array. Y is a 1- or 2-element integer vector. $Y[1]$ is a negative integer that specifies the tie number of a native file. The optional second element $Y[2]$ specifies the data type to which the array X is to be converted before it is written to the file.

The shy file index result returned is the position within the file of the end of the record, which is also the start of the following one.

Name Classification: $R\leftarrow \Box NC \;\; Y$

Y must be a simple character scalar, vector or matrix. A scalar or vector is treated as a one-row matrix. Each (implied) row of Y is taken to be an APL name. R is a simple integer vector of shape $1\uparrow ^-2\uparrow 1\;\; 1,\rho Y$.

Each element of R is the name class of the active referent of the object named in the corresponding row of Y. A name class may be:

$^-1$	invalid name
0	unused name
1	label
2	variable
3	function
4	operator
9	GUI object/namespace

Example

```
      A←ι5

      □NC 2 3ρ'A  FOO'
2 3
```

Symbols which represent primitive functions or operators, or system constants, variables, namespaces and functions, are classified as invalid names.

Examples

```
      □NC 3 1ρ'∼+*'
¯1 ¯1 ¯1

      □NC 3 3ρ'□CR□DL□TS'
¯1 ¯1 ¯1
```

The name class of an external variable is 2 as soon as it is associated by `⎕XT` even though it may not have a value.

Example

```
      ⎕NC'V'
0

      'EXT/ARRAY' ⎕XT 'V'

      ⎕NC'V'
2

      V
VALUE ERROR
      V
      ^
```

The name class of the left argument in an ambivalent function, may be 0 or 2 dependent on whether the function is called monadically or dyadically.

Example

```
      ∇ R←{A} PLUS B
[1]    →L1ρ⍨0≠⎕NC'A' ◇ A←0
[2]   L1:R←A+B
      ∇

      PLUS 2
2

      2 PLUS 2
4
```

Native File Create: $\{R\}\leftarrow X$ $\square NCREATE$ Y

This function creates a new file. Under DOS/Windows the file is opened in compatibility mode. The name of the new file is specified by the left argument X which must be a simple character vector or scalar containing a valid pathname for the file. Y is 0 or a negative integer value that specifies an (unused) tie number by which the file may subsequently be referred.

The shy result of $\square NCREATE$ is the tie number of the new file.

Automatic Tie Number Allocation

A tie number of 0 as argument to a create or tie operation, allocates, and returns as an explicit result, the first (closest to zero) available tie number. This allows you to simplify code. For example:

from:

```
tie←¯1+⌊/0,□NNUMS    ⍝ With next available number,
file □NCREATE tie    ⍝ ... create file.
```

to:

```
tie←file □NCREATE 0 ⍝ Create with first available no
```

.

Native File Erase: $\{R\}\leftarrow X$ $\square NERASE$ Y

This function erases (deletes) a native file. Y is a negative integer tie number associated with a tied native file. X is a simple character vector or scalar containing the name of the same file and must be **identical** to the name used when it was opened by $\square NCREATE$ or $\square NTIE$.

The shy result of $\square NERASE$ is the tie number that the erased file had.

Example

```
file □nerase file □ntie 0
```

Name List: $R \leftarrow \{X\} \square NL \ Y$

Y must be a simple integer scalar or vector containing only 1, 2, 3, 4 or 9. X is optional. If present, it must be a simple character scalar or vector. R is a simple character matrix of names of active objects whose name class is included in Y (see the system function $\square NC$) in standard sorted order.

If X is supplied, R contains only those names which begin with any character of X. Standard sorted order is the collation order of $\square AV$.

Examples

```
      □NL 2 3
A
FAST
FIND
FOO
V

      'AV' □NL 2 3
A
V
```

Native File Lock: $\{R\} \leftarrow X \ \square NLOCK \ Y$

This function assists the controlled update of shared native files by locking a range of bytes.

Locking enables controlled update of native files by co-operating users. A process requesting a lock on a region of a file will be *blocked* until that region becomes available. A *write-lock* is exclusive, whereas a *read-lock* is shared. In other words, any byte in a file may be in one of only three states:

- Unlocked
- Write-locked by exactly one process.
- Read-locked by any number of processes.

Y must be a simple integer scalar or vector containing 1, 2 or 3 items namely:
1. Tie number
2. Offset (from 0) of first byte of region. Defaults to 0
3. Number of bytes to lock. Defaults to maximum possible file size

X is optional. If present, it must be a simple integer scalar or vector containing 1 or 2 items, namely:
1. Type: 0: Unlock, 1:Read lock, 2:Write lock.
2. Timeout: Number of seconds to wait for lock until TIMEOUT error. Defaults to indefinite wait.

The shy result *R* is *Y*. To unlock the file, this value should subsequently be supplied in the right argument to 0⎕*NLOCK*.

Examples:

```
2 ⎕NLOCK ¯1        ⍝ write-lock whole file
0 ⎕NLOCK ¯1        ⍝ unlock whole file.
1 ⎕NLOCK ¯1        ⍝ read (share) lock whole file.
2 ⎕NLOCK¨⎕NNUMS    ⍝ write-lock all files.
0 ⎕NLOCK¨⎕NNUMS    ⍝ unlock all files.

1 ⎕NLOCK ¯1 12 1   ⍝ read-lock byte 12.
1 ⎕NLOCK ¯1 0 10   ⍝ read-lock first 10 bytes.
2 ⎕NLOCK ¯1 20     ⍝ write-lock from byte 20 onwards.
2 ⎕NLOCK ¯1 10 2   ⍝ write-lock 2 bytes from byte 10
0 ⎕NLOCK ¯1 12 1   ⍝ remove lock from byte 12.
```

To lock the region immediately beyond the end of the file prior extending it:

```
region←2 ⎕NLOCK ¯1, ⎕NSIZE ¯1  ⍝ write-lock from EOF.
... ⎕NAPPEND ¯1                ⍝ append bytes to file
... ⎕NAPPEND ¯1                ⍝ append bytes to file
0 ⎕NLOCK ¯1,region             ⍝ release lock.
```

The left argument may have a second optional item that specifies a *timeout* value. If a lock has not been acquired within this number of seconds, the acquisition is abandoned and a *TIMEOUT* error reported.

```
2 10 ⎕nlock ¯1     ⍝ wait up to 10 seconds for lock.
```

Notes:

There is no *per-byte* cost associated with region locking. It takes the same time to lock/unlock a region, irrespective of that region's size.

Different file servers implement locks in slightly different ways. For example on some systems, locks are *advisory*. This means that a write lock on a region precludes other locks intersecting that region, but doesn't stop reads or writes across the region. On the other hand, *mandatory* locks block both other locks *and* read/write operations. ⎕*NLOCK* will just pass the server's functionality along to the APL programmer without trying to standardise it across different systems.

All locks on a file will be removed by *⎕NUNTIE*.

Blocked locking requests can be freed by a strong interrupt. Under Windows, this operation is performed from the Dyalog APL pop-up menu in the system tray.

Errors

In this release, an attempt to unlock a region that contains bytes that have not been locked results in a *DOMAIN* error.

A *LIMIT ERROR* results if the operating system lock daemon has insufficient resources to honour the locking request.

Some systems support only write locks. In this case an attempt to set a read lock will generate a DOMAIN ERROR, and it may be appropriate for the APL programmer to trap the error and apply a write lock.

No attempt will be made to detect deadlock. Some servers do this and if such a condition is detected, a DEADLOCK error (1008) will be reported.

Native File Names: *R←⎕NNAMES*

This niladic function reports the names of all currently open native files. *R* is a character matrix. Each row contains the name of a tied native file padded if necessary with blanks. The names are **identical** to those that were given when opening the files with *⎕NCREATE* or *⎕NTIE*. The rows of the result are in the order in which the files were tied.

Native File Numbers: *R←⎕NNUMS*

This niladic function reports the tie numbers associated with all currently open native files. *R* is an integer vector of negative tie numbers. The elements of the result are in the order in which the files were tied.

Enqueue Event: $\{R\}\leftarrow\{X\}\square NQ \ Y$

This system function generates an event or invokes a method.

While APL is executing, events occur "naturally" as a result of user action or of communication with other applications. These events are added to the event queue as and when they occur, and are subsequently removed and processed one by one by $\square DQ$. $\square NQ$ provides an "artificial" means to generate an event and is analogous to $\square SIGNAL$.

If the left argument X is omitted or is 0, $\square NQ$ adds the event specified by Y to the bottom of the event queue. The event will subsequently be processed by $\square DQ$ when it reaches the top of the queue.

If X is 1, the event is actioned **immediately** by $\square NQ$ itself and is processed in exactly the same way as it would be processed by $\square DQ$. For example, if the event has a callback function attached, $\square NQ$ will invoke it directly.   See $\square DQ$ for further details.

Note that it is not possible for one thread to use 1 $\square NQ$ to send an event to another thread.

If X is 2 and the name supplied is the name of an event, $\square NQ$ performs the default processing for the event immediately, but does **not** invoke a callback function if there is one attached.

If X is 2 and the name supplied is the name of a (Dyalog APL) method, $\square NQ$ invokes the method. Its (shy) result is the result produced by the method.

If X is 3, $\square NQ$ invokes a method in an OLE Control. The (shy) result of $\square NQ$ is the result produced by the method.

If X is 4, $\square NQ$ signals an event from an ActiveXControl object to its host application. The (shy) result of $\square NQ$ is the result returned by the host application and depends upon the syntax of the event. This case is only applicable to ActiveXControl objects.

Y is a nested vector containing an event message. The first two elements of Y are :

$Y[1]$: Object name -   a character vector

$Y[2]$: Event Type -   a numeric scalar or character vector which specifies an event or method.

Y[1] must contain the name of an existing object. If not, *⎕NQ* terminates with a
VALUE ERROR. If *Y*[2] specifies a standard event type, subsequent elements must
conform to the structure defined for that event type. If not, *⎕NQ* terminates with a
SYNTAX ERROR. If *Y*[2] specifies a non-standard event type, *Y*[3] onwards (if
present) may contain arbitrary information. Although any event type not listed herein
may be used, numbers in the range 0-1000 are reserved for future extensions.

 If *⎕NQ* is used monadically, or with a left argument of 0, its (shy) result is always an
empty character vector. If a left argument of 1 is specified, *⎕NQ* returns *Y* unchanged
or a modified *Y* if the callback function returns its modified argument as a result.
 If the left argument is 2, *⎕NQ* returns either the value 1 or a value that is appropriate.

Examples

```
      ⍝ Send a keystroke ("A") to an Edit Field
      ⎕NQ 'TEST.ED' 'KeyPress' 'A'

      ⍝ Iconify all top-level Forms
      {⎕NQ ⍵ 'StateChange' 1}¨'Form'⎕WN'.'

      ⍝ Set the focus to a particular field
      ⎕NQ 'TEST.ED3' 40

      ⍝ Throw a new page on a printer
      1 ⎕NQ 'PR1' 'NewPage'


      ⍝ Terminate ⎕DQ under program control

      'TEST'⎕WC 'Form' ... ('Event' 1001 1)
      ...
      ⎕DQ 'TEST'
      ...
      ⎕NQ 'TEST' 1001  ⍝ From a callback

      ⍝ Call GetItemState method for a TreeView 'F.TV'
      +2 ⎕NQ'F.TV' 'GetItemState' 6
96

      +2 ⎕NQ'.' 'GetEnvironment' 'Dyalog'
c:\Z\2\dyalog82
```

Nested Representation: $R \leftarrow \Box NR \ Y$

Y must be a simple character scalar or vector which represents the name of a function or a defined operator.

If Y is a name of a defined function or defined operator, R is a vector of text vectors. The first element of R contains the text of the function or operator header. Subsequent elements contain lines of the function or operator. Elements of R contain no unnecessary blanks, except for leading indentation of control structures and the blanks which precede comments.

If Y is the name of a variable, a locked function or operator, an external function or a namespace, or is undefined, R is an empty vector.

Example

```
      ∇R←MEAN X      ⍝ Average
[1]   R←(+/X)÷ρX
      ∇

      +F←□NR'MEAN'
  R←MEAN X     ⍝Average   R←(+/X)÷ρX

      ρF
2
      DISPLAY F
.→--------------------------------------.
| .→--------------------.  .→----------. |
| | R←MEAN X    ⍝ Average|  | R←(+/X)÷ρX| |
| '--------------------'  '----------' |
'∊--------------------------------------'
```

The definition of $\Box NR$ has been extended to names assigned to functions by specification ($\leftarrow$), and to local names of functions used as operands to defined operators. In these cases, the result of $\Box NR$ is identical to that of $\Box CR$ except that the representation of defined functions and operators is as described above.

Example

```
    AVG←MEAN∘,

    +F←⎕NR'AVG'
  R←MEAN X      ⍝ Average   R←(+/X)÷ρX   ∘,

    ρF
3

    DISPLAY F
.→-------------------------------------------------.
| .→---------------------------------------.       |
| | .→--------------------.  .→----------.  | ∘ , |
| | | R←MEAN X    ⍝ Average|  | R←(+/X)÷ρX|  | - - |
| | '--------------------'  '----------'  |       |
| '∊-------------------------------------'       |
'∊-------------------------------------------------'
```

Native File Read: $R←⎕NREAD\ Y$

This monadic function reads data from a native file. `Y` is a 3- or 4-element integer vector whose elements are as follows:

[1]   negative tie number,
[2]   conversion code (see below),
[3]   count,
[4]   start byte, counting from 0.

`Y[2]` specifies conversion to an APL internal form as follows

Value	Number of bytes read	Result Type	Result shape
11	*count*	1 bit boolean	8 × *count*
82	*count*	8 bits character	*count*
83	***count***	**8 bits integer**	***count***
163	**2 × *count***	**16 bits integer**	***count***
323	4 × *count*	32 bits integer	*count*
645	8 × *count*	64bits floating	*count*

Example

```
    DATA←⎕NREAD ¯1 82 (⎕NSIZE ¯1) 0
```

Native File Rename: $\{R\} \leftarrow X$ $\square NRENAME$ Y

$\square NRENAME$ is used to rename a native file.

Y is a negative integer tie number associated with a tied native file. X is a simple character vector or scalar containing a valid (and unused) file name.

The shy result of $\square NRENAME$ is the tie number of the renamed file.

Native File Replace: $\{R\} \leftarrow X$ $\square NREPLACE$ Y

$\square NREPLACE$ is used to write data to a native file, replacing data which is already there.

X must be a simple homogeneous APL array containing the data to be written.

Y is a 2- or 3-element integer vector whose elements are as follows:

[1]   negative tie number,
[2]   start byte, counting from 0, at which the data is to be written,
[3]   conversion code (optional).

See $\square NREAD$ for a list of valid conversion codes.

The shy result is the position within the file of the end of the record, or, equivalently, the start of the following one. Used, for example, in:

```
A Replace sequentially from indx.
{α ⎕NREPLACE tie ω}/vec,indx
```

Native File Resize: $\{R\} \leftarrow X$ $\square NRESIZE$ Y

This function changes the size of a native file.

Y is a negative integer tie number associated with a tied native file.

X is a single integer value that specifies the new size of the file in bytes. If X is smaller than the current file size, the file is truncated. If X is larger than the current file size, the file is extended and the value of additional bytes is undefined.

The shy result of $\square NRESIZE$ is the tie number of the resized file.

Create Namespace: $\{R\}\leftarrow\{X\}\square NS\ Y$

If specified, X must be a simple character scalar or vector identifying the name of a namespace.

Y is either a character array which represents a list of names of objects to be copied into the namespace, or is an array produced by the $\square OR$ of an existing namespace. In the first case, Y must be a simple character scalar, vector, matrix or a nested vector of character vectors identifying zero or more workspace objects to be copied into the namespace X. The identifiers in X and Y may be simple names or compound names separated by '.' and including the names of the special namespaces '#', '##' and ''.

The namespace X is created if it doesn't already exist. If the name is already in use for an object other than a namespace, APL issues a $DOMAIN\ ERROR$.

If X is omitted, an unnamed namespace is created.

The objects identified in the list Y are copied into the namespace X.

If X is specified, the result R is the full name (starting $\#$. or $\square SE$.) of the namespace X. If X is omitted, the result R is a namespace reference, or *ref*, to an unnamed namespace.

Examples

```
      +'X'□NS''              ⍝ Create namespace X.
#.X
      +'X'□NS'VEC' 'UTIL.DISP'⍝ Copy VEC and DISP to X.
#.X
      )CS X                  ⍝ Change to namespace X.
#.X
      +'Y'□NS'#.MAT' '##.VEC' ⍝ Create #.X.Y &copy into it
#.X.Y
      +'#.UTIL'□NS'Y.MAT'    ⍝ Copy MAT from Y to #.UTIL.
#.UTIL
      +'#'□NS'Y'             ⍝ Copy namespace Y to root.
#
      +''□NS'#.MAT'          ⍝ Copy MAT to current space.
#.X
      +''□NS''               ⍝ Display current namespace.
#.X
      +'Z'□NS □OR'Y'         ⍝ Create namespace from □OR.
#.X.Z
```

```
        NONAME←⎕NS ''              ⍝ Create unnamed namespace
        NONAME
#.[Namespace]

        DATA←⎕NS¨4⍴⊂''             ⍝ Create 4-element vector of
                                   ⍝ distinct unnamed namespace
s
        DATA
 #.[Namespace]  #.[Namespace]  #.[Namespace]  #.[Namespace
]
```

Namespace Indicator: $R←⎕NSI$

R is a nested vector of character vectors containing the names of the spaces from which functions in the state indicator were called (ρ⎕*NSI*↔ρ⎕*SI*).

Note that ⎕*NSI* contains the names of spaces *from which* functions were called not those *in which* they are currently running.

Example

```
        )OBJECTS
XX      YY

        ⎕VR 'YY.FOO'
     ∇ FOO
[1]    ⎕SE.GOO
     ∇
        ⎕VR'⎕SE.GOO'
     ∇ GOO
[1]    ⎕SI,[1.5]⎕NSI
     ∇


        )CS XX
#.XX
        #.YY.FOO
 GOO  #.YY
 FOO  #.XX
```

Native File Size: $R←⎕NSIZE Y$

This reports the size of a native file.

Y is a negative integer tie number associated with a tied native file. The result *R* is the size of the file in bytes.

Native File Tie: $\quad\quad\quad\quad\quad\quad\quad\quad$ *{R}←X ⎕NTIE Y*

⎕*NTIE* opens a native file.

X is a simple character vector or scalar containing a valid pathname for an existing native file.

Y is a 1- or 2-element vector. *Y[1]* is a negative integer value that specifies an (unused) tie number by which the file may subsequently be referred. *Y[2]* is optional and specifies the mode in which the file is to be opened. This is an integer value calculated as the sum of 2 codes. The first code refers to the type of access needed from users who have already tied the native file. The second code refers to the type of access you wish to grant to users who subsequently try to open the file while you have it open.

Needed from existing users		Granted to subsequent users	
0	read access	0	compatibility mode
1	write access	16	no access (exclusive)
2	read and write access	32	read access
		48	write access
		64	read and write access

On Unix systems, the first code (16|*mode*) is passed to the open(2) call as the access parameter. See include file fcntl.h for details.

Automatic Tie Number Allocation

A tie number of 0 as argument to a create or tie operation, allocates, and returns as an explicit result, the first (closest to zero) available tie number. This allows you to simplify code. For example:

from:

```
        tie←¯1+⌊/0,⎕NNUMS     ⍝ With next available number,
        file ⎕NTIE tie         ⍝ ... tie file.
```

to:

```
        tie←file ⎕NTIE 0      ⍝ Tie with first available no.
```

Example

```
ntie←{                       ⍝ tie file and return tie no.
    α←2+64                   ⍝ default all access.
    ω ⎕ntie 0 α              ⍝ return new tie no.
```

```
}
```

Null Item: $R \leftarrow \Box NULL$

This is a reference to a null item, such as may be returned across the COM interface to represent a null value. An example might be the value of an empty cell in a spreadsheet.

$\Box NULL$ may be used in any context that accepts a namespace reference, in particular:

- As the argument to a defined function
- As an item of an array.
- As the argument to those primitive functions that take character data arguments, for example: =, ≠, ≡, ≠ , ρ, ⊃, ⊂, ...

Example

```
    'EX'⎕WC'OLEClient' 'Excel.Application'
    WB←EX.Workbooks.Open 'simple.xls'

    (WB.Sheets.Item 1).UsedRange.Value2
[Null]  [Null]  [Null]  [Null]  [Null]
[Null]    Year  [Null]  [Null]  [Null]
[Null]    1999    2000    2001    2002
[Null]  [Null]  [Null]  [Null]  [Null]
Sales     100      76     120     150
[Null]  [Null]  [Null]  [Null]  [Null]
Costs      80      60     100     110
[Null]  [Null]  [Null]  [Null]  [Null]
Margin     20      16      20      40
```

To determine which of the cells are filled , you can compare the array with $\Box NULL$.

```
    ⎕NULL≠¨(WB.Sheets.Item 1).UsedRange.Value2
0 0 0 0 0
0 1 0 0 0
0 1 1 1 1
0 0 0 0 0
1 1 1 1 1
0 0 0 0 0
1 1 1 1 1
0 0 0 0 0
1 1 1 1 1
```

Native File Untie: $\{R\}\leftarrow\square NUNTIE\ Y$

This closes one or more native files. *Y* is a scalar or vector of negative integer tie numbers. The files associated with elements of *Y* are closed. Native file untie with a zero length argument (*□NUNTIE* ⍬) flushes all file buffers to disk - see *□FUNTIE* for more explanation.

The shy result of *□NUNTIE* is a vector of tie numbers of the files **actually untied**.

Native File Translate: $\{R\}\leftarrow\{X\}\square NXLATE\ Y$

This associates a character translation vector with a native file or, if *Y* is 0, with the use by *□DR*.

A translate vector is a 256-element vector of integers from 0-255. Each element maps the corresponding *□AV* position onto an ANSI character code.

For example, to map *□AV[17+□IO]* onto ANSI 'a' (code 97), element 17 of the translate vector is set to 97.

Y is either a negative integer tie number associated with a tied native file or 0. If *Y* is negative, monadic *□NXLATE* returns the current translation vector associated with the corresponding native file. If specified, the left argument *X* is a 256-element vector of integers that specifies a new translate vector. In this case, the old translate vector is returned as a shy result. If *Y* is 0, it refers to the translate vector used by *□DR* to convert to and from character data.

The system treats a translate vector with value (⍳256)-*□IO* as meaning *no translation* and thus provides raw input/output bypassing the whole translation process.

The default translation vector established at *□NTIE* or *□NCREATE* time, maps *□AV* characters to their corresponding ANSI positions and is derived from the mapping defined in the current output translation table (normally WIN.DOT)

Between them, ANSI and RAW translations should cater for most uses.

Sign Off APL: $\square OFF$

This niladic system function terminates the APL session, returning to the shell command level . The active workspace does not replace the last continuation workspace.

Object Representation: $R \leftarrow \Box OR \ Y$

$\Box OR$ converts a function, operator or namespace to a special form, described as its *object representation*, that may be assigned to a variable and/or stored on a component file.

Taking the $\Box OR$ of a function or operator is an extremely fast operation as it simply changes the type information in the object's header, leaving its internal structure unaltered. Converting the object representation back to an executable function or operator using $\Box FX$ is also very fast. $\Box OR$ is therefore the recommended form for storing functions and operators on component files and is significantly faster than using $\Box CR$, $\Box VR$ or $\Box NR$.

$\Box OR$ may also be used to convert a namespace (either a plain namespace or a GUI object) into a form that can be stored in a variable or on a component file. The namespace may be reconstructed using $\Box NS$ or $\Box WC$ with its original name or with a new one. $\Box OR$ may therefore be used to *clone* a namespace or GUI object.

Y must be a simple character scalar or vector which contains the name of an APL object.

If Y is the name of a variable, the result R is its value. In this case, $R \leftarrow \Box OR \ Y$ is identical to $R \leftarrow \mathbf{±} Y$.

Otherwise, R is a special form of the name Y, re-classified as a variable. The rank of R is 0 (R is scalar), and the depth of R is 1. These unique characteristics distinguish the result of $\Box OR$ from any other object. The type of R ($\in R$) is itself. Note that although R is scalar, it may not be index assigned to an element of an array unless it is enclosed.

If Y is the name of a function or operator, R is in the domain of the monadic functions Depth ($\equiv$), Disclose ($\supset$), Enclose ($\subset$), Rotate(ϕ), Transpose($\otimes$), Index ing($[\ ]$), Format ($\bar{\ast}$), Identity ($+$), Shape (ρ), Type ($\in$) and Unique ($\cup$), of the dyadic functions Assignment ($\leftarrow$), Without ($\sim$), Index Of (ι), Intersection ($\cap$), Match ($\equiv$), Membership ($\in$), Not Match ($\neq$) and Union ($\cup$), and of the monadic system functions Canonical Representation ($\Box CR$), Cross-Reference ($\Box REFS$), Fix ($\Box FX$), Format ($\Box FMT$), Nested Representation ($\Box NR$) and Vector Representation ($\Box VR$).

Nested arrays which include the object representations of functions and operators are in the domain of many mixed functions which do not use the values of items of the arrays.

Note that a $\Box OR$ object can be transmitted through an 'APL-style' TCP socket. This technique may be used to transfer objects including namespaces between APL sessions.

The object representation forms of namespaces produced by `⎕OR` may not be used as arguments to any primitive functions. The only operations permitted for such objects (or arrays containing such objects) are `⎕EX`, `⎕FAPPEND`, `⎕FREPLACE`, `⎕NS`, and `⎕WC`.

Example

```
      F←⎕OR ⎕FX'R←FOO' 'R←10'

      ρF

      ρρF
0
      ≡F
1
      F≡∈F
1
```

The display of the `⎕OR` form of a function or operator is a listing of the function or operator. If the `⎕OR` form of a function or operator has been enclosed, then the result will display as the operator name preceded by the symbol ∇. It is permitted to apply `⎕OR` to a locked function or operator. In this instance the result will display as for the enclosed form.

Examples

```
       F
       ∇  R←FOO
[1]      R←10
       ∇

      ⊂F
 ∇FOO

      ⎕LOCK'FOO'

      ⎕OR'FOO'
∇FOO

      A←ι5

      A[3]←⊂F

      A
1 2  ∇FOO  4 5
```

For the `⎕OR` forms of two functions or operators to be considered identical, their unlocked display forms must be the same, they must either both be locked or unlocked, and any monitors, trace and stop vectors must be the same.

Example

```
      F←⎕OR ⎕FX 'R←A PLUS B' 'R←A+B'

      F≡⎕OR 'PLUS'
1

      1 ⎕STOP 'PLUS'

      F≡⎕OR 'PLUS'
0
```

Namespace Examples

The following example sets up a namespace called *UTILS*, copies into it the contents of the *UTIL* workspace, then writes it to a component file:

```
      )CLEAR
clear ws
      )NS UTILS
#.UTILS
      )CS UTILS
#.UTILS
      )COPY UTIL
C:\WDYALOG\WS\UTIL saved Fri Mar 17 12:48:06 1995
      )CS
#
      'ORTEST' ⎕FCREATE 1
      (⎕OR'UTILS')⎕FAPPEND 1
```

The namespace can be restored with `⎕NS`, using either the original name or a new one:

```
      )CLEAR
clear ws
      'UTILS' ⎕NS ⎕FREAD 1 1
#.UTILS
      )CLEAR
clear ws
      'NEWUTILS' ⎕NS ⎕FREAD 1 1
#.NEWUTILS
```

This example illustrates how `⎕OR` can be used to clone a GUI object; in this case a Group containing some Button objects. Note that `⎕WC` will accept **only** a `⎕OR` object as its argument (or preceded by the "Type" keyword). You may not specify any other properties in the same `⎕WC` statement, but you must instead use `⎕WS` to reset them afterwards.

```
'F'⎕WC'Form'
'F.G1' ⎕WC 'Group' '&One' (10 10)(80 30)
'F.G1.B2'⎕WC'Button' '&Blue' (40 10)('Style' 'Radio')
'F.G1.B3'⎕WC'Button' '&Green' (60 10)('Style' 'Radio')
'F.G1.B1'⎕WC'Button' '&Red' (20 10)('Style' 'Radio')
'F.G2' ⎕WC ⎕OR 'F.G1'
'F.G2' ⎕WS ('Caption' 'Two')('Posn' 10 60)
```

Note too that `⎕WC` and `⎕NS` may be used interchangeably to rebuild *pure* namespaces or GUI namespaces from a `⎕OR` object. You may therefore use `⎕NS` to rebuild a Form or use `⎕WC` to rebuild a pure namespace that has no GUI components.

Search Path: `⎕PATH`

`⎕PATH` is a simple character vector representing a blank-separated list of namespaces. It is approximately analogous to the PATH variable in MS-DOS or UNIX

The `⎕PATH` variable can be used to identify a namespace in which commonly used utility functions reside. Functions or operators (**NOT** variables) which are copied into this namespace and *exported* (see `⎕EXPORT`) can then be used directly from anywhere in the workspace without giving their full path names.

Example

To make the `DISPLAY` function available directly from within any namespace.

```
⍝ Create and reference utility namespace.
⎕PATH←'⎕se.util'⎕ns''
⍝ Copy DISPLAY function from UTIL into it.
'DISPLAY'⎕se.util.⎕cy'UTIL'
⍝ (Remember to save the session to file).
```

In detail, `⎕PATH` works as follows:

When a reference to a name cannot be found in the current namespace, the system searches for it from left to right in the list of namespaces indicated by □*PATH*. In each namespace, if the name references a defined function (or operator) *and* the export type of that function is non-zero (see □*EXPORT*), then it is used to satisfy the reference. If the search exhausts all the namespaces in □*PATH* without finding a qualifying reference, the system issues a VALUE ERROR in the normal manner.

The special character ↑ stands for the list of namespace ancestors:

> *## ##.## ##.##.## ...*

In other words, the search is conducted upwards through enclosing namespaces, emulating the static scope rule inherent in modern block-structured languages.

Note that the □*PATH* mechanism is used ONLY if the function reference cannot be satisfied in the current namespace. This is analogous to the case when the MS_DOS or UNIX PATH variable begins with a '.'.

Examples

□*PATH*	*Search in ...*
1. '□*se.util*'	*Current space, then* *□se.util, then* *VALUE ERROR*
2. '↑'	*Current space* *Parent space: ##* *Parent's parent space: ##.##* *...* *Root: # (or □se if current space* * was inside □se)* *VALUE ERROR*
3. '*util* ↑ □*se.util*'	*Current space* *util (relative to current space)* *Parent space: ##* *...* *Root: # or □se* *□se.util* *VALUE ERROR*

Note that □*PATH* is a *session* variable. This means that it is workspace-wide and survives)*LOAD* and)*CLEAR*. It can of course, be localised by a defined function or operator.

Program Function Key: $R \leftarrow \{X\} \square PFKEY \ Y$

$\square PFKEY$ is a system function that sets or queries the programmable function keys. $\square PFKEY$ associates a sequence of keystrokes with a function key. When the user subsequently presses the key, it is as if he had typed the associated keystrokes one by one.

Y is an integer scalar in the range 0-255 specifying a programmable function key. If X is omitted the result R is the current setting of the key. If the key has not been defined previously, the result is an empty character vector.

If X is specified it is a simple or nested character vector defining the new setting of the key. The value of X is returned in the result R.

The elements of X are either character scalars or 2-element character vectors which specify Input Translate Table codes.

Programmable function keys are recognised in any of the three types of window (SESSION, EDIT and TRACE) provided by the Dyalog APL development environment. $\square SR$ operates with the 'raw' function keys and ignores programmed settings.

Note that key definitions can reference other function keys.

The size of the buffer associated with $\square PFKEY$ is specified by the *pfkey_size* parameter.

Examples

```
      (')FNS',⊂'ER') □PFKEY 1
 )FNS  ER


      DISPLAY □PFKEY 1
.→-------------.
|          .→-. |
| )  F  N  S |ER| |
| -  -  -  - '--' |
'∊-------------'


      (')VARS',⊂'ER') □PFKEY 2
 )VARS  ER
      'F1' 'F2' □PFKEY 3    ⍝ Does )FNS and )VARS
 F1  F2
```

Print Precision: `⎕PP`

`⎕PP` is the number of significant digits in the display of numeric output.

`⎕PP` may be assigned any integer value in the range 1 to 16. The value in a clear workspace is 10.

`⎕PP` is used to format numbers displayed directly. It is an implicit argument of monadic function Format (`⍕`), `⎕CR`, `⎕NR`, `⎕VR`, `⎕DM`, `⎕ED`, monadic `⎕FMT` and display of numbers via `⎕` and `⍞` output. `⎕PP` is ignored for the display of integers.

Examples

```
      ⎕PP←10

      ÷3 6
0.3333333333 0.1666666667

      ⎕PP←3

      ÷3 6
0.333 0.167
```

Print Width: `⎕PW`

`⎕PW` is the maximum number of output characters per line before folding the display.

`⎕PW` may be assigned any integer value in the range 30 to 32767.

If an attempt is made to display a line wider than `⎕PW`, then the display will be folded at or before the `⎕PW` width and the folded portions indented 6 spaces. The display of a simple numeric array may be folded at a width less than `⎕PW` so that individual numbers are not split.

`⎕PW` only affects output, either direct or through `⎕` output. It does not affect the result of the function Format (`⍕`), of the system function `⎕FMT`, or output through the system functions `⎕ARBOUT` and `⎕ARBIN`, or output through `⍞`.

Note that if the auto_pw parameter is set to 1, `⎕PW` is reset dynamically whenever the Session window is resized. In these circumstances, a value assigned to `⎕PW` will only be effective until the Session Window is next resized.

Examples

```
      ⎕PW←30


      ⎕←3ρ÷3
0.3333333333 0.3333333333
      0.3333333333
```

Cross References: *R←⎕REFS Y*

Y must be a simple character scalar or vector, identifying the name of a function or operator, or the object representation form of a function or operator (see *⎕OR*). *R* is a simple character matrix, with one name per row, of identified names in the function or operator in *Y* excluding distinguished names of system constants, variables or functions.

Example

```
      ⎕VR'OPTIONS'
    ∇ OPTIONS;OPTS;INP
[1]   ⍝ REQUESTS AND EXECUTES AN OPTION
[2]    OPTS ←'INPUT' 'REPORT' 'END'
[3]   IN:INP←ASK'OPTION:'
[4]    →EXρ̈(⊂INP)∈OPTS
[5]    'INVALID OPTION. SELECT FROM',OPTS ◇ →IN
[6]   EX:→EX+OPTSι⊂INP
[7]    INPUT ◇ →IN
[8]    REPORT ◇ →IN
[9]   END:
    ∇

      ⎕REFS'OPTIONS'
ASK
END
EX
IN
INP
INPUT
OPTIONS
OPTS
REPORT
```

If *Y* is locked or is an External Function, *R* contains its name only. For example:

```
      ⎕LOCK 'OPTIONS' ◇ ⎕REFS 'OPTIONS'
OPTIONS
```

If *Y* is the name of a primitive, external or derived function, *R* is an empty matrix with shape 0 0.

Random Link: `⎕RL`

`⎕RL` is used or set to establish a base for generating random numbers.

`⎕RL` may be assigned any integer value in the range 1 to 2147483646. The value in a clear workspace is 16807.

Repeatable results can be obtained from Roll or Deal if `⎕RL` is set to a particular value first.

`⎕RL` is used and set implicitly by the functions Roll and Deal (`?`).

Examples

```
      ⎕RL
16807

      ?9 9 9
2 7 5

      ?9
7

      ⎕RL
984943658

      ⎕RL←16807
      ?9 9 9
2 7 5
```

Response Time Limit: `⎕RTL`

A non-zero value in `⎕RTL` places a time limit, in seconds, for input requested via `⎕` , `⎕ARBIN` , and `⎕SR`. `⎕RTL` may be assigned any integer in the range 0 to 32767. The value is a clear workspace is 0.

Example

```
      ⎕RTL←5 ◇ ⎕←'FUEL QUANTITY?' ◇ R←⎕
FUEL QUANTITY?
TIMEOUT
      ⎕RTL←5 ◇ ⎕←'FUEL QUANTITY?' ◇ R←⎕
```

Save Workspace: $\quad\{R\}\leftarrow\{X\}\Box SAVE\ Y$

Y must be a simple character scalar or vector, identifying a workspace name. *R* is a simple logical scalar. The active workspace is saved with the given name in *Y*. In the active workspace, the value 1 is returned. The result is suppressed if not used or assigned.

The optional left argument *X* is either 0 or 1. If *X* is omitted or 1, the saved version of the workspace has execution suspended at the point of exit from the $\Box SAVE$ function. If the saved workspace is subsequently loaded by $\Box LOAD$, execution is resumed, and the value 0 is returned if the result is used or assigned, or otherwise the result is suppressed. In this case, the latent expression value ($\Box LX$) is ignored.

If *X* is 0, the workspace is saved without any State Indicator in effect. The effect is the same as if you first executed $)RESET$ and then $)SAVE$. In this case, when the workspace is subsequently loaded, the value of the latent expression ($\Box LX$) is honoured if applicable.

A *DOMAIN ERROR* is reported if the name in *Y* is not a valid workspace name, or the reference is to an unauthorised directory.

Note that the values of all system variables (including $\Box SM$) and all GUI objects are saved.

Example

```
      (⊃'SAVED' 'ACTIVE' [⎕IO+⎕SAVE'TEMP']),' WS'
ACTIVE WS
      ⎕LOAD 'TEMP'
SAVED WS
```

Screen Dimensions: $\quad R\leftarrow\Box SD$

$\Box SD$ is a 2-element integer vector containing the number of rows and columns on the screen, or in the USER window.

Under DOS, $\Box SD$ always contains (25 80). For asynchronous terminals under UNIX, the screen size is taken from the terminal database **terminfo** or **termcap**.

In window implementations of Dyalog APL, $\Box SD$ reports the current size (in characters) of the USER window or the current size of the SM object, whichever is appropriate.

Session Namespace: *⎕SE*

⎕SE is a system namespace. Its GUI components (MenuBar, ToolBar, and so forth) define the appearance and behaviour of the APL Session window and may be customised to suit individual requirements.

⎕SE is maintained separately from the active workspace and is not affected by *)LOAD* or *)CLEAR*. It is therefore useful for containing utility functions. The contents of *⎕SE* may be saved in and loaded from a .DSE file.

See *User Guide* for further details.

Execute (UNIX) Command: *{R}←⎕SH Y*

⎕SH executes a UNIX shell command or a DOS command. *⎕SH* is a synonym of *⎕CMD*. Either function may be used in either environment (UNIX or DOS) with exactly the same effect. *⎕SH* is probably more natural for the UNIX user. This section describes the behaviour of *⎕SH* and *⎕CMD* under UNIX. See *⎕CMD* for a discussion of the behaviour of these system functions under DOS.

Y must be a simple character scalar or vector representing a UNIX shell command. *R* is a nested vector of character vectors.

Y may be any acceptable UNIX command. It could cause another process to be entered, such as sed or vi. If the command does not return a result, *R* is ⊂'' but the result is suppressed if not explicitly used or assigned. If the command has a non-zero exit code, then APL will signal a *DOMAIN ERROR*. If the command returns a result and has a zero exit code, then each element of *R* will be a line from the standard output (stdout) of the command. Output from standard error (stderr) is not captured unless redirected to stdout.

Examples

```
      ⎕SH'ls'
FILES WS temp

      ⎕SH 'rm WS/TEST'

      ⎕SH 'grep bin /etc/passwd ; exit 0'
bin:!:2:2::/bin:

      ⎕SH 'apl MYWS <inputfile >out1 2>out2 &'
```

See *User Guide* for further information.

Start (UNIX) Auxiliary Processor: *X* □*SH* *Y*

Used dyadically, □*SH* starts an Auxiliary Processor. The effect, as far as the APL user is concerned, is identical under both DOS and UNIX although there are differences in the method of implementation. □*SH* is a synonym of □*CMD*. Either function may be used in either environment (UNIX or DOS) with exactly the same effect. □*SH* is probably more natural for the UNIX user. This section describes the behaviour of □*SH* and □*CMD* under UNIX. See □*CMD* for a discussion of the behaviour of these system functions under DOS.

X must be a simple character vector. *Y* may be a simple character scalar or vector, or a nested character vector.

□*SH* loads the Auxiliary Processor from the file named by *X* using a search-path defined by the environment variable WSPATH.

The effect of starting an AP is that one or more **external functions** are defined in the workspace. These appear as locked functions and may be used in exactly the same way as regular defined functions.

When an external function is used in an expression, the argument(s) (if any) are **piped** to the AP for processing. If the function returns a result, APL halts while the AP is processing and waits for the result. If not it continues processing in parallel.

The syntax of dyadic □*SH* is similar to the UNIX execl(2) system call, where '*taskname*' is the name of the auxiliary processor to be executed and *arg0* through *argn* are the parameters of the calling line to be passed to the task, viz.

        '*taskname*' □*SH* '*arg0*' '*arg1*' ... '*argn*'

See *User Guide* for further information.

Examples

        '*xutils*' □*SH* '*xutils*' '*ss*' '*dbr*'

        '*/bin/sh*' □*SH* '*sh*' '*-c*' '*adb test*'

Shadow Name: ⎕SHADOW Y

Y must be a simple character scalar, vector or matrix identifying one or more APL names. For a vector Y, names are separated by one or more blanks. For a matrix Y, each row is taken to be a single name.

Each valid name in Y is shadowed in the most recently invoked defined function or operator, as though it were included in the list of local names in the function or operator header. The class of the name becomes 0 (undefined). The name ceases to be shadowed when execution of the shadowing function or operator is completed. Shadow has no effect when the state indicator is empty.

If a name is ill-formed, or if it is the name of a system constant or system function, *DOMAIN ERROR* is reported.

If the name of a top-level GUI object is shadowed, it is made inactive.

Example

```
      ⎕VR'RUN'
    ∇ NAME RUN FN
[1]    ⍝ RUNS FUNCTION NAMED <NAME> DEFINED
[2]    ⍝ FROM REPRESENTATION FORM <FN>
[3]     ⎕SHADOW NAME
[4]    ⍎⎕FX FN
    ∇

      0 ⎕STOP 'RUN'

      'FOO' RUN 'R←FOO' 'R←10'
10

RUN[0]

      )SINL
RUN[0]*        FOO      FN      NAME

      →⎕LC

      FOO
VALUE ERROR
      FOO
      ∧
```

State Indicator: $R \leftarrow \Box SI$

R is a nested vector of vectors giving the names of the functions or operators in the execution stack.

Example

```
      )SI
PLUS[2]*
.
MATDIV[4]
FOO[1]*
⍠
```

```
      ⎕SI
 PLUS  MATDIV  FOO
```

```
      (⍴⎕LC)=⍴⎕SI
1
```

If execution stops in a callback function, $\Box DQ$ will appear on the stack, and may occur more than once

```
      )SI
ERRFN[7]*
⎕DQ
CALC
⎕DQ
MAIN
```

To edit the function on the top of the stack:

```
      ⎕ED ⊃⎕SI
```

The name of the function which called this one:

```
      ⊃1↓⎕SI
```

To check if the function ΔN is pendent:

```
      ((⊂∆N)∊1↓⎕SI)/'Warning : ',∆N,' is pendent'
```

See also $\Box XSI$.

Signal Event: $\{X\}\square SIGNAL\ Y$

Y must be a single positive integer scalar or a null vector. *X* is optional. If present, *X* must be a simple character scalar or vector, or an object reference. If *Y* is null, nothing is signalled.

Y is taken to be an event number in the range 1-999. *X* is an optional text message. If omitted, the standard event message shown for the corresponding event number in Figure 7(i) or 7(ii) is assumed. If there is no standard message, a message of the form *ERROR NUMBER n* is composed, where *n* is the event number in *Y*.

The effect of the system function is to interrupt execution. The state indicator is cut back to exit from the function or operator containing the line that invoked $\square SIGNAL$ or the Execute (⍎) expression that invoked $\square SIGNAL$, and an error is then generated.

An error interrupt may be trapped if the system variable $\square TRAP$ is set to intercept the event. Otherwise, the standard system action is taken (which may involve cutting back the state indicator further if there are locked functions or operators in the state indicator). The standard event message is replaced by the text given in *X*, if present.

Example

```
      ⎕VR'DIVIDE'
    ∇ R←A DIVIDE B;⎕TRAP
[1]    ⎕TRAP←11 'E' '→ERR'
[2]    R←A÷B ◇ →0
[3]  ERR:'DIVISION ERROR' ⎕SIGNAL 11
    ∇

      2 4 6 DIVIDE 0
DIVISION ERROR
      2 4 6 DIVIDE 0
      ∧
```

If you are using the Microsoft .Net Framework, you may use $\square SIGNAL$ to throw an exception by specifying a value of 90 in *Y*. In this case, if you specify the optional left argument *X* , it must be a reference to a .Net object that is or derives from the Microsoft .Net class System.Exception. The following example illustrates a *constructor* function *CTOR* that expects to be called with a value for $\square IO$ (0 or 1)

```
    ∇ CTOR IO;EX
[1]    :If IO∊0 1
[2]        ⎕IO←IO
[3]    :Else
[4]        EX←ArgumentException.New'IO must be 0 or 1'
[5]        EX ⎕SIGNAL 90
```

```
      [6]      :EndIf
          ∇
```

Size of Object: $R \leftarrow \Box SIZE\ Y$

Y must be a simple character scalar, vector or matrix. A scalar or vector is treated as a single row matrix. Each row is taken to be an APL name. *R* is a simple integer vector of non-negative elements and of shape $1 \uparrow \bar{}2 \uparrow 1\ 1, \rho Y$.

If the name in a row of *Y* identifies an object with an active referent, the workspace required in bytes by that object is returned in the corresponding element of *R*. Otherwise, 0 is returned in that element of *R*.

The result returned for an external variable is the space required to store the external array. The result for a system constant, variable or function is 0. The result returned for a GUI object gives the amount of workspace needed to store it, but excludes the space required for its children.

Note: Wherever possible, Dyalog APL 'shares' the whole or part of a workspace object rather than generates a separate copy.

Examples
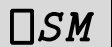
```
      ⎕VR 'FOO'
    ∇ R←FOO
[1]   R←10
    ∇

     A←⍳10

     'EXT/ARRAY' ⎕XT'E' ◇ E←⍳20

     ⎕SIZE ↑'A' 'FOO' 'E' 'UND'
28 76 120 0
```

Screen Map: *⎕SM*

⎕SM is a system variable that defines a character-based user interface (as opposed to a graphical user interface). In versions of Dyalog APL for DOS/386 and in those that support asynchronous terminals, *⎕SM* defines a **form** that is displayed on the **USER SCREEN**. The implementation of *⎕SM* in "window" environments is compatible with these versions. In Dyalog APL/X, *⎕SM* occupies its own separate window on the display, but is otherwise equivalent. In versions of Dyalog APL with GUI support, *⎕SM* either occupies its own separate window (as in Dyalog APL/X) or, if it exists, uses the window assigned to the SM object. This allows *⎕SM* to be used in a GUI application in conjunction with other GUI components.

In general *⎕SM* is a nested matrix containing between 3 and 13 columns. Each row of *⎕SM* represents a **field**; each column a **field attribute**.

⎕SM is discussed in detail in the *User Guide*

The columns have the following meanings :

Column	Description	Default
1	Field Contents	*N/A*
2	Field Position - Top Row	*N/A*
3	Field Position - Left Column	*N/A*
4	Window Size - Rows	0
5	Window Size - Columns	0
6	Field Type	0
7	Behaviour	0
8	Video Attributes	0
9	Active Video Attributes	⁻1
10	Home Element - Row	1
11	Home Element - Column	1
12	Scrolling Group - Vertical	0
13	Scrolling Group - Horizontal	0

With the exception of columns 1 and 8, all elements in *⎕SM* are integer scalar values.

Elements in column 1 (Field Contents) may be :

a)   A numeric scalar

b)   A numeric vector

c)   A 1-column numeric matrix

d)   A character scalar

e)   A character vector

f)   A character matrix (rank 2)

g)   A nested matrix defining a sub-form whose structure and contents must
         conform to that defined for ⎕*SM* as a whole. This definition is recursive. Note
         however that a sub-form must be a matrix - a vector is not allowed.

Elements in column 8 (Video Attributes) may be :

a)   An integer scalar that specifies the appearance of the entire field.

b)   An integer array of the same shape as the field contents. Each element
         specifies the appearance of the corresponding element in the field contents.

Screen Management (DOS & Async Terminals)

Dyalog APL for DOS/386 and UNIX systems (Async terminals) manages two screens;
the SESSION screen and the USER screen. If the SESSION screen is current, an
assignment to ⎕*SM* causes the display to switch to the USER screen and show the form
defined by ⎕*SM*.

If the USER screen is current, any change in the value of ⎕*SM* is immediately reflected
by a corresponding change in the appearance of the display. However, an assignment
to ⎕*SM* that leaves its value unchanged has no effect.

Dyalog APL automatically switches to the SESSION screen for default output, if it enters immediate input mode (6-space prompt), or through use of ⎕ or ⍞. This means that typing

>   ⎕SM ← expression

in the APL session will cause the screen to switch first to the USER screen, display the form defined by ⎕SM, and then switch back to the SESSION screen to issue the 6-space prompt. This normally happens so quickly that all the user sees is a flash on the screen. To retain the USER screen in view it is necessary to issue a call to ⎕SR or for APL to continue processing
e.g.

>   ⎕SM ← expression   ◇   ⎕SR 1

or

>   ⎕SM ← expression   ◇   ⎕DL 5

Screen Management (Window Versions)

In Dyalog APL/X, and optionally in Dyalog APL/W, ⎕SM is displayed in a separate **USER WINDOW** on the screen. In an end-user application this may be the only Dyalog APL window. However, during development, there will be a SESSION window, and perhaps EDIT and TRACE windows too.

The USER Window will only accept input during execution of ⎕SR. It is otherwise "output-only". Furthermore, during the execution of ⎕SR it is the only active window, and the SESSION, EDIT and TRACE Windows will not respond to user input.

Screen Management (GUI Versions)

In versions of Dyalog APL that provide GUI support, there is a special SM object that defines the position and size of the window to be associated with ⎕SM. This allows character-mode applications developed for previous versions of Dyalog APL to be migrated to and integrated with GUI environments without the need for a total re-write.

Effect of Localisation

Like all system variables (with the exception of $\square TRAP$) $\square SM$ is subject to "pass-through localisation". This means that a localised $\square SM$ assumes its value from the calling environment. The localisation of $\square SM$ does not, of itself therefore, affect the appearance of the display. However, reassignment of a localised $\square SM$ causes the new form to overlay rather than replace whatever forms are defined further down the stack. The localisation of $\square SM$ thus provides a simple method of defining pop-up forms, help messages, etc.

The user may edit the form defined by $\square SM$ using the system function $\square SR$. Under the control of $\square SR$ the user may change the following elements in $\square SM$ which may afterwards be referenced to obtain the new values.

Column 1 :   Field Contents
Column 10 :   Home Element - Row (by scrolling vertically)
Column 11 :   Home Element - Column (by scrolling horizontally)

See *User Guide* for a full description of the behaviour and usage of $\square SM$ and $\square SR$.

Screen Read: $R \leftarrow \{X\} \square SR \ \ Y$

$\square SR$ is a system function that allows the user to edit or otherwise interact with the form defined by $\square SM$.

In Dyalog APL for DOS or in versions that support asynchronous terminals, if the current screen is the SESSION screen, $\square SR$ immediately switches to the USER SCREEN and displays the form defined by $\square SM$.

In Dyalog APL/X, $\square SR$ causes the input cursor to be positioned in the USER window. During execution of $\square SR$, only the USER Window defined by $\square SM$ will accept input and respond to the keyboard or mouse. The SESSION and any EDIT and TRACE Windows that may appear on the display are dormant.

In versions of Dyalog APL with GUI support, a single SM object may be defined. This object defines the size and position of the $\square SM$ window, and allows $\square SM$ to be used in conjunctions with other GUI components. In these versions, $\square SR$ acts as a superset of $\square DQ$ (see $\square DQ$) but additionally controls the character-based user interface defined by $\square SM$.

Y is an integer vector that specifies the fields which the user may visit. In versions with GUI support, *Y* may additionally contain the names of GUI objects with which the user may also interact.

If specified, *X* may be an enclosed vector of character vectors defining *EXIT_KEYS* or a 2-element nested vector defining *EXIT_KEYS* and the *INITIAL_CONTEXT*.

The result *R* is the *EXIT_CONTEXT*.

Thus the 3 uses of *⎕SR* are :

```
        EXIT_CONTEXT ← ⎕SR FIELDS

        EXIT_CONTEXT ← (⊂EXIT_KEYS)⎕SR FIELDS

        EXIT_CONTEXT ← (EXIT_KEYS)(INITIAL_CONTEXT)⎕SR FIELD
S
```

FIELDS

If an element of *Y* is an integer scalar, it specifies a field as the index of a row in *⎕SM* (if *⎕SM* is a vector it is regarded as having 1 row).

If an element of *Y* is an integer vector, it specifies a sub-field. The first element in *Y* specifies the top-level field as above. The next element is used to index a row in the form defined by ⊃*⎕SM[Y[1];1]* and so forth.

If an element of *Y* is a character scalar or vector, it specifies the name of a top-level GUI object with which the user may also interact. Such an object must be a "top-level" object, i.e. the *Root* object ('.') or a *Form* or pop-up *Menu*. This feature is implemented ONLY in versions of Dyalog APL with GUI support.

EXIT_KEYS

Each element of *EXIT_KEYS* is a 2-character code from the Input Translate Table for the keyboard. If the user presses one of these keys, *⎕SR* will terminate and return a result.

If *EXIT_KEYS* is not specified, it defaults to :

```
        'ER' 'EP' 'QT'
```

which (normally) specifies <Enter>, <Esc> and <Shift+Esc>.

INITIAL_CONTEXT

This is a vector of between 3 and 6 elements with the following meanings and defaults:

Element	Description	Default
1	Initial Field	*N/A*
2	Initial Cursor Position - Row	*N/A*
3	Initial Cursor Position - Col	*N/A*
4	Initial Keystroke	''
5	(ignored)	*N/A*
6	Changed Field Flags	0

Structure of INITIAL_CONTEXT

INITIAL_CONTEXT[1] specifies the field in which the cursor is to be placed. It is an integer scalar or vector, and must be a member of *Y*. It must not specify a field which has *BUTTON* behaviour (64), as the cursor is not allowed to enter such a field.

INITIAL_CONTEXT[2 3] are integer scalars which specify the initial cursor position within the field in terms of row and column numbers.

INITIAL_CONTEXT[4] is either empty, or a 2-element character vector specifying the initial keystroke as a code from the Input Translate Table for the keyboard.

INITIAL_CONTEXT[5] is ignored. It is included so that the *EXIT_CONTEXT* result of one call to □*SR* can be used as the *INITIAL_CONTEXT* to a subsequent call.

INITIAL_CONTEXT[6] is a boolean scalar or vector the same length as *Y*. It specifies which of the fields in *Y* has been modified by the user.

EXIT_CONTEXT

The result *EXIT_CONTEXT* is a 6 or 9-element vector whose first 6 elements have the same structure as the *INITIAL_CONTEXT*. Elements 7-9 **only** apply to those versions of Dyalog APL that provide mouse support.

Element	Description
1	Final Field
2	Final Cursor Position - Row
3	Final Cursor Position - Col
4	Terminating Keystroke
5	Event Code
6	Changed Field Flags
7	Pointer Field
8	Pointer Position - Row
9	Pointer Position - Col

Structure of the Result of □*SR*

EXIT_CONTEXT[1] contains the field in which the cursor was when □*SR* terminated due to the user pressing an exit key or due to an event occurring. It is an integer scalar or vector, and a member of *Y*.

EXIT_CONTEXT[2 3] are integer scalars which specify the row and column position of the cursor within the field *EXIT_CONTEXT*[1] when □*SR* terminated.

EXIT_CONTEXT[4] is a 2-element character vector specifying the last keystroke pressed by the user before □*SR* terminated. Unless □*SR* terminated due to an event, *EXIT_CONTEXT*[4] will contain one of the exit keys defined by *X*. The keystroke is defined in terms of an Input Translate Table code.

EXIT_CONTEXT[5] contains the **sum** of the event codes that caused □*SR* to terminate. For example, if the user pressed a mouse button on a *BUTTON* field (event code 64) **and** the current field has *MODIFIED* behaviour (event code 2) *EXIT_CONTEXT*[5] will have the value 66.

EXIT_CONTEXT[6] is a boolean scalar or vector the same length as *Y*. It specifies which of the fields in *Y* has been modified by the user during **this** □*SR*, ORed with *INITIAL_CONTEXT*[6]. Thus if the *EXIT_CONTEXT* of one call to □*SR* is fed back as the *INITIAL_CONTEXT* of the next, *EXIT_CONTEXT*[6] records the fields changed since the start of the process.

EXIT_CONTEXT (Window Versions)

□*SR* returns a 9-element result **ONLY** if it is terminated by the user pressing a mouse button. In this case :

EXIT_CONTEXT[7] contains the field over which the mouse pointer was positioned when the user pressed a button. It is an integer scalar or vector, and a member of *Y*.

EXIT_CONTEXT[8 9] are integer scalars which specify the row and column position of the mouse pointer within the field *EXIT_CONTEXT*[7] when □*SR* terminated.

See *User Guide* for a full description of the behaviour and usage of □*SR* and □*SM*.

State Indicator Stack: *R←□STACK*

R is a two-column matrix, with one row per entry in the State Indicator.

Column 1 :   □*OR* form of user defined functions or operators on the State Indicator. Null for entries that are not user defined functions or operators.

Column 2 :   Indication of the type of the item on the stack.

space	user defined function or operator
⍎	execute level
□	evaluated input
⋆	desk calculator level
□*DQ*	in callback function
other	primitive operator

Example

```
        )SI
PLUS[2]*
.
MATDIV[4]
FOO[1]*
⍋

        ⎕STACK
           *
∇PLUS
           .
∇MATDIV
           *
∇FOO
           ⍒
           *

        ρ⎕STACK
8 2

        (ρ⎕LC)=ρ⎕STACK
0
```

Pendent defined functions and operators may be edited in Dyalog APL with no
resulting SI damage. However, only the visible definition is changed; the pendent
version on the stack is retained until its execution is complete. When the function or
operator is displayed, only the visible version is seen. Hence ⎕STACK is a tool which
allows the user to display the form of the actual function or operator being executed.

Example

To display the version of MATDIV currently pendent on the stack:

```
        ⊃⎕STACK[4;1]
      ∇ R←A MATDIV B
[1]   ⍝ Divide matrix A by matrix B
[2]    C←A⌹B
[3]   ⍝ Check accuracy
[4]    D←⌊0.5+A PLUS.TIMES B
      ∇
```

State of Object: $R \leftarrow \square STATE\ Y$

Y must be a simple character scalar or vector which is taken to be the name of an APL object. The result returned is a nested vector of 4 elements as described below. $\square STATE$ supplies information about shadowed or localised objects that is otherwise unobtainable.

$1 \supset R$ Boolean vector, element set to 1 if and only if this level shadows Y.

Note: $(\rho 1 \supset R) = \rho \square LC$

$2 \supset R$ Numeric vector giving the stack state of this name as it entered this level:

0   not on stack
1   suspended
2   pendent (may also be suspended)
3   active (may also be pendent or suspended)

Note: $(\rho 2 \supset R) = \rho \square LC$

$3 \supset R$ Numeric vector giving the name classification of Y as it entered this level.

Note: $(\rho 3 \supset R) = +/1 \supset R$

$4 \supset R$ Vector giving the contents of Y before it was shadowed at this level.

Note: $(\rho 4 \supset R) = +/0 \neq 3 \supset R$

Example

```
      ⎕FMT∘⎕OR¨'FN1' 'FN2' 'FN3'
    ∇ FN1;A;B;C        ∇ FN2;A;C              ∇ FN3;A
[1]    A←1        [1]    A←'HELLO'        [1]    A←100
[2]    B←2        [2]    B←'EVERYONE'      [2]    ∘
[3]    C←3        [3]    C←'HOW ARE YOU?'        ∇
[4]    FN2        [4]    FN3
    ∇                  ∇


      )SI
FN3[2]*
FN2[4]
FN1[4]


      ⎕STATE 'A'
 1 1 1  0 0 0  2 2 0   HELLO  1


      R←⎕STATE '⎕TRAP'
```

Set Stop: $\{R\}←X \ \Box STOP \ Y$

Y must be a simple character scalar or vector which is taken to be the name of a visible defined function or operator. X must be a simple non-negative integer scalar or vector. R is a simple integer vector of non-negative elements. X identifies the numbers of lines in the function or operator named by Y on which a stop control is to be placed. Numbers outside the range of line numbers in the function or operator (other than 0) are ignored. The number 0 indicates that a stop control is to be placed immediately prior to exit from the function or operator. If X is empty, all existing stop controls are cancelled. The value of X is independent of $\Box IO$.

R is a vector of the line numbers on which a stop control has been placed in ascending order. The result is suppressed unless it is explicitly used or assigned.

Examples

```
      +(0,ι10) ⎕STOP 'FOO'
0 1
```

Existing stop controls in the function or operator named by Y are cancelled before new stop controls are set:

```
      +1 ⎕STOP 'FOO'
1
```

All stop controls may be cancelled by giving *X* an empty vector:

```
      ρ'' ⎕STOP 'FOO'
0


      ρθ ⎕STOP 'FOO'
0
```

Attempts to set stop controls in a locked function or operator are ignored.

```
      ⎕LOCK'FOO'

      +0 1 ⎕STOP'FOO'
```

The effect of ⎕*STOP* when a function or operator is invoked is to suspend execution at the beginning of any line in the function or operator on which a stop control is placed immediately before that line is executed, and immediately before exiting from the function or operator if a stop control of 0 is set. Execution may be resumed by a branch expression. A stop control interrupt (1001) may also be trapped - see the ⎕*TRAP* system variable.

Example

```
      ⎕FX'R←FOO' 'R←10'

      0 1 ⎕STOP'FOO'

      FOO
FOO[1]

      R
VALUE ERROR
      R
      ^

      →1
FOO[0]

      R
10

      →⎕LC
10
```

Query Stop: $R \leftarrow \Box STOP\ Y$

Y must be a simple character scalar or vector which is taken to be the name of a visible defined function or operator. R is a simple non-negative integer vector of the line numbers of the function or operator named by Y on which stop controls are set, shown in ascending order. The value 0 in R indicates that a stop control is set immediately prior to exit from the function or operator.

Example

```
      ☐STOP'FOO'
0 1
```

Set Access Control: $R \leftarrow X\ \ \Box SVC\ \ Y$

This system function sets access control on one or more shared variables.

Y is a character scalar, vector, or matrix containing names of shared variables. Each name may optionally be paired with its surrogate. If so, the surrogate must be separated from the name by at least one space.

X may be a 4-element boolean vector which specifies the access control to be applied to all of the shared variables named in Y. Alternatively, X may be a 4-column boolean matrix whose rows specify the access control for the corresponding name in Y. X may also be a scalar or a 1-element vector. If so, it treated as if it were a 4-element vector with the same value in each element.

Each shared variable has a current access control vector which is a 4-element boolean vector. A 1 in each of the four positions has the following impact :

[1]   You cannot **set** a new value for the shared variable until after an intervening use or set by your partner.

[2]   Your partner cannot **set** a new value for the shared variable until after an intervening use or set by you.

[3]   You cannot **use** the value of the shared variable until after an intervening set by your partner.

[4]   Your partner cannot **use** the value of the shared variable until after an intervening set by you.

The effect of ⎕*SVC* is to reset the access control vectors for each of the shared variables named in *Y* by OR-ing the values most recently specified by your partner with the values in *X*. This means that you cannot reset elements of the control vector which your partner has set to 1.

Note that the initial value of your partner's access control vector is normally 0 0 0 0. However, if it is a non-APL client application that has established a hot DDE link, its access control vector is defined to be 1 0 0 1. This inhibits either partner from setting the value of the shared variable twice, without an intervening use (or set) by the other. This prevents loss of data which is deemed to be desirable from the nature of the link. (An application that requests a hot link is assumed to require every value of the shared variable, and not to miss any). Note that APL's way of inhibiting another application from setting the value twice (without an intervening use) is to delay the acknowledgement of the DDE message containing the second value until the variable has been used by the APL workspace. An application that waits for an acknowledgement will therefore hang until this happens. An application that does not wait will carry on obliviously.

The result *R* is a boolean vector or matrix, corresponding to the structure of *X*, which contains the new access control settings. If *Y* refers to a name which is not a shared variable, or if the surrogate name is mis-spelt, the corresponding value in *R* is 4⍴0.

Examples

```
      1 0 0 1 ⎕SVC 'X'
1 0 0 1

      1 ⎕SVC 'X EXTNAME'
1 1 1 1

      (2 4⍴1 0 0 1 0 1 1 0) ⎕SVC ↑'ONE' 'TWO'
1 1 1 1
0 1 1 0
```

Query Access Control: $R \leftarrow \square SVC\ Y$

This system function queries the access control on one or more shared variables.

Y is a character scalar, vector, or matrix containing names of shared variables. Each name may optionally be paired with its surrogate. If so, the surrogate must be separated from the name by at least one space.

If Y specifies a single name, the result R is a boolean vector containing the current effective access control vector. If Y is a matrix of names, R is a boolean matrix whose rows contain the current effective access control vectors for the corresponding row in Y.

For further information, see the preceding section on setting the access control vector.

Example

```
      □SVC 'X'
0 0 0 0
```

Shared Variable Offer: $R \leftarrow X\ \square SVO\ Y$

This system function offers to share one or more variables with another APL workspace or with another application. Shared variables are implemented using Dynamic Data Exchange (**DDE**) and may be used to communicate with any other application that supports this protocol. See *Interface Guide* for further details.

Y is a character scalar, vector or matrix. If it is a vector it contains a name and optionally an external name or surrogate. The first name is the name used internally in the current workspace. The external name is the name used to make the connection with the partner and, if specified, must be separated from the internal name by one or more blanks. If the partner is another application, the external name corresponds to the DDE **item** specified by that application. If the external name is omitted, the internal name is used instead. The internal name must be a valid APL name and be either undefined or be the name of a variable. There are no such restrictions on the content of the external name.

Instead of an external name, Y may contain the special symbol '♣' separated from the (internal) name by a blank. This is used to implement a mechanism for sending DDE_EXECUTE messages, and is described at the end of this section.

If Y is a scalar, it specifies a single 1-character name. If Y is a matrix, each row of Y specifies a name and an optional external name as for the vector case.

The left argument *X* is a character vector or matrix. If it is a vector, it contains a string that defines the **protocol**, the **application** to which the shared variable is to be connected, and the **topic** of the conversation. These three components are separated by the characters ':' and '|' respectively. The protocol is currently always '*DDE*', but future implementations of Dyalog APL may support additional communications protocols if applicable. If *Y* specifies more than one name, *X* may be a vector or a matrix with one row per row in *Y*.

If the shared variable offer is a general one (server), *X*, or the corresponding row of *X*, should contain '*DDE:*'.

The result *R* is a numeric scalar or vector with one element for each name in *Y* and indicates the "degree of coupling". A value of 2 indicates that the variable is fully coupled (via a warm or hot DDE link) with a shared variable in another APL workspace, or with a DDE item in another application. A value of 1 indicates that there is no connection, or that the second application rejected a warm link. In this case, a transfer of data may have taken place (via a cold link) but the connection is no longer open. Effectively, APL treats an application that insists on a cold link as if it immediately retracts the sharing after setting or using the value, whichever is appropriate.

Examples

```
      'DDE:' ⎕SVO 'X'
1

      'DDE:' ⎕SVO 'X SALES_92'
1

      'DDE:' ⎕SVO ↑'X SALES_92' 'COSTS_92'
1 1

      'DDE:DYALOG|SERV_WS' ⎕SVO 'X'
2

      'DDE:EXCEL|SHEET1' ⎕SVO 'DATA R1C1:R10C12'
2
```

A special syntax is used to provide a mechanism for sending DDE_EXECUTE messages to another application. This case is identified by specifying the '⍎' symbol in place of the external name. The subsequent assignment of a character vector to a variable shared with the external name of '⍎' causes the value of the variable to be transmitted in the form of a DDE_EXECUTE message. The value of the variable is then reset to 1 or 0 corresponding to a positive or negative acknowledgement from the partner. In most (if not all) applications, commands transmitted in DDE_EXECUTE messages must

be enclosed in square brackets [ ]. For details, see the relevant documentation for the external application.

Examples :

```
      'DDE:EXCEL|SYSTEM' ⎕SVO 'X ⍀'
2


      X←'[OPEN("c:\mydir\mysheet.xls")]'
      X
1


      X←'[SELECT("R1C1:R5C10")]'
      X
1
```

Query Degree of Coupling: $R←⎕SVO\ Y$

This system function returns the current degree of coupling for one or more shared variables.

Y is a character scalar, vector or matrix. If it is a vector it contains a shared variable name and optionally its external name or surrogate separated from it by one of more blanks.

If Y is a scalar, it specifies a single 1-character name. If Y is a matrix, each row of Y specifies a name and an optional external name as for the vector case.

If Y specifies a single name, the result R is a 1-element vector whose value 0, 1 or 2 indicates its current degree of coupling. If Y specifies more than one name, R is a vector whose elements indicate the current degree of coupling of the variable specified by the corresponding row in Y. A value of 2 indicates that the variable is fully coupled (via a warm or hot DDE link) with a shared variable in another APL workspace, or with a DDE item in another application. A value of 1 indicates that you have offered the variable but there is no such connection, or that the second application rejected a warm link. In this case, a transfer of data may have taken place (via a cold link) but the connection is no longer open. A value of 0 indicates that the name is not a shared variable.

Examples

```
      ⎕SVO 'X'
2
      ⎕SVO ↑'X SALES' 'Y' 'JUNK'
2 1 0
```

Shared Variable Query: $R \leftarrow \Box SVQ \ Y$

This system function is implemented for compatibility with other versions of APL but currently performs no useful function. Its purpose is to obtain a list of outstanding shared variable offers made to you, to which you have not yet responded.

Using DDE as the communication protocol, it is not possible to implement $\Box SVQ$ effectively.

Shared Variable Retract Offer: $R \leftarrow \Box SVR \ Y$

This system function terminates communication via one or more shared variables, or aborts shared variable offers that have not yet been accepted.

Y is a character scalar, vector or matrix. If it is a vector it contains a shared variable name and optionally its external name or surrogate separated from it by one of more blanks. If Y is a scalar, it specifies a single 1-character name. If Y is a matrix, each row of Y specifies a name and an optional external name as for the vector case.

If Y specifies a single name, the result R is a 4-element vector indicating the state of the variable prior to retraction. If Y specifies more than one name, R is a matrix whose rows indicate the previous state of the variable specified by the corresponding row in Y.

See $\Box SVS$ for further information on the possible states of a shared variable.

Shared Variable State: $R \leftarrow \Box SVS \ Y$

This system function returns the current state of one or more shared variables.

Y is a character scalar, vector or matrix. If it is a vector it contains a shared variable name and optionally its external name or surrogate separated from it by one of more blanks. If Y is a scalar, it specifies a single 1-character name. If Y is a matrix, each row of Y specifies a name and an optional external name as for the vector case.

If Y specifies a single name, the result R is a 4-element vector indicating its current state. If Y specifies more than one name, R is a matrix whose rows indicate the current state of the variable specified by the corresponding row in Y.

There are four possible shared variable states :

0 0 1 1 :   means that you and your partner are both aware of the current value, and
                 neither has since reset it. This is also the initial value of the state when
                 the link is first established.

1 0 1 0 :   means that you have reset the shared variable and your partner has not
                 yet used it. This state can only occur if both partners are APL
                 workspaces.

0 1 0 1 :   means that your partner has reset the shared variable but that you have
                 not yet used it.

0 0 0 0 :   the name is not that of a shared variable.

Examples

```
      ⎕SVS 'X'
0 1 0 1

      ⎕SVS ↑'X SALES' 'Y' 'JUNK'
0 0 1 1
1 0 1 0
0 0 0 0
```

Terminal Control: (⎕ML)   R←⎕TC

This is a simple three element vector. If ⎕ML < 3 this is ordered as follows:

⎕TC[1] - Backspace
⎕TC[2] - Linefeed
⎕TC[3] - Newline

Note that ⎕TC ≡ ⎕AV[⎕IO+ι3] for ⎕ML< 3 .

If ⎕ML ≥ 3 the order of the elements of ⎕TC is instead compatible with IBM's APL2:

⎕TC[1] - Backspace
⎕TC[2] - Newline
⎕TC[3] - Linefeed

Elements of ⎕TC beyond 3 are not defined but are reserved.

Thread Child Numbers: $R \leftarrow \Box TCNUMS \ Y$

Y must be a simple array of integers representing thread numbers.

The result R is a simple integer vector of the child threads of each thread of Y.

Examples

```
      ⎕TCNUMS 0
2 3

      ⎕TCNUMS 2 3
4 5 6 7 8 9
```

Current Thread Identity: $R \leftarrow \Box TID$

R is a simple integer scalar whose value is the number of the current thread.

Examples

```
      ⎕TID     ⍝ Base thread number
0

      ⎕&'⎕TID' ⍝ Thread number of async ⎕.
1
```

Kill Thread: $\{R\}\leftarrow\{X\}\square TKILL\ Y$

Y must be a simple array of integers representing thread numbers to be terminated. X is a boolean single, defaulting to 1, which indicates that all descendant threads should also be terminated.

The shy result R is a vector of the numbers of all threads that have been terminated.

The **base thread** 0 is always excluded from the cull.

Examples

```
      ⎕TKILL 0          ⍝ Kill background threads.

      ⎕TKILL ⎕TID        ⍝ Kill self and descendants.

    0 ⎕TKILL ⎕TID        ⍝ Kill self only.

      ⎕TKILL ⎕TCNUMS ⎕TID ⍝ Kill descendants.
```

Thread Numbers: $R\leftarrow\square TNUMS$

$\square TNUMS$ reports the numbers of all current threads.

R is a simple integer vector of the base thread and all its living descendants.

Example

```
      ⎕TNUMS
0 2 4 5 6 3 7 8 9
```

Set Trace: $\{R\}\leftarrow X \ \Box TRACE \ Y$

Y must be a simple character scalar or vector which is taken to be the name of a visible defined function or operator. X must be a simple non-negative integer scalar or vector. R is a simple integer vector of non-negative elements.

X identifies the numbers of lines in the function or operator named by Y on which a trace control is to be placed. Numbers outside the range of line numbers in the function or operator (other than 0) are ignored. The number 0 indicates that a trace control is to be placed immediately prior to exit from the function or operator. The value of X is independent of $\Box IO$.

Example

```
      +(0,ι10) □TRACE'FOO'
0 1
```

Existing trace controls in the function or operator named by Y are cancelled before new trace controls are set:

```
      + 1 □TRACE'FOO'
1
```

All trace controls may be cancelled by giving X an empty vector:

```
      ρθ □TRACE 'FOO'
0
```

Attempts to set trace controls in a locked function or operator are ignored.

```
      □LOCK 'FOO'
      +1 □TRACE 'FOO'
```

The effect of trace controls when a function or operator is invoked is to display the result of each complete expression for lines with trace controls as they are executed, and the result of the function if trace control 0 is set. If a line contains expressions separated by ◇, the result of each complete expression is displayed for that line after execution.

The result of a complete expression is displayed even where the result would normally be suppressed. In particular:

1. the result of a branch statement is displayed;
2. the result ("pass-through" value) of assignment is displayed;
3. the result of a function whose result would normally be suppressed is displayed;

For each traced line, the output from $\Box TRACE$ is displayed as a two element vector, the first element of which contains the function or operator name and line number, and the second element of which takes one of two forms.

1. The result of the line, displayed as in standard output.
2. → followed by a line number.

Example

```
      ⎕VR 'DSL'
     ∇ R←DSL SKIP;A;B;C;D
[1]    A←2×3+4
[2]    B←(2 3ρ'ABCDEF')A
[3]    →NEXT×ιSKIP
[4]    'SKIPPED LINE'
[5]   NEXT:C←'one' ◇ D←'two'
[6]   END:R←C D
     ∇


      (0,ι6) ⎕TRACE 'DSL'


      DSL 1
 DSL[1]  14
 DSL[2]   ABC    14
          DEF
 DSL[3]  →5
 DSL[5]  one
 DSL[5]  two
 DSL[6]   one    two
 DSL[0]   one    two
 one  two
```

Query Trace: $R←\Box TRACE\ Y$

Y must be a simple character scalar or vector which is taken to be the name of a visible defined function or operator. R is a simple non-negative integer vector of the line numbers of the function or operator named by Y on which trace controls are set, shown in ascending order. The value 0 in R indicates that a trace control is set to display the result of the function or operator immediately prior to exit.

Example

```
      ⎕TRACE'DSL'
0 1 2 3 4 5 6
```

Trap Event: $\Box TRAP$

This is a non-simple vector. An item of $\Box TRAP$ specifies an action to be taken when one of a set of events occurs. An item of $\Box TRAP$ is a 2 or 3 element vector whose items are simple scalars or vectors in the following order:

1. an integer vector whose value is one or more event codes selected from the list in the Figure on the following two pages.
2. a character scalar whose value is an action code selected from the letters C, E, N or S.
3. if element 2 is the letter C or E, this item is a character vector forming a valid APL expression or series of expressions separated by $\diamond$. Otherwise, this element is omitted.

An EVENT may be an APL execution error, an interrupt by the user or the system, a control interrupt caused by the $\Box STOP$ system function, or an event generated by the $\Box SIGNAL$ system function.

When an event occurs, the system searches for a trap definition for that event. The most local $\Box TRAP$ value is searched first, followed by successive shadowed values of $\Box TRAP$, and finally the global $\Box TRAP$ value. Separate actions defined in a single $\Box TRAP$ value are searched from **left to right**. If a trap definition for the event is found, the defined action is taken. Otherwise, the normal system action is followed.

The ACTION code identifies the nature of the action to be taken when an associated event occurs. Permitted codes are interpreted as follows:

C Cutback   The state indicator is 'cut back' to the environment in which the $\Box TRAP$ is locally defined (or to immediate execution level). The APL expression in element 3 of the same $\Box TRAP$ item is then executed.

E Execute   The APL expression in element 3 of the same $\Box TRAP$ item is executed in the environment in which the event occurred.

N Next   The event is excluded from the current $\Box TRAP$ definition. The search will continue through further localised definitions of $\Box TRAP$.

S Stop   Stops the search and causes the normal APL action to be taken in the environment in which the event occurred.

Code	Event
0	**Any event** in range 1-999
1	WS FULL
2	SYNTAX ERROR
3	INDEX ERROR
4	RANK ERROR
5	LENGTH ERROR
6	VALUE ERROR
7	FORMAT ERROR
10	LIMIT ERROR
11	DOMAIN ERROR
12	HOLD ERROR
16	NONCE ERROR
18	FILE TIE ERROR
19	FILE ACCESS ERROR
20	FILE INDEX ERROR
21	FILE FULL
22	FILE NAME ERROR
23	FILE DAMAGED
24	FILE TIED
25	FILE TIED REMOTELY
26	FILE SYSTEM ERROR
28	FILE SYSTEM NOT AVAILABLE
30	FILE SYSTEM TIES USED UP
31	FILE TIE QUOTA USED UP
32	FILE NAME QUOTA USED UP
34	FILE SYSTEM NO SPACE
35	FILE ACCESS ERROR - CONVERTING FILE
38	FILE COMPONENT DAMAGED

Figure 5(i): $\square TRAP$ **Event Codes**

See *Chapter 7* for further details.

Code	Event
52	FIELD CONTENTS RANK ERROR
53	FIELD CONTENTS TOO MANY COLUMNS
54	FIELD POSITION ERROR
55	FIELD SIZE ERROR
56	FIELD CONTENTS/TYPE MISMATCH
57	FIELD TYPE/BEHAVIOUR UNRECOGNISED
58	FIELD ATTRIBUTES RANK ERROR
59	FIELD ATTRIBUTES LENGTH ERROR
60	FULL-SCREEN ERROR
61	KEY CODE UNRECOGNISED
62	KEY CODE RANK ERROR
63	KEY CODE TYPE ERROR
70	FORMAT FILE ACCESS ERROR
71	FORMAT FILE ERROR
72	NO PIPES
76	PROCESSOR TABLE FULL
84	TRAP ERROR
90	EXCEPTION
200-499	Reserved for distributed auxiliary processors
500-999	User-defined events
1000	**Any event** in range 1001-1006
1001	Stop vector
1002	Weak interrupt
1003	INTERRUPT
1005	EOF INTERRUPT
1006	TIMEOUT
1007	RESIZE (Dyalog APL/X, Dyalog APL/W)
1008	DEADLOCK

Figure 5(i): □*TRAP* Event Codes (Continued)

See :*Trap* as an alternative 'control structured' error trapping mechanism.

Examples

```
      ⎕TRAP←⊂(3 4 5) 'E' 'ERROR' ◇ ρ⎕TRAP
1


      ⎕TRAP
 3 4 5 E ERROR
```

Items may be specified as scalars. If there is only a single trap definition, it need not be enclosed. However, the value of ⎕*TRAP* will be rigorously correct:

```
      ⎕TRAP←11 'E' '→LAB'

      ⎕TRAP
 11 E →ERR

      ρ⎕TRAP
1
```

The value of ⎕*TRAP* in a clear workspace is an empty vector whose prototype is 0ρ(θ '' ''). A convenient way of cancelling a ⎕*TRAP* definition is:

```
      ⎕TRAP←0ρ⎕TRAP
```

Event codes 0 and 1000 allow all events in the respective ranges 1-999 and 1000-1006 to be trapped. Specific event codes may be excluded by the *N* action (which must precede the general event action):

```
      ⎕TRAP←(1 'N')(0 'E' '→GENERR')
```

The 'stop' action is a useful mechanism for cancelling trap definitions during development of applications.

The 'cut-back' action is useful for returning control to a known point in the application system when errors occur. The following example shows a function that selects and executes an option with a general trap to return control to the function when an untrapped event occurs:

```
      ∇ SELECT;OPT;□TRAP
[1]    ⍝ Option selection and execution
[2]    ⍝ A general cut-back trap
[3]     □TRAP←(0 1000)'C' '→ERR'
[4]    INP:□←'OPTION : ' ◇ OPT←(OPT≠' ')/OPT←9↓□
[5]     →EXρ⍨(⊂OPT)∊OPTIONS ◇ 'INVALID OPTION' ◇ →INP
[6]    EX:⍕OPT ◇ →INP
[7]    ERR:ERRORΔACTION ◇ →INP
[8]    END:
      ∇
```

User-defined events may be signalled through the □*SIGNAL* system function. A user-defined event (in the range 500-999) may be trapped explicitly, or implicitly by the event code 0.

Example

```
      □TRAP←500 'E' '''USER EVENT 500 - TRAPPED'''

      □SIGNAL 500
USER EVENT 500 - TRAPPED
```

Time Stamp: $R←□TS$

This is a seven element vector which identifies the clock time set on the particular installation as follows:

> □*TS*[1]   - Year
> □*TS*[2]   - Month
> □*TS*[3]   - Day
> □*TS*[4]   - Hour
> □*TS*[5]   - Minute
> □*TS*[6]   - Second
> □*TS*[7]   - Millisecond

Example

```
      □TS
1989 7 11 10 42 59 123
```

Note that on some systems, where time is maintained only to the nearest second, a zero is returned for the seventh (millisecond) field.

Wait for Threads to Terminate: $R \leftarrow \Box TSYNC \ Y$

Y must be a simple array of thread numbers.

If *Y* is a simple scalar, *R* is an array, the result (if any) of the thread.

If *Y* is a simple non-scalar, *R* has the same shape as *Y*, and result is an array of enclosed thread results.

The interpreter detects a potential deadlock if a number of threads wait for each other in a cyclic dependency. In this case, the thread that attempts to cause the deadlock issues error number 1008: *DEADLOCK*.

Examples

```
      dup←{ω ω}          ⍝ Duplicate

      ⎕←dup&88           ⍝ Show thread number
11
88 88

      ⎕TSYNC dup&88     ⍝ Wait for result
88 88

      ⎕TSYNC,dup&88
 88 88

      ⎕TSYNC dup&1 2 3
 1 2 3  1 2 3

      ⎕TSYNC dup&¨1 2 3
 1 1  2 2  3 3

      ⎕TSYNC ⎕TID       ⍝ Wait for self
DEADLOCK
      ⎕TSYNC ⎕TID
      ∧

      ⎕EN
1008
```

Using (Microsoft .Net Search Path):   `⎕USING`

`⎕USING` specifies a list of Microsoft .Net Namespaces that are to be searched for a reference to a .Net class.

`⎕USING` is a vector of character vectors, each element of which specifies the name of a .Net Namespace followed optionally by a comma (,) and the Assembly in which it is to be found.

If the Assembly is defined in the *global assembly cache*, you need only specify its name. If not, you must specify a full or relative pathname.

If the Microsoft .Net Framework is installed, the System namespace in `mscorlib.dll` is automatically loaded when Dyalog APL starts. To access this namespace, it is not necessary to specify the name of the Assembly.

`⎕USING` has *namespace scope*. If the local value of `⎕USING` is anything other than empty, and you reference a name that would otherwise generate a *VALUE ERROR*, APL searches the list of .Net Namespaces and Assemblies specified by `⎕USING` for a class of that name. If it is found, an entry for the class is added to the symbol table in the current space and the class is used as specified. Note that subsequent references to that class in the current space will be identified immediately.

If `⎕USING` is empty (its default value in a *CLEAR WS*) no such search is performed.

Note that when you assign a value to `⎕USING`, you may specify a simple character vector or a vector of character vectors.

Examples:

```
  ⎕USING←'System'
  DISPLAY ⎕USING
.→---------.
| .→-----. |
| |System| |
| '------' |
'∊---------'

  ⎕USING,←⊂'System.Windows.Forms,System.Windows.Forms.dll'
  ⎕USING,←⊂'System.Drawing,System.Drawing.dll'
```

An Assembly may contain top-level classes which are not packaged into .Net Namespaces. In this case, you omit the Namespace name. For example:

```
⎕USING←,⊂'.\LoanService.dll'
```

Verify & Fix Input: $R \leftarrow \{X\} \square VFI \ Y$

Y must be a simple character scalar or vector. X is optional. If present, X must be a
simple character scalar or vector. R is a nested vector of length two whose first item is
a simple logical vector and whose second item is a simple numeric vector of the same
length as the first item of R.

Y is the character representation of a series of numeric constants. If X is omitted,
adjacent numeric strings are separated by one or more blanks. Leading and trailing
blanks and separating blanks in excess of one are redundant and ignored. If X is
present, X specifies one or more alternative separating characters. Blanks in leading
and trailing positions in Y and between numeric strings separated also by the
character(s) in X are redundant and ignored. Leading, trailing and adjacent occurrences
of the character(s) in X are not redundant. The character 0 is implied in Y before a
leading character, after a trailing character, and between each adjacent pair of
characters specified by X.

The length of the items of R is the same as the number of identifiable strings (or
implied strings) in Y separated by blank or the value of X. An element of the first item
of R is 1 where the corresponding string in Y is a valid numeric representation, or 0
otherwise. An element of the second item of R is the numeric value of the
corresponding string in Y if it is a valid numeric representation, or 0 otherwise.

Examples

```
      □VFI '12.1 1E1 1A1 ¯10'
 1 1 0 1   12.1 10 0 ¯10

      ⊃(//□VFI'12.1 1E1 1A1 ¯10')
12.1 10 ¯10

      ',' □VFI'3.9,2.4,,76,'
 1 1 1 1 1   3.9 2.4 0 76 0

      '◇' □VFI'1 ◇ 2 3 ◇ 4 '
 1 0 1  1 0 4
      θθ≡□VFI''
1
```

Vector Representation: $R \leftarrow \Box VR \ Y$

Y must be a simple character scalar or vector which represents the name of a function or defined operator.

If *Y* is the name of a defined function or defined operator, *R* is a simple character vector containing a character representation of the function or operator with each line except the last terminated by the newline character ($\Box TC[3]$). Its display form is as follows:

1. the header line starts at column 8 with the ∇ symbol in column 6,

2. the line number for each line of the function starts in column 1,

3. the statement contained in each line starts at column 8 except for labelled lines or lines beginning with ⍝ which start at column 7,

4. the header line and statements contain no redundant blanks beyond column 7 except that the ◇ separator is surrounded by single blanks, control structure indentation is preserved and comments retain embedded blanks as originally defined,

5. the last line shows only the ∇ character in column 6.

If *Y* is the name of a variable, a locked function or operator, an external function, or is undefined, *R* is an empty vector.

Example

```
      ρV←□VR'PLUS'
1 2 8

      V
    ∇ R←{A}PLUS B
[1]   ⍝ MONADIC OR DYADIC +
[2]    →DYADICρ⍨2=□NC'A'  ◇  R←B  ◇  →END
[3]   DYADIC:R←A+B  ◇  →END
[4]   END:
    ∇
```

The definition of ⎕*VR* has been extended to names assigned to functions by specification (←), and to local names of functions used as operands to defined operators. In these cases, the result of ⎕*VR* is identical to that of ⎕*CR* except that the representation of defined functions and operators is as described above.

Example

```
      AVG←MEAN∘,

      +F←⎕VR'AVG'
      ∇ R←MEAN X      ⍝ Arithmetic mean
[1]    R←(+/X)÷ρX
      ∇ ∘,

      ρF
3

      DISPLAY F
.→-----------------------------------------.
| .→---------------------------------.     |
| |     ∇ R←MEAN X    ⍝ Arithmetic mean| ∘ , |
| |[1]    R←(+/X)÷ρX                  | - - |
| |     ∇                            |     |
| '---------------------------------'     |
'∈-----------------------------------------'
```

Workspace Available: *R←⎕WA*

This is a simple integer scalar. It identifies the total available space in the active workspace area given as the number of bytes it could hold.

A side effect of using ⎕*WA* is an internal reorganisation of the workspace; a process know as COMPACTION when all objects are shuffled to one end. After a compaction there is only one free block rather than many small ones.

Example

```
      ⎕WA
261412
```

Windows Create Object: $\qquad\qquad \{R\}\leftarrow\{X\}\Box WC\ Y$

This system function creates a GUI **object**. Y is either a vector which specifies **properties** that determine the new object's appearance and behaviour, or the $\Box OR$ of a GUI object that exists or previously existed. X is a character vector which specifies the name of the new object, and its position in the object hierarchy.

If X is omitted, $\Box WC$ attaches a GUI component to the current namespace, retaining any functions, variables and other namespaces that it may contain. Monadic $\Box WC$ is discussed in detail at the end of this section.

If Y is a nested vector each element specifies a property. The $Type$ property (which specifies the class of the object) **must** be specified. Most other properties take default values and need not be explicitly stated. Properties (including $Type$) may be declared either positionally or with a keyword followed by a value. Note that $Type$ must always be the first property specified. Properties are specified positionally by placing their values in Y in the order prescribed for an object of that type.

If Y is a result of $\Box OR$, the new object is a complete copy of the one from which the $\Box OR$ was made, including any child objects, namespaces, functions and variables that it contained at that time.

The shy result R is the full name (starting #. or $\Box SE$.) of the namespace X.

An object's name is specified by giving its full pathname in the object hierarchy. At the top of the hierarchy is the $Root$ object whose name is ".". Below "." there may be one or more "top-level" objects. The names of these objects follow the standard rules for other APL objects as described in *Chapter 1*.

Names for sub-objects follow the same rules except that the character "." is used as a delimiter to indicate parent/child relationships.

The following are examples of legal and illegal names :

Legal	Illegal
*FORM*1	*FORM* 1
*form_*23	*form#*1
*Form*1.*Gp*1	1_*Form*
*F*1.*g*2.*b*34	*Form*+1

If *X* refers to the name of an APL variable, label, function, or operator, a *DOMAIN ERROR* is reported. If *X* refers to the name of an existing GUI object or namespace, the existing one is replaced by the new one. The effect is the same as if it were deleted first.

If *Y* refers to a non-existent property, or to a property that is not defined for the type of object *X*, a *DOMAIN ERROR* is reported. A *DOMAIN ERROR* is also reported if a value is given that is inconsistent with the corresponding property. This can occur for example, if *Y* specifies values positionally and in the wrong order.

A "top-level" object created by *⎕WC* whose name is localised in a function/operator header, is deleted on exit from the function/operator. All objects, including sub-objects, can be deleted using *⎕EX*.

GUI objects are named **relative** to the current namespace, so the following examples are equivalent:

```
      'F1.B1' ⎕WC 'Button'
```

is equivalent to :

```
      )CS F1
#.F1
      'B1' ⎕WC 'Button'
      )CS
#
```

is equivalent to :

```
      'B1' F1.⎕WC 'Button'
```

Examples

```
⍝ Create a default Form called F1

      'F1' ⎕WC 'Form'

⍝ Create a Form with specified properties (by position)
⍝   Caption = "My Application"  (Title)
⍝   Posn    = 10 30  (10% down, 30% across)
⍝   Size    = 80 60  (80% high, 60% wide)

      'F1' ⎕WC 'Form' 'My Application' (10 30)(80 60)

⍝ Create a Form with specified properties (by keyword)
⍝   Caption = "My Application"  (Title)
⍝   Posn    = 10 30  (10% down, 30% across)
⍝   Size    = 80 60  (80% high, 60% wide)

      PROPS←⊂'Type' 'Form'
      PROPS,←⊂'Caption' 'My Application'
      PROPS,←⊂'Posn' 10 30
      PROPS,←⊂'Size' 80 60
      'F1' ⎕WC PROPS

⍝ Create a default Button (a pushbutton) in the Form F1

      'F1.BTN' ⎕WC 'Button'

⍝ Create a pushbutton labelled "Ok"
⍝ 10% down and 10% across from the start of the FORM
⍝ with callback function FOO associated with EVENT 30
⍝ (this event occurs when the user presses the button)

      'F1.BTN'⎕WC'Button' '&Ok' (10 10)('Event' 30 'FOO')
```

Monadic *⎕WC* is used to *attach* a GUI component to an existing object. The existing object must be a pure namespace or a GUI object. The operation may be performed by changing space to the object or by running *⎕WC inside* the object using the *dot* syntax. For example, the following statements are equivalent.

```
      )CS F
#.F
      ⎕WC 'Form'  ⍝ Attach a Form to this namespace
```

```
        )CS
#
        F.⎕WC'Form' ⍝ Attach a Form to namespace F
```

Windows Get Property: $R \leftarrow \{X\} \square WG \ Y$

This system function returns property values for a GUI object.

X is a namespace reference or a character vector containing the name of the object. Y is a character vector or a vector of character vectors containing the name(s) of the properties whose values are required. The result R contains the current values of the specified properties. If Y specifies a single property name, a single property value is returned. If Y specifies more than one property, R is a vector with one element per name in Y.

If X refers to a non-existent GUI name, a *VALUE ERROR* is reported. If Y refers to a non-existent property, or to a property that is not defined for the type of object X, a *DOMAIN ERROR* is reported.

GUI objects are named **relative** to the current namespace. A null value of X (referring to the namespace in which the function is being evaluated) may be omitted. The following examples are equivalent:

```
'F1.B1' ⎕WG 'Caption'
'B1' F1.⎕WG 'Caption'
'' F1.B1.⎕WG 'Caption'
F1.B1.⎕WG 'Caption'
```

Examples

```
'F1' ⎕WC 'Form' 'TEST'

'F1' ⎕WG 'Caption'
TEST

'F1' ⎕WG 'MaxButton'
1

'F1' ⎕WG 'Size'
50 50

DISPLAY 'F1' ⎕WG 'Caption' 'MaxButton' 'Size'
.→-----------------.
| .→---.    .→----. |
| |TEST| 1 |50 50| |
| '----'    '~----' |
'∈-----------------'
```

Windows Child Names: $R \leftarrow \{X\} \square WN \quad Y$

This system function reports the names of the objects whose parent is Y. This is a character scalar or vector containing the name of an existing GUI object.

The optional left argument X is a character vector which specifies the $Type$ of object to be reported.

The result R is a vector of character vectors containing the names of the direct children of Y which are of $Type$ X, or all of them if X is not specified. The names of objects further down the tree are not returned, but can be obtained by recursive use of $\square WN$.

If Y refers to a non-existent GUI name, a $VALUE\ ERROR$ is reported.

Note that $\square WN$ reports **only** those child objects visible from the current thread.

GUI objects are named **relative** to the current namespace. The following examples are equivalent:

```
     ⎕WN 'F1.B1'
     F1.⎕WN 'B1'
     F1.B1.⎕WN ''
```

Examples

```
     ⎕WN'.'
 FORM1   MENU1

     ⎕WN'FORM1'
 CANCEL   GROUP1   MENUBAR   OK

     'Button'⎕WN'FORM1'
 CANCEL   OK
```

Windows Set Property: {*X*}⎕*WS Y*

This system function resets property values for a GUI object.

X is a namespace reference or a character vector containing the name of the object. *Y* defines the property or properties to be changed and the new value or values. If a single property is to be changed, *Y* is a vector whose first element *Y*[1] is a character vector containing the property name. If *Y* is of length 2, *Y*[2] contains the corresponding property value. However, if the property value is itself a numeric or nested vector, its elements may be may be specified in *Y*[2 3 4 ...] instead of as a single nested element in *Y*[2]. If *Y* specifies more than one property, they may be declared either positionally or with a keyword followed by a value. Properties are specified positionally by placing their values in *Y* in the order prescribed for an object of that type. Note that the first property in *Y* must always be specified with a keyword because the *Type* property (which is expected first) may not be changed using ⎕*WS*.

If *X* refers to a non-existent GUI name, a *VALUE ERROR* is reported. If *Y* refers to a non-existent property, or to a property that is not defined for the type of object *X*, or to a property whose value may not be changed by ⎕*WS*, a *DOMAIN ERROR* is reported.

GUI objects are named **relative** to the current namespace. A null value of *X* (referring to the namespace in which the function is being evaluated) may be omitted. The following examples are equivalent:

```
'F1.B1' ⎕WS 'Caption' '&Ok'
'B1' F1.⎕WS 'Caption' '&Ok'
'' F1.B1.⎕WS 'Caption' '&Ok'
F1.B1.⎕WS 'Caption' '&Ok'
```

Examples

```
'F1' ⎕WC 'Form'   ⍝ A default Form

'F1' ⎕WS 'Active' 0

'F1' ⎕WS 'Caption' 'My Application'

'F1' ⎕WS 'Posn' 0 0

'F1' ⎕WS ('Active' 1)('Event' 'Configure' 'FOO')

'F1' ⎕WS 'Junk' 10
DOMAIN ERROR

'F1' ⎕WS 'MaxButton' 0
DOMAIN ERROR
```

Workspace Identification: $\Box WSID$

This is a simple character vector. It contains the identification name of the active workspace. If a new name is assigned, that name becomes the identification name of the active workspace, provided that it is a correctly formed name.

See page 1 for workspace naming conventions.

It is useful, though not essential, to associate workspaces with a specific directory in order to distinguish workspaces from other files.

The value of $\Box WSID$ in a clear workspace is $'CLEAR WS'$.

Example

```
      ⎕WSID
CLEAR WS

      ⎕WSID←'WS/MYWORK'      (UNIX)

      ⎕WSID←'B:\WS\MYWORK'   (DOS)
```

Window Expose: $\Box WX$

$\Box WX$ is a Boolean system variable that determines whether or not the names of properties, methods and events provided by a Dyalog APL GUI object are exposed.

If $\Box WX$ is 1, the names of properties, methods and events are exposed as reserved names in GUI namespaces and can be accessed directly by name. This means that the same names may not be used for global variables in GUI namespaces.

The value of $\Box WX$ in a clear workspace is defined by the default_wx parameter (see User Guide) which itself defaults to 1.

$\Box WX$ has namespace scope and may be localised in a function header. This allows you to create a utility namespace or utility function in which the exposure of GUI objects is known and determined, regardless of its global value in the workspace.

Extended State Indicator: $R \leftarrow \Box XSI$

R is a nested vector of character vectors giving the full path names of the functions or operators in the execution stack. Note that if a function has changed space, its original (home) space is reported, rather than its current one.

Example

In the following, function *foo* in namespace *x* has called *goo* in namespace *y*. Function *goo* has then changed space ($\Box CS$) to namespace *z* where it has been suspended:

```
      )si
[z] y.goo[2]*
x.foo[1]
```

$\Box XSI$ reports the full path name of each function:

```
      ⎕xsi
 #.y.goo  #.x.foo
```

This can be used for example, to edit all functions in the stack, irrespective of the current namespace by typing: $\Box ed \ \Box xsi$

See also $\Box SI$.

Set External Variable: *X* □*XT* *Y*

Y must be a simple character scalar or vector which is taken to be a variable name. *X* must be a simple character scalar or vector which is taken to be a file reference. The name given by *Y* is identified as an EXTERNAL VARIABLE associated with an EXTERNAL ARRAY whose value may be stored in file identified by *X*. See *User Guide* for file naming conventions under DOS and UNIX.

If *Y* is the name of a defined function or operator, a label or a namespace in the active workspace, a *DOMAIN ERROR* is reported.

Example

```
        'EXT\ARRAY' □XT 'V'
```

If the file reference does not exist, the external variable has no value until a value is assigned:

```
        V
VALUE ERROR
        V
        ^
```

A value assigned to an external variable is stored in file space, not within the workspace:

```
        □WA
2261186

        V←ι100000

        □WA
2261186
```

There are no specific restrictions placed on the use of external variables. They must conform to the normal requirements when used as arguments of functions or as operands of operators. The essential difference between a variable and an external variable is that an external variable requires only temporary workspace for an operation to accommodate (usually) a part of its value.

Examples

```
      +/V
15

      V[3]←⊂'ABC'

      V
1 2  ABC  4 5

      ρ¨V
      3
```

Assignment allows the structure or the value of an external variable to be changed without fully defining the external array in the workspace.

Examples

```
      V,←⊂2 4ρι8

      ρV
6

      V[6]
1 2 3 4
5 6 7 8

      V[1 2 4 5 6]×←10

      V
10 20  ABC  40 50  10 20 30 40
                   50 60 70 80
```

An external array is (usually) preserved in file space when the name of the external variable is disassociated from the file. It may be re-associated with any valid variable name.

Example

```
      ⎕EX'V'

      'EXT\ARRAY'⎕XT'F'

      F
10 20  ABC  40 50  10 20 30 40
                   50 60 70 80
```

In Unix versions, if X is an empty vector, the external array is associated with a temporary file which is erased when the array is disassociated.

Example

```
      ''⎕XT'TEMP'

      TEMP←ι10

      +/TEMP×TEMP
385

      ⎕EX'TEMP'
```

An external array may be erased using the native file function: *⎕NERASE*.

In a multi-user environment (UNIX or a DOS LAN) a new file associated with an external array is created with access permission for owner read/write. An existing file is opened for exclusive use (by the owner) if the permissions remain at this level. If the access permissions allow any other users to read and write to the file, the file is opened for shared use. In Unix versions, access permissions may be modified using the appropriate Operating System command, or in Windows using the supplied function *XVAR* from the UTIL workspace.

Query External Variable: *R←⎕XT Y*

Y must be a simple character scalar or vector which is taken to be a variable name. *R* is a simple character vector containing the file reference of the external array associated with the variable named by *Y*, or the null vector if there is no associated external array.

Example

```
      ⎕XT'V'
EXT\ARRAY

      ρ⎕XT'G'
0
```

CHAPTER 6

System Commands

System commands are **not** executable APL expressions. They provide services or information associated with the workspace and the **external environment**.

Command	Description
)*CLEAR*	Clear the workspace
)*CMD Y*	Execute a (DOS) Command
)*CONTINUE*	Save a Continue workspace and terminate APL
)*COPY {Y}*	Copy objects from another workspace
)*CS {Y}*	Change current namespace
)*DROP {Y}*	Drop named workspace
)*ED Y*	Edit object(s)
)*ERASE Y*	Erase object(s)
)*EVENTS*	List events of GUI namespace or object
)*FNS {Y}*	List user defined Functions
)*HOLDS*	Display Held tokens
)*LIB {Y}*	List workspaces in a directory
)*LOAD {Y}*	Load a workspace
)*NS {Y}*	Create a global Namespace
)*METHODS*	List methods in GUI namespace or object
)*OBJECTS {Y}*	List global namespaces
)*OBS {Y}*	List global namespaces (alternative form)
)*OFF*	Terminate the APL session
)*OPS {Y}*	List user defined Operators
)*PCOPY {Y}*	Perform Protected Copy of objects
)*PROPS*	List properties of GUI namespace or object
)*RESET*	Reset the state indicator
)*SAVE {Y}*	Save the workspace
)*SH {Y}*	Execute a (UNIX) Shell command
)*SI*	State Indicator
)*SINL*	State Indicator with local Name Lists
)*TID {Y}*	Switch current Thread Identity
)*VARS {Y}*	List user defined global Variables
)*WSID {Y}*	Workspace Identification
)*XLOAD Y*	Load a workspace; do not execute $\square LX$

{ } indicates that the parameter(s) denoted by *Y* are optional.

Figure 6(i) : System Commands

Command Presentation

System commands may be entered from immediate execution mode or in response to the prompt ⎕: within evaluated input. All system commands begin with the symbol `)`, known as a right parenthesis. All system commands may be entered in upper or lower case.   Each command is described in alphabetical order in this chapter.

Clear Workspace: `)CLEAR`

This command clears the active workspace and gives the report "`clear ws`". The active workspace is lost. The name of a clear workspace is `CLEAR WS`. System variables are initialised with their default values as described in *Chapter 5*.

In GUI implementations of Dyalog APL, `)CLEAR` expunges all GUI objects, discards any unprocessed events in the event queue and resets the properties of the `Root` object `'.'` to their default values.

Example

```
        )CLEAR
clear ws
```

Execute (DOS) Command: `)CMD cmd`

This command allows DOS or UNIX shell commands to be given from APL. `)CMD` is a synonym of `)SH`. Either command may be given in either environment (DOS or UNIX) with exactly the same effect. `)CMD` is probably more natural for the DOS user. This section describes the behaviour of `)CMD` and `)SH` under DOS/Windows. See `)SH` for a discussion of the behaviour of these commands under UNIX.

The system functions `⎕CMD` and `⎕SH` provide similar facilities but may be executed from within APL code.

The implementation of `)CMD` differs between Dyalog APL Version 7 and Version 8. Specifically, Version 8 takes advantage of Windows 95 and Windows NT facilities that are not available in Windows 3.x. The two different implementations are described below.

Version 7 Implementation

Under DOS/Windows, `)CMD` creates a temporary .PIF file called QCMD1234.PIF in the root directory of your default drive (usually C:). This file is initially a copy of WDYALOG\QUADCMD.PIF which is then dynamically changed to run *your* DOS shell together with the particular command that you have asked to be run; i.e. the right argument you supplied to `)CMD`. The DOS shell to be run is obtained from your COMSPEC environment variable. If this is not defined, it defaults to C:\COMMAND.COM.

The temporary .PIF file is then executed in a new (but **invisible**) window. The result of the command (if any) is captured by redirecting output to a temporary file and then reading and deleting the file on completion of the command. The result is then displayed in the APL session. The temporary file is created in the directory specified by your TMP environment variable, or in the root directory of the current drive if TMP is undefined. The file is called POPEN.TMP prefixed by your Task ID and the value of the **aplnid** entry in APL.INI.

Note that under MS-Windows, you may not execute `)CMD` without a command. If you wish to, you can easily open a new DOS window outside APL.

Example

```
     )CMD DIR

 Volume in drive C has no label
 Directory of C:\PETE\WS

.              <DIR>      5-07-94   3.02p
..             <DIR>      5-07-94   3.02p
SALES     DWS    110092 5-07-94   3.29p
EXPENSES DWS    154207 5-07-94   3.29p
```

If **cmd** issues prompts and expects user input, it is **ESSENTIAL** to explicitly redirect output to the console. If this is done, APL detects the presence of a " > " in the command line and runs COMMAND.COM in a visible window and does not direct output to the temporary file. If you fail to do this your system will appear to hang because there is no mechanism for you to receive or respond to the prompt.

Example

```
     )CMD DATE>CON
```
(DOS window appears)
```
Current date is Wed 19-07-1995
Enter new date (dd-mm-yy): 20-07-95
```
(DOS window disappears)

Errors

cannot create temp file
APL was unable to create a temporary .PIF file or could not find a shell to run.

could not execute command
APL has successfully started the .PIF file but has timed out waiting for the DOS
window to be created. The time-out is defined by the `qcmd_timeout=` parameter in
APL.INI (default 5000 msecs).

not enough memory
not enough memory or corrupt shell
invalid shell

Version 8 Implementation

Under Version 8, the argument of `)CMD` is simply passed to the appropriate command
processor for execution and its output is received using an *unnamed pipe*.

By default, `)CMD` will execute the string `('command.com /c',Y)` under
Windows 95 and the string `('cmd.exe /c',Y)` under Windows NT; where `Y` is
the argument given to `)CMD`. However, the implementation permits the use of
alternative command processors as follows.

Before execution, the argument is prefixed and postfixed with strings defined by the
APL parameters CMD_PREFIX and CMD_POSTFIX. The former specifies the name
of your command processor and any parameters that it requires. The latter specifies a
string which may be required. If CMD_PREFIX is not defined, it defaults to the name
defined by the environment variable COMSPEC followed by " /c". If COMSPEC is
not defined, it defaults to COMMAND.COM or CMD.EXE as appropriate. If
CMD_POSTFIX is not defined, it defaults to an empty vector.

Save Continuation:)*CONTINUE*

This command saves the active workspace under the name *CONTINUE* and ends the Dyalog APL session.

When you subsequently start another Dyalog APL session, the *CONTINUE* workspace is loaded automatically. When a *CONTINUE* workspace is loaded, the latent expression (if any) is NOT executed.

Note that the values of all system variables (including ⎕*SM*) and GUI objects are also saved in *CONTINUE*.

Copy Workspace:)*COPY {ws {nms}}*

This command brings all or selected global objects *nms* from a stored workspace with the given name. A stored workspace is one which has previously been saved with the system command)*SAVE* or the system function ⎕*SAVE*. See page 1 for the rules for specifying a workspace name.

)*COPY* does NOT copy ⎕*SM*.

If the list of names is excluded, all defined objects (including namespaces and GUI objects) are copied. Copied objects are defined at the global level in the active workspace. Existing global objects in the active workspace with the same name as a copied object are replaced. If the copied object replaces either a function in the state indicator, or an object that is an operand of an operator in the state indicator, or a function whose left argument is being executed, the original object remains defined until its execution is completed or it is no longer referenced by an operator in the state indicator. If the workspace name is not valid or does not exist or if access to the workspace is not authorised, the system reports "*ws not found*".

You may copy an object from a namespace by specifying its full pathname. The object will be copied to the current namespace in the active workspace, losing its original parent and gaining a new one in the process. You may only copy a GUI object into a namespace that is a suitable parent for that object. For example, you could only copy a Group object from a saved workspace if the current namespace in the active workspace is itself a Form, SubForm or Group.

If the workspace name identifies a file that is not a workspace, the system reports "*bad ws*".

If "*ws*" is omitted, the file open dialog box is displayed and all objects copied from the selected workspace.

If the workspace name identifies a valid, readable workspace, the system reports the
workspace name, "*saved*" and the date and time when the workspace was last saved.

If the list of names is included, the names of system variables may also be included and
copied into the active workspace. The global referents will be copied.

If an object is not found in the stored workspace, the system reports "*not found*"
followed by the name of the object.

Examples

```
     )COPY WS/UTILITY
WS/UTILITY saved Mon Nov  1 13:11:19 1992

     )COPY TEMP ⎕LX FOO X A.B.C
./TEMP saved Mon Nov  1 14:20:47 1992
not found X
```

Implementation Note

`)COPY` is implemented in a way that is designed to provide optimum performance
under most circumstances. The `)COPY` process `)LOAD`s the target workspace into
memory. Some or all of the contents of the target workspace are then copied from the
secondary task into the primary task.

During the copy, all processing of the two symbol tables, stacks, and internal pointers
occurs in memory without further (random) disk accesses. In general this design
makes `)COPY` (which is traditionally a slow operation) extremely fast. If however the
target workspace is very large it can cause problems with excessive paging. If the
source workspace is too large to be loaded into the second APL task, the system reports
"ws too large".

Change Space: `)CS {nm}`

`)CS` changes the current space to the **global** namespace *nm*.

If no *name* is given, the system changes to the top level (Root) namespace.

If *name* is not the name of a global namespace, the system reports the error message *Namespace does not exist.*

name may be either a simple name or a compound name separated by '.', including one of the special names '#' (Root) or '##' (Parent).

Examples

```
      )CS
#
      )CS X
#.X
      )CS Y.Z
#.X.Y.Z
      )CS ##
#.X.Y
      )CS #.UTIL
#.UTIL
```

Drop Workspace: `)DROP {ws}`

This command removes the specified workspace from disk storage. See page 1 for information regarding the rules for specifying a workspace name.

If "`ws`" is omitted, a file open dialog box is displayed to elicit the workspace name.

Example

```
      )DROP WS/TEMP
Thu Sep 17 10:32:18 1998
```

Edit Object: `)ED nms`

`)ED` invokes the Dyalog APL editor and opens an Edit window for each of the objects specified in `nms`.

If a name specifies a new symbol it is taken to be a function/operator. However, if a name is localised in a suspended function/operator but is otherwise undefined, it is assumed to be a vector of character vectors.

The type of a new object may be specified explicitly by preceding its name with an appropriate symbol as follows :

"∇"   - function/operator
"∈"   - vector of character vectors
"-"   - character matrix
"→"   - character vector

The first object named becomes the top window on the stack. See *User Guide* for details.

`)ED` ignores names which specify GUI objects.

Examples

```
      )ED MYFUNCTION

      )ED ∇FOO -MAT ∈VECVEC
```

Erase Object: `)ERASE nms`

This command erases named global defined objects (functions, operators, variables, namespaces and GUI objects) from the active workspace or current namespace.

If a named object is a function or operator in the state indicator, or the object is an operand of an operator in the state indicator, or the object is a function whose left argument is being executed, the object remains defined until its execution is completed or it is no longer referenced by an operator in the state indicator. However, the name is available immediately for other uses.

If a named object is a GUI object, the object and all its children are deleted and removed from the screen.

If an object is not erased for any reason, the system reports `not found` followed by the name of the object.

Erasing objects such as external functions may have other implications: see the entry for expunge object (`⎕EX`) for details.

Example

```
      )ERASE FOO A ⎕IO
not found ⎕IO
```

List Events:) *EVENTS*

The) *EVENTS* system command lists the Events that may be generated by the object associated with the current space.

For example:

```
    □CS 'BB'□WC'BrowseBox'

    )EVENTS
Close    Create   FileBoxCancel   FileBoxOK
```

) *EVENTS* produces no output when executed in a pure (non-GUI) namespace, for example:

```
    □CS 'X' □NS ''
    )EVENTS
```

List Global Defined Functions:) *FNS {nm}*

This command displays the names of global defined functions in the active workspace or current namespace. Names are displayed in □ *AV* collation order. If a name is included after the command, only those names starting at or after the given name in collation order are displayed.

Examples

```
    )FNS
ASK DISPLAY GET PUT ZILCH
    )FNS G
GET PUT ZILCH
```

Display Held Tokens:)*HOLDS*

System command)*HOLDS* displays a list of tokens which have been acquired or requested by the :*Hold* control structure.

Each line of the display is of the form:

token: acq   req   req ...

Where *acq* is the number of *the* thread that has acquired the token, and *req* is the number of *a* thread which is requesting it. For a token to appear in the display, a thread (and only one thread) must have acquired it, whereas any number of threads can be requesting it.

Example

Thread 300's attempt to acquire token '*blue*' results in a deadlock:

```
300:DEADLOCK
Sema4[1] :Hold 'blue'
        ^

      )HOLDS
blue:    100
green:   200      100
red:     300      200      100
```

- *Blue* has been acquired by thread 100.
- *Green* has been acquired by 200 and requested by 100.
- *Red* has been acquired by 300 and requested by 200 and 100.

The following cycle of dependencies has caused the deadlock:

Thread 300 attempts to acquire *blue*,	300 → *blue*
which is owned by 100,	↑ ↓
which is waiting for *red*,	*red* ← 100
which is owned by 300.	

List Workspace Library:)*LIB {dir}*

This command lists the names of Dyalog APL workspaces contained in the given directory.

Example

```
      )LIB WS
MYWORK TEMP
```

If a directory is not given, the workspaces on the user's APL workspace path (WSPATH)are listed. In this case, the listing is divided into sections identifying the directories concerned. The current directory is identified as ".".

Example

```
      )LIB
.
      PDTEMP  WORK   GRAPHICS
C:\DYALOG\WS
      DISPLAY GROUPS
```

Load Workspace:)*LOAD* {*ws*}

This command causes the named stored workspace to be loaded. The current active workspace is lost.

If "*ws*" is a full or relative pathname, only the specified directory is examined. If not, the APL workspace path (WSPATH as specified in APL.INI) is traversed in search of the named workspace. A stored workspace is one which has previously been saved with the system command)*SAVE* or the system function ⎕*SAVE*. If "*ws*" is omitted, the File Open dialog box is displayed.

If the workspace name is not valid or does not exist or if access to the workspace is not authorised, the system reports "*ws not found*". If the workspace name identifies a file or directory that is not a workspace, the system reports workspace name "*is not a ws*". If successfully loaded, the system reports workspace name "*saved*", followed by the date and time when the workspace was last saved. If the workspace is too large to be loaded into the APL session, the system reports "*ws too large*". After loading the workspace, the latent expression (⎕*LX*) is executed unless APL was invoked with the -x option.

If the workspace contains any GUI objects whose *Visible* property is 1, these objects will be displayed. If the workspace contains a non-empty ⎕*SM* but does not contain an SM GUI object, the form defined by ⎕*SM* will be displayed in a window on the screen.

Holding the Ctrl key down while entering a)*LOAD* command or selecting a workspace from the session file menu now causes the incoming latent expression to be *traced*.

Holding the Shift key down while selecting a workspace from the session file menu will *prevent* execution of the latent expression.

Example

```
      )LOAD SMDEMO
/usr/dyalog/WS/SMDEMO saved Wed Sep 6 21:46:27 1989
Type HOWDEMO for help
```

List Methods:)*METHODS*

The)*METHODS* system command lists the Methods that apply to the object associated with the current space.

For example:

```
      ⎕CS 'F'⎕WC'Form'
      )METHODS
Animate ChooseFont      Detach  GetFocus        GetTextSiz
e
Wait
```

)*METHODS* produces no output when executed in a pure (non-GUI) namespace, for example:

```
      ⎕CS 'X' ⎕NS ''
      )METHODS
```

Create Namespace: `)NS {nm}`

`)NS` creates a **global** namespace and displays its full name, *nm*.

nm may be either a simple name or a compound name separated by `'.'`, including one of the special names `'#'` (Root) or `'##'` (Parent).

If *name* does not start with the special Root space identifier `'#'`, the new namespace is created relative to the current one.

If *name* is already in use for a workspace object other than a namespace, the command fails and displays the error message *Name already exists*.

If *name* is an existing namespace, no change occurs.

`)NS` with no *nm* specification displays the current namespace.

Examples

```
      )NS
#

      )NS W.X
#.W.X

      )CS W.X
#.W.X

      )NS Y.Z
#.W.X.Y.Z

      )NS
#.W.X
```

List Global Namespaces: *)OBJECTS {nm}*

This command displays the names of global **namespaces** in the active workspace. Names are displayed in the *⎕AV* collating order. If a name is included after the command, only those names starting at or after the given name in collating order are displayed. Namespaces are objects created using *⎕NS*, *)NS* or *⎕WC* and have name class 9.

Note: *)OBS* can be used as an **alternative** to *)OBJECTS*

Examples

```
      )OBJECTS
FORM1    UTIL    WSDOC    XREF

      )OBS W
WSDOC   XREF
```

List Global Namespaces: *)OBS {nm}*

This command is the same as the *)OBJECTS* command above.

Sign Off APL: *)OFF*

This command terminates the APL session, returning to the Operating System command processor or shell.

List Global Defined Operators: *)OPS {nm}*

This command displays the names of global defined operators in the active workspace or current namespace. Names are displayed in *⎕AV* collation order. If a name is included after the command, only those names starting at or after the given name in collation order are displayed.

Examples

```
      )OPS
AND DOIF DUAL ELSE POWER

      )OPS E
ELSE POWER
```

Protected Copy:)*PCOPY {ws {nms}}*

This command brings all or selected global objects from a stored workspace with the given name provided that there is no existing global usage of the name in the active workspace. A stored workspace is one which has previously been saved with the system command)*SAVE* or the system function □*SAVE*.

)*PCOPY* does not copy □*SM*. This restriction may be removed in a later release.

If the workspace name is not valid or does not exist or if access to the workspace is not authorised, the system reports "*ws not found*". If the workspace name identifies a file that is not a workspace, or is a workspace with an invalid version number (one that is greater than the version of the current APL) the system reports "*bad ws*". See page 1 for the rules for specifying a workspace name.

If the workspace name is the name of a valid, readable workspace, the system reports the workspace name, "*saved*", and the date and time that the workspace was last saved.

If the list of names is excluded, all global defined objects (functions and variables) are copied. If an object is not found in the stored workspace, the system reports "*not found*" followed by the name of the object. If an object cannot be copied into the active workspace because there is an existing referent, the system reports "*not copied*" followed by the name of the object.

For further information, see the implementation note under)*COPY*.

Examples

```
      )PCOPY WS/UTILITY
WS/UTILITY saved Mon Nov  1 13:11:19 1993
not copied COPIED IF
not copied COPIED JOIN

      )PCOPY TEMP FOO X
./TEMP saved Mon Nov  1 14:20:47 1993
not found X
```

List Properties: ` )PROPS`

The `)PROPS` system command lists the Properties of the object associated with the current space.

For example:

```
      ⎕CS 'BB'⎕WC'BrowseBox'

      )PROPS
BrowseFor        Caption ChildList        Data    Event
EventList        HasEdit KeepOnClose      MethodList
PropList         StartIn Target  Translate        Type
```

`)PROPS` produces no output when executed in a pure (non GUI) namespace, for example:

```
      ⎕CS 'X' ⎕NS ''
      )PROPS
```

Reset State Indicator: ` )RESET`

This command cancels all suspensions recorded in the state indicator and discards any unprocessed events in the event queue.

Example

```
      )SI
FOO[1]*
⍋
FOO[1]*

      )RESET

      )SI
```

Save Workspace:)SAVE {ws}

This command saves the active workspace after compaction of its contents. The workspace is saved with its state of execution intact. A stored workspace may subsequently be loaded with the system command)LOAD or the system function □LOAD, and objects may be copied from a stored workspace with the system commands)COPY or)PCOPY or the system function □CY.

This command may fail with one of the following error messages:

a)   unacceptable char
        The given workspace name was ill-formed.

b)   not saved this ws is WSID
        An attempt was made to change the name of the workspace for the save, and that workspace already existed.

c)   not saved this ws is CLEAR WS
        The active workspace was CLEAR WS and no attempt was made to change the name.

d)   cannot create
        The user does not have access to create the file OR the workspace name conflicts with an existing non-workspace file.

e)   cannot save with windows open
        A workspace may not be saved if trace or edit windows are open.

An existing stored workspace with the same name will be replaced. The active workspace may be renamed by the system command)WSID or the system function □WSID.

After a successful save, the system reports the workspace name, "saved", followed by the time and date.

Example

```
    )SAVE MYWORK
./MYWORK saved Thu Sep 17 10:32:20 1998
```

Execute (UNIX) Command: `)SH {cmd}`

This command allows DOS or UNIX shell commands to be given from APL. `)SH` is a synonym of `)CMD`. Either command may be given in either environment (DOS or UNIX) with exactly the same effect. `)SH` is probably more natural for the UNIX user. This section describes the behaviour of `)SH` and `)CMD` under UNIX. See `)CMD` for a discussion of their behaviour under DOS.

The system commands `⎕SH` and `⎕CMD` provide similar facilities but may be executed from within APL code.

`)SH` allows UNIX shell commands to be given from APL. The argument must be entered in the appropriate case (usually lower-case). The result of the command, if any, is displayed.

In Dyalog APL/X, `)SH` without an argument opens a new window. In non-X versions, `)SH` without an argument causes APL to fork a (Bourne) shell and wait. When the shell is closed ("exit" or Ctrl-d), control returns to APL. See *User Guide* for further information.

Example

```
      )SH ls
EXT
FILES
```

State Indicator:)*SI*

This command displays the contents of the state indicator in the active workspace. The state indicator identifies those operations which are suspended or pendent for each suspension.

The list consists of a line for each suspended or pendent operation beginning with the most recently suspended function or operator. Each line may be:

1. The name of a defined function or operator, followed by the line number at which the operation is halted, and followed by the * symbol if the operation is suspended. The name of a function or operator defined in a namespace other than the root is displayed relative to the root. For example, *UTIL.PRINT*. In addition, the display of a function or operator which has dynamically changed space away from its origin, is prefixed with its current space. For example, [□*SE*] *TRAV*.

2. A primitive operator symbol.

3. The Execute function symbol (⍎).

4. The Evaluated Input symbol (□).

5. The System Function □*DQ* or □*SR* (occurs when executing a callback function).

Examples

```
      )SI
PLUS[2]*
.
MATDIV[4]
FOO[1]*
⍎
```

This example indicates that at some point function *FOO* was executed and suspended on line 1. Subsequently, function *MATDIV* was invoked, with a function derived from the Inner Product or Outer Product operator (.) having defined function *PLUS* as an operand.

In the following, function *foo* in namespace *x* has called *goo* in namespace *y*. Function *goo* has then changed space (□*CS*) to namespace *z* where it has been suspended:

```
      )si
[z] y.goo[2]*
x.foo[1]
```

MultiThreading

In a multithreading application, where parent threads spawn child threads, the state
indicator assumes the structure of a branching tree. Branches of the tree are represented
by indenting lines belonging to child threads. For example:

```
      )SI
·    Calc[1]
&5
·    ·     DivSub[1]
·    &7
·    ·     DivSub[1]
·    &6
·    Div[2]*
&4
Sub[3]
Main[4]
```

Here, *Main* has called *Sub*, which has spawned threads 4 and 5 with functions: *Div*
and *Calc*. Function *Div*, after spawning *DivSub* in each of threads 6 and 7, has been
suspended at line [2].

State Indicator & Name List:) *SINL*

This command displays the contents of the state indicator together with local names.
The display is the same as for) *SI* (see above) except that a list of local names is
appended to each defined function or operator line.

Example

```
      )SINL
PLUS[2]*          B        A        R        DYADIC  END
·
MATDIV[4]         R        END      I        J       ▯TRAP
FOO[1]* R
♠
```

Thread Identity: *)TID {tid}*

Only a single thread can be suspended at any one time. You can switch the suspension between threads by using *)TID*.

If no thread number is given, *)TID* reports the number of the current thread.

Only those threads which are paused between lines of code, can be suspended. Attempting to switch to a thread that is, for example, pendent or waiting for a *:Hold*, will fail.

Examples

```
      ⍝State indicator
      )si
·   print[1]
&3
·   ·   sub_calc[2]*
·   &2
·   calc[1]
&1

      ⍝Current thread
      )tid
is 2

      ⍝Switch suspension to thread 3
      )tid 3
was 2

      ⍝State indicator
      )si
·   print[1]*
&3
·   ·   sub_calc[2]
·   &2
·   calc[1]
&1

      ⍝Attempt to switch to pendent thread 1
      )tid 1
Can't switch, this thread is 3
```

List Global Defined Variables: `)VARS {nm}`

This command displays the names of global defined variables in the active workspace or current namespace. Names are displayed in `⎕AV` collation order. If a name is included after the command, only those names starting at or after the given name in collation order are displayed.

Examples

```
      )VARS
A   B   F   TEMP VAR

      )VARS F
F TEMP VAR
```

Workspace Identification: `)WSID {ws}`

This command displays or sets the name of the active workspace.

If a workspace name is not specified, `)WSID` reports the name of the current active workspace. The name reported is the full path name, including directory references.

If a workspace name is given, the current active workspace is renamed accordingly. The previous name of the active workspace (excluding directory references) is reported. See page 1 for the rules for specifying a workspace name.

Examples

```
      )LOAD WS/TEMP
WS/TEMP saved Thu Sep 17 10:32:19 1998

      )WSID
is WS/TEMP

      )WSID WS/KEEP
was WS/TEMP

      )WSID
WS/KEEP
```

Load without Latent Expression: `)XLOAD {ws}`

This command causes the named stored workspace to be loaded. The current active workspace is lost.

`)XLOAD` is identical in effect to `)LOAD` except that `)XLOAD` does **not** cause the expression defined by the latent expression `⎕LX` in the saved workspace to be executed.

C H A P T E R   7

Error Messages

The error messages reported by APL are described in this chapter. Standard APL messages that provide information or report error conditions are summarised in Figure 7(i) and described in alphabetical order.

APL also reports messages originating from the Operating System (DOS or UNIX) which are summarised in Figure 7(ii). Only those Operating System error messages that might occur through normal usage of APL operations are described in this manual. Other messages could occur as a direct or indirect consequence of using the Operating System interface functions $\square CMD$ and $\square SH$ or system commands $)CMD$ and $)SH$, or when a non-standard device is specified for the system functions $\square ARBIN$ or $\square ARBOUT$. Refer to the DOS or UNIX reference manual for further information about these messages.

Most errors may be trapped using the system variable $\square TRAP$, thereby retaining control and inhibiting the standard system action and error report. For convenience, Figures 7(i) and 7(ii) identify the error code for trappable errors. The error code is also identified in the heading block for each error message when applicable.

See *User Guide* for a full description of the Error Handling facilities in Dyalog APL.

Standard Error Action

The standard system action in the event of an error or interrupt whilst executing an expression is to suspend execution and display an error report. If necessary, the state indicator is cut back to a statement such that there is no halted locked function visible in the state indicator.

The error report consists of up to three lines

1. The error message, preceded by the symbol ⋭ if the error occurred while evaluating the Execute function.

2. The statement in which the error occurred (or expression being evaluated by the Execute function), preceded by the name of the function and line number where execution is suspended unless the state indicator has been cut back to immediate execution mode. If the state indicator has been cut back because of a locked function in execution, the displayed statement is that from which the locked function was invoked.

3. The symbol ∧ under the last referenced symbol or name when the error occurred. All code to the right of the ∧ symbol in the expression will have been evaluated.

Examples

```
      X PLUS U
VALUE ERROR
      X PLUS U
             ∧
      FOO
INDEX ERROR
FOO[2] X←X+A[I]
           ∧

      CALC
⋭DOMAIN ERROR
CALC[5] ÷0
        ∧
```

Error Code	Report
	bad ws
	cannot create <u>*name*</u>
	clear ws
1008	*copy incomplete*
	DEADLOCK
11	*defn error*
1005	*DOMAIN ERROR*
90	*EOF INTERRUPT*
53	*EXCEPTION*
54	*FIELD CONTENTS RANK ERROR*
55	*FIELD CONTENTS TOO MANY COLUMNS*
56	*FIELD POSITION ERROR*
57	*FIELD SIZE ERROR*
58	*FIELD CONTENTS/TYPE MISMATCH*
59	*FIELD TYPE/BEHAVIOUR UNRECOGNISED*
60	*FIELD ATTRIBUTES RANK ERROR*
61	*FIELD ATTRIBUTES LENGTH ERROR*
62	*FULL-SCREEN ERROR*
63	*KEY CODE UNRECOGNISED*
70	*KEY CODE RANK ERROR*
71	*KEY CODE TYPE ERROR*
19	*FORMAT FILE ACCESS ERROR*
35	*FORMAT FILE ERROR*
38	*FILE ACCESS ERROR*
23	*FILE ACCESS ERROR - CONVERTING FILE*
21	*FILE COMPONENT DAMAGED*
20	*FILE DAMAGED*
22	*FILE FULL*
32	*FILE INDEX ERROR*
26	*FILE NAME ERROR*
34	*FILE NAME QUOTA USED UP*
28	*FILE SYSTEM ERROR*
30	*FILE SYSTEM NO SPACE*
18	*FILE SYSTEM NOT AVAILABLE*
24	*FILE SYSTEM TIES USED UP*
25	*FILE TIE ERROR*
31	*FILE TIED*
	FILE TIED REMOTELY
	FILE TIE QUOTA USED UP

Figure 7(i) : APL Error Messages

Error Code	Report
7	*FORMAT ERROR*
	incorrect command
12	*HOLD ERROR*
3	*INDEX ERROR*
	insufficient resources
1003	*INTERRUPT*
	is <u>name</u>
5	*LENGTH ERROR*
10	*LIMIT ERROR*
16	*NONCE ERROR*
72	*NO PIPES*
	<u>name</u> is not a ws
	Name already exists
	Namespace does not exist
	not copied <u>name</u>
	not found <u>name</u>
	not saved this ws is <u>name</u>
76	*PROCESSOR TABLE FULL*
4	*RANK ERROR*
1007	*RESIZE*
	<u>name</u> saved <u>date/time</u>
2	*SYNTAX ERROR*
	sys error <u>number</u>
1006	
	TIMEOUT
84	*too many names*
6	*TRAP ERROR*
	VALUE ERROR
	warning duplicate label
	warning duplicate name
	warning label name present in line 0
	warning pendent operation
	warning unmatched brackets
	warning unmatched parentheses
1	*was <u>name</u>*
	WS FULL
	ws not found
	ws too large

Figure 7(i) : APL Error Messages (Continued)

Error Code	Report
101	*FILE ERROR 1 Not owner*
102	*FILE ERROR 2 No such file or directory*
103	*FILE ERROR 3 No such process*
104	*FILE ERROR 4 Interrupted system call*
105	*FILE ERROR 5 I/O error*
106	*FILE ERROR 6 No such device or address*
107	*FILE ERROR 7 Arg list too long*
108	*FILE ERROR 8 Exec format error*
109	*FILE ERROR 9 Bad file number*
110	*FILE ERROR 10 No children*
111	*FILE ERROR 11 No more processes*
112	*FILE ERROR 12 Not enough code*
113	*FILE ERROR 13 Permission denied*
114	*FILE ERROR 14 Bad address*
115	*FILE ERROR 15 Block device required*
116	*FILE ERROR 16 Mount device busy*
117	*FILE ERROR 17 File exists*
118	*FILE ERROR 18 Cross-device link*
119	*FILE ERROR 19 No such device*
120	*FILE ERROR 20 Not a directory*
121	*FILE ERROR 21 Is a directory*
122	*FILE ERROR 22 Invalid argument*
123	*FILE ERROR 23 File table overflow*
124	*FILE ERROR 24 Too many open files*
125	*FILE ERROR 25 Not a typewriter*
126	*FILE ERROR 26 Text file busy*
127	*FILE ERROR 27 File too large*
128	*FILE ERROR 28 No space left on device*
129	*FILE ERROR 29 Illegal seek*
130	*FILE ERROR 30 Read-only file system*
131	*FILE ERROR 31 Too many links*
132	*FILE ERROR 32 Broken pipe*
133	*FILE ERROR 33 Math argument*
134	*FILE ERROR 34 Result too large*

Figure 7(ii) : Typical Operating System Error Messages

Note that operating system error messages **vary** between different operating systems. Dyalog APL reports the error code associated with the operating system error message that it received. Figure 7(iii) gives the messages recieved under Windows 95.

Error Code	Report
101	*FILE ERROR 1 No such file or directory*
102	*FILE ERROR 2 No such file or directory*
103	*FILE ERROR 3 Exec format error*
105	*FILE ERROR 5 Not enough memory*
106	*FILE ERROR 6 Permission denied*
107	*FILE ERROR 7 Argument list too big*
108	*FILE ERROR 8 Exec format error*
109	*FILE ERROR 9 Bad file number*
	FILE ERROR 11 Too many open files
111	*FILE ERROR 12 Not enough memory*
112	*FILE ERROR 13 Permission denied*
113	*FILE ERROR 14 Result too large*
114	*FILE ERROR 15 Resource deadlock would occur*
115	*FILE ERROR 17 File exists*
117	*FILE ERROR 18 Cross-device link*
118	*FILE ERROR 22 Invalid argument*
122	*FILE ERROR 23 File table overflow*
123	*FILE ERROR 24 Too many open files*
124	*FILE ERROR 33 Argument too large*
133	*FILE ERROR 34 Result too large*
134	*FILE ERROR 45 Resource deadlock would occur*
145	

Figure 7(iii) : Windows 95 Operating System Error Messages

APL Error Messages

There follows an alphabetical list of error messages reported from within Dyalog APL.

bad ws

This report is given when an attempt is made to)*COPY* or)*PCOPY* from a file that is not a valid workspace file. Invalid files include workspaces that were created by a version of Dyalog APL later than the version currently being used.

cannot create name

This report is given when an attempt is made to)*SAVE* a workspace with a name that is either the name of an existing, non-workspace file, or the name of a workspace that the user does not have permission to overwrite or create.

clear ws

This message is displayed when the system command)*CLEAR* is issued.

Example

```
      )CLEAR
clear ws
```

copy incomplete

This report is given when an attempted)*COPY* or)*PCOPY* fails to complete. Reasons include:

* Failure to identify the incoming file as a workspace.

* Not enough active workspace to accommodate the copy.

DEADLOCK   **1008**

If two threads succeed in acquiring a hold of two different tokens, and then each asks to hold the other token, they will both stop and wait for the other to release its token. The interpreter detects such cases and issues an error (1008) *DEADLOCK*.

This condition can also occur if calls to `⎕TSYNC` cause a cyclic dependency.

defn error

This report is given when either:

- The system editor is invoked in order to edit a function that does not exist, or the named function is pendent or locked, or the given name is an object other than a function.

- The system editor is invoked to define a new function whose name is already active.

- The header line of a function is replaced or edited in definition mode with a line whose syntax is incompatible with that of a header line. The original header line is re-displayed by the system editor with the cursor placed at the end of the line. Back-spacing to the beginning of the line followed by line-feed restores the original header line.

Examples

```
      X←1
      ∇X
defn error

      ∇FOO[0⎕]
[0]   R←FOO
[0]   R←FOO:X
defn error
[0]   R←FOO:X

      ⎕LOCK'FOO'
      ∇FOO[⎕]
defn error
```

DOMAIN ERROR   11

This report is given when either:

- An argument of a function is not of the correct type or its numeric value is outside the range of permitted values or its character value does not constitute valid name(s) in the context.

- An array operand of an operator is not an array, or it is not of the correct type, or its numeric value is outside the range of permitted values. A function operand of an operator is not one of a prescribed set of functions.

- A value assigned to a system variable is not of the correct type, or its numeric value is outside the range of permitted values

- The result produced by a function includes numeric elements which cannot be fully represented.

Examples

```
      1÷0
DOMAIN ERROR
      1÷0
      ∧

      (×∘'CAT')2 4 6
DOMAIN ERROR
      (×∘'CAT')2 4 6
      ∧

      ⎕IO←5
DOMAIN ERROR
      ⎕IO←5
      ∧
```

EOF INTERRUPT   1005

This report is given on encountering the end-of-file when reading input from a file. This condition could occur when an input to APL is from a file.

EXCEPTION 90

This report is given when a Micorosft .Net object throws an exception. Details are available in *⎕EXCEPTION*.

FIELD CONTENTS RANK ERROR 52

This report is given if a field content of rank greater than 2 is assigned to ⎕*SM*.

FIELD CONTENTS TOO MANY COLUMNS 53

This report is given if the content of a numeric or date field assigned to ⎕*SM* has more than one column.

FIELD POSITION ERROR 54

This report is given if the location of the field assigned to ⎕*SM* is outside the screen.

FIELD CONTENTS/TYPE MISMATCH 56

This report is given if the field contents assigned to ⎕*SM* does not conform with the given field type eg. character content with numeric type.

FIELD TYPE/BEHAVIOUR UNRECOGNISED 57

This report is given if the field type or behaviour code assigned to ⎕*SM* is invalid.

FIELD ATTRIBUTES RANK ERROR   58

This report is given if the current video attribute assigned to $\Box SM$ is non-scalar but its rank does not match that of the field contents.

FIELD ATTRIBUTES LENGTH ERROR   59

This report is given if the current video attribute assigned to $\Box SM$ is non-scalar but its dimensions do not match those of the field contents.

FULL-SCREEN ERROR   60

This report is given if the required full screen capabilities are not available to $\Box SM$. This report is only generated in UNIX environments.

KEY CODE UNRECOGNISED   61

This report is given if a key code supplied to $\Box SR$ or $\Box PFKEY$ is not recognised as a valid code.

KEY CODE RANK ERROR   62

This report is given if a key code supplied to $\Box SR$ or $\Box PFKEY$ is not a scalar or a vector.

KEY CODE TYPE ERROR   63

This report is given if a key code supplied to $\Box SR$ or $\Box PFKEY$ is numeric or nested; ie. is not a valid key code.

FORMAT FILE ACCESS ERROR   **70**

This report is given if the date format file to be used by □*SM* does not exist or cannot be accessed.

FORMAT FILE ERROR   **71**

This report is given if the date format file to be used by □*SM* is ill-formed.

FILE ACCESS ERROR   **19**

This report is given when the user attempts to execute a file system function for which the user is not authorised, or has supplied the wrong passnumber. It also occurs if the file specified as the argument to □*FERASE* or □*FRENAME* is not exclusively tied.

Examples

```
      'SALES' □FSTIE 1

      □FRDAC 1
0 4121 0
0 4137 99

      X □FREPLACE 1
FILE ACCESS ERROR
      X □FREPLACE 1
      ^

      'SALES' □FERASE 1
FILE ACCESS ERROR
      'SALES' □FERASE 1
      ^
```

FILE ACCESS ERROR – CONVERTING FILE   35

When a new version of Dyalog APL is used, it may be that improvements to the component file system demand that the internal structure of component files must alter. This alteration is performed by the interpreter on the first occasion that the file is accessed. If the operating system file permissions deny the ability to perform such a restructure, this report is given.

FILE COMPONENT DAMAGED   38

This report is given if an attempt is made to access a component that is not a valid APL object. This will rarely occur, but may happen as a result of a previous computer system failure. Components files may be checked using **qfsck**. (See *User Guide.*)

FILE DAMAGED   23

This report is given if a component file becomes damaged. This rarely occurs but may result from a computer system failure. Components files may be checked using **qfsck**. (See *User Guide.*)

FILE FULL   21

This report is given if the file operation would cause the file to exceed its file size limit.

FILE INDEX ERROR   20

This report is given when an attempt is made to reference a non-existent component.

Example

```
      ⎕FSIZE 1
1 21 16578 4294967295

      ⎕FREAD 1 34
FILE INDEX ERROR
      ⎕FREAD 1 34
      ∧
      ⎕FDROP 1 50
FILE INDEX ERROR
      ⎕FDROP 1 50
      ∧
```

FILE NAME ERROR   **22**

This report is given if:

- the user attempts to *⎕FCREATE* using the name of an existing file.

- the user attempts to *⎕FTIE* or *⎕FSTIE* a non-existent file, or a file that is not a component file.

- the user attempts to *⎕FERASE* a component file with a name other than the EXACT name that was used when the file was tied.

FILE NAME QUOTA USED UP   **32**

This report is given when the user attempts to execute a file system command that would result in the User's File Name Quota (see *User Guide*) being exceeded.

This can occur with *⎕FCREATE*, *⎕FTIE*, *⎕FSTIE* or *⎕FRENAME* .

FILE SYSTEM ERROR 26

This report is given if the File System Control Block (FSCB) is removed or altered while files are tied.

Contact the System Administrator. If this occurs when a file is share-tied, the file may be damaged. It is therefore advisable to check the integrity of all such files using `qfsck`.

See *User Guide* for details.

FILE SYSTEM NO SPACE 34

This report is given if the user attempts a file operation that cannot be completed because there is insufficient disk space.

FILE SYSTEM NOT AVAILABLE 28

This report is given if the File System Control Block (FSCB) is missing or inaccessible. See *User Guide* for details.

FILE SYSTEM TIES USED UP 30

This report is given if the File System Control Block (FSCB) is full. See *User Guide* for details.

FILE TIE ERROR   18

This report is given when the argument to a file system function contains a file tie
number used as if it were tied when it is not or as if it were available when it is already
tied. It also occurs if the argument to ⎕*FHOLD* contains the names of non-existent
external variables.

Examples

```
      ⎕FNAMES,⎕FNUMS
SALES  1
COSTS  2
PROFIT 3

      X ⎕FAPPEND 4
FILE TIE ERROR
      X ⎕FAPPEND 4
      ^
      'NEWSALES' ⎕FCREATE 2
FILE TIE ERROR
      'NEWSALES' ⎕FCREATE 2
      ^

      'EXTVFILE' ⎕XT'BIGMAT'
      ⎕FHOLD 'BIGMAT'
FILE TIE ERROR
      ⎕FHOLD 'BIGMAT'
      ^
      ⎕FHOLD⊂'BIGMAT'
```

FILE TIED   24

This report is given if the user attempts to tie a file that is exclusively tied by another
task, or attempts to exclusively tie a file that is already share-tied by another task.

FILE TIED REMOTELY   25

This report is given if the user attempts to tie a file that is exclusively tied by another
task, or attempts to exclusively tie a file that is already share-tied by another task; and
that task is running on other than the user's processor.

FILE TIE QUOTA USED UP 31

This report is given if an attempt is made to □*FTIE*, □*FSTIE* or □*FCREATE* a file when the user already has the maximum number of files tied. (See File Tie Quota , *User Guide*)

FORMAT ERROR 7

This report is given when the format specification in the left argument of system function □*FMT* is ill-formed.

Example

```
      'A1,1X,I5'□FMT CODE NUMBER
FORMAT ERROR
      'A1,1X,I5'□FMT CODE NUMBER
       ^
```

(The correct specification should be '*A1,X1,I5*' .)

HOLD ERROR 12

This report is given when an attempt is made to save a workspace using the system function □*SAVE* if any external arrays or component files are currently held (as a result of a prior use of the system function □*FHOLD*).

Example

```
      ∇HOLD∆SAVE
[1]    □FHOLD 1
[2]    □SAVE 'TEST'
      ∇

      'FILE' □FSTIE 1

      HOLD∆SAVE
HOLD ERROR
HOLD∆SAVE[2] □SAVE'TEST'
              ^
```

incorrect command

This report is given when an unrecognised system command is entered.

Example

```
      )CLERA
incorrect command
```

INDEX ERROR   3

This report is given when either:

*   The value of an index, whilst being within comparison tolerance of an integer, is outside the range of values defined by the index vector along an axis of the array being indexed. The permitted range is dependent on the value of *⎕IO*.

*   The value specified for an axis, whilst being within comparison tolerance of an integer for a derived function requiring an integer axis value or a non-integer for a derived function requiring a non-integer, is outside the range of values compatible with the rank(s) of the array argument(s) of the derived function. Axis is dependent on the value of *⎕IO*.

Examples

```
      A
1  2  3
4  5  6

      A[1;4]
INDEX ERROR
      A[1;4]
      ^

      ↑[2]'ABC' 'DEF'
INDEX ERROR
      ↑[2]'ABC' 'DEF'
      ^
```

insufficient resources

This report only occurs in Version 7 and is generated by *)COPY* or *)LOAD*. It occurs when the operating system runs out of GUI memory while attempting to recreate GUI

objects which have been saved in the workspace. Close some applications and try again.

INTERRUPT 1003

This report is given when execution is suspended by entering a hard interrupt. A hard interrupt causes execution to suspend as soon as possible without leaving the environment in a damaged state.

Example

```
      1 1 2 ⍴(2 100ρι200)∘.|?1000ρ200
(Hard interrupt)
INTERRUPT
      1 1 2 ⍴(2 100ρι200)∘.|?1000ρ200
            ^
```

is name

This report is given in response to the system command *)WSID* when used without a parameter. *name* is the name of the active workspace including directory references given when loaded or named. If the workspace has not been named, the system reports is *CLEAR WS*.

Example

```
      )WSID
is WS/UTILITY
```

LENGTH ERROR 5

This report is given when the shape of the arguments of a function do not conform, but the ranks do conform.

Example

```
      2 3+4 5 6
LENGTH ERROR
      2 3+4 5 6
      ^
```

LIMIT ERROR   10

This report is given when a system limit is exceeded. System limits are installation dependent.

Example

```
      (16ρ1)ρ1
LIMIT ERROR
      (16ρ1)ρ1
      ^
```

NONCE ERROR   16

This report is given when a system function or piece of syntax is not currently implemented but is reserved for future use.

NO PIPES   72

This message applies to the UNIX environment ONLY.

This message is given when the limit on the number of pipes communicating between tasks is exceeded. An installation-set quota is assigned for each task. An associated task may require more than one pipe. The message occurs on attempting to exceed the account's quota when either:

- An APL session is started

- A non-APL task is started by the system function *⎕SH*

- An external variable is used.

It is necessary to release pipes by terminating sufficient tasks before proceeding with the required activity. In practice, the error is most likely to occur when using the system function *⎕SH*.

Examples

```
      'via' ⎕SH 'via'
NO PIPES
      'via' ⎕SH 'via'
      ^

      'EXT/ARRAY' ⎕XT 'EXVAR'
NO PIPES
      'EXT/ARRAY' ⎕XT 'EXVAR'
                       ^
```

name is not a ws

This report is given when the name specified as the parameter of the system commands
)*LOAD*,)*COPY* or)*PCOPY* is a reference to an existing file or directory that is not
identified as a workspace.

This will also occur if an attempt is made to)*LOAD* a workspace that was)*SAVE*'d
using a later version of Dyalog APL.

Example

```
      )LOAD EXT\ARRAY
EXT\ARRAY is not a ws
```

Name already exists

This report is given when an)*NS* command is issued with a name which is already in
use for a workspace object other than a namespace.

Namespace does not exist

This report is given when a)*CS* command is issued with a name which is not the name
of a global namespace.

not copied *name*

This report is given for each object named or implied in the parameter list of the system command `)PCOPY` which was not copied because of an existing global referent to that name in the active workspace.

Example

```
     )PCOPY WS/UTILITY A FOO Z
WS/UTILITY saved Mon Nov 1 13:11:19 1993
not copied Z
```

not found *name*

This report is given when either:

- An object named in the parameter list of the system command `)ERASE` is not erased because it was not found or it is not eligible to be erased.

- An object named in the parameter list (or implied list) of names to be copied from a saved workspace for the system commands `)COPY` or `)PCOPY` is not copied because it was not found in the saved workspace.

Examples

```
     )ERASE □IO
not found □IO

     )COPY WS/UTILITY UND
WS/UTILITY saved Mon Nov 1 13:11:19 1993
not found UND
```

not saved this ws is name

This report is given in the following situations:

- When the system command)*SAVE* is used without a name, and the workspace is not named. In this case the system reports not saved this ws is *CLEAR WS*.

- When the system command)*SAVE* is used with a name, and that name is not the current name of the workspace, but is the name of an existing file.

In neither case is the workspace renamed.

Examples

```
      )CLEAR
      )SAVE
not saved this ws is CLEAR WS

      )WSID JOHND
      )SAVE
      )WSID ANDYS
      )SAVE JOHND
not saved this ws is ANDYS
```

PROCESSOR TABLE FULL   *76*

This report can only occur in a UNIX environment.

This report is given when the limit on the number of processes (tasks) that the computer system can support would be exceeded. The limit is installation dependent. The report is given when an attempt is made to initiate a further process, occurring when an APL session is started.

It is necessary to wait until active processes are completed before the required task may proceed. If the condition should occur frequently, the solution is to increase the limit on the number of processes for the computer system.

Example

```
      'prefect' ⎕SH 'prefect'
PROCESSOR TABLE FULL
      'prefect' ⎕SH 'prefect'
      ^
```

RANK ERROR   4

This report is given when the rank of an argument or operand does not conform to the requirements of the function or operator, or the ranks of the arguments of a function do not conform.

Example

```
      2 3 + 2 2ρ10 11 12 13
RANK ERROR
      2 3 + 2 2ρ10 11 12 13
      ^
```

RESIZE   1007

This report is given when the user resizes the ⎕SM window. It is only applicable to Dyalog APL/X and Dyalog APL/W.

name saved *date/time*

This report is given when a workspace is saved, loaded or copied.

date/time is the date and time at which the workspace was most recently saved.

Examples

```
      )LOAD WS/UTILITY
WS/UTILITY saved Fri Sep 11 10:34:35 1998

      )COPY SPACES GEOFF JOHND VINCE
./SPACES saved Wed Sep 30 16:12:56 1998
```

SYNTAX ERROR   2

This report is given when a line of characters does not constitute a meaningful statement. This condition occurs when either:

- An illegal symbol is found in an expression.

- Brackets, parentheses or quotes in an expression are not matched.

- Parentheses in an expression are not matched.

- Quotes in an expression are not matched.

- A value is assigned to a function, label, constant or system constant.

- A strictly dyadic function (or derived function) is used monadically.

- A monadic function (or derived function) is used dyadically.

- A monadic or dyadic function (or derived function) is used without any arguments.

- The operand of an operator is not an array when an array is required.

- The operand of an operator is not a function (or derived function) when a function is required.

- The operand of an operator is a function (or derived function) with incorrect valency.

- A dyadic operator is used with only a single operand.

- An operator is used without any operands.

Examples

```
      (A>10/A
SYNTAX ERROR
      (A>10/A
      ^

      ⊤2 4 8
SYNTAX ERROR
      ⊤2 4 8
      ^

      A.+1 2 3
SYNTAX ERROR
      A.+1 2 3
      ^
```

sys error number

This report is given when an internal error occurs in Dyalog APL.

Under UNIX it may be necessary to enter a hard interrupt to obtain the UNIX command prompt, or even to **kill** your processes from another screen. Under DOS it may be necessary to reboot your PC.

If this error occurs, please submit a fault report to your Dyalog APL distributor.

TIMEOUT   1006

This report is given when the time limit specified by the system variable *⎕RTL* is exceeded while awaiting input through character input (⎕) or *⎕SR*.

It is usual for this error to be trapped.

Example

```
      ⎕RTL←5 ◇ ⎕←'RESPOND WITHIN 5 SECONDS: ' ◇ R←⎕
RESPOND WITHIN 5 SECONDS:
TIMEOUT
      ⎕RTL←5 ◇ ⎕←'RESPOND WITHIN 5 SECONDS: ' ◇ R←⎕
                                                   ^
```

TRAP ERROR   84

This report is given when a workspace full condition occurs whilst searching for a definition set for the system variable *⎕TRAP* after a trappable error has occurred. It does not occur when an expression in a *⎕TRAP* definition is being executed.

`too many names`

This report is given by the function editor when the number of distinct names (other than distinguished names beginning with the symbol ⎕) referenced in a defined function exceeds the system limit of 4096.

VALUE ERROR   6

This report is given when either:

- There is no active definition for a name encountered in an expression.

- A function does not return a result in a context where a result is required.

Examples

```
      )CLEAR

      X
VALUE ERROR
      X
      ^

      ∇ HELLO
[1]    'HI THERE'
[2]    ∇

      2+HELLO
HI THERE
VALUE ERROR
      2+HELLO
       ^
```

warning duplicate label

This warning message is reported on closing definition mode when one or more labels are duplicated in the body of the defined function. This does not prevent the definition of the function in the active workspace. The value of a duplicated label is the lowest of the line-numbers in which the labels occur.

warning duplicate name

This warning message is reported on closing definition mode when one or more names are duplicated in the header line of the function. This may be perfectly valid. Definition of the function in the active workspace is not prevented. The order in which values are associated with names in the header line is described in *Chapter 2*.

warning pendent operation

This report is given on opening and closing definition mode when attempting to edit a pendant function or operator.

Example

```
[0]    ∇FOO
[1]     GOO
[2]    ∇

[0]    ∇GOO
[1]     ∘
[2]    ∇


      FOO
SYNTAX ERROR
GOO[1] ∘
      ^

      ∇FOO
warning pendent operation
[0]    ∇FOO
[1]     GOO
[2]    ∇
warning pendent operation
```

warning label name present in line 0

This warning message is reported on closing definition mode when one or more label names also occur in the header line of the function. This does not prevent definition of the function in the active workspace. The order in which values are associated with names is described in *Chapter 2*.

warning unmatched brackets

This report is given after adding or editing a function line in definition mode when it is found that there is not an opening bracket to match a closing bracket, or vice versa, in an expression. This is a warning message only. The function line will be accepted even though syntactically incorrect.

Example

```
[3]    A[;B[;2]←0
warning unmatched brackets
[4]
```

warning unmatched parentheses

This report is given after adding or editing a function line in definition mode when it is found that there is not an opening parenthesis to match a closing parenthesis, or vice versa, in an expression. This is a warning message only. The function line will be accepted even though syntactically incorrect.

Example

```
[4]    X←(E>2)∧E<10)≠A
warning unmatched parentheses
[5]
```

was name

This report is given when the system command)*WSID* is used with a parameter specifying the name of a workspace. The message identifies the former *name* of the workspace. If the workspace was not named, the given report is was *CLEAR WS*.

Example

```
      )WSID TEMP
was UTILITY
```

WS FULL   1

This report is given when there is insufficient workspace in which to perform an operation. Workspace available is identified by the system constant *⎕WA*.

The maximum workspace size allowed is defined by the environment variable MAXWS. See *User Guide* for details.

Example

```
      ⎕WA ρ 1.2
WS FULL
      ⎕WA ρ 1.2
      ^
```

ws not found

This report is given when a workspace named by the system commands)*LOAD*,)*COPY* or)*PCOPY* does not exist as a file, or when the user does not have read access authorisation for the file.

Examples

```
      )LOAD NOWS
ws not found

      )COPY NOWS A FOO X
ws not found
```

ws too large

This report is given when:

- the user attempts to `)LOAD` a workspace that needs a greater work area than the maximum that the user is currently permitted.

- the user attempts to `)COPY` or `)PCOPY` from a workspace that would require a greater work area than the user is currently permitted if the workspace were to be loaded.

The maximum work area permitted is set using the environment variable MAXWS.

Operating System Error Messages

There follows a numerically sorted list of error messages emanating from a typical operating system and reported through Dyalog APL.

`FILE ERROR 1 Not owner`   `101`

This report is given when an attempt is made to modify a file in a way which is forbidden except to the owner or super-user, or in some instances only to a super-user.

`FILE ERROR 2 No such file or directory 102`

This report is given when a file (which should exist) does not exist, or when a directory in a path name does not exist.

`FILE ERROR 5 I/O error`   `105`

This report is given when a physical I/O error occurred whilst reading from or writing to a device, indicating a hardware fault on the device being accessed.

`FILE ERROR 6 No such device or address 106`

This report is given when a device does not exist or the device is addressed beyond its limits. Examples are a tape which has not been mounted or a tape which is being accessed beyond the end of the tape.

`FILE ERROR 13 Permission denied`   `113`

This report is given when an attempt is made to access a file in a way forbidden to the account.

FILE ERROR 20 *Not a directory*   120

This report is given when the request assumes that a directory name is required but the name specifies a file or is not a legal name.

FILE ERROR 21 *Is a directory*   121

This report is given when an attempt is made to write into a directory.

FILE ERROR 23 *File table overflow*   123

This report is given when the system limit on the number of open files is full and a request is made to open another file. It is necessary to wait until the number of open files is reduced. If this error occurs frequently, the system limit should be increased.

FILE ERROR 24 *Too many open files*   124

This report is given when the task limit on the number of open files is exceeded. It may occur when an APL session is started or when a shell command is issued to start an external process through the system command □*SH*. It is necessary to reduce the number of open files. It may be necessary to increase the limit on the number of open files to overcome the problem.

FILE ERROR 26 *Text file busy*   126

This report is given when an attempt is made to write a file which is a load module currently in use. This situation could occur on assigning a value to an external variable whose associated external file name conflicts with an existing load module's name.

FILE ERROR 27 *File too large*   127

This report is given when a write to a file would cause the system limit on file size to be exceeded.

FILE ERROR 28 *No space left on device* 128

This report is given when a write to a file would exceed the capacity of the device containing the file.

FILE ERROR 30 *Read-only file system*   130

This report is given when an attempt is made to write to a device which can only be read from. This would occur with a write-protected tape.

C H A P T E R   8

APL Character Sets

The mapping between the Dyalog APL/W atomic vector and the symbols displayed on your screen or printer depends upon two factors. The first is the Output Translate Table in use (normally WIN.DOT) which maps $\Box AV$ to font positions. The second factor is the location of the symbols in your font. Both are to a large extent user configurable.

Dyalog APL/W provides two separate font layouts. The first includes traditional APL underscored characters $\underline{A}$-$\underline{Z}$ at the expense of certain (less commonly used) National Language symbols. This mapping is shown in Figure 8(i). The $\Box AV$ obtained using the second font layout contains a richer set of National Language symbols but no underscored characters and is shown in Figure 8(ii).

Three characters have no display but have the following effect when sent to a terminal:

$\Box AV[1+\Box IO]$	BS	Back-Space by one character.
$\Box AV[2+\Box IO]$	LF	Line-Feed to next line with the cursor remaining in the same column position.
$\Box AV[3+\Box IO]$	NL	New Line, causing the cursor to be placed at the beginning of the next line.

These characters are also contained in the system constant $\Box TC$. Other control characters found early in the atomic vector include, for example,

$\Box AV[9+\Box IO]$	HT	Horizontal Tab character.

Dyalog APL Atomic Vector

Standard Characters

nul	bs	lf	cr	sp	ff	(06)	bel	esc	tab	⍳	+	%	'	α	ω
_	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	(01)	(02)	¯	.	θ
0	1	2	3	4	5	6	7	8	9	(03)	⊢	¥	$	£	¢
∆	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	(04)	(05)	ý	.	(7f)
_∆_	_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_	{	∈	}	⊣	□
¨	À	Ä	Å	Æ	˜	É	Ñ	Ö	Ø	Ü	β	à	á	â	ä
å	æ	ç	è	é	ê	ë	í	î	ï	ñ	[	/	≠	\	⍀
<	≤	=	≥	>	≠	∨	∧	−	+	÷	×	?	∈	ρ	~
↑	↓	ι	○	*	⌈	⌊	∇	∘	(	⊂	⊃	∩	∪	⊥	⊤
|	;	,	⍲	⍱	⍢	⍔	⍥	⌽	⊖	⊛	⌹	!	⍕	⍎	⍪
⍮	≡	≠	ó	ô	ö	ø	"	#	(1e)	&	´	⌋	⌝	⌜	⌞
⊥	−	⊦	⊣	⊥	⊤	|	@	ù	ú	û	^	ü	`	¦	¶
:	⊆	¿	¡	◊	←	→	⍺	)	]	(1f)	(a0)	⍳	□	□	∵

Figure 8(i) : Dyalog Std ⎕AV (with underscores)

Note that ASCII symbols ˆ (caret) and ¦ (pipe) are interpreted as APL functions ∧ (And) and │ (Magnitude or Residue).

There are currently 9 unused character positions. (xx) means that the unused $\square AV$ character at this position is mapped to font position (hex) xx , and to this byte value for $\square NAPPEND$ and $\square NREPLACE$.

Alternative Characters

nul	bs	lf	cr	sp	ff	(06)	bel	esc	tab	I	+	%	'	α	ω
_	a	b	c	d	e	f	g	ħ	i	j	k	l	m	n	o
p	q	r	s	t	u	v	w	x	y	z	(01)	(02)	‾	.	θ
0	1	2	3	4	5	6	7	8	9	(03)	⊢	¥	$	£	¢
Δ	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	(04)	(05)	ý	.	(7f)
Δ̲	Á	Â	Ã	Ç	È	Ê	Ë	Ì	Í	Î	Ï	Đ	Ò	Ó	Ô
Õ	Ù	Ú	Û	Ý	Þ	ã	ì	ð	ò	õ	{	€	}	⊣	▯
¨	À	Ä	Å	Æ	≈	É	Ñ	Ö	Ø	Ü	ß	à	á	â	ä
å	æ	ç	è	é	ê	ë	í	î	ï	ñ	[	/	≠	\	↟
<	≤	=	≥	>	≠	∨	∧	−	+	÷	×	?	∊	ρ	~
↑	↓	⍳	○	*	⌈	⌊	∇	∘	(	⊂	⊃	∩	∪	⊥	⊤
|	;	,	⍦	⍋	⍢	⍟	⍣	⌽	⊖	⊛	⌷	!	⍤	⍚	⍫
⍭	≡	≢	ó	ô	ö	⌀	"	#	(1e)	&	´	⌐	⌐	⌙	⌞
+	−	⊢	⊣	⊥	⊤	|	@	ù	ú	û	^	ü	`	¦	¶
:	⊆	¿	¡	◇	←	→	Ɐ	)	]	(1f)	(a0)	I	▯	▯	⍨

Figure 8(ii) : Dyalog Alt $\square AV$ (**without** underscores)

Dyalog APL Font Layout

Standard Characters

nul	(2b)	(2c)	(3a)	(5b)	(5c)	(06)	bel	bs	tab	lf	I̲	ff	cr	╪	^
⌋	⌐	⌐	∟	┼	─	├	┤	⊥	⊤	│	esc	´	¦	(d9)	(fa)
sp	!	"	#	$	%	&	'	(	)	*	+	,	-	.	/
0	1	2	3	4	5	6	7	8	9	:	;	<	=	>	?
@	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	[	\	]	∧	_
`	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	{	|	}	~	(5f)
€	⊣	⊤	⊥	←	→	↑	↓	≤	≥	⍒	⍲	▯	⍞	⌶	Δ̲
⍢	Δ	∇	⍋	⍷	⍥	⍦	⌈	⌊	⍀	≠	⊂	⊃	∩	∪	∨
(fb)	¡	¢	£	⊢	¥	▯	⌶	¨	⍀	◊	θ	≠	≡	̦	‾
∘	○	φ	⍤	⊖	⊛	¶	·	α	∈	∊̲	≢	⍳	ρ	ω	¿
À	A̲	B̲	C̲	Ä	Å	Æ	D̲	E̲	É	F̲	G̲	H̲	I̲	J̲	K̲
L̲	Ñ	M̲	N̲	O̲	P̲	Ö	×	Ø	Q̲	R̲	S̲	Ü	T̲	U̲	β
à	á	â	V̲	ä	å	æ	ç	è	é	ê	ë	W̲	í	î	ï
X̲	ñ	Y̲	ó	ô	Z̲	ö	÷	ø	ù	ú	û	ü	ý	̈~	̈*

Figure 8(iii) : Dyalog Std Font Layout (<u>with</u> underscores)

Figure 8(iii) illustrates the Dyalog Std font layout which includes underscored characters (A̲ - Z̲). Figure 8(iv) illustrates the Dyalog Alt layout which does **not** include the underscored characters.

Chapter 8: APL Character Sets 533

Note: (xx) means that the unused ⎕AV character at (hex) xx is mapped to this font position, and that ⎕NREAD of this character is mapped into ⎕AV[xx].

Alternative Characters

nul	(2b)	(2c)	(3a)	(5b)	(5c)	(06)	bel	bs	tab	lf	⊤	ff	cr	┴	^
⌋	⌐	⌐	∟	┼	─	├	┤	⊥	⊤	│	esc	´	¦	(d9)	(fa)
sp	!	"	#	$	%	&	'	(	)	*	+	,	-	.	/
0	1	2	3	4	5	6	7	8	9	:	;	<	=	>	?
@	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	[	\	]	^	_
`	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	{	|	}	~	(5f)
€	⊣	⊤	⊥	←	→	↑	↓	≤	≥	⋆	�	⎕	⍞	⍠	∆
∇̈	∆	∇	⍋	⍒	⌽	⍉	⌈	⌊	⍀	≠	⊂	⊃	∩	∪	∨
(fb)	¡	¢	£	⊢	¥	⌷	I	¨	A	◇	θ	≠	≡	⍮	¯
∘	○	φ	⍇	⊖	⊛	¶	·	α	∈	⊆	≢	ι	ρ	ω	¿
À	Á	Â	Ã	Ä	Å	Æ	Ç	È	É	Ê	Ë	Ì	Í	Î	Ï
Đ	Ñ	Ò	Ó	Ô	Õ	Ö	×	Ø	Ù	Ú	Û	Ü	Ý	Þ	β
à	á	â	ã	ä	å	æ	ç	è	é	ê	ë	ì	í	î	ï
ð	ñ	ò	ó	ô	õ	ö	÷	ø	ù	ú	û	ü	ý	∼̈	⋆̈

Figure 8(iv) : Dyalog Alt Font Layout (<u>without</u> underscores)

Note that the APL quote resides in the ASCII quote position, which means that when reading text with embedded quotes from a plain ASCII file, the quotes will be converted by default (see ⎕NXLATE) into APL quotes.

Note also that most printer drivers have trouble printing the characters to be found in the first two rows of the above font tables because they are often used as control characters.

Overstrike Combinations

You may enter certain composite characters by typing two symbols (in either order) with an intervening left cursor movement. This applies only if the cursor is in **replace mode** and does not apply if it is in insert mode. Permissible overstrike combinations are listed in Figure 8(v).

$\underline{A}$	A _	$\underline{V}$	V _	⏀	○ |
$\underline{B}$	B _			⊖	○ −
$\underline{C}$	C _	$\underline{W}$	W _		
$\underline{D}$	D _			⊛	○ ⋆
$\underline{E}$	E _	$\underline{X}$	X _		
$\underline{F}$	F _			⊡	⎕ ÷
$\underline{G}$	G _	$\underline{Y}$	Y _		
$\underline{H}$	H _			!	' .
$\underline{I}$	I _	$\underline{Z}$	Z _		
$\underline{J}$	J _			⍕	⊤ ○
$\underline{K}$	K _	$\underline{\Delta}$	Δ _		
$\underline{L}$	L _			⍚	⊥ ○
$\underline{M}$	M _	⌶	⊤ ⊥		
$\underline{N}$	N _			⍠	, −
$\underline{O}$	O _	┼	| −		
$\underline{P}$	P _			≡	= _
$\underline{Q}$	Q _	⊖	0 ~		
$\underline{R}$	R _			≠	≠ _
$\underline{S}$	S _	\$	S |		
$\underline{T}$	T _			⍞	⎕ '
$\underline{U}$	U _	£	L −		
		⍨	~ ¨	⍴	∩ ○
		¢	⊂ |	{	[○
		⍘	[ ]	}	] ○
		⌿	− /	◇	∧ ∨

⊀	– \	◇	< >
⊻	∨ ~	⍒	∇ ~
⋆	∧ ~	⊣	) –
⍒	| ∇	∊	(–
⍋	| ∆	∈	∈ _
⍉	○ \		

Figure 8(v) : Overstrike Combinations

Symbolic Index

Alphabetic Index