

## 8. Arrays

An *array* is a Lisp object that consists of a group of cells, each of which may contain an object. The individual cells are selected by numerical *subscripts*.

The *rank* of an array (the number of dimensions which the array has) is the number of subscripts used to refer to one of the elements of the array. The rank may be any integer from zero to seven, inclusively.

The lowest value for any subscript is zero; the highest value is a property of the array. Each dimension has a size, which is the lowest number which is too great to be used as a subscript. For example, in a one-dimensional array of five elements, the size of the one and only dimension is five, and the acceptable values of the subscript are zero, one, two, three, and four.

The most basic primitive functions for handling arrays are: `make-array`, which is used for the creation of arrays, `aref`, which is used for examining the contents of arrays, and `aset`, which is used for storing into arrays.

An array is a regular Lisp object, and it is common for an array to be the binding of a symbol, or the `car` or `cdr` of a cons, or, in fact, an element of an array. There are many functions, described in this chapter, which take arrays as arguments and perform useful operations on them.

Another way of handling arrays, inherited from Maclisp, is to treat them as functions. In this case each array has a name, which is a symbol whose function definition is the array. Zetalisp supports this style by allowing an array to be *applied* to arguments, as if it were a function. The arguments are treated as subscripts and the array is referenced appropriately. The `store` special form (see page 142) is also supported. This kind of array referencing is considered to be obsolete and is slower than the usual kind. It should not be used in new programs.

There are many types of arrays. Some types of arrays can hold Lisp objects of any type; the other types of arrays can only hold fixnums or flonums. The array types are known by a set of symbols whose names begin with "art-" (for ARray Type).

The most commonly used type is called `art-q`. An `art-q` array simply holds Lisp objects of any type.

Similar to the `art-q` type is the `art-q-list`. Like the `art-q`, its elements may be any Lisp object. The difference is that the `art-q-list` array "doubles" as a list; the function `g-l-p` will take an `art-q-list` array and return a list whose elements are those of the array, and whose actual substance is that of the array. If you `rplaca` elements of the list, the corresponding element of the array will change, and if you store into the array, the corresponding element of the list will change the same way. An attempt to `rplacd` the list will cause a `sys:rplacd-wrong-representation-type` error, since arrays cannot implement that operation.

There is a set of types called `art-1b`, `art-2b`, `art-4b`, `art-8b`, and `art-16b`; these names are short for "1 bit", "2 bits", and so on. Each element of an `art-nb` array is a non-negative fixnum, and only the least significant *n* bits are remembered in the array; all of the others are

discarded. Thus **art-1b** arrays store only 0 and 1, and if you store a 5 into an **art-2b** array and look at it later, you will find a 1 rather than a 5.

These arrays are used when it is known beforehand that the fixnums which will be stored are non-negative and limited in size to a certain number of bits. Their advantage over the **art-q** array is that they occupy less storage, because more than one element of the array is kept in a single machine word. (For example, 32 elements of an **art-1b** array or 2 elements of an **art-16b** array will fit into one word).

There are also **art-32b** arrays which have 32 bits per element. Since fixnums only have 24 bits anyway, these are the same as **art-q** arrays except that they only hold fixnums. They are not compatible with the other "bit" array types and generally should not be used.

Character strings are implemented by the **art-string** array type. This type acts similarly to the **art-8b**; its elements must be fixnums, of which only the least significant eight bits are stored. However, many important system functions, including **read**, **print**, and **eval**, treat **art-string** arrays very differently from the other kinds of arrays. These arrays are usually called *strings*, and chapter 9 of this manual deals with functions that manipulate them.

An **art-fat-string** array is a character string with wider characters, containing 16 bits rather than 8 bits. The extra bits are ignored by string operations, such as comparison, on these strings; typically they are used to hold font information.

An **art-half-fix** array contains half-size fixnums. Each element of the array is a signed 16-bit integer; the range is from -32768 to 32767 inclusive.

The **art-float** array type is a special-purpose type whose elements are flonums. When storing into such an array the value (any kind of number) will be converted to a flonum, using the **float** function (see page 113). The advantage of storing flonums in an **art-float** array rather than an **art-q** array is that the numbers in an **art-float** array are not true Lisp objects. Instead the array remembers the numerical value, and when it is **aref**'ed creates a Lisp object (a flonum) to hold the value. Because the system does special storage management for bignums and flonums that are intermediate results, the use of **art-float** arrays can save a lot of work for the garbage collector and hence greatly increase performance. An intermediate result is a Lisp object passed as an argument, stored in a local variable, or returned as the value of a function, but not stored into a special variable, a non-**art-float** array, or list structure. **art-float** arrays also provide a locality of reference advantage over **art-q** arrays containing flonums, since the flonums are contained in the array rather than being separate objects probably on different pages of memory.

The **art-fps-float** array type is another special-purpose type whose elements are flonums. The internal format of this array is compatible with the PDP-11/VAX single-precision floating-point format. The primary purpose of this array type is to interface with the FPS array processor, which can transfer data directly in and out of such an array.

Any type of number may be stored into an **art-fps-float** array, but it will in effect be converted to a flonum, and then rounded off to the 24-bit precision of the PDP-11. If the magnitude of the number is too large, the largest valid floating-point number will be stored. If the magnitude is too small, zero will be stored.

When an element of an `art-fps-float` array is read, a new flonum is created containing the value, just as with an `art-float` array.

The `art-complex` array type is a special purpose type whose elements are arbitrary numbers, which may be complex numbers. (Most of the numeric array types can only hold real numbers.) As compared with an ordinary `art-q` array, `art-complex` provides an advantage in garbage collection similar to what `art-float` provides for floating point numbers.

The `art-complex-float` array type is a special purpose type whose elements are numbers (real or complex) whose real and imaginary parts are both floating point numbers. (If you store a non-floating-point number into the array, its real and imaginary parts are converted to floating point.) This provides maximum advantage in garbage collection if all the elements you wish to store in the array are numbers with floating point real and imaginary parts.

The `art-complex-fps-float` array type is similar to `art-complex-float` but each real or imaginary part is stored in the form used by the FPS array processor. Each element occupies two words, the first being the real part and the second being the imaginary part.

There are three types of arrays which exist only for the implementation of *stack groups*; these types are called `art-stack-group-head`, `art-special-pdl`, and `art-reg-pdl`. Their elements may be any Lisp object; their use is explained in the section on stack groups (see chapter 12, page 186).

### **array-types**

*Variable*

The value of `array-types` is a list of all of the array type symbols such as `art-q`, `art-4b`, `art-string` and so on. The values of these symbols are internal array type code numbers for the corresponding type.

### **array-types** *array-type-code*

Given an internal numeric array-type code, returns the symbolic name of that type.

### **array-elements-per-q**

*Variable*

`array-elements-per-q` is an association list (see page 78) which associates each array type symbol with the number of array elements stored in one word, for an array of that type. If the value is negative, it is instead the number of words per array element, for arrays whose elements are more than one word long.

### **array-elements-per-q** *array-type-code*

Given the internal array-type code number, returns the number of array elements stored in one word, for an array of that type. If the value is negative, it is instead the number of words per array element, for arrays whose elements are more than one word long.

### **array-bits-per-element**

*Variable*

The value of `array-bits-per-element` is an association list (see page 78) which associates each array type symbol with the number of bits of unsigned number it can hold, or `nil` if it can hold Lisp objects. This can be used to tell whether an array can hold Lisp objects or not.

**array-bits-per-element** *array-type-code*

Given the internal array-type code numbers, returns the number of bits per cell for unsigned numeric arrays, or nil for a type of array that can contain Lisp objects.

**array-element-size** *array*

Given an array, returns the number of bits that fit in an element of that array. For arrays that can hold general Lisp objects, the result is 24., assuming you will be storing unsigned fixnums in the array.

## 8.1 Extra Features of Arrays

Any array may have an *array leader*. An array leader is like a one-dimensional **art-q** array which is attached to the main array. So an array which has a leader acts like two arrays joined together. The leader can be stored into and examined by a special set of functions, different from those used for the main array: **array-leader** and **store-array-leader**. The leader is always one-dimensional, and always can hold any kind of Lisp object, regardless of the type or rank of the main part of the array.

Very often the main part of an array will be a homogeneous set of objects, while the leader will be used to remember a few associated non-homogeneous pieces of data. In this case the leader is not used like an array; each slot is used differently from the others. Explicit numeric subscripts should not be used for the leader elements of such an array; instead the leader should be described by a **defstruct** (see page 300).

By convention, element 0 of the array leader of an array is used to hold the number of elements in the array that are "active" in some sense. When the zeroth element is used this way, it is called a *fill pointer*. Many array-processing functions recognize the fill pointer. For instance, if a string (an array of type **art-string**) has seven elements, but its fill pointer contains the value five, then only elements zero through four of the string are considered to be "active"; the string's printed representation will be five characters long, string-searching functions will stop after the fifth element, etc.

**fill-pointer** *array*

Returns the fill pointer of *array*, or nil if it does not have one. This function can be used with **self** to set the array's fill pointer.

The system does not provide a way to turn off the fill-pointer convention; any array that has a leader must reserve element 0 for the fill pointer or avoid using many of the array functions.

Leader element 1 is used in conjunction with the "named structure" feature to associate a "data type" with the array; see page 312. Element 1 is only treated specially if the array is flagged as a named structure.

### 8.1.1 Displaced Arrays

The following explanation of *displaced arrays* is probably not of interest to a beginner; the section may be passed over without losing the continuity of the manual.

Normally, an array is represented as a small amount of header information, followed by the contents of the array. However, sometimes it is desirable to have the header information removed from the actual contents. One such occasion is when the contents of the array must be located in a special part of the Lisp Machine's address space, such as the area used for the control of input/output devices, or the bitmap memory which generates the TV image. Displaced arrays are also used to reference certain special system tables, which are at fixed addresses so the microcode can access them easily.

If you give `make-array` a fixnum or a locative as the value of the `:displaced-to` option, it will create a displaced array referring to that location of virtual memory and its successors. References to elements of the displaced array will access that part of storage, and return the contents; the regular `aref` and `aset` functions are used. If the array is one whose elements are Lisp objects, caution should be used: if the region of address space does not contain typed Lisp objects, the integrity of the storage system and the garbage collector could be damaged. If the array is one whose elements are bytes (such as an `art-4b` type), then there is no problem. It is important to know, in this case, that the elements of such arrays are allocated from the right to the left within the 32-bit words.

It is also possible to have an array whose contents, instead of being located at a fixed place in virtual memory, are defined to be those of another array. Such an array is called an *indirect array*, and is created by giving `make-array` an array as the value of the `:displaced-to` option. The effects of this are simple if both arrays have the same type; the two arrays share all elements. An object stored in a certain element of one can be retrieved from the corresponding element of the other. This, by itself, is not very useful. However, if the arrays have different rank, the manner of accessing the elements differs. Thus, creating a one-dimensional array of nine elements, indirected to a second, two-dimensional array of three elements by three, allows access to the elements in either a one-dimensional or a two-dimensional manner. Weird effects can be produced if the new array is of a different type than the old array; this is not generally recommended. Indirecting an `art-mb` array to an `art-nb` array will do the obvious thing. For instance, if  $m$  is 4 and  $n$  is 1, each element of the first array will contain four bits from the second array, in right-to-left order.

It is also possible to create an indirect array in such a way that when an attempt is made to reference it or store into it, a constant number is added to the subscript given. This number is called the *index-offset*. It is specified at the time the indirect array is created, by giving a fixnum to `make-array` as the value of the `:displaced-index-offset` option. The length of the indirect array need not be the full length of the array it indirects to; it can be smaller. Thus the indirect array can cover just a subrange of the original array. The `nsubstring` function (see page 146) creates such arrays. When using index offsets with multi-dimensional arrays, there is only one index offset; it is added in to the "linearized" subscript which is the result of multiplying each subscript by an appropriate coefficient and adding them together.

## 8.2 Basic Array Functions

**make-array** *dimensions* &rest *options*.

This is the primitive function for making arrays. *dimensions* should be a list of fixnums which are the dimensions of the array; the length of the list will be the rank of the array. For convenience when making a one-dimensional array, the single dimension may be provided as a fixnum rather than a list of one fixnum.

*options* are alternating keywords and values. The keywords may be any of the following:

- :area**           The value specifies in which area (see chapter 15, page 223) the list should be created. It should be either an area number (a fixnum), or `nil` to mean the default area.
- :type**           The value should be a symbolic name of an array type; the most common of these is `art-q`, which is the default. The elements of the array are initialized according to the type: if the array is of a type whose elements may only be fixnums or flonums, then every element of the array will initially be 0 or 0.0; otherwise, every element will initially be `nil`. See the description of array types on page 121. The value of the option may also be the value of a symbol which is an array type name (that is, an internal numeric array type code).
- :initial-value**   This specifies the value to be stored in each element of the new array. If it is not specified, it is `nil` for arrays that can hold arbitrary objects, or 0 or 0.0 for numeric arrays.
- :displaced-to**   If this is not `nil`, then the array will be a *displaced* array. If the value is a fixnum or a locative, `make-array` will create a regular displaced array which refers to the specified section of virtual address space. If the value is an array, `make-array` will create an indirect array (see page 125).
- :leader-length**   The value should be a fixnum. The array will have a leader with that many elements. The elements of the leader will be initialized to `nil` unless the `:leader-list` option is given (see below).
- :leader-list**     The value should be a list. Call the number of elements in the list *n*. The first *n* elements of the leader will be initialized from successive elements of this list. If the `:leader-length` option is not specified, then the length of the leader will be *n*. If the `:leader-length` option is given, and its value is greater than *n*, then the *n*th and following leader elements will be initialized to `nil`. If its value is less than *n*, an error is signalled. The leader elements are filled in forward order; that is, the `car` of the list will be stored in leader element 0, the `cadr` in element 1, and so on.
- :fill-pointer**    The value should be a fixnum. The array will have a leader with at least one element, and this fixnum will go in that element.

Using the `:fill-pointer` option is equivalent to using `:leader-list` with a list one element long. It avoids consing the list, and is also compatible with Common Lisp.

**:displaced-index-offset**

If this is present, the value of the `:displaced-to` option should be an array, and the value should be a non-negative fixnum; it is made to be the `index-offset` of the created indirect array. (See page 125.)

**:named-structure-symbol**

If this is not nil, it is a symbol to be stored in the `named-structure` cell of the array. The array will be tagged as a named structure (see page 312.) If the array has a leader, then this symbol will be stored in leader element 1 regardless of the value of the `:leader-list` option. If the array does not have a leader, then this symbol will be stored in array element zero.

**Examples:**

```

;; Create a one-dimensional array of five elements.
(make-array 5)
;; Create a two-dimensional array,
;; three by four, with four-bit elements.
(make-array '(3 4) ':type 'art-4b)
;; Create an array with a three-element leader.
(make-array 5 ':leader-length 3)
;; Create an array containing 5 t's,
;; and a fill pointer saying the array is full.
(make-array 5 ':initial-value t ':fill-pointer 5)
;; Create a named-structure with five leader
;; elements, initializing some of them.
(setq b (make-array 20 ':leader-length 5
                    ':leader-list '(0 nil foo)
                    ':named-structure-symbol 'bar))

(array-leader b 0) => 0
(array-leader b 1) => bar
(array-leader b 2) => foo
(array-leader b 3) => nil
(array-leader b 4) => nil

```

`make-array` returns the newly-created array, and also returns, as a second value, the number of words allocated in the process of creating the array, i.e. the `%structure-total-size` of the array.

When `make-array` was originally implemented, it took its arguments in the following fixed pattern:

```

(make-array area type dimensions
           &optional displaced-to leader
                   displaced-index-offset
                   named-structure-symbol)

```

*leader* was a combination of the `:leader-length` and `:leader-list` options, and the list was in reverse order. This obsolete form is still supported so that old programs will continue to work, but the new keyword-argument form is preferred.

**aref** *array &rest subscripts*

Returns the element of *array* selected by the *subscripts*. The *subscripts* must be fixnums and their number must match the rank of *array*.

**ar-1** *array i*

**ar-2** *array i j*

**ar-3** *array i j k*

These are obsolete versions of **aref** that only work for one-, two-, or three-dimensional arrays, respectively. There is no reason ever to use them.

**aset** *x array &rest subscripts*

Stores *x* into the element of *array* selected by the *subscripts*. The *subscripts* must be fixnums and their number must match the rank of *array*. The returned value is *x*.

**as-1** *x array i*

**as-2** *x array i j*

**as-3** *x array i j k*

These are obsolete versions of **aset** that only work for one-, two-, or three-dimensional arrays, respectively. There is no reason ever to use them.

**aloc** *array &rest subscripts*

Returns a locative pointer to the element-cell of *array* selected by the *subscripts*. The *subscripts* must be fixnums and their number must match the rank of *array*. See the explanation of locatives in chapter 13, page 197.

**ap-1** *array i*

**ap-2** *array i j*

**ap-3** *array i j k*

These are obsolete versions of **aloc** that only work for one-, two-, or three-dimensional arrays, respectively. There is no reason ever to use them.

The compiler turns **aref** into **ar-1**, **ar-2**, etc. according to the number of subscripts specified, turns **aset** into **as-1**, **as-2**, etc., and turns **aloc** into **ap-1**, **ap-2**, etc. For arrays with more than three dimensions the compiler uses the slightly less efficient form since the special routines only exist for one, two and three dimensions. There is no reason for any program to call **ar-1**, **as-1**, **ar-2**, etc. explicitly; they are documented because there used to be such a reason, and many old programs use these functions. New programs should use **aref**, **aset**, and **aloc**.

A related function, provided only for Maclisp compatibility, is **arraycall** (page 142).

**ar-1-force** *array i*

**as-1-force** *value array i*

**ap-1-force** *array i*

These functions access an array with a single subscript regardless of how many dimensions the array has. They may be useful for manipulating arrays of varying rank, as an alternative to maintaining and updating lists of subscripts or to creating one-dimensional indirect arrays.



In using these functions, you must pay attention to the order in which the array elements are actually stored. See section 8.8, page 137.

**array-leader** *array i*

*array* should be an array with a leader, and *i* should be a fixnum. This returns the *i*th element of *array*'s leader. This is analogous to **aref**.

**store-array-leader** *x array i*

*array* should be an array with a leader, and *i* should be a fixnum. *x* may be any object. *x* is stored in the *i*th element of *array*'s leader. **store-array-leader** returns *x*. This is analogous to **aset**.

**ap-leader** *array i*

*array* should be an array with a leader, and *i* should be a fixnum. This returns a locative pointer to the *i*th element of *array*'s leader. See the explanation of locatives, chapter 13, page 197. This is analogous to **aloc**.

Here are the conditions signaled for various errors in accessing arrays.

**sys:array-has-no-leader** (sys:bad-array-mixin error) *Condition*

This is signaled on a reference to the leader of an array that doesn't have one. The condition instance supports the **:array** operation, which returns the array that was used.

The **:new-array** proceed-type is provided.

**sys:bad-array-mixin** *Condition Flavor*

This mixin is used in the conditions signaled by several kinds of problems pertaining to arrays. It defines prompting for the **:new-array** proceed type.

**sys:array-wrong-number-of-dimensions** (sys:bad-array-mixin error) *Condition*

This is signaled when an array is referenced (either reading or writing) with the wrong number of subscripts; for example, (**aref** "foo" 1 2).

The **:array** operation on the condition instance returns the array that was used. The **:subscripts-used** operation returns the list of subscripts used.

The **:new-array** proceed type is provided. It expects one argument, an array to use instead of the original one.

**sys:subscript-out-of-bounds** (error) *Condition*

This is signaled when there are the right number of subscripts but their values specify an element that falls outside the bounds of the array. The same condition is used by **sys:%instance-ref**, etc., when the index is out of bounds in the instance.

The condition instance supports the operations **:object** and **:subscripts-used**, which return the array or instance and the list of subscripts.

The `:new-subscript` proceed type is provided. It takes an appropriate number of subscripts as arguments. You should provide as many subscripts as there originally were.

**sys:number-array-not-allowed** (sys:bad-array-mixin error) *Condition*

This is signaled by an attempt to use `alloc` on a numeric array. The `:array` operation and the `:new-array` proceed type are available.

### 8.3 Getting Information About an Array

**array-type** *array*

Returns the symbolic type of *array*.

Example:

```
(setq a (make-array '(3 5)))
(array-type a) => art-q
```

**array-length** *array*

*array* may be any array. This returns the total number of elements in *array*. For a one-dimensional array, this is one greater than the maximum allowable subscript. (But if fill pointers are being used, you may want to use `array-active-length`.)

Example:

```
(array-length (make-array 3)) => 3
(array-length (make-array '(3 5)))
=> 17 ;octal, which is 15. decimal
```

**array-active-length** *array*

If *array* does not have a fill pointer, then this returns whatever `(array-length array)` would have. If *array* does have a fill pointer, `array-active-length` returns it. See the general explanation of the use of fill pointers on page 124.

**array-rank** *array*

Returns the number of dimensions of *array*.

Example:

```
(array-rank (make-array '(3 5))) => 2
```

**array-dimension** *array n*

Returns the length of dimension *n* of *array*, or nil if *n* is negative or not less than the rank of *array*.

```
(setq a (make-array '(2 3)))
(array-dimension a 0) => 2
(array-dimension a 1) => 3
```

**array-dimension-n** *n array*

*array* may be any kind of array, and *n* should be a fixnum. If *n* is between 1 and the rank of *array*, this returns the *n*th dimension of *array*. If *n* is 0, this returns the length of the leader of *array*; if *array* has no leader it returns nil. If *n* is any other value, this returns nil.

This function is obsolete.

Examples:

```
(setq a (make-array '(3 5) ':leader-length 7))
(array-dimension-n 1 a) => 3
(array-dimension-n 2 a) => 5
(array-dimension-n 3 a) => nil
(array-dimension-n 0 a) => 7
```

### **array-dimensions** *array*

`array-dimensions` returns a list whose elements are the dimensions of *array*.

Example:

```
(setq a (make-array '(3 5)))
(array-dimensions a) => (3 5)
```

Note: the list returned by `(array-dimensions x)` is equal to the `cdr` of the list returned by `(arraydims x)`.

### **arraydims** *array*

*array* may be any array; it also may be a symbol whose function cell contains an array, for Maclisp compatibility (see section 8.11, page 141). `arraydims` returns a list whose first element is the symbolic name of the type of *array*, and whose remaining elements are its dimensions.

Example:

```
(setq a (make-array '(3 5)))
(arraydims a) => (art-q 3 5)
```

### **array-in-bounds-p** *array* &rest *subscripts*

This function checks whether *subscripts* is a legal set of subscripts for *array*, and returns `t` if they are; otherwise it returns `nil`.

### **array-displaced-p** *array*

*array* may be any kind of array. This predicate returns `t` if *array* is any kind of displaced array (including an indirect array). Otherwise it returns `nil`.

### **array-indirect-p** *array*

*array* may be any kind of array. This predicate returns `t` if *array* is an indirect array. Otherwise it returns `nil`.

### **array-indexed-p** *array*

*array* may be any kind of array. This predicate returns `t` if *array* is an indirect array with an index-offset. Otherwise it returns `nil`.

### **array-index-offset** *array*

*array* may be any kind of array. This returns the index offset of *array* if it is an indirect array which has an index offset. Otherwise it returns `nil`.

**array-has-leader-p** *array*

*array* may be any array. This predicate returns **t** if *array* has a leader; otherwise it returns **nil**.

**array-leader-length** *array*

*array* may be any array. This returns the length of *array*'s leader if it has one, or **nil** if it does not.

## 8.4 Changing the Size of an Array

**adjust-array-size** *array new-size*

If *array* is a one-dimensional array, its size is changed to be *new-size*. If *array* has more than one dimension, its size (**array-length**) is changed to *new-size* by changing only the last dimension.

If *array* is made smaller, the extra elements are lost; if *array* is made bigger, the new elements are initialized in the same fashion as **make-array** (see page 126) would initialize them: either to **nil** or **0**, depending on the type of array.

Example:

```
(setq a (make-array 5))
(aset 'foo a 4)
(aref a 4) => foo
(adjust-array-size a 2)
(aref a 4) => an error occurs
```

If the size of the array is being increased, **adjust-array-size** may have to allocate a new array somewhere. In that case, it alters *array* so that references to it will be made to the new array instead, by means of "invisible pointers" (see **structure-forward**, page 203). **adjust-array-size** will return this new array if it creates one, and otherwise it will return *array*. Be careful to be consistent about using the returned result of **adjust-array-size**, because you may end up holding two arrays which are not the same (i.e. not **eq**), but which share the same contents.

**array-grow** *array &rest dimensions*

**array-grow** creates a new array of the same type as *array*, with the specified dimensions. Those elements of *array* that are still in bounds are copied into the new array. The elements of the new array that are not in the bounds of *array* are initialized to **nil** or **0** as appropriate. If *array* has a leader, the new array will have a copy of it. **array-grow** returns the new array and also forwards *array* to it, like **adjust-array-size**.

Unlike **adjust-array-size**, **array-grow** always creates a new array rather than growing or shrinking the array in place. But **array-grow** of a multi-dimensional array can change all the subscripts and move the elements around in memory to keep each element at the same logical place in the array.

**si:change-indirect-array** *array type dimlist displaced-p index-offset*

Change an indirect array *array*'s type, size, or target pointed at. *type* specifies the new array type, *dimlist* its new dimensions, *displaced-p* the target it should point to (an array, locative or fixnum), and *index-offset* the new offset in the new target.

*array* is returned.

## 8.5 Arrays Overlaid With Lists

These functions manipulate **art-q-list** arrays, which were introduced on page 121.

**g-1-p** *array*

*array* should be an **art-q-list** array. This returns a list which shares the storage of *array*.

Example:

```
(setq a (make-array 4 ':type 'art-q-list))
(aref a 0) => nil
(setq b (g-1-p a)) => (nil nil nil nil)
(rplaca b t)
b => (t nil nil nil)
(aref a 0) => t
(aset 30 a 2)
b => (t nil 30 nil)
```

The following two functions work strangely, in the same way that **store** does, and should not be used in new programs.

**get-list-pointer-into-array** *array-ref*

The argument *array-ref* is ignored, but should be a reference to an **art-q-list** array by applying the array to subscripts (rather than by **aref**). This returns a list object which is a portion of the "list" of the array, beginning with the last element of the last array which has been called as a function.

**get-locative-pointer-into-array** *array-ref*

**get-locative-pointer-into-array** is similar to **get-list-pointer-into-array**, except that it returns a locative, and doesn't require the array to be **art-q-list**. Use **aloc** instead of this function in new programs.

## 8.6 Adding to the End of an Array

**array-push** *array x*

*array* must be a one-dimensional array which has a fill pointer and *x* may be any object. **array-push** attempts to store *x* in the element of the array designated by the fill pointer, and increase the fill pointer by one. If the fill pointer does not designate an element of the array (specifically, when it gets too big), it is unaffected and **array-push** returns **nil**; otherwise, the two actions (storing and incrementing) happen uninterruptibly, and **array-push** returns the *former* value of the fill pointer, i.e. the array index in which it stored *x*. If the array is of type **art-q-list**, an operation similar to **nconc** has taken place, in that

the element has been added to the list by changing the cdr of the formerly last element. The cdr coding is updated to ensure this.

**array-push-extend** *array x &optional extension*

**array-push-extend** is just like **array-push** except that if the fill pointer gets too large, the array is grown to fit the new element; i.e. it never "fails" the way **array-push** does, and so never returns **nil**. *extension* is the number of elements to be added to the array if it needs to be grown. It defaults to something reasonable, based on the size of the array.

**array-pop** *array*

*array* must be a one-dimensional array which has a fill pointer. The fill pointer is decreased by one and the array element designated by the new value of the fill pointer is returned. If the new value does not designate any element of the array (specifically, if it had already reached zero), an error is caused. The two operations (decrementing and array referencing) happen uninterruptibly. If the array is of type **art-q-list**, an operation similar to **nbutlast** has taken place. The cdr coding is updated to ensure this.

**sys:fill-pointer-not-fixnum** (**sys:bad-array-mixin error**) *Condition*

This is signaled when one of the functions in this section is used with an array whose leader element zero is not a fixnum. Most other array accessing operations will simply assume that the array has no fill pointer in such a case, but these cannot be performed without a fill pointer.

The **:array** operation on the condition instance returns the array that was used. The **:new-array** proceed type is supported, with one argument, an array.

## 8.7 Copying an Array

**array-initialize** *array value &optional start end*

Stores *value* into all or part of *array*. *start* and *end* are optional indices which delimit the part of *array* to be initialized. They default to the beginning and end of the array.

This function is by far the fastest way to do the job.

**fillarray** *array x*

*array* may be any type of array, or, for Maclisp compatibility, a symbol whose function cell contains an array. It can also be **nil**, in which case an array of type **art-q** is created. There are two forms of this function, depending on the type of *x*.

If *x* is a list, then **fillarray** fills up *array* with the elements of *list*. If *x* is too short to fill up all of *array*, then the last element of *x* is used to fill the remaining elements of *array*. If *x* is too long, the extra elements are ignored. If *x* is **nil** (the empty list), *array* is filled with the default initial value for its array type (**nil** or **0**).

If *x* is an array (or, for Maclisp compatibility, a symbol whose function cell contains an array), then the elements of *array* are filled up from the elements of *x*. If *x* is too small, then the extra elements of *array* are not affected.

If *array* is multi-dimensional, the elements are accessed in row-major order: the last subscript varies the most quickly. The same is true of *x* if it is an array.

**fillarray** returns *array*; or, if *array* was *nil*, the newly created array.

**listarray** *array* &optional *limit*

*array* may be any type of array or, for Maclisp compatibility, a symbol whose function cell contains an array. **listarray** creates and returns a list whose elements are those of *array*. If *limit* is present, it should be a fixnum, and only the first *limit* (if there are more than that many) elements of *array* are used, and so the maximum length of the returned list is *limit*.

If *array* is multi-dimensional, the elements are accessed in row-major order: the last subscript varies the most quickly.

**list-array-leader** *array* &optional *limit*

*array* may be any type of array, or, for Maclisp compatibility, a symbol whose function cell contains an array. **list-array-leader** creates and returns a list whose elements are those of *array*'s leader. If *limit* is present, it should be a fixnum, and only the first *limit* (if there are more than that many) elements of *array*'s leader are used, and so the maximum length of the returned list is *limit*. If *array* has no leader, *nil* is returned.

**copy-array-contents** *from to*

*from* and *to* must be arrays. The contents of *from* is copied into the contents of *to*, element by element. If *to* is shorter than *from*, the rest of *from* is ignored. If *from* is shorter than *to*, the rest of *to* is filled with *nil* if it is a q-type array, or 0 if it is a numeric array or a string, or 0.0 if it is a flonum array. This function always returns *t*.

Note that even if *from* or *to* has a leader, the whole array is used; the convention that leader element 0 is the "active" length of the array is not used by this function. The leader itself is not copied.

**copy-array-contents** works on multi-dimensional arrays. *from* and *to* are "linearized" subscripts, and elements are taken in the order they are stored in memory. See **array-index-order**, page 137.

**copy-array-contents-and-leader** *from to*

This is just like **copy-array-contents**, but the leader of *from* (if any) is also copied into *to*. **copy-array-contents** copies only the main part of the array.

**copy-array-portion** *from-array from-start from-end to-array to-start to-end*

The portion of the array *from-array* with indices greater than or equal to *from-start* and less than *from-end* is copied into the portion of the array *to-array* with indices greater than or equal to *to-start* and less than *to-end*, element by element. If there are more elements in the selected portion of *to-array* than in the selected portion of *from-array*, the extra elements are filled with the default value as by **copy-array-contents**. If there are more elements in the selected portion of *from-array*, the extra ones are ignored. Multi-dimensional arrays are treated the same way as **copy-array-contents** treats them. This function always returns *t*.

**bitblt** *alu width height from-array from-x from-y to-array to-x to-y*

*from-array* and *to-array* must be two-dimensional arrays of bits or bytes (*art-1b*, *art-2b*, *art-4b*, *art-8b*, *art-16b*, or *art-32b*). **bitblt** copies a rectangular portion of *from-array* into a rectangular portion of *to-array*. The value stored can be a Boolean function of the new value and the value already there, under the control of *alu* (see below). This function is most commonly used in connection with raster images for TV displays.

The top-left corner of the source rectangle is (*ar-2-reverse from-array from-x from-y*). The top-left corner of the destination rectangle is (*ar-2-reverse to-array to-x to-y*). *width* and *height* are the dimensions of both rectangles. If *width* or *height* is zero, **bitblt** does nothing. The *x* coordinates and *width* may be used as the first or the second dimension of the array, according to *array-index-order*. The choice is made so that *x* will come out as horizontal on the TV screen. See page 137.

*from-array* and *to-array* are allowed to be the same array. **bitblt** normally traverses the arrays in increasing order of *x* and *y* subscripts. If *width* is negative, then (*abs width*) is used as the width, but the processing of the *x* direction is done backwards, starting with the highest value of *x* and working down. If *height* is negative it is treated analogously. When **bitblt**'ing an array to itself, when the two rectangles overlap, it may be necessary to work backwards to achieve effects such as shifting the entire array downwards by a certain number of rows. Note that negativity of *width* or *height* does not affect the (*x,y*) coordinates specified by the arguments, which are still the top-left corner even if **bitblt** starts at some other corner.

If the two arrays are of different types, **bitblt** works bit-wise and not element-wise. That is, if you **bitblt** from an *art-2b* array into an *art-4b* array, then two elements of the *from-array* will correspond to one element of the *to-array*.

If **bitblt** goes outside the bounds of the source array, it wraps around. This allows such operations as the replication of a small stipple pattern through a large array. If **bitblt** goes outside the bounds of the destination array, it signals an error.

If *src* is an element of the source rectangle, and *dst* is the corresponding element of the destination rectangle, then **bitblt** changes the value of *dst* to (*boole alu src dst*). See the *boole* function (page 114). There are symbolic names for some of the most useful *alu* functions; they are *tv:alu-seta* (plain copy), *tv:alu-ior* (inclusive or), *tv:alu-xor* (exclusive or), and *tv:alu-andca* (and with complement of source).

**bitblt** is written in highly-optimized microcode and goes very much faster than the same thing written with ordinary *aref* and *aset* operations would. Unfortunately this causes **bitblt** to have a couple of strange restrictions. Wrap-around does not work correctly if *from-array* is an indirect array with an *index-offset*. **bitblt** will signal an error if the first dimensions of *from-array* and *to-array* are not both integral multiples of the machine word length. For *art-1b* arrays, the first dimension must be a multiple of 32., for *art-2b* arrays it must be a multiple of 16., etc.

**%blt** (page 206) is often useful for copying parts of arrays. It can be used to shift a part of an array either up or down.



## 8.8 Order of Array Elements

Currently, multi-dimensional arrays are stored in column-major order rather than row-major order as in Maclisp. Row-major order means that successive memory locations differ in the last subscript, while column-major order means that successive memory locations differ in the first subscript. However, as part of the adoption of Common Lisp, the Lisp machine will switch to the more standard row-major order. This change will take place a few months from now.

### **array-index-order**

*Variable*

If this variable is non-nil, arrays are stored in row-major order (last subscript varies fastest). Otherwise, they are stored in column-major order. The default value of this symbol is nil.

Most user code has no need to know about which order array elements are stored in. There are three known reasons to care: use of multidimensional indirect arrays; paging efficiency (if you want to reference every element in a multi-dimensional array and move linearly through memory to improve locality of reference, you must vary the first subscript fastest rather than the last, in column-major order); and access to the tv screen or to arrays of pixels copied to or from the screen with `bitblt`. The latter is the most important one.

The bits on the screen are actually stored in rows, which means that the dimension that varies fastest has to be the horizontal position. As a result, if arrays are stored in row-major order, the horizontal position must be the second subscript, but if arrays are stored in column-major order, the horizontal position must be the first subscript. To ease the conversion of code that uses arrays of pixels, several bridging functions are provided:

### **make-pixel-array** *width height &rest options*

This is like `make-array` except that the dimensions of the array are *width* and *height*, in whichever order is correct. *width* will be the dimension in the subscript that varies fastest in memory, and *height* will be the other dimension. *options* are passed along to `make-array` to specify everything but the size of the array.

### **pixel-array-width** *array*

Returns the extent of *array*, a two-dimensional array, in the dimension that varies faster through memory. For a screen array, this will always be the width.

### **pixel-array-height** *array*

Returns the extent of *array*, a two-dimensional array, in the dimension that varies slower through memory. For a screen array, this will always be the height.

### **ar-2-reverse** *array horizontal-index vertical-index*

Returns the element of *array* at *horizontal-index* and *vertical-index*. *horizontal-index* will be used as the subscript in whichever dimension varies faster through memory.

### **as-2-reverse** *newvalue array horizontal-index vertical-index*

Stores *newvalue* into the element of *array* at *horizontal-index* and *vertical-index*. *horizontal-index* will be used as the subscript in whichever dimension varies faster through memory.

By replacing calls to `make-array`, `array-dimension-n`, `aref` and `aset` with these functions, you can trivially change your code (without changing the order of arguments) so that it will work no matter what order array elements are stored in, and still display properly on the screen. When the Big Change happens, you will not even need to recompile your program.

## 8.9 Matrices and Systems of Linear Equations

The functions in this section perform some useful matrix operations. The matrices are represented as two-dimensional Lisp arrays. These functions are part of the mathematics package rather than the kernel array system, hence the "math:" in the names.

**math:multiply-matrices** *matrix-1 matrix-2* &optional *matrix-3*

Multiplies *matrix-1* by *matrix-2*. If *matrix-3* is supplied, `multiply-matrices` stores the results into *matrix-3* and returns *matrix-3*; otherwise it creates an array to contain the answer and returns that. All matrices must be two-dimensional arrays, and the first dimension of *matrix-2* must equal the second dimension of *matrix-1*.

**math:invert-matrix** *matrix* &optional *into-matrix*

Computes the inverse of *matrix*. If *into-matrix* is supplied, stores the result into it and returns it; otherwise it creates an array to hold the result and returns that. *matrix* must be two-dimensional and square. The Gauss-Jordan algorithm with partial pivoting is used. Note: if you want to solve a set of simultaneous equations, you should not use this function; use `math:decompose` and `math:solve` (see below).

**math:transpose-matrix** *matrix* &optional *into-matrix*

Transposes *matrix*. If *into-matrix* is supplied, stores the result into it and returns it; otherwise it creates an array to hold the result and returns that. *matrix* must be a two-dimensional array. *into-matrix*, if provided, must be two-dimensional and have sufficient dimensions to hold the transpose of *matrix*.

**math:determinant** *matrix*

Returns the determinant of *matrix*. *matrix* must be a two-dimensional square matrix.

The next two functions are used to solve sets of simultaneous linear equations. `math:decompose` takes a matrix holding the coefficients of the equations and produces the LU decomposition; this decomposition can then be passed to `math:solve` along with a vector of right-hand sides to get the values of the variables. If you want to solve the same equations for many different sets of right-hand side values, you only need to call `math:decompose` once. In terms of the argument names used below, these two functions exist to solve the vector equation  $Ax = b$  for  $x$ .  $A$  is a matrix.  $b$  and  $x$  are vectors.

**math:decompose** *a* &optional *lu ps*

Computes the LU decomposition of matrix *a*. If *lu* is non-nil, stores the result into it and returns it; otherwise it creates an array to hold the result, and returns that. The lower triangle of *lu*, with ones added along the diagonal, is L, and the upper triangle of *lu* is U, such that the product of L and U is *a*. Gaussian elimination with partial pivoting is used. The *lu* array is permuted by rows according to the permutation array *ps*, which is also produced by this function; if the argument *ps* is supplied, the permutation

array is stored into it; otherwise, an array is created to hold it. This function returns two values, the LU decomposition and the permutation array.

**math:solve** *lu ps b* &optional *x*

This function takes the LU decomposition and associated permutation array produced by **math:decompose** and solves the set of simultaneous equations defined by the original matrix *a* and the right-hand sides in the vector *b*. If *x* is supplied, the solutions are stored into it and it is returned; otherwise an array is created to hold the solutions and that is returned. *b* must be a one-dimensional array.

**math:list-2d-array** *array*

Returns a list of lists containing the values in *array*, which must be a two-dimensional array. There is one element for each row; each element is a list of the values in that row.

**math:fill-2d-array** *array list*

This is the opposite of **math:list-2d-array**. *list* should be a list of lists, with each element being a list corresponding to a row. *array*'s elements are stored from the list. Unlike **fillarray** (see page 134), if *list* is not long enough, **math:fill-2d-array** "wraps around", starting over at the beginning. The lists which are elements of *list* also work this way.

**math:singular-matrix** (sys:arithmetic-error error)

*Condition*

This is signaled when any of the matrix manipulation functions in this section has trouble because of a singular matrix. (In some functions, such as **math:determinant**, a singular matrix is not a problem.)

The **:matrix** operation on the condition instance returns the matrix which is singular.

## 8.10 Planes

A *plane* is effectively an array whose bounds, in each dimension, are plus-infinity and minus-infinity; all integers are legal as indices. Planes may be of any rank. When you create a plane, you do not need to specify any size, just the rank. You also specify a default value. At that moment, every component of the plane has that value. As you can't ever change more than a finite number of components, only a finite region of the plane need actually be stored. When you refer to an element for which space has not actually been allocated, you just get the default value.

The regular array accessing functions don't work on planes. You can use **make-plane** to create a plane, **plane-aref** or **plane-ref** to get the value of a component, and **plane-aset** or **plane-store** to store into a component. **array-rank** will work on a plane.

A plane is actually stored as an array with a leader. The array corresponds to a rectangular, aligned region of the plane, containing all the components in which a **plane-store** has been done (and, usually, others which have never been altered). The lowest-coordinate corner of that rectangular region is given by the **plane-origin** in the array leader. The highest-coordinate corner can be found by adding the **plane-origin** to the **array-dimensions** of the array. The **plane-**

**default** is the contents of all the elements of the plane that are not actually stored in the array. The **plane-extension** is the amount to extend a plane by in any direction when the plane needs to be extended. The default is 32.

If you never use any negative indices, then the **plane-origin** will be all zeroes and you can use regular array functions, such as **aref** and **aset**, to access the portion of the plane that is actually stored. This can be useful to speed up certain algorithms. In this case you can even use the **bitblt** function on a two-dimensional plane of bits or bytes, provided you don't change the **plane-extension** to a number that is not a multiple of 32.

**make-plane** *rank &rest options*

Creates and returns a plane. *rank* is the number of dimensions. *options* is a list of alternating keyword symbols and values. The allowed keywords are:

**:type**           The array type symbol (e.g. **art-1b**) specifying the type of the array out of which the plane is made.

**:default-value** The default component value as explained above.

**:extension**     The amount by which to extend the plane, as explained above.

Example:

```
(make-plane 2 ':type 'art-4b ':default-value 3)
```

creates a two-dimensional plane of type **art-4b**, with default value 3.

**plane-origin** *plane*

A list of numbers, giving the lowest coordinate values actually stored.

**plane-default** *plane*

This is the contents of the infinite number of plane elements that are not actually stored.

**plane-extension** *plane*

The amount to extend the plane by, in any direction, when **plane-store** is done outside of the currently-stored portion.

**plane-aref** *plane &rest subscripts*

**plane-ref** *plane subscripts*

These two functions return the contents of a specified element of a plane. They differ only in the way they take their arguments; **plane-aref** wants the subscripts as arguments, while **plane-ref** wants a list of subscripts.

**plane-aset** *datum plane &rest subscripts*

**plane-store** *datum plane subscripts*

These two functions store *datum* into the specified element of a plane, extending it if necessary, and return *datum*. They differ only in the way they take their arguments; **plane-aset** wants the subscripts as arguments, while **plane-store** wants a list of subscripts.

## 8.11 Maclisp Array Compatibility

The functions in this section are provided only for Maclisp compatibility and should not be used in new programs.

Fixnum arrays do not exist (however, see Zetalisp's small-positive-integer arrays). Flonum arrays exist but you do not use them in the same way; no declarations are required or allowed. "Un-garbage-collected" arrays do not exist. Readtables and obarrays are represented as arrays, but Zetalisp does not use special array types for them. See the descriptions of `read` (page 383) and `intern` (page 512) for information about readtables and obarrays (packages). There are no "dead" arrays, nor are Multics "external" arrays provided.

The `arraycall` function exists for compatibility but should not be used (see `aref`, page 128.)

Subscripts are always checked for validity, regardless of the value of `*rset` and whether the code is compiled or not. However, in a multi-dimensional array, an error is caused only if the subscripts would have resulted in a reference to storage outside of the array. For example, if you have a 2 by 7 array and refer to an element with subscripts 3 and 1, no error will be caused despite the fact that the reference is invalid; but if you refer to element 1 by 100, an error will be caused. In other words, subscript errors will be caught if and only if they refer to storage outside the array; some errors are undetected, but they will only clobber (alter randomly) some other element of the same array, not something completely unpredictable.

Currently, multi-dimensional arrays are stored in column-major order rather than row-major order as in Maclisp. See section 8.8, page 137 for further discussion of this issue.

`loadarrays` and `dumparrays` are not provided. However, arrays can be put into QFASL files; see section 16.8, page 245.

The `*rearray` function is not provided, since not all of its functionality is available in Zetalisp. Its most common uses are implemented by `adjust-array-size`.

In Maclisp, arrays are usually kept on the `array` property of symbols, and the symbols are used instead of the arrays. In order to provide some degree of compatibility for this manner of using arrays, the `array`, `*array`, and `store` functions are provided, and when arrays are applied to arguments, the arguments are treated as subscripts and `apply` returns the corresponding element of the array.

**array** *&quote symbol type &eval &rest dims*

This creates an `art-q` type array in `default-array-area` with the given dimensions. (That is, *dims* is given to `make-array` as its first argument.) *type* is ignored. If *symbol* is `nil`, the array is returned; otherwise, the array is put in the function cell of *symbol*, and *symbol* is returned.

**\*array** *symbol type &rest dims*

This is just like `array`, except that all of the arguments are evaluated.

**store** *array-ref* *x*

*Special Form*

**store** stores *x* into the specified array element. *array-ref* should be a form which references an array by calling it as a function (**aref** forms are not acceptable). First *x* is evaluated, then *array-ref* is evaluated, and then the value of *x* is stored into the array cell last referenced by a function call, presumably the one in *array-ref*.

**xstore** *x* *array-ref*

This is just like **store**, but it is not a special form; this is because the arguments are in the other order. This function only exists for the compiler to compile the **store** special form into, and should never be used by programs.

**arraycall** *ignored array &rest subscripts*

(**arraycall** *t array sub1 sub2...*) is the same as (**aref** *array sub1 sub2...*). It exists for Maclisp compatibility.