Annex C

(informative)

Extended notes

C.1 Section 4 notes

C.1.1 Intrinsic and derived data types (4.3, 4.4)

FORTRAN 77 provided only data types explicitly defined in the standard (logical, integer, real, double precision, complex, and character). This standard provides those intrinsic types and provides derived types to allow the creation of new data types. A derived-type definition specifies a data structure consisting of components of intrinsic types and of derived types. Such a type definition does not represent a data object, but rather, a template for declaring named objects of that derived type. For example, the definition

```
TYPE POINT
   INTEGER X_COORD
   INTEGER Y_COORD
END TYPE POINT
```

specifies a new derived type named POINT which is composed of two components of intrinsic type integer (X_COORD and Y_COORD). The statement TYPE (POINT) FIRST, LAST declares two data objects, FIRST and LAST, that can hold values of type POINT.

FORTRAN 77 provided REAL and DOUBLE PRECISION intrinsic types as approximations to mathematical real numbers. This standard generalizes REAL as an intrinsic type with a type parameter that selects the approximation method. The type parameter is named kind and has values that are processor dependent. DOUBLE PRECISION is treated as a synonym for REAL (k), where k is the implementation-defined kind type parameter value KIND (0.0D0).

Real literal constants may be specified with a kind type parameter to ensure that they have a particular kind type parameter value (4.3.1.2).

For example, with the specifications

```
INTEGER Q
PARAMETER (Q = 8)
REAL (Q) B
```

the literal constant 10.93_Q has the same precision as the variable B.

FORTRAN 77 did not allow zero-length character strings. They are permitted by this standard (4.3.2.1).

Objects are of different derived type if they are declared using different derived-type definitions. For example,

```
TYPE APPLES
   INTEGER NUMBER
END TYPE APPLES
TYPE ORANGES
   INTEGER NUMBER
END TYPE ORANGES
TYPE (APPLES) COUNT1
TYPE (ORANGES) COUNT2
```

```
COUNT1 = COUNT2 ! Erroneous statement mixing apples and oranges
```

Even though all components of objects of type APPLES and objects of type ORANGES have identical intrinsic types, the objects are of different types.

The distinction between the ultimate and direct components of a derived type can be made clear with an example.

```
TYPE PERSON
  CHARACTER (30) :: NAME
  INTEGER :: AGE
END TYPE PERSON

TYPE HOUSEHOLD
  INTEGER :: NUM_PEOPLE
  TYPE (PERSON), POINTER :: PEOPLE(:)
END TYPE HOUSEHOLD

TYPE COMMUNITY
  INTEGER :: NUM_HOUSES
  TYPE (HOUSEHOLD) :: HOUSES (100)
END TYPE COMMUNITY
```

The ultimate and direct components of PERSON are NAME and AGE. The ultimate and direct components of HOUSEHOLD are NUM_PEOPLE and PEOPLE. The ultimate components of COMMUNITY are NUM_HOUSES, NUM_PEOPLE, and PEOPLE. However, the direct components of COMMUNITY are NUM_HOUSES, HOUSES, NUM_PEOPLE, and PEOPLE. Note that the components of PERSON are neither ultimate nor direct components of HOUSEHOLD or COMMUNITY because the component PEOPLE has the POINTER attribute. The direct components include all the ultimate components and any components of derived type that lead to those ultimate components.

C.1.2 Selection of the approximation methods (4.3.1.2)

This standard permits the selection of the real approximation method for an entire program to be parameterized through the use of the parameterized real data type and module. This is accomplished by defining a named integer constant, say FLOAT, to have a specific kind type parameter value, and to use that named constant in all real, complex, and derived-type declarations. For example, the specification statements

```
INTEGER FLOAT
PARAMETER (FLOAT = 8)
REAL (FLOAT) X, Y
COMPLEX (FLOAT) Z
```

specify that the approximation method corresponding to a kind type parameter value of 8 is supplied for the data objects X, Y, and Z in the program unit. The kind type parameter value FLOAT can be made available to an entire program by placing the INTEGER and PARAMETER specification statements in a module and accessing the named constant FLOAT with a USE statement. Note that by changing 8 to 4 once in the module, a different approximation method is selected.

To avoid the use of the processor-dependent values 4 or 8, replace 8 by KIND (0.0) or KIND (0.0D0). Another way to avoid these processor-dependent values is to select the kind value using the intrinsic inquiry function SELECTED_REAL_KIND. This function, given integer arguments P and R specifying minimum requirements for decimal precision and decimal exponent range, respectively, returns the kind type parameter value of the approximation method that has at least P decimal digits of precision and at least a range for positive numbers of 10^{-R} to 10^{R}. In the above specification statement, the 8 may be replaced by, for instance, SELECTED_REAL_KIND (10, 50), which requires an approximation method to be selected with at least 10 decimal digits of precision and an exponent range from 10^{-50} to 10^{50}. There are no magnitude or ordering constraints placed on kind values, in order that implementors may have flexibility in assigning such values and may add new kinds without changing previously assigned kind values.

As kind values have no portable meaning, a good practice is to use them in programs only through named constants as described above (for example, SINGLE, IEEE_SINGLE, DOUBLE, and QUAD), rather than using the kind values directly.

C.1.3 Pointers (4.4.1)

Pointers are names that can change dynamically their association with a target object. In a sense, a normal variable is a name with a fixed association with a specific object. A normal variable name refers to the same storage space throughout the lifetime of a variable. A pointer name may refer to different storage space, or even no storage space, at different times. A variable may be considered to be a descriptor for space to hold values of the appropriate type, type parameters, and array rank such that the values stored in the descriptor are fixed when the variable is created by its declaration. A pointer also may be considered to be a descriptor, but one whose values may be changed dynamically so as to describe different pieces of storage. When a pointer is declared, space to hold the descriptor is created, but the space for the target object is not created.

A derived type may have one or more components that are defined to be pointers. It may have a component that is a pointer to an object of the same derived type. This "recursive" data definition allows dynamic data structures such as linked lists, trees, and graphs to be constructed. For example:

```
TYPE NODE          ! Define a "recursive" type
   INTEGER :: VALUE = 0
   TYPE (NODE), POINTER :: NEXT_NODE => NULL ( )
END TYPE NODE

TYPE (NODE), TARGET :: HEAD       ! Automatically initialized
TYPE (NODE), POINTER :: CURRENT, TEMP  ! Declare pointers
INTEGER :: IOEM, K

CURRENT => HEAD                   ! CURRENT points to head of list
DO
   READ (*, *, IOSTAT = IOEM) K  ! Read next value, if any
   IF (IOEM /= 0) EXIT
   ALLOCATE (TEMP)                ! Create new cell each iteration
   TEMP % VALUE = K                ! Assign value to cell
    CURRENT % NEXT_NODE => TEMP   ! Attach new cell to list
   CURRENT => TEMP                ! CURRENT points to new end of list
END DO
```

A list is now constructed and the last linked cell contains a disassociated pointer. A loop can be used to "walk through" the list.

```
CURRENT => HEAD
DO
   IF (.NOT. ASSOCIATED (CURRENT % NEXT_NODE)) EXIT
   CURRENT => CURRENT % NEXT_CELL
   WRITE (*, *) CURRENT % VALUE
END DO
```

C.2 Section 5 notes

C.2.1 The POINTER attribute (5.1.2.7)

The POINTER attribute shall be specified to declare a pointer. The type, type parameters, and rank, which may be specified in the same statement or with one or more attribute specification statements, determine the characteristics of the target objects that may be associated with the pointers declared in the statement. An obvious model for interpreting declarations of pointers is that such declarations create for each name a descriptor. Such a descriptor includes all the data necessary to describe fully and locate in memory an object and all subobjects of the type, type parameters, and rank specified. The descriptor is created empty; it does not contain values describing how to access an actual memory space. These descriptor values will be filled in when the pointer is associated with actual target space.

The following example illustrates the use of pointers in an iterative algorithm:

```
PROGRAM DYNAM_ITER
   REAL, DIMENSION (:, :), POINTER :: A, B, SWAP  ! Declare pointers
   ...
   READ (*, *) N, M
   ALLOCATE (A (N, M), B (N, M))  ! Allocate target arrays
   ! Read values into A
   ...
   ITER: DO
      ...
      ! Apply  transformation of values in A to produce values in B
      ...
      IF (CONVERGED) EXIT ITER
      ! Swap A and B
      SWAP => A; A => B; B => SWAP
   END DO ITER
   ...
END PROGRAM DYNAM_ITER
```

C.2.2 The TARGET attribute (5.1.2.8)

The TARGET attribute shall be specified for any nonpointer object that may, during the execution of the program, become associated with a pointer. This attribute is defined primarily for optimization purposes. It allows the processor to assume that any nonpointer object not explicitly declared as a target may be referred to only by way of its original declared name. It also means that implicitly-declared objects shall not be used as pointer targets. This will allow a processor to perform optimizations that otherwise would not be possible in the presence of certain pointers.

The following example illustrates the use of the TARGET attribute in an iterative algorithm:

```
PROGRAM ITER
   REAL, DIMENSION (1000, 1000), TARGET :: A, B
   REAL, DIMENSION (:, :), POINTER       :: IN, OUT, SWAP
   ...
   ! Read values into A
   ...
   IN => A            ! Associate IN with target A
   OUT => B           ! Associate OUT with target B
   ...
   ITER:DO
      ...
      ! Apply transformation of IN values to produce OUT
      ...
      IF (CONVERGED) EXIT ITER
```

```
         ! Swap IN and OUT
         SWAP => IN; IN => OUT; OUT => SWAP
      END DO ITER
      . . .
   END PROGRAM ITER
```

C.3 Section 6 notes

C.3.1 Structure components (6.1.2)

Components of a structure are referenced by writing the components of successive levels of the structure hierarchy until the desired component is described. For example,

```
TYPE ID_NUMBERS
   INTEGER SSN
   INTEGER EMPLOYEE_NUMBER
END TYPE ID_NUMBERS

TYPE PERSON_ID
   CHARACTER (LEN=30) LAST_NAME
   CHARACTER (LEN=1) MIDDLE_INITIAL
   CHARACTER (LEN=30) FIRST_NAME
   TYPE (ID_NUMBERS) NUMBER
END TYPE PERSON_ID

TYPE PERSON
   INTEGER AGE
   TYPE (PERSON_ID) ID
END TYPE PERSON

TYPE (PERSON) GEORGE, MARY

PRINT *, GEORGE % AGE              ! Print the AGE component
PRINT *, MARY % ID % LAST_NAME     ! Print LAST_NAME of MARY
PRINT *, MARY % ID % NUMBER % SSN ! Print SSN of MARY
PRINT *, GEORGE % ID % NUMBER  ! Print SSN and EMPLOYEE_NUMBER of GEORGE
```

A structure component may be a data object of intrinsic type as in the case of GEORGE % AGE or it may be of derived type as in the case of GEORGE % ID % NUMBER. The resultant component may be a scalar or an array of intrinsic or derived type.

```
TYPE LARGE
   INTEGER ELT (10)
   INTEGER VAL
END TYPE LARGE

TYPE (LARGE) A (5)        ! 5 element array, each of whose elements
                          ! includes a 10 element array ELT and
                          ! a scalar VAL.
PRINT *, A (1)            ! Prints 10 element array ELT and scalar VAL.
PRINT *, A (1) % ELT (3) ! Prints scalar element 3
                          ! of array element 1 of A.
PRINT *, A (2:4) % VAL    ! Prints scalar VAL for array elements
                          ! 2 to 4 of A.
```

C.3.2 Pointer allocation and association

The effect of ALLOCATE, DEALLOCATE, NULLIFY, and pointer assignment is that they are interpreted as changing the values in the descriptor that is the pointer. An ALLOCATE is assumed to create space for a suitable object and to "assign" to the pointer the values necessary to describe that space. A NULLIFY breaks the association of the pointer with the space. A DEALLOCATE breaks the association and releases the space. Depending on the implementation, it could be seen as setting a flag in the pointer that indicates whether the values in the descriptor are valid, or it could clear the descriptor values to some (say zero) value indicative of the pointer not currently pointing to anything. A pointer assignment copies the values necessary to describe the space occupied by the target into the descriptor that is the pointer. Descriptors are copied, values of objects are not.

If PA and PB are both pointers and PB currently is associated with a target, then

```
PA => PB
```

results in PA being associated with the same target as PB. If PB was disassociated, then PA becomes disassociated.

The standard is specified so that such associations are direct and independent. A subsequent statement

```
PB => D
```

or

```
ALLOCATE (PB)
```

has no effect on the association of PA with its target. A statement

```
DEALLOCATE (PB)
```

leaves PA as a "dangling pointer" to space that has been released. The program shall not use PA again until it becomes associated via pointer assignment or an ALLOCATE statement.

DEALLOCATE should only be used to release space that was created by a previous ALLOCATE. Thus the following is invalid:

```
REAL, TARGET :: T
REAL, POINTER :: P
   ...
P => T
DEALLOCATE (P) ! Not allowed: P's target was not allocated
```

The basic principle is that ALLOCATE, NULLIFY, and pointer assignment primarily affect the pointer rather than the target. ALLOCATE creates a new target but, other than breaking its connection with the specified pointer, it has no effect on the old target. Neither NULLIFY nor pointer assignment has any effect on targets. A given piece of memory that was allocated and associated with a pointer will become inaccessible to a program if the pointer is nullified and no other pointer was associated with this piece of memory. Such pieces of memory may be reused by the processor if this is expedient. However, whether such inaccessible memory is in fact reused is entirely processor dependent.

C.4 Section 7 notes

C.4.1 Character assignment

The FORTRAN 77 restriction that none of the character positions being defined in the character assignment statement may be referenced in the expression has been removed (7.5.1.5).

C.4.2 Evaluation of function references

If more than one function reference appears in a statement, they may be executed in any order (subject to a function result being evaluated after the evaluation of its arguments) and their values shall not depend on the order of execution. This lack of dependence on order of evaluation permits parallel execution of the function references (7.1.7.1).

C.4.3 Pointers in expressions

A pointer is basically considered to be like any other variable when it is used as a primary in an expression. If a pointer is used as an operand to an operator that expects a value, the pointer will automatically deliver the value stored in the space currently described by the pointer, that is, the value of the target object currently associated with the pointer. In value-demanding expression contexts, pointers are dereferenced.

C.4.4 Pointers on the left side of an assignment

A pointer that appears on the left of an intrinsic assignment statement also is dereferenced and is taken to be referring to the space that is its current target. Therefore, the assignment statement specifies the normal copying of the value of the right-hand expression into this target space. All the normal rules of intrinsic assignment hold; the type and type parameters of the expression and the pointer target shall agree and the shapes shall be conformable.

For intrinsic assignment of derived types, nonpointer components are assigned and pointer components are pointer assigned. Dereferencing is applied only to entire scalar objects, not selectively to pointer subobjects.

For example, suppose a type such as

```
TYPE CELL
   INTEGER :: VAL
   TYPE (CELL), POINTER :: NEXT_CELL
END TYPE CELL
```

is defined and objects such as HEAD and CURRENT are declared using

```
TYPE (CELL), TARGET :: HEAD
TYPE (CELL), POINTER :: CURRENT
```

If a linked list has been created and attached to HEAD and the pointer CURRENT has been allocated space, statements such as

```
CURRENT = HEAD
CURRENT = CURRENT % NEXT_CELL
```

cause the contents of the cells referenced on the right to be copied to the cell referred to by CURRENT. In particular, the right-hand side of the second statement causes the pointer component in the cell, CURRENT, to be selected. This pointer is dereferenced because it is in an expression context to produce the target's integer value and a pointer to a cell that is in the target's NEXT_CELL component. The left-hand side causes the pointer CURRENT to be dereferenced to produce its present target, namely space to hold a cell (an integer and a cell pointer). The integer value on the right is copied to the integer space on the left and the pointer components are pointer assigned (the descriptor on the right is copied into the space for a descriptor on the left). When a statement such as

```
CURRENT => CURRENT % NEXT_CELL
```

is executed, the descriptor value in CURRENT % NEXT_CELL is copied to the descriptor named CURRENT. In this case, CURRENT is made to point at a different target.

In the intrinsic assignment statement, the space associated with the current pointer does not change but the values stored in that space do. In the pointer assignment, the current pointer is made to associate with different space. Using the intrinsic assignment causes a linked list of cells to be moved up through the current "window"; the pointer assignment causes the current pointer to be moved down through the list of cells.

C.4.5 An example of a FORALL construct containing a WHERE construct

```
INTEGER :: A(5,5)
...
FORALL (I = 1:5)
   WHERE (A(I,:) .EQ. 0)
      A(:,I) = I
   ELSEWHERE (A(I,:) > 2)
      A(I,:) = 6
   END WHERE
END FORALL
```

If prior to execution of the FORALL, A has the value

```
A =      1  0  0  0  0
         2  1  1  1  0
         1  2  2  0  2
         2  1  0  2  3
         1  0  0  0  0
```

After execution of the assignment statements following the WHERE statement A has the value A′. The mask created from row one is used to mask the assignments to column one; the mask from row two is used to mask assignments to column two; etc.

```
A′ =     1  0  0  0  0
         1  1  1  1  5
         1  2  2  4  5
         1  1  3  2  5
         1  2  0  0  5
```

The masks created for assignments following the ELSEWHERE statement are

```
.NOT. (A(I,:) .EQ. 0) .AND. (A′(I,:) > 2)
```

Thus the only elements affected by the assignments following the ELSEWHERE statement are A(3, 5) and A(4, 5). After execution of the FORALL construct, A has the value

```
A =      1  0  0  0  0
         1  1  1  1  5
         1  2  2  4  6
         1  1  3  2  6
         1  2  0  0  5
```

C.4.6 Examples of FORALL statements

Example 1:

```
FORALL (J=1:M,  K=1:N)  X(K, J) = Y(J, K)
FORALL (K=1:N) X(K, 1:M) = Y(1:M, K)
```

These statements both copy columns 1 through N of array Y into rows 1 through N of array X. They are equivalent to

```
X(1:N, 1:M) = TRANSPOSE (Y(1:M, 1:N) )
```

Example 2:

The following FORALL statement computes five partial sums of subarrays of J.

```
J = (/ 1, 2, 3, 4, 5 /)
FORALL (K = 1:5) J(K) = SUM (J(1:K) )
```

SUM is allowed in a FORALL because intrinsic functions are pure (12.6). After execution of the FORALL statement, J = (/ 1, 3, 6, 10, 15 /).

Example 3:

```
FORALL (I = 2:N-1) X(I) = (X(I-1) + 2*X(I) + X(I+1) ) / 4
```

has the same effect as

```
X(2:N-1) = (X(1:N-2) + 2*X(2:N-1) + X(3:N+1) ) / 4
```

C.5 Section 8 notes

C.5.1 Loop control

Fortran provides several forms of loop control:

(1) With an iteration count and a DO variable. This is the classic Fortran DO loop.
(2) Test a logical condition before each execution of the loop (DO WHILE).
(3) DO "forever".

C.5.2 The CASE construct

At most one case block is selected for execution within a CASE construct, and there is no fall-through from one block into another block within a CASE construct. Thus there is no requirement for the user to exit explicitly from a block.

C.5.3 Additional examples of DO constructs

The following are all valid examples of block DO constructs.

Example 1:

```
      SUM = 0.0
      READ (IUN) N
      OUTER: DO L = 1, N           ! A DO with a construct name
         READ (IUN) IQUAL, M, ARRAY (1:M)
         IF (IQUAL < IQUAL_MIN) CYCLE OUTER    ! Skip inner loop
         INNER: DO 40 I = 1, M     ! A DO with a label and a name
            CALL CALCULATE (ARRAY (I), RESULT)
            IF (RESULT < 0.0) CYCLE
            SUM = SUM + RESULT
            IF (SUM > SUM_MAX) EXIT OUTER
   40    END DO INNER
      END DO OUTER
```

The outer loop has an iteration count of MAX (N, 0), and will execute that number of times or until SUM exceeds SUM_MAX, in which case the EXIT OUTER statement terminates both loops. The inner loop is skipped by the first CYCLE statement if the quality flag, IQUAL, is too low. If CALCULATE returns a negative RESULT, the second CYCLE statement prevents it from being summed. Both loops have construct names and the inner loop also has a label. A construct name is required in the EXIT statement in order to terminate both loops, but is optional in the CYCLE statements because each belongs to its innermost loop.

Example 2:

```
      N = 0
      DO 50, I = 1, 10
         J = I
         DO K = 1, 5
            L = K
            N = N + 1  ! This statement executes 50 times
         END DO        ! Nonlabeled DO inside a labeled DO
   50 CONTINUE
```

After execution of the above program fragment, I = 11, J = 10, K = 6, L = 5, and N = 50.

Example 3:

```
      N = 0
      DO I = 1, 10
         J = I
         DO 60, K = 5, 1  ! This inner loop is never executed
            L = K
            N = N + 1
   60    CONTINUE         ! Labeled DO inside a nonlabeled DO
      END DO
```

After execution of the above program fragment, I = 11, J = 10, K = 5, N = 0, and L is not defined by these statements.

The following are all valid examples of nonblock DO constructs:

Example 4:

```
      DO 70
         READ (IUN, '(1X, G14.7)', IOSTAT = IOS) X
         IF (IOS .NE. 0) EXIT
         IF (X < 0.) GOTO 70
         CALL SUBA (X)
         CALL SUBB (X)
              ...
         CALL SUBY (X)
         CYCLE
   70    CALL SUBNEG (X)  ! SUBNEG called only when X < 0.
```

This is not a block DO construct because it ends with a statement other than END DO or CONTINUE. The loop will continue to execute until an end-of-file condition or input/output error occurs.

Example 5:

```
      SUM = 0.0
      READ (IUN) N
      DO 80, L = 1, N
         READ (IUN) IQUAL, M, ARRAY (1:M)
         IF (IQUAL < IQUAL_MIN) M = 0  ! Skip inner loop
         DO 80 I = 1, M
            CALL CALCULATE (ARRAY (I), RESULT)
            IF (RESULT < 0.) CYCLE
            SUM = SUM + RESULT
            IF (SUM > SUM_MAX) GOTO 81
   80    CONTINUE  ! This CONTINUE is shared by both loops
   81 CONTINUE
```

This example is similar to Example 1 above, except that the two loops are not block DO constructs because they share the CONTINUE statement with the label 80. The terminal construct of the outer DO is the entire inner DO construct. The inner loop is skipped by forcing M to zero. If SUM grows too large, both loops are terminated by branching to the CONTINUE statement labeled 81. The CYCLE statement in the inner loop is used to skip negative values of RESULT.

Example 6:

```
      N = 0
      DO 100 I = 1, 10
         J = I
         DO 100 K = 1, 5
            L = K
  100       N = N + 1  ! This statement executes 50 times
```

In this example, the two loops share an assignment statement. After execution of this program fragment, I = 11, J = 10, K = 6, L = 5, and N = 50.

Example 7:

```
      N = 0
      DO 200 I = 1, 10
         J = I
         DO 200 K = 5, 1  ! This inner loop is never executed
            L = K
  200       N = N + 1
```

This example is very similar to the previous one, except that the inner loop is never executed. After execution of this program fragment, I = 11, J = 10, K = 5, N = 0, and L is not defined by these statements.

C.5.4 Examples of invalid DO constructs

The following are all examples of invalid skeleton DO constructs:

Example 1:

```
DO I = 1, 10
   ...
END DO LOOP   ! No matching construct name
```

Example 2:

```
LOOP: DO 1000 I = 1, 10   ! No matching construct name
   ...
1000  CONTINUE
```

Example 3:

```
LOOP1: DO
   ...
END DO LOOP2   ! Construct names don't match
```

Example 4:

```
DO I = 1, 10   ! Label required or ...
   ...
1010 CONTINUE  ! ... END DO required
```

Example 5:

```
DO 1020 I = 1, 10
   ...
1021 END DO        ! Labels don't match
```

Example 6:

```
FIRST: DO I = 1, 10
   SECOND: DO J = 1, 5
      ...
   END DO FIRST   ! Improperly nested DOs
END DO SECOND
```

C.6 Section 9 notes

C.6.1 Files (9.2)

This standard accommodates, but does not require, file cataloging. To do this, several concepts are introduced.

C.6.1.1 File connection (9.3)

Before any input/output may be performed on a file, it shall be connected to a unit. The unit then serves as a designator for that file as long as it is connected. To be connected does not imply that "buffers" have or have not been allocated, that "file-control tables" have or have not been filled out, or that any other method of implementation has been used. Connection means that (barring some other fault) a READ or WRITE statement may be executed on the unit, hence on the file. Without a connection, a READ or WRITE statement shall not be executed.

C.6.1.2 File existence (9.2.1.1)

Totally independent of the connection state is the property of existence, this being a file property. The processor "knows" of a set of files that exist at a given time for a given program. This set would include tapes ready to read, files in a catalog, a keyboard, a printer, etc. The set may exclude files inaccessible to the program because of security, because they are already in use by another program, etc. This standard does not specify which files exist, hence wide latitude is available to a processor to implement security, locks, privilege techniques, etc. Existence is a convenient concept to designate all of the files that a program can potentially process.

All four combinations of connection and existence may occur:

Connect	Exist	Examples
Yes	Yes	A card reader loaded and ready to be read
Yes	No	A printer before the first line is written
No	Yes	A file named 'JOAN' in the catalog
No	No	A file on a reel of tape, not known to the processor

Means are provided to create, delete, connect, and disconnect files.

C.6.1.3 File names (9.3.4.1)

A file may have a name. The form of a file name is not specified. If a system does not have some form of cataloging or tape labeling for at least some of its files, all file names will disappear at the termination of execution. This is a valid implementation. Nowhere does this standard require names to survive for any period of time longer than the execution time span of a program. Therefore, this standard does not impose cataloging as a prerequisite. The naming feature is intended to allow use of a cataloging system where one exists.

C.6.1.4 File access (9.2.1.2)

This standard does not address problems of security, protection, locking, and many other concepts that may be part of the concept of "right of access". Such concepts are considered to be in the province of an operating system.

The OPEN and INQUIRE statements can be extended naturally to consider these things.

Possible access methods for a file are: sequential and direct. The processor may implement two different types of files, each with its own access method. It might also implement one type of file with two different access methods.

Direct access to files is of a simple and commonly available type, that is, fixed-length records. The key is a positive integer.

C.6.1.5 Nonadvancing input/output (9.2.1.3.1)

Data transfer statements affect the positioning of an external file. In FORTRAN 77, if no error or end-of-file condition exists, the file is positioned after the record just read or written and that record becomes the preceding record. This standard contains the record positioning ADVANCE= specifier in a data transfer statement that provides the capability of maintaining a position within the current record from one formatted data transfer statement to the next data transfer statement.

The value NO provides this capability. The value YES positions the file after the record just read or written. The default is YES.

The tab edit descriptor and the slash are still appropriate for use with this type of record access but the tab will not reposition before the left tab limit.

A BACKSPACE of a file that is currently positioned within a record causes the specified unit to be positioned before the current record.

If the last data transfer statement was WRITE and the file is currently positioned within a record, the file will be positioned implicitly after the current record before an ENDFILE record is written to the file, that is, a REWIND, BACKSPACE, or ENDFILE statement following a nonadvancing WRITE statement causes the file to be positioned at the end of the current output record before the endfile record is written to the file.

This standard provides a SIZE= specifier to be used with nonadvancing data transfer statements. The variable in the SIZE= specifier will contain the count of the number of characters that make up the sequence of values read by the data edit descriptors in this input statement.

The count is especially helpful if there is only one list item in the input list since it will contain the number of characters that were present for the item.

The EOR= specifier is provided to indicate when an end-of-record condition has been encountered during a nonadvancing data transfer statement. The end-of-record condition is not an error condition. If this specifier is present, the current input list item that required more characters than the record contained will be padded with blanks if PAD= 'YES' is in effect. This means that the iolist item was successfully completed. The file will then be positioned after the current record. The IOSTAT= specifier, if present, will be defined with a processor-dependent negative value and the data transfer statement will be terminated. Program execution will continue with the statement specified in the EOR= specifier. The EOR= specifier gives the capability of taking control of execution when the end-of-record has been found. Implied-DO variables retain their last defined value and any remaining items in the *iolist* retain their definition status when an end-of-record condition occurs. The SIZE= specifier, if present, will contain the number of characters read with the data edit descriptors during this READ statement.

For nonadvancing input, the processor is not required to read partial records. The processor may read the entire record into an internal buffer and make successive portions of the record available to successive input statements.

In an implementation of nonadvancing I/O in which a nonadvancing write to a terminal device causes immediate display of the output, such a write can be employed as a mechanism to output a prompt. In this case, the statement

```
WRITE (*, FMT='(A)', ADVANCE='NO') 'CONTINUE?(Y/N): '
```

would result in the prompt

```
CONTINUE?(Y/N):
```

being displayed with no subsequent line feed.

The response, which might be read by a statement of the form

```
READ (*, FMT='(A)') ANSWER
```

can then be entered on the same line as the prompt as in

```
CONTINUE?(Y/N): Y
```

The standard does not require that an implementation of nonadvancing I/O operate in this manner. For example, an implementation of nonadvancing output in which the display of the output is deferred until the current record is complete is also standard conforming. Such an implementation will not, however, allow a prompting mechanism of this kind to operate.

C.6.2 OPEN statement (9.3.4)

A file may become connected to a unit in either of two ways: preconnection or execution of an OPEN statement. Preconnection is performed prior to the beginning of execution of a program by means external to Fortran. For example, it may be done by job control action or by processor-established defaults. Execution of an OPEN statement is not required to access preconnected files (9.3.3).

The OPEN statement provides a means to access existing files that are not preconnected. An OPEN statement may be used in either of two ways: with a file name (open-by-name) and without a file name (open-by-unit). A unit is given in either case. Open-by-name connects the specified file to the specified unit. Open-by-unit connects a processor-determined default file to the specified unit. (The default file may or may not have a name.)

Therefore, there are three ways a file may become connected and hence processed: preconnection, open-by-name, and open-by-unit. Once a file is connected, there is no means in standard Fortran to determine how it became connected.

An OPEN statement may also be used to create a new file. In fact, any of the foregoing three connection methods may be performed on a file that does not exist. When a unit is preconnected, writing the first record creates the file. With the other two methods, execution of the OPEN statement creates the file.

When an OPEN statement is executed, the unit specified in the OPEN may or may not already be connected to a file. If it is already connected to a file (either through preconnection or by a prior OPEN), then omitting the FILE= specifier in the OPEN statement implies that the file is to remain connected to the unit. Such an OPEN statement may be used to change the values of the BLANK=, DELIM=, or PAD= specifiers.

Note that, since an OPEN that specifies STATUS = 'SCRATCH' is not allowed to have a FILE= specifier, such an OPEN always attempts to retain any connection that the specified unit may have. If the unit were already connected to a file, and if that connection did not have a STATUS of SCRATCH, then the OPEN would be illegal because the value of the STATUS= specifier shall not be changed by the OPEN.

If the value of the ACTION= specifier is WRITE, then READ statements shall not refer to this connection. ACTION = 'WRITE' does not restrict positioning by a BACKSPACE statement or positioning specified by the POSITION= specifier with the value APPEND. However, a BACKSPACE statement or an OPEN statement containing POSITION = 'APPEND' may fail if the processor requires reading of the file to achieve the positioning.

The following examples illustrate these rules. In the first example, unit 10 is preconnected to a SCRATCH file; the OPEN statement changes the value of PAD= to YES.

```
CHARACTER (LEN = 20) CH1
WRITE (10, '(A)') 'THIS IS RECORD 1'
OPEN (UNIT = 10, STATUS = 'SCRATCH', PAD = 'YES')
REWIND 10
READ (10, '(A20)') CH1   ! CH1 now has the value
                         ! 'THIS IS RECORD 1    '
```

In the next example, unit 12 is first connected to a file named FRED, with a status of OLD. The second OPEN statement then opens unit 12 again, retaining the connection to the file FRED, but changing the value of the DELIM= specifier to QUOTE.

```
CHARACTER (LEN = 25) CH2, CH3
OPEN (12, FILE = 'FRED', STATUS = 'OLD', DELIM = 'NONE')
CH2 = '"THIS STRING HAS QUOTES."'
              ! Quotes in string CH2
WRITE (12, *) CH2          ! Written with no delimiters
```

```
OPEN (12, DELIM = 'QUOTE')  ! Now quote is the delimiter
REWIND 12
READ (12, *) CH3  ! CH3 now has the value
                  ! 'THIS STRING HAS QUOTES.  '
```

The next example is invalid because it attempts to change the value of the STATUS= specifier.

```
OPEN (10, FILE = 'FRED', STATUS = 'OLD')
WRITE (10, *) A, B, C
OPEN (10, STATUS = 'SCRATCH')   ! Attempts to make FRED
                                ! a SCRATCH file
```

The previous example could be made valid by closing the unit first, as in the next example.

```
OPEN (10, FILE = 'FRED', STATUS = 'OLD')
WRITE (10, *) A, B, C
CLOSE (10)
OPEN (10, STATUS = 'SCRATCH')   ! Opens a different
                                ! SCRATCH file
```

C.6.3 Connection properties (9.3.2)

When a unit becomes connected to a file, either by execution of an OPEN statement or by preconnection, the following connection properties may be established:

(1) An access method, which is sequential or direct, is established for the connection (9.3.4.3).

(2) A form, which is formatted or unformatted, is established for a connection to a file that exists or is created by the connection. For a connection that results from execution of an OPEN statement, a default form (which depends on the access method, as described in 9.2.1.2) is established if no form is specified. For a preconnected file that exists, a form is established by preconnection. For a preconnected file that does not exist, a form may be established, or the establishment of a form may be delayed until the file is created (for example, by execution of a formatted or unformatted WRITE statement) (9.3.4.4).

(3) A record length may be established. If the access method is direct, the connection establishes a record length that specifies the length of each record of the file. An existing file with records that are not all of equal length shall not be connected for direct access.

If the access method is sequential, records of varying lengths are permitted. In this case, the record length established specifies the maximum length of a record in the file (9.3.4.5).

(4) A blank significance property, which is ZERO or NULL, is established for a connection for which the form is formatted. This property has no effect on output. For a connection that results from execution of an OPEN statement, the blank significance property is NULL by default if no blank significance property is specified. For a preconnected file, the property is NULL. The blank significance property of the connection is effective at the beginning of each formatted input statement. During execution of the statement, any BN or BZ edit descriptors encountered may temporarily change the effect of embedded and trailing blanks (9.3.4.6).

FORTRAN 77 did not define default values for the blank significance properties of internal and preconnected files. This standard defines the default values for these files to be NULL, matching that of files connected by the OPEN statement.

A processor has wide latitude in adapting these concepts and actions to its own cataloging and job control conventions. Some processors may require job control action to specify the set of files that exist or that will be created by a program. Some processors may require no job control action prior to execution. This standard enables processors to perform dynamic open, close, or file creation operations, but it does not require such capabilities of the processor.

The meaning of "open" in contexts other than Fortran may include such things as mounting a tape, console messages, spooling, label checking, security checking, etc. These actions may occur upon job control action external to Fortran, upon execution of an OPEN statement, or upon execution of the first read or write of the file. The OPEN statement describes properties of the connection to the file and may or may not cause physical activities to take place. It is a place for an implementation to define properties of a file beyond those required in standard Fortran.

C.6.4 CLOSE statement (9.3.5)

Similarly, the actions of dismounting a tape, protection, etc. of a "close" may be implicit at the end of a run. The CLOSE statement may or may not cause such actions to occur. This is another place to extend file properties beyond those of standard Fortran. Note, however, that the execution of a CLOSE statement on a unit followed by an OPEN statement on the same unit to the same file or to a different file is a permissible sequence of events. The processor shall not deny this sequence solely because the implementation chooses to do the physical act of closing the file at the termination of execution of the program.

C.6.5 INQUIRE statement (9.6)

Table C.1 indicates the values assigned to the INQUIRE statement specifier variables when no error condition is encountered during execution of the INQUIRE statement.

Table C.1 Values assigned to INQUIRE specifier variables

Specifier	INQUIRE by file		INQUIRE by unit	
	Unconnected	Connected	Connected	Unconnected
ACCESS=	UNDEFINED	SEQUENTIAL or DIRECT		UNDEFINED
ACTION=	UNDEFINED	READ, WRITE, or READWRITE		UNDEFINED
BLANK=	UNDEFINED	NULL, ZERO, or UNDEFINED		UNDEFINED
DELIM=	UNDEFINED	APOSTROPHE, QUOTE, NONE, or UNDEFINED		UNDEFINED
DIRECT=	UNKNOWN	YES, NO, or UNKNOWN		UNKNOWN
EXIST=	.TRUE. if file exists, .FALSE. otherwise		.TRUE. if unit exists, .FALSE. otherwise	
FORM=	UNDEFINED	FORMATTED or UNFORMATTED		UNDEFINED
FORMATTED=	UNKNOWN	YES, NO, or UNKNOWN		UNKNOWN
IOLENGTH=	RECL= value for *output-item-list*			
IOSTAT=	0 for no error, a positive integer for an error			
NAME=	Filename (may not be same as FILE= value)		Filename if named, else undefined	UNDEFINED
NAMED=	.TRUE.		.TRUE. if file named, .FALSE. otherwise	.FALSE.
NEXTREC=	Undefined	If direct access, next record #; else undefined		Undefined
NUMBER=	–1	Unit number		–1
OPENED=	.FALSE.	.TRUE.		.FALSE.
PAD=	YES	YES or NO		YES
POSITION=	UNDEFINED	REWIND, APPEND, ASIS, UNDEFINED, or a processor-dependent value		UNDEFINED
READ=	UNKNOWN	YES, NO, or UNKNOWN		UNKNOWN
READWRITE=	UNKNOWN	YES, NO, or UNKNOWN		UNKNOWN
RECL=	Undefined	If direct access, record length; else maximum record length		Undefined
SEQUENTIAL=	UNKNOWN	YES, NO, or UNKNOWN		UNKNOWN
UNFORMATTED=	UNKNOWN	YES, NO, or UNKNOWN		UNKNOWN
WRITE=	UNKNOWN	YES, NO, or UNKNOWN		UNKNOWN

C.7 Section 10 notes

C.7.1 Number of records (10.3, 10.4, 10.6.2)

The number of records read by an explicitly formatted advancing input statement can be determined from the following rule: a record is read at the beginning of the format scan (even if the input list is empty), at each slash edit descriptor encountered in the format, and when a format rescan occurs at the end of the format.

The number of records written by an explicitly formatted advancing output statement can be determined from the following rule: a record is written when a slash edit descriptor is encountered in the format, when a format rescan occurs at the end of the format, and at completion of execution of the output statement (even if the output list is empty). Thus, the occurrence of n successive slashes between two other edit descriptors causes $n - 1$ blank lines if the records are printed. The occurrence of n slashes at the beginning or end of a complete format specification causes n blank lines if the records are printed. However, a complete format specification containing n slashes ($n > 0$) and no other edit descriptors causes $n + 1$ blank lines if the records are printed. For example, the statements

```
  PRINT 3
3 FORMAT (/)
```

will write two records that cause two blank lines if the records are printed.

C.7.2 List-directed input (10.8.1)

The following examples illustrate list-directed input. A blank character is represented by b.

Example 1:

Program:

```
J = 3
READ *, I
READ *, J
```

Sequential input file:

```
record 1:  b1b,4bbbbb
record 2:  ,2bbbbbbbb
```

Result: I = 1, J = 3.

Explanation: The second READ statement reads the second record. The initial comma in the record designates a null value; therefore, J is not redefined.

Example 2:

Program:

```
CHARACTER A *8, B *1
READ *, A, B
```

Sequential input file:

```
record 1:  'bbbbbbbb'
record 2:  'QXY'b'Z'
```

Result: A = 'bbbbbbbb', B = 'Q'

Explanation: In the first record, the rightmost apostrophe is interpreted as delimiting the constant (it cannot be the first of a pair of embedded apostrophes representing a single apostrophe because this would involve the prohibited "splitting" of the pair by the end of a record); therefore, A is assigned the character constant 'bbbbbbbb'. The end of a record acts as a blank, which in this case is a value separator because it occurs between two constants.

C.8 Section 11 notes

C.8.1 Main program and block data program unit (11.1, 11.4)

The name of the main program or of a block data program unit has no explicit use within the Fortran language. It is available for documentation and for possible use by a processor.

A processor may implement an unnamed main program or unnamed block data program unit by assigning it a default name. However, this name shall not conflict with any other global name in a standard-conforming program. This might be done by making the default name one which is not permitted in a standard-conforming program (for example, by including a character not normally allowed in names) or by providing some external mechanism such that for any given program the default name can be changed to one that is otherwise unused.

C.8.2 Dependent compilation (11.3)

This standard, like its predecessors, is intended to permit the implementation of conforming processors in which a program can be broken into multiple units, each of which can be separately translated in preparation for execution. Such processors are commonly described as supporting separate compilation. There is an important difference between the way separate compilation can be implemented under this standard and the way it could be implemented under the previous standards. Under the previous standards, any information required to translate a program unit was specified in that program unit. Each translation was thus totally independent of all others. Under this standard, a program unit can use information that was specified in a separate module and thus may be dependent on that module. The implementation of this dependency in a processor may be that the translation of a program unit may depend on the results of translating one or more modules. Processors implementing the dependency this way are commonly described as supporting dependent compilation.

The dependencies involved here are new only in the sense that the Fortran processor is now aware of them. The same information dependencies existed under the previous standards, but it was the programmer's responsibility to transport the information necessary to resolve them by making redundant specifications of the information in multiple program units. The availability of separate but dependent compilation offers several potential advantages over the redundant textual specification of information:

(1) Specifying information at a single place in the program ensures that different program units using that information will be translated consistently. Redundant specification leaves the possibility that different information will erroneously be specified. Even if some kind of textual inclusion facility is used to ensure that the text of the specifications is identical in all involved program units, the presence of other specifications (for example, an IMPLICIT statement) may change the interpretation of that text.

(2) During the revision of a program, it is possible for a processor to assist in determining whether different program units have been translated using different (incompatible) versions of a module, although there is no requirement that a processor provide such assistance. Inconsistencies in redundant textual specification of information, on the other hand, tend to be much more difficult to detect.

(3) Putting information in a module provides a way of packaging it. Without modules, redundant specifications frequently shall be interleaved with other specifications in a program unit, making convenient packaging of such information difficult.

(4) Because a processor may be implemented such that the specifications in a module are translated once and then repeatedly referenced, there is the potential for greater efficiency than when the processor shall translate redundant specifications of information in multiple program units.

The exact meaning of the requirement that the public portions of a module be available at the time of reference is processor dependent. For example, a processor could consider a module to be available only after it has been compiled and require that if the module has been compiled separately, the result of that compilation shall be identified to the compiler when compiling program units that use it.

C.8.2.1 USE statement and dependent compilation (11.3.2)

Another benefit of the USE statement is its enhanced facilities for name management. If one needs to use only selected entities in a module, one can do so without having to worry about the names of all the other entities in that module. If one needs to use two different modules that happen to contain entities with the same name, there are several ways to deal with the conflict. If none of the entities with the same name are to be used, they can simply be ignored. If the name happens to refer to the same entity in both modules (for example, if both modules obtained it from a third module), then there is no confusion about what the name denotes and the name can be freely used. If the entities are different and one or both is to be used, the local renaming facility in the USE statement makes it possible to give those entities different names in the program unit containing the USE statements.

A typical implementation of dependent but separate compilation may involve storing the result of translating a module in a file (or file element) whose name is derived from the name of the module. Note, however, that the name of a module is limited only by the Fortran rules and not by the names allowed in the file system. Thus the processor may have to provide a mapping between Fortran names and file system names.

The result of translating a module could reasonably either contain only the information textually specified in the module (with "pointers" to information originally textually specified in other modules) or contain all information specified in the module (including copies of information originally specified in other modules). Although the former approach would appear to save on storage space, the latter approach can greatly simplify the logic necessary to process a USE statement and can avoid the necessity of imposing a limit on the logical "nesting" of modules via the USE statement.

Variables declared in a module retain their definition status on much the same basis as variables in a common block. That is, saved variables retain their definition status throughout the execution of a program, while variables that are not saved retain their definition status only during the execution of scoping units that reference the module. In some cases, it may be appropriate to put a USE statement such as

```
USE MY_MODULE, ONLY:
```

in a scoping unit in order to assure that other procedures that it references can communicate through the module. In such a case, the scoping unit would not access any entities from the module, but the variables not saved in the module would retain their definition status throughout the execution of the scoping unit.

There is an increased potential for undetected errors in a scoping unit that uses both implicit typing and the USE statement. For example, in the program fragment

```
SUBROUTINE SUB
   USE MY_MODULE
   IMPLICIT INTEGER (I-N), REAL (A-H, O-Z)
   X = F (B)
   A = G (X) + H (X + 1)
END SUBROUTINE SUB
```

X could be either an implicitly typed real variable or a variable obtained from the module MY_MODULE and might change from one to the other because of changes in MY_MODULE unrelated to the action performed by SUB. Logic errors resulting from this kind of situation can be extremely difficult to locate. Thus, the use of these features together is discouraged.

C.8.2.2 Accessibility attributes (11.3.1)

The PUBLIC and PRIVATE attributes, which can be declared only in modules, divide the entities in a module into those which are actually relevant to a scoping unit referencing the module and those that are not. This information may be used to improve the performance of a Fortran processor. For example, it may be possible to discard much of the information on the private entities once a module has been translated, thus saving on both storage and the time to search it. Similarly, it may be possible to recognize that two versions of a module differ only in the private entities they contain and avoid retranslating program units that use that module when switching from one version of the module to the other.

C.8.3 Examples of the use of modules

C.8.3.1 Identical common blocks

A common block and all its associated specification statements may be placed in a module named, for example, MY_COMMON and accessed by a USE statement of the form

```
USE MY_COMMON
```

that accesses the whole module without any renaming. This ensures that all instances of the common block are identical. Module MY_COMMON could contain more than one common block.

C.8.3.2 Global data

A module may contain only data objects, for example:

```
MODULE DATA_MODULE
   SAVE
   REAL A (10), B, C (20,20)
   INTEGER :: I=0
   INTEGER, PARAMETER :: J=10
   COMPLEX D (J,J)
END MODULE DATA_MODULE
```

Data objects made global in this manner may have any combination of data types.

Access to some of these may be made by a USE statement with the ONLY option, such as:

```
USE DATA_MODULE, ONLY: A, B, D
```

and access to all of them may be made by the following USE statement:

```
USE DATA_MODULE
```

Access to all of them with some renaming to avoid name conflicts may be made by:

```
USE DATA_MODULE, AMODULE => A, DMODULE => D
```

C.8.3.3 Derived types

A derived type may be defined in a module and accessed in a number of program units. For example:

```
MODULE SPARSE
   TYPE NONZERO
      REAL A
      INTEGER  I, J
   END TYPE NONZERO
```

```
END MODULE SPARSE
```

defines a type consisting of a real component and two integer components for holding the numerical value of a nonzero matrix element and its row and column indices.

C.8.3.4 Global allocatable arrays

Many programs need large global allocatable arrays whose sizes are not known before program execution. A simple form for such a program is:

```
PROGRAM GLOBAL_WORK
   CALL CONFIGURE_ARRAYS      ! Perform the appropriate allocations
   CALL COMPUTE               ! Use the arrays in computations
END PROGRAM GLOBAL_WORK
MODULE WORK_ARRAYS            ! An example set of work arrays
   INTEGER  N
   REAL, ALLOCATABLE, SAVE :: A (:), B (:, :), C (:, :, :)
END MODULE WORK_ARRAYS
SUBROUTINE CONFIGURE_ARRAYS   ! Process to set up work arrays
   USE WORK_ARRAYS
   READ (*, *)  N
   ALLOCATE (A (N), B (N, N), C (N, N, 2 * N))
END SUBROUTINE CONFIGURE_ARRAYS
SUBROUTINE COMPUTE
   USE WORK_ARRAYS
   ...  !  Computations involving arrays A, B, and C
END SUBROUTINE COMPUTE
```

Typically, many subprograms need access to the work arrays, and all such subprograms would contain the statement

```
USE WORK_ARRAYS
```

C.8.3.5 Procedure libraries

Interface blocks for external procedures in a library may be gathered into a module. This permits the use of argument keywords and optional arguments, and allows static checking of the references. Different versions may be constructed for different applications, using argument keywords in common use in each application.

An example is the following library module:

```
MODULE LIBRARY_LLS
   INTERFACE
      SUBROUTINE LLS (X, A, F, FLAG)
         REAL X (:, :)
         ! The SIZE in the next statement is an intrinsic function
         REAL, DIMENSION (SIZE (X, 2)) :: A, F
         INTEGER FLAG
      END SUBROUTINE LLS
         ...
   END INTERFACE
      ...
END MODULE LIBRARY_LLS
```

This module allows the subroutine LLS to be invoked:

```
USE LIBRARY_LLS
   ...
CALL LLS (X = ABC, A = D, F = XX, FLAG = IFLAG)
   ...
```

C.8.3.6 Operator extensions

In order to extend an intrinsic operator symbol to have an additional meaning, an interface block specifying that operator symbol in the OPERATOR option of the INTERFACE statement may be placed in a module.

For example, // may be extended to perform concatenation of two derived-type objects serving as varying length character strings and + may be extended to specify matrix addition for type MATRIX or interval arithmetic addition for type INTERVAL.

A module might contain several such interface blocks. An operator may be defined by an external function (either in Fortran or some other language) and its procedure interface placed in the module.

C.8.3.7 Data abstraction

In addition to providing a portable means of avoiding the redundant specification of information in multiple program units, a module provides a convenient means of "packaging" related entities, such as the definitions of the representation and operations of an abstract data type. The following example of a module defines a data abstraction for a SET data type where the elements of each set are of type integer. The standard set operations of UNION, INTERSECTION, and DIFFERENCE are provided. The CARDINALITY function returns the cardinality of (number of elements in) its set argument. Two functions returning logical values are included, ELEMENT and SUBSET. ELEMENT defines the operator `.IN.` and SUBSET extends the operator <=. ELEMENT determines if a given scalar integer value is an element of a given set, and SUBSET determines if a given set is a subset of another given set. (Two sets may be checked for equality by comparing cardinality and checking that one is a subset of the other, or checking to see if each is a subset of the other.)

The transfer function SETF converts a vector of integer values to the corresponding set, with duplicate values removed. Thus, a vector of constant values can be used as set constants. An inverse transfer function VECTOR returns the elements of a set as a vector of values in ascending order. In this SET implementation, set data objects have a maximum cardinality of 200.

```
MODULE INTEGER_SETS
! This module is intended to illustrate use of the module facility
! to define a new data type, along with suitable operators.

INTEGER, PARAMETER :: MAX_SET_CARD = 200

TYPE SET                              ! Define SET data type
   PRIVATE
   INTEGER CARD
   INTEGER ELEMENT (MAX_SET_CARD)
END TYPE SET

INTERFACE OPERATOR (.IN.)
   MODULE PROCEDURE ELEMENT
END INTERFACE OPERATOR (.IN.)

INTERFACE OPERATOR (<=)
   MODULE PROCEDURE SUBSET
END INTERFACE OPERATOR (<=)

INTERFACE OPERATOR (+)
   MODULE PROCEDURE UNION
END INTERFACE OPERATOR (+)

INTERFACE OPERATOR (-)
   MODULE PROCEDURE DIFFERENCE
```

```
END INTERFACE OPERATOR (-)

INTERFACE OPERATOR (*)
   MODULE PROCEDURE INTERSECTION
END INTERFACE OPERATOR (*)

CONTAINS

INTEGER FUNCTION CARDINALITY (A)   ! Returns cardinality of set A
   TYPE (SET), INTENT (IN) :: A
   CARDINALITY = A % CARD
END FUNCTION CARDINALITY

LOGICAL FUNCTION ELEMENT (X, A)          ! Determines if
   INTEGER, INTENT(IN) :: X              ! element X is in set A
   TYPE (SET), INTENT(IN) :: A
   ELEMENT = ANY (A % ELEMENT (1 : A % CARD) .EQ. X)
END FUNCTION ELEMENT

FUNCTION UNION (A, B)                    ! Union of sets A and B
   TYPE (SET) UNION
   TYPE (SET), INTENT(IN) :: A, B
   INTEGER J
   UNION = A
   DO J = 1, B % CARD
      IF (.NOT. (B % ELEMENT (J) .IN. A)) THEN
         IF (UNION % CARD < MAX_SET_CARD) THEN
            UNION % CARD = UNION % CARD + 1
            UNION % ELEMENT (UNION % CARD) = &
                  B % ELEMENT (J)
         ELSE
            ! Maximum set size exceeded . . .
         END IF
      END IF
   END DO
END FUNCTION UNION

FUNCTION DIFFERENCE (A, B)          ! Difference of sets A and B
   TYPE (SET) DIFFERENCE
   TYPE (SET), INTENT(IN) :: A, B
   INTEGER J, X
   DIFFERENCE % CARD = 0            ! The empty set
   DO J = 1, A % CARD
      X = A % ELEMENT (J)
      IF (.NOT. (X .IN. B)) DIFFERENCE = DIFFERENCE + SET (1, X)
   END DO
END FUNCTION DIFFERENCE

FUNCTION INTERSECTION (A, B)       ! Intersection of sets A and B
   TYPE (SET) INTERSECTION
   TYPE (SET), INTENT(IN) :: A, B
   INTERSECTION = A - (A - B)
END FUNCTION INTERSECTION

LOGICAL FUNCTION SUBSET (A, B)          ! Determines if set A is
   TYPE (SET), INTENT(IN) :: A, B       ! a subset of set B
   INTEGER I
   SUBSET = A % CARD <= B % CARD
   IF (.NOT. SUBSET) RETURN             ! For efficiency
```

```
   DO I = 1, A % CARD
      SUBSET = SUBSET .AND. (A % ELEMENT (I) .IN. B)
   END DO
END FUNCTION SUBSET

TYPE (SET) FUNCTION SETF (V)      ! Transfer function between a vector
   INTEGER V (:)                  ! of elements and a set of elements
   INTEGER J                      ! removing duplicate elements
   SETF % CARD = 0
   DO J = 1, SIZE (V)
      IF (.NOT. (V (J) .IN. SETF)) THEN
         IF (SETF % CARD < MAX_SET_CARD) THEN
            SETF % CARD = SETF % CARD + 1
            SETF % ELEMENT (SETF % CARD) = V (J)
         ELSE
            ! Maximum set size exceeded . . .
         END IF
      END IF
   END DO
END FUNCTION SETF

FUNCTION VECTOR (A)               ! Transfer the values of set A
   TYPE (SET), INTENT (IN) :: A   ! into a vector in ascending order
   INTEGER, POINTER :: VECTOR (:)
   INTEGER I, J, K
   ALLOCATE (VECTOR (A % CARD))
   VECTOR = A % ELEMENT (1 : A % CARD)
   DO I = 1, A % CARD - 1          ! Use a better sort if
      DO J = I + 1, A % CARD       ! A % CARD is large
         IF (VECTOR (I) > VECTOR (J)) THEN
            K = VECTOR (J); VECTOR (J) = VECTOR (I); VECTOR (I) = K
         END IF
      END DO
   END DO
END FUNCTION VECTOR

END MODULE INTEGER_SETS
```

Examples of using INTEGER_SETS (A, B, and C are variables of type SET; X is an integer variable):

```
! Check to see if A has more than 10 elements
IF (CARDINALITY (A) > 10) ...

! Check for X an element of A but not of B
IF (X .IN. (A - B)) ...

! C is the union of A and the result of B intersected
! with the integers 1 to 100
C = A + B * SETF ((/ (I, I = 1, 100) /))

! Does A have any even numbers in the range 1:100?
IF (CARDINALITY (A * SETF ((/ (I, I = 2, 100, 2) /))) > 0) ...

PRINT *, VECTOR (B) ! Print out the elements of set B, in ascending order
```

C.8.3.8 Public entities renamed

At times it may be necessary to rename entities that are accessed with USE statements. Care should be taken if the referenced modules also contain USE statements.

The following example illustrates renaming features of the USE statement.

```
MODULE J; REAL JX, JY, JZ; END MODULE J
MODULE K
   USE J, ONLY : KX => JX, KY => JY
   ! KX and KY are local names to module K
   REAL KZ       ! KZ is local name to module K
   REAL JZ       ! JZ is local name to module K
END MODULE K
PROGRAM RENAME
   USE J; USE K
   ! Module J's entity JX is accessible under names JX and KX
   ! Module J's entity JY is accessible under names JY and KY
   ! Module K's entity KZ is accessible under name KZ
   ! Module J's entity JZ and K's entity JZ are different entities
   ! and shall not be referenced
   ...
END PROGRAM RENAME
```

C.9 Section 12 notes

C.9.1 Portability problems with external procedures (12.3.2.2)

There is a potential portability problem in a scoping unit that references an external procedure without declaring it in either an EXTERNAL statement or a procedure interface block. On a different processor, the name of that procedure may be the name of a nonstandard intrinsic procedure and the processor would be permitted to interpret those procedure references as references to that intrinsic procedure. (On that processor, the program would also be viewed as not conforming to the standard because of the references to the nonstandard intrinsic procedure.) Declaration in an EXTERNAL statement or a procedure interface block causes the references to be to the external procedure regardless of the availability of an intrinsic procedure with the same name. Note that declaration of the type of a procedure is not enough to make it external, even if the type is inconsistent with the type of the result of an intrinsic of the same name.

C.9.2 Procedures defined by means other than Fortran (12.5.3)

A processor is not required to provide any means other than Fortran for defining external procedures. Among the means that might be supported are the machine assembly language, other high level languages, the Fortran language extended with nonstandard features, and the Fortran language as supported by another Fortran processor (for example, a previously existing FORTRAN 77 processor).

Procedures defined by means other than Fortran are considered external procedures because their definitions are not in a Fortran program unit and because they are referenced using global names. The use of the term external should not be construed as any kind of restriction on the way in which these procedures may be defined. For example, if the means other than Fortran has its own facilities for internal and external procedures, it is permissible to use them. If the means other than Fortran can create an "internal" procedure with a global name, it is permissible for such an "internal" procedure to be considered by Fortran to be an external procedure. The means other than Fortran for defining external procedures, including any restrictions on the structure for organization of those procedures, are entirely processor dependent.

A Fortran processor may limit its support of procedures defined by means other than Fortran such that these procedures may affect entities in the Fortran environment only on the same basis as procedures written in Fortran. For example, it might prohibit the value of a local variable from being changed by a procedure reference unless that variable were one of the arguments to the procedure.

C.9.3 Procedure interfaces (12.3)

In FORTRAN 77, the interface to an external procedure was always deduced from the form of references to that procedure and any declarations of the procedure name in the referencing program unit. In this standard, features such as argument keywords and optional arguments make it impossible to deduce sufficient information about the dummy arguments from the nature of the actual arguments to be associated with them, and features such as array-valued function results and pointer function results make necessary extensions to the declaration of a procedure that cannot be done in a way that would be analogous with the handling of such declarations in FORTRAN 77. Hence, mechanisms are provided through which all the information about a procedure's interface may be made available in a scoping unit that references it. A procedure whose interface shall be deduced as in FORTRAN 77 is described as having an implicit interface. A procedure whose interface is fully known is described as having an explicit interface.

A scoping unit is allowed to contain a procedure interface block for procedures that do not exist in the program, provided the procedure described is never referenced. The purpose of this rule is to allow implementations in which the use of a module providing procedure interface blocks describing the interface of every routine in a library would not automatically cause each of those library routines to be a part of the program referencing the module. Instead, only those library procedures actually referenced would be a part of the program. (In implementation terms, the mere presence of a procedure interface block would not generate an external reference in such an implementation.)

C.9.4 Argument association and evaluation (12.4.1.1)

There is a significant difference between the argument association allowed in this standard and that supported by FORTRAN 77 and FORTRAN 66. In FORTRAN 77 and 66, actual arguments were limited to consecutive storage units. With the exception of assumed length character dummy arguments, the structure imposed on that sequence of storage units was always determined in the invoked procedure and not taken from the actual argument. Thus it was possible to implement FORTRAN 66 and FORTRAN 77 argument association by supplying only the location of the first storage unit (except for character arguments, where the length would also have to be supplied). However, this standard allows arguments that do not reside in consecutive storage locations (for example, an array section), and dummy arguments that assume additional structural information from the actual argument (for example, assumed-shape dummy arguments). Thus, the mechanism to implement the argument association allowed in this standard needs to be more general.

Because there are practical advantages to a processor that can support references to and from procedures defined by a FORTRAN 77 processor, requirements for explicit interfaces have been added to make it possible to determine whether a simple (FORTRAN 66/FORTRAN 77) argument association implementation mechanism is sufficient or whether the more general mechanism is necessary (12.3.1.1). Thus a processor can be implemented whose procedures expect the simple mechanism to be used whenever the procedure's interface is one which uses only FORTRAN 77 features and which expects the more general mechanism otherwise (for example, if there are assumed-shape or optional arguments). At the point of reference, the appropriate mechanism can be determined from the interface if it is explicit and can be assumed to be the simple mechanism if it is not. Note that if the simple mechanism is determined to be what the procedure expects, it may be necessary for the processor to allocate consecutive temporary storage for the actual argument, copy the actual argument to the temporary storage, reference the procedure using the temporary storage in place of the actual argument, copy the contents of temporary storage back to the actual argument, and deallocate the temporary storage.

Note that while this is the specific implementation method these rules were designed to support, it is not the only one possible. For example, on some processors, it may be possible to implement the general argument association in such a way that the information involved in FORTRAN 77 argument association may be found in the same places and the "extra" information is placed so it does not disturb a procedure expecting only FORTRAN 77 argument association. With such an implementation, argument association could be translated without regard to whether the interface is explicit or implicit. Alternatively, it would be possible to disallow discontiguous arguments when calling procedures defined by the FORTRAN 77 processor and let any copying to and from contiguous storage be done explicitly in the program. Yet another possibility would be not to allow references to procedures defined by a FORTRAN 77 processor.

The provisions for expression evaluation give the processor considerable flexibility for obtaining expression values in the most efficient way possible. This includes not evaluating or only partially evaluating an operand, for example, if the value of the expression can be determined otherwise (7.1.7.1). This flexibility applies to function argument evaluation, including the order of argument evaluation, delaying argument evaluation, and omitting argument evaluation. A processor may delay the evaluation of an argument in a procedure reference until the execution of the procedure refers to the value of that argument, provided delaying the evaluation of the argument does not otherwise affect the results of the program. The processor may, with similar restrictions, entirely omit the evaluation of an argument not referenced in the execution of the procedure. This gives processors latitude for optimization (for example, for parallel processing).

C.9.5 Pointers and targets as arguments (12.4.1.1)

If a dummy argument is declared to be a pointer, it may be matched only by an actual argument that also is a pointer, and the characteristics of both arguments shall agree. A model for such an association is that descriptor values of the actual pointer are copied to the dummy pointer. If the actual pointer has an associated target, this target becomes accessible via the dummy pointer. If the dummy pointer becomes associated with a different target during execution of the procedure, this target will be accessible via the actual pointer after the procedure completes execution. If the dummy pointer becomes associated with a local target that ceases to exist when the procedure completes, the actual pointer will be left dangling in an undefined state. Such dangling pointers shall not be used.

When execution of a procedure completes, any pointer that remains defined and that is associated with a dummy argument that has the TARGET attribute and is either a scalar or an assumed-shape array, remains associated with the corresponding actual argument if the actual argument has the TARGET attribute and is not an array section with a vector subscript.

```
REAL, POINTER      :: PBEST
REAL, TARGET       :: B (10000)
CALL BEST (PBEST, B)            ! Upon return PBEST is associated
   ...                          ! with the "best" element of B
CONTAINS
  SUBROUTINE BEST (P, A)
    REAL, POINTER     :: P
    REAL, TARGET      :: A (:)
      ...                       ! Find the "best" element A(I)
    P => A (I)
  RETURN
  END SUBROUTINE BEST
END
```

When procedure BEST completes, the pointer PBEST is associated with an element of B.

An actual argument without the TARGET attribute can become associated with a dummy argument with the TARGET attribute. This permits pointers to become associated with the dummy argument during execution of the procedure that contains the dummy argument. For example:

```
INTEGER LARGE(100,100)
CALL SUB (LARGE)
   ...
CALL SUB ()
CONTAINS
  SUBROUTINE SUB(ARG)
    INTEGER, TARGET, OPTIONAL :: ARG(100,100)
    INTEGER, POINTER, DIMENSION(:,:) :: PARG
    IF (PRESENT(ARG)) THEN
      PARG => ARG
    ELSE
      ALLOCATE (PARG(100,100))
      PARG = 0
    ENDIF
       ...  ! Code with lots of references to PARG
    IF (.NOT. PRESENT(ARG)) DEALLOCATE(PARG)
  END SUBROUTINE SUB
END
```

Within subroutine SUB the pointer PARG is either associated with the dummy argument ARG or it is associated with an allocated target. The bulk of the code can reference PARG without further calls to the PRESENT intrinsic.

C.10 Section 14 notes

C.10.1 Examples of host association (14.6.1.3)

The first two examples are examples of valid host association. The third example is an example of invalid host association.

Example 1:

```
PROGRAM A
   INTEGER I, J
   ...
CONTAINS
   SUBROUTINE B
      INTEGER I  ! Declaration of I hides
                 ! program A's declaration of I
         ...
      I = J      ! Use of variable J from program A
                 ! through host association
   END SUBROUTINE B
END PROGRAM A
```

Example 2:

```
PROGRAM A
   TYPE T
      ...
   END TYPE T
   ...
CONTAINS
   SUBROUTINE B
      IMPLICIT TYPE (T) (C)   ! Refers to type T declared below
                              ! in subroutine B, not type T
                              ! declared above in program A
```

```
         ...
      TYPE T
         ...
      END TYPE T
         ...
   END SUBROUTINE B
END PROGRAM A
```

Example 3:

```
PROGRAM Q
   REAL (KIND = 1) :: C
      ...
CONTAINS
   SUBROUTINE R
      REAL (KIND = KIND (C)) :: D  ! Invalid declaration
                                   ! See below
      REAL (KIND = 2) :: C
         ...
   END SUBROUTINE R
END PROGRAM Q
```

In the declaration of D in subroutine R, the use of C would refer to the declaration of C in subroutine R, not program Q. However, it is invalid because the declaration of C shall occur before it is used in the declaration of D (7.1.6.1).

C.11 Array feature notes

C.11.1 Summary of features

This section is a summary of the principal array features.

C.11.1.1 Whole array expressions and assignments (7.5.1.2, 7.5.1.5)

An important new feature is that whole array expressions and assignments are permitted. For example, the statement

```
A = B + C * SIN (D)
```

where A, B, C, and D are arrays of the same shape, is permitted. It is interpreted element-by-element; that is, the sine function is taken on each element of D, each result is multiplied by the corresponding element of C, added to the corresponding element of B, and assigned to the corresponding element of A. Functions, including user-written functions, may be array valued and may be generic with scalar versions. All arrays in an expression or across an assignment shall conform; that is, have exactly the same shape (number of dimensions and set of lengths in each dimension), but scalars may be included freely and these are interpreted as being broadcast to a conforming array. Expressions are evaluated before any assignment takes place.

C.11.1.2 Array sections (2.4.5, 6.2.2.3)

Whenever whole arrays may be used, it is also possible to use subarrays called "sections". For example:

```
A (:, 1:N, 2, 3:1:-1)
```

consists of a subarray containing the whole of the first dimension, positions 1 to N of the second dimension, position 2 of the third dimension and positions 1 to 3 in reverse order of the fourth dimension. This is an artificial example chosen to illustrate the different forms. Of course, a common use may be to select a row or column of an array, for example:

```
A (:, J)
```

C.11.1.3 WHERE statement (7.5.3)

The WHERE statement applies a conforming logical array as a mask on the individual operations in the expression and in the assignment. For example:

```
WHERE (A .GT. 0) B = LOG (A)
```

takes the logarithm only for positive components of A and makes assignments only in these positions.

The WHERE statement also has a block form (WHERE construct).

C.11.1.4 Automatic and allocatable arrays (5.1, 5.1.2.4.3)

A major advance for writing modular software is the presence of automatic arrays, created on entry to a subprogram and destroyed on return, and allocatable arrays whose rank is fixed but whose actual size and lifetime is fully under the programmer's control through explicit ALLOCATE and DEALLOCATE statements. The declarations

```
SUBROUTINE X (N, A, B)
   REAL WORK (N, N); REAL, ALLOCATABLE :: HEAP (:, :)
```

specify an automatic array WORK and an allocatable array HEAP. Note that a stack is an adequate storage mechanism for the implementation of automatic arrays, but a heap will be needed for allocatable arrays.

C.11.1.5 Array constructors (4.5)

Arrays, and in particular array constants, may be constructed with array constructors exemplified by:

```
(/ 1.0, 3.0, 7.2 /)
```

which is a rank-one array of size 3,

```
(/ (1.3, 2.7, L = 1, 10), 7.1 /)
```

which is a rank-one array of size 21 and contains the pair of real constants 1.3 and 2.7 repeated 10 times followed by 7.1, and

```
(/ (I, I = 1, N) /)
```

which contains the integers 1, 2, ..., N. Only rank-one arrays may be constructed in this way, but higher dimensional arrays may be made from them by means of the intrinsic function RESHAPE.

C.11.2 Examples

The array features have the potential to simplify the way that almost any array-using program is conceived and written. Many algorithms involving arrays can now be written conveniently as a series of computations with whole arrays.

C.11.2.1 Unconditional array computations

At the simplest level, statements such as

```
A = B + C
```

or

```
S = SUM (A)
```

can take the place of entire DO loops. The loops were required to perform array addition or to sum all the elements of an array.

Further examples of unconditional operations on arrays that are simple to write are:

matrix multiply	`P = MATMUL (Q, R)`
largest array element	`L = MAXVAL (P)`
factorial N	`F = PRODUCT ((/ (K, K = 2, N) /))`

The Fourier sum $F = \sum_{i=1}^{N} a_i \times \cos x_i$ may also be computed without writing a DO loop if one makes use of the element-by-element definition of array expressions as described in Section 7. Thus, we can write

```
F = SUM (A * COS (X))
```

The successive stages of calculation of F would then involve the arrays:

A = (/ A (1), ..., A (N) /)
X = (/ X (1), ..., X (N) /)
COS (X) = (/ COS (X (1)), ..., COS (X (N)) /)
A * COS (X) = (/ A (1) * COS (X (1)), ..., A (N) * COS (X (N)) /)

The final scalar result is obtained simply by summing the elements of the last of these arrays. Thus, the processor is dealing with arrays at every step of the calculation.

C.11.2.2 Conditional array computations

Suppose we wish to compute the Fourier sum in the above example, but to include only those terms *a(i) cos x(i)* that satisfy the condition that the coefficient *a(i)* is less than 0.01 in absolute value. More precisely, we are now interested in evaluating the conditional Fourier sum

$$CF = \sum_{|a_i| < 0.01} a_i \times \cos x_i$$

where the index runs from 1 to N as before.

This can be done by using the MASK parameter of the SUM function, which restricts the summation of the elements of the array A ∗ COS (X) to those elements that correspond to true elements of MASK. Clearly, the mask required is the logical array expression ABS (A) .LT. 0.01. Note that the stages of evaluation of this expression are:

A = (/ A (1), ..., A (N) /)
ABS (A) = (/ ABS (A (1)), ..., ABS (A (N)) /)
ABS (A) .LT. 0.01 = (/ ABS (A (1)) .LT. 0.01, ..., ABS (A (N)) .LT. 0.01 /)

The conditional Fourier sum we arrive at is:

```
CF = SUM (A * COS (X), MASK = ABS (A) .LT. 0.01)
```

If the mask is all false, the value of CF is zero.

The use of a mask to define a subset of an array is crucial to the action of the WHERE statement. Thus for example, to zero an entire array, we may write simply A = 0; but to set only the negative elements to zero, we need to write the conditional assignment

```
WHERE (A .LT. 0)  A = 0
```

The WHERE statement complements ordinary array assignment by providing array assignment to any subset of an array that can be restricted by a logical expression.

In the Ising model described below, the WHERE statement predominates in use over the ordinary array assignment statement.

C.11.2.3 A simple program: the Ising model

The Ising model is a well-known Monte Carlo simulation in 3-dimensional Euclidean space which is useful in certain physical studies. We will consider in some detail how this might be programmed. The model may be described in terms of a logical array of shape N by N by N. Each gridpoint is a single logical variable which is to be interpreted as either an up-spin (true) or a down-spin (false).

The Ising model operates by passing through many successive states. The transition to the next state is governed by a local probabilistic process. At each transition, all gridpoints change state simultaneously. Every spin either flips to its opposite state or not according to a rule that depends only on the states of its 6 nearest neighbors in the surrounding grid. The neighbors of gridpoints on the boundary faces of the model cube are defined by assuming cubic periodicity. In effect, this extends the grid periodically by replicating it in all directions throughout space.

The rule states that a spin is flipped to its opposite parity for certain gridpoints where a mere 3 or fewer of the 6 nearest neighbors currently have the same parity as it does. Also, the flip is executed only with probability P (4), P (5), or P (6) if as many as 4, 5, or 6 of them have the same parity as it does. (The rule seems to promote neighborhood alignments that may presumably lead to equilibrium in the long run.)

C.11.2.3.1 Problems to be solved

Some of the programming problems that we will need to solve in order to translate the Ising model into Fortran statements using entire arrays are

(1) Counting nearest neighbors that have the same spin;

(2) Providing an array-valued function to return an array of random numbers; and

(3) Determining which gridpoints are to be flipped.

C.11.2.3.2 Solutions in Fortran

The arrays needed are:

```
LOGICAL ISING (N, N, N), FLIPS (N, N, N)
INTEGER ONES (N, N, N), COUNT (N, N, N)
REAL THRESHOLD (N, N, N)
```

The array-valued function needed is:

```
FUNCTION RAND (N)
REAL RAND (N, N, N)
```

The transition probabilities are specified in the array

```
REAL P (6)
```

The first task is to count the number of nearest neighbors of each gridpoint *g* that have the same spin as *g*.

Assuming that ISING is given to us, the statements

```
ONES = 0
WHERE (ISING) ONES = 1
```

make the array ONES into an exact analog of ISING in which 1 stands for an up-spin and 0 for a down-spin.

The next array we construct, COUNT, will record for every gridpoint of ISING the number of spins to be found among the 6 nearest neighbors of that gridpoint. COUNT will be computed by adding together 6 arrays, one for each of the 6 relative positions in which a nearest neighbor is found. Each of the 6 arrays is obtained from the ONES array by shifting the ONES array one place circularly along one of its dimensions. This use of circular shifting imparts the cubic periodicity.

```
COUNT = CSHIFT (ONES, SHIFT = -1, DIM = 1)  &
      + CSHIFT (ONES, SHIFT =  1, DIM = 1)  &
      + CSHIFT (ONES, SHIFT = -1, DIM = 2)  &
      + CSHIFT (ONES, SHIFT =  1, DIM = 2)  &
      + CSHIFT (ONES, SHIFT = -1, DIM = 3)  &
      + CSHIFT (ONES, SHIFT =  1, DIM = 3)
```

At this point, COUNT contains the count of nearest neighbor up-spins even at the gridpoints where the Ising model has a down-spin. But we want a count of down-spins at those gridpoints, so we correct COUNT at the down (false) points of ISING by writing:

```
WHERE (.NOT. ISING)  COUNT = 6 - COUNT
```

Our object now is to use these counts of what may be called the "like-minded nearest neighbors" to decide which gridpoints are to be flipped. This decision will be recorded as the true elements of an array FLIP. The decision to flip will be based on the use of uniformly distributed random numbers from the interval $0 \le p < 1$. These will be provided at each gridpoint by the array-valued function RAND. The flip will occur at a given point if and only if the random number at that point is less than a certain threshold value. In particular, by making the threshold value equal to 1 at the points where there are 3 or fewer like-minded nearest neighbors, we guarantee that a flip occurs at those points (because p is always less than 1). Similarly, the threshold values corresponding to counts of 4, 5, and 6 are assigned P (4), P (5), and P (6) in order to achieve the desired probabilities of a flip at those points (P (4), P (5), and P (6) are input parameters in the range 0 to 1).

The thresholds are established by the statements:

```
THRESHOLD = 1.0
WHERE (COUNT .EQ. 4) THRESHOLD = P (4)
WHERE (COUNT .EQ. 5) THRESHOLD = P (5)
WHERE (COUNT .EQ. 6) THRESHOLD = P (6)
```

and the spins that are to be flipped are located by the statement:

```
FLIPS = RAND (N) .LE. THRESHOLD
```

All that remains to complete one transition to the next state of the ISING model is to r spins in ISING wherever FLIPS is true:

```
WHERE (FLIPS)  ISING = .NOT. ISING
```

C.11.2.3.3 The complete Fortran subroutine

The complete code, enclosed in a subroutine that performs a sequence of transitions, is as follows:

```
SUBROUTINE TRANSITION (N, ISING, ITERATIONS, P)

   LOGICAL ISING (N, N, N), FLIPS (N, N, N)
   INTEGER ONES (N, N, N), COUNT (N, N, N)
   REAL THRESHOLD (N, N, N), P (6)

   DO I = 1, ITERATIONS
      ONES = 0
      WHERE (ISING)  ONES = 1
```

```
      COUNT = CSHIFT (ONES, -1, 1) + CSHIFT (ONES, 1, 1) &
            + CSHIFT (ONES, -1, 2) + CSHIFT (ONES, 1, 2) &
            + CSHIFT (ONES, -1, 3) + CSHIFT (ONES, 1, 3)
      WHERE (.NOT. ISING)  COUNT = 6 - COUNT
      THRESHOLD = 1.0
      WHERE (COUNT .EQ. 4)  THRESHOLD = P (4)
      WHERE (COUNT .EQ. 5)  THRESHOLD = P (5)
      WHERE (COUNT .EQ. 6)  THRESHOLD = P (6)
      FLIPS = RAND (N) .LE. THRESHOLD
      WHERE (FLIPS)  ISING = .NOT. ISING
   END DO

CONTAINS
   FUNCTION RAND (N)
      REAL RAND (N, N, N)
      CALL RANDOM_NUMBER (HARVEST = RAND)
      RETURN
   END FUNCTION RAND
END
```

C.11.2.3.4 Reduction of storage

The array ISING could be removed (at some loss of clarity) by representing the model in ONES all the time. The array FLIPS can be avoided by combining the two statements that use it as:

```
WHERE (RAND (N) .LE. THRESHOLD)  ISING = .NOT. ISING
```

but an extra temporary array would probably be needed. Thus, the scope for saving storage while performing whole array operations is limited. If N is small, this will not matter and the use of whole array operations is likely to lead to good execution speed. If N is large, storage may be very important and adequate efficiency will probably be available by performing the operations plane by plane. The resulting code is not as elegant, but all the arrays except ISING will have size of order N^2 instead of N^3.

C.11.3 FORmula TRANslation and array processing

Many mathematical formulas can be translated directly into Fortran by use of the array processing features.

We assume the following array declarations:

```
REAL X (N), A (M, N)
```

Some examples of mathematical formulas and corresponding Fortran expressions follow.

C.11.3.1 A sum of products

The expression $\sum_{j=1}^{N} \prod_{i=1}^{M} a_{ij}$

can be formed using the Fortran expression

```
SUM (PRODUCT (A, DIM=1))
```

The argument DIM=1 means that the product is to be computed down each column of A. If A had the value $\begin{bmatrix} B & C & D \\ E & F & G \end{bmatrix}$ the result of this expression is BE + CF + DG.

C.11.3.2 A product of sums

The expression $\prod_{i=1}^{M} \sum_{j=1}^{N} a_{ij}$ can be formed using the Fortran expression

```
PRODUCT (SUM (A, DIM = 2))
```

The argument DIM = 2 means that the sum is to be computed along each row of A. If A had the value $\begin{bmatrix} \mathrm{B} & \mathrm{C} & \mathrm{D} \\ \mathrm{E} & \mathrm{F} & \mathrm{G} \end{bmatrix}$ the result of this expression is (B+C+D)(E+F+G).

C.11.3.3 Addition of selected elements

The expression $\sum_{x_i > 0.0} x_i$ can be formed using the Fortran expression

```
SUM (X, MASK = X .GT. 0.0)
```

The mask locates the positive elements of the array of rank one. If X has the vector value (0.0, –0.1, 0.2, 0.3, 0.2, –0.1, 0.0), the result of this expression is 0.7.

C.11.4 Sum of squared residuals

The expression $\sum_{i=1}^{N} (x_i - x_{\mathrm{mean}})^2$ can be formed using the Fortran statements

```
XMEAN = SUM (X) / SIZE (X)
SS = SUM ((X - XMEAN) ** 2)
```

Thus, SS is the sum of the squared residuals.

C.11.5 Vector norms: infinity-norm and one-norm

The infinity-norm of vector X = (X (1), ..., X (N)) is defined as the largest of the numbers ABS (X (1)), ..., ABS (X (N)) and therefore has the value MAXVAL (ABS (X)).

The one-norm of vector X is defined as the *sum* of the numbers ABS (X (1)), ..., ABS (X (N)) and therefore has the value SUM (ABS (X)).

C.11.6 Matrix norms: infinity-norm and one-norm

The infinity-norm of the matrix A = (A (I, J)) is the largest row-sum of the matrix ABS (A (I, J)) and therefore has the value MAXVAL (SUM (ABS (A), DIM = 2)).

The one-norm of the matrix A = (A (I, J)) is the largest column-sum of the matrix ABS (A (I, J)) and therefore has the value MAXVAL (SUM (ABS (A), DIM = 1)).

C.11.7 Logical queries

The intrinsic functions allow quite complicated questions about tabular data to be answered without use of loops or conditional constructs. Consider, for example, the questions asked below about a simple tabulation of students' test scores.

Suppose the rectangular table T (M, N) contains the test scores of M students who have taken N different tests. T is an integer matrix with entries in the range 0 to 100.

Example: The scores on 4 tests made by 3 students are held as the table

$$T = \begin{bmatrix} 85 & 76 & 90 & 60 \\ 71 & 45 & 50 & 80 \\ 66 & 45 & 21 & 55 \end{bmatrix}$$

Question: What is each student's top score?

Answer: MAXVAL (T, DIM = 2); in the example: [90, 80, 66].

Question: What is the average of all the scores?

Answer: SUM (T) / SIZE (T); in the example: 62.

Question: How many of the scores in the table are above average?

Answer: ABOVE = T **.**GT**.** SUM (T) / SIZE (T); N = COUNT (ABOVE); in the example: ABOVE is the logical array (t = true, . = false): $\begin{bmatrix} t & t & t & . \\ t & . & . & t \\ t & . & . & . \end{bmatrix}$ and COUNT (ABOVE) is 6.

Question: What was the lowest score in the above-average group of scores?

Answer: MINVAL (T, MASK = ABOVE), where ABOVE is as defined previously; in the example: 66.

Question: Was there a student whose scores were all above average?

Answer: With ABOVE as previously defined, the answer is yes or no according as the value of the expression ANY (ALL (ABOVE, DIM = 2)) is true or false; in the example, the answer is no.

C.11.8 Parallel computations

The most straightforward kind of parallel processing is to do the same thing at the same time to many operands. Matrix addition is a good example of this very simple form of parallel processing. Thus, the array assignment A = B + C specifies that corresponding elements of the identically-shaped arrays B and C be added together in parallel and that the resulting sums be assigned in parallel to the array A.

The process being done in parallel in the example of matrix addition is of course the process of addition; the array feature that implements matrix addition as a parallel process is the element-by-element evaluation of array expressions.

These observations lead us to look to element-by-element computation as a means of implementing other simple parallel processing algorithms.

C.11.9 Example of element-by-element computation

Several polynomials of the same degree may be evaluated at the same point by arranging their coefficients as the rows of a matrix and applying Horner's method for polynomial evaluation to the columns of the matrix so formed.

The procedure is illustrated by the code to evaluate the three cubic polynomials

$$P(t) = 1 + 2t - 3t^2 + 4t^3$$

$$Q(t) = 2 - 3t + 4t^2 - 5t^3$$

$$R(t) = 3 + 4t - 5t^2 + 6t^3$$

in parallel at the point t = X and to place the resulting vector of numbers [P (X), Q (X), R (X)] in the real array RESULT (3).

The code to compute RESULT is just the one statement

```
RESULT = M (:, 1) + X * (M (:, 2) + X * (M (:, 3) + X * M (:, 4)))
```

where M represents the matrix M (3, 4) with value $\begin{bmatrix} 1 & 2 & -3 & 4 \\ 2 & -3 & 4 & -5 \\ 3 & 4 & -5 & 6 \end{bmatrix}$.

C.11.10 Bit manipulation and inquiry procedures

The procedures IOR, IAND, NOT, IEOR, ISHFT, ISHFTC, IBITS, MVBITS, BTEST, IBSET, and IBCLR are defined by MIL-STD 1753 for scalar arguments and are extended in this standard to accept array arguments and to return array-valued results.