

# Vector Fortran for Numerical Problems on CRAY-1

Paul Swarztrauber  
Guest Editor

W. P. PETERSEN *Bell Laboratories*

Wesley Petersen's current research interest include numerical integration of stochastic differential equations, Monte Carlo methods, and parallel/vector computers.

Author's Present Address:  
W. P. Petersen, Bell  
Laboratories, 600 Mountain  
Avenue, Room 2F-211,  
Murray Hill, NJ 07974.

## 1. INTRODUCTION

Computations in vector mode on CRAY-1 can be an order of magnitude faster than in scalar mode. This paper deals with the basics: vector hardware operations and how the CFT (Cray Fortran) compiler makes use of them. A little understanding of the vector hardware is useful because some operations on CRAY-1 are more efficient than others. Despite the hardware motivation, a modular "each block does a vector computation" method usually works very well on any large computer.

### 1.1 Segmentation and Unrolling Loops

To CRAY-1, vectors are regularly spaced arrays of data that can be processed by segments. Regularly spaced data means that each element is the same number of memory locations from its predecessor. For example, the elements  $\{A(N-1), A(N-3), A(N-5), A(N-7), \dots\}$  are regularly spaced, while  $\{A(1), A(2), A(4), A(8), \dots\}$  are not.

In CFT, the principle engines of vector operations are DO loops. If there are  $n$  repetitions of the loop, executing as many as 64 at a time<sup>1</sup> gives

$$n = rsl + 64q$$

where  $rsl \leq 64$  is the number in the residual segment, which is processed first, and  $q$  is the number of additional segments of length 64. Machine instructions generated by CFT for vector DO loops calculate  $q$  and  $rsl$  to "unroll" the loop into segments of length  $\leq 64$ . A vector length register VL [1] is set to the number per segment. All unrolling of loops is transparent to the user, with CFT doing all segmenting and appropriate addressing.

### 1.2 Identical and Independent Operations

In writing vectorized codes, it is important to understand that vector elements really must be independent of one another, but must be treated identically. Vectorized conditional calculations have superfluous operations. Consider the summation<sup>2</sup> of  $N$  elements in an array  $A$  which skips the addition of null values, as shown in Figure 1.

**ABSTRACT:** *This is a practical guide to vector Fortran for programming numerical problems on CRAY-1. The intent is to illustrate those constructions which effectively use the hardware through familiar and useful examples.*

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission. © 1983 ACM 0001-0782/83/1100-1000 75¢

<sup>1</sup> Vector registers having 64 words each are discussed in Section 1.

<sup>2</sup> See Sections 2.4 and 4.2.

```

SUM = 0
DO 1 I = 1,N
  IF (A(I).EQ.0)GOTO 1
  SUM = SUM + A(I)
1 CONTINUE

```

FIGURE 1. Non-Vector Summation.

```

DO 1 I = 2,N
  A(I) = A(I-1)
2 CONTINUE

```

FIGURE 2. Non-Vector Dependency Case.




FIGURE 3. Register S1 Added to V0.

Plucking out the IF statement allows CFT to compile vector hardware instructions. Addition of zero elements may be superfluous, but should be done anyway. The idea is that each  $A(I)$  must be treated exactly like every other. Conditional statements IF and computed GOTO imply that some data are very different from others and are treated accordingly. Branching, which includes IF, GOTO, and CALL statements, inhibits the use of the CRAY-1 vector hardware.

Furthermore, CFT considers an array to be a vector only if it is clear that no element of that array depends upon the previous computation of another. For example, in Figure 2, the natural order of  $I = 2, 3, 4, \dots$  requires that  $A(I-1)$  be replaced by  $A(I-2)$  for  $A(I)$  to be properly set. Thus, the  $A(I)$ s in Figure 2 must be set one at a time rather than by segments.

Indexing arrays so that CFT compiles vector machine instructions is generally straightforward, but there are subtleties. Section 3 of this document, the CFT reference manual [2, Part 3, Section 2], and the paper of Higbie [3] deal with indexing in some detail.

## 2. CRAY-1 VECTOR HARDWARE

Only about one-fourth of the CRAY-1 machine instructions use vector registers. The following notation will be used to

describe machine instructions [4]:

$$V1 \leftarrow S1 + V0$$

Here, the content of scalar register S1 is added element by element to the contents of vector register V0 and the results stored in vector register V1.

Figure 3 is a pictorial representation of this instruction. Paths to each register are represented by arrows. The end of each path is a pointer indicating only one element at a time. Once this instruction is issued and addition begins, the pointer in V0 is incremented each clock period ( $12.5 \times 10^{-9}$  seconds), until the operands are exhausted. A similar pointer in result V1 is incremented when results begin to emerge from the adder, three clock periods later in this case.

The integer adder, which may contain three separate pairs of operands concurrently in distinct stages of processing, is called a pipelined functional unit. It has some analogy to a short piece of pipe into which marbles are being pushed; until the pipe is full, no marbles emerge from the other end. Even though the first result does not emerge from the pipe until several clock periods later, successive results arrive one clock period apart. One 64-bit integer addition takes three clock cycles, but pushing a segment of 64 additions through only takes 67 cycles. Effectively, this is only slightly more than one cycle per addition—a factor of three faster than one at a time. Longer pipelines show even greater improvements—typically six- or seven-fold for floating point operations. Further, with chaining and overlap, described in the next section, several operations may run concurrently.

Machine instructions generated by CFT to execute a vectorized Do-loop control the segmentation (Section 1.1), and involve some of eight vector registers V0, V1, ..., V7 of 64 words each. Scalar registers S0, S1, ..., S7 may also be used as operands in some vector operations, as in Figure 3 for example. Vector merging, that is, selecting vector elements word-by-word is implemented by a correspondence between the 64 bits of the S and VM registers and the 64 words of the V registers (see Section 3.2). Operations may run concurrently if certain independence criteria are satisfied. In particular, each of seven functional units may run independently to perform the operations shown in Table I.

<sup>1</sup>Division uses a 30-bit reciprocal approximation and one Newton iteration [5, pp. 5-53].

TABLE I. Vector Functional Units

| Unit    | operations | Purpose                                |
|---------|------------|--|
| memory  | load       | load register from memory              |
|         | store      | store to memory from register          |
|         | F          | f.p. (truncated) multiply              |
|         | *R         | f.p. (rounded) multiply                |
|         | *I         | f.p. (iterative) multiply <sup>3</sup> |
| /       | /H         | reciprocal approximation <sup>3</sup>  |
|         | + / -      | f.p. add / f.p. subtract               |
| logical | &          | logical .and.                          |
|         | !          | logical .or.                           |
|         | \          | exclusive .or.                         |
|         | VM         | form vector mask                       |
| merge   |            | vector merge                           |
|         | + / -      | integer add / integer subtract         |
| shift   | >          | right shift                            |
|         | <          | left shift                             |

|    |   |        |                    |
|----|---|--------|--------------------|
| V0 | ← | B      | b/s                |
| V1 | ← | /HV0   | half precision 1/b |
| V2 | ← | V0*V1  | 2.0 - b/Hb         |
| V3 | ← | V1*RV2 | full precision 1/b |
| V4 | ← | A      | a/s                |
| V5 | ← | V4*RV3 | a/b/s              |

FIGURE 4. Division by Reciprocal Approximation.

Each vector functional unit is independent, and results from one unit may be fed into another as operands—hence, operations may chain together.

Notice that there are vector hardware provisions only for processing of data, not for creation. In particular, the vector hardware will not create an array of integers. This means that arrays may not be generated directly from DO-loop index variables in vector mode: A(1) = FLOAT(1), for example.

### 2.1 Chaining and Functional Unit Overlap

If successive operations use different functional units, they may run concurrently. Division using a Newton iteration is an interesting example; we want  $a/b$ , and use the operations from Table I in Figure 4.

In a sense, there are only four separate operations in division, with the first pair linking together to form a chain, and instructions four and five overlapping. Briefly, separation occurs as follows:

1. Memory access is independent of reciprocal approximation; so, as soon as the first  $b_i$  arrives in V0 the second instruction is issued. As the  $b_i$ s are stored into V0, they are immediately copied by the reciprocal unit and start through that pipe, eventually emerging to be stored into V1. This linking to form a short chain is shown in Figure 5.
2. Although multiplication is independent of reciprocal approximation, the pointer in V0 has moved down before the first result arrives in V1. Thus, V0 is "busy" as an operand, and the third instruction must wait until the first pair are finished.

3. Instruction four uses the same functional unit (multiply) as the third. This unit is busy until the third is finished. Instruction five (load A) may issue after four begins, and will run concurrently. Operations begun by the issue of instructions four and five do not chain together, but "overlap" and run concurrently.
4. Instruction six must wait until four (running concurrently with five) is finished because it uses the multiply unit again.

### 2.2 Memory Access and Timing

With some exceptions (division for example), the number of sequences that run concurrently is approximately the number of memory access instructions. Since there is only one port to memory, and in Fortran *all symbols are in memory*, this is usually the critical resource. A lower bound on any loop timing (in seconds) is

$$T \geq (12.5 \times 10^{-9}) \times 8/7 \times (\text{number of memory references})$$

where  $8/7$  accounts for pipeline overhead. Indexing, that is, segment counting, array offsets, setting VL, etc., is independent of vector operations and runs concurrently with the last store into memory from the loop. Indexing by segments adds little to loop overhead and is transparent to the CFT user. If there are no function calls, the bound is a reasonable timing approximation, roughly within a factor of two.

Memory is also the critical resource in Input/Output requests. Loading vector registers locks up memory. In principle, I/O runs concurrently with computation, but vector loading keeps memory pretty busy. Thus, for a particular user job, overlapping I/O with vector loops is not as helpful as might be expected. To make things more complicated, unless the job is running by itself, the operating system (COS—Cray Operating System) will start up another job while waiting for the I/O request to finish. However, a few things will help.

1. FORMATTed records should be avoided wherever possible. FORMAT statements are executed in an interpretative manner character-by-character, and are very slow.
2. Use of BUFFER IN and BUFFER OUT to transfer datasets will allow overlapping the I/O with the computation, which may help sometimes.




FIGURE 5. Chaining Pictorial.

3 Usually, sequential I/O using READ and WRITE statements works best because it's very simple and lets the operating system take care of all the concurrency headaches. It is best to use long fixed-length records—somewhere around 500 words. Every dataset transfer involves at least one sector of 512 words, which includes some block and record control words. It is also advisable to try to avoid array references having multiple-of-eight increments which cause memory bank conflicts.

Sequential array elements are stored in sequential banks of memory. There are 16 banks, each with a cycle time of four clock periods [1, Section 6-21]. Fetching or storing by multiple-of-eight increments conflicts with this cycle. Memory bank conflicts also break chaining with independent operations. These are important considerations for FFTs and Cyclic Reduction (which are discussed in Sections 4.4 and 4.5.3).

### 2.3 Reduction Operations

Operations with a result vector having the same number of elements as the operands are usually more efficient than, say, dot-products—which reduce dimensionality. Consider a DO-loop that computes the product of N elements of array A. Result PROD is of dimension unity, but A is an N-dimensional operand, (Figure 6) Through CFT version 1.09, the reductions used a curious property of the vector hardware. After version 1.10, a somewhat slower procedure is used, which is compatible with the CRAY X-MP. In what follows the pre-1.10 method is discussed.

To do reductions like this, three steps are necessary. For simplicity, let  $N = 64q$ , and divide A into q segments, each of length 64. In the first step, the following accumulation happens. Segments 1 and 2 are multiplied together, element-by-element, to yield a 64-element result. That result is now multiplied by the 3rd segment—still a 64-element result, then the 4th segment, and so on until all q segments are used.

A second step reduces this 64-element partial result by a recursive hardware operation [1, pp. 3-14] similar to the following, if V0 contains the 64-element partial result,

$$V1 \leftarrow V1 * RV0$$

where the first element of V1 is set to 1.0. Result register V1 is also an operand. In this case, the pointer in V1 cannot advance until functional unit time + 2 clock periods ( $7 + 2 = 9$ ) later, when the first result is ready. When completed, the last nine elements of V1 are the product of the first element of V0 and every ninth element of V0. Use of identical registers for operands and results is deliberate and recursive, producing a useful reduction of the 64 products. An example of this recursive mechanism is shown in [I], but is too involved to reproduce here.

Finally, in the last of three steps, f.u. time + 2 elements in V1 must be pulled out and the reduction completed. This remaining step, which uses the S-registers, can be scheduled to take about 45 clock periods, somewhat less than a vector segment operation on 64 elements. Such a macro is invoked

```

PROD = 1.0
DO 1 I=1,N
  PROD = PROD*A(I)
1 CONTINUE
    
```

FIGURE 6. Product of Vector Elements.

by CFT to do reductions of sums and products only. If N is large, two additional steps after the initial accumulation add little extra time, since each is executed only once. Compared to the q accumulation operations in the first step, the last two become unimportant. However, for moderate-length vectors (less than 100), reductions are less efficient than vector → vector operations. Sections 4.1 and 4.2 give examples that illustrate the point.

### 3. PROGRAMMING THE LOOPS

The appearance of

```
VECTOR BLOCK BEGINS AT SEQ. NO.      n, P =  addr
```

at the end of a listing of a compiled subprogram means that CFT has generated a vector DO-loop in a particular block of code. A block is a basic unit of code which is locally optimized by CFT and is demarcated by register usage. Subroutine or user-defined function calls, GOTO statements, and inner DO-loops force blocking. CFT only vectorizes inner loops! Since blocking does not usually start exactly at the DO statement, to identify the vectorized loops in a compiled listing look for the first inner loop following the sequence number beginning the vector block. This will be the vectorized loop.

Many CFT intrinsic functions, SIN, EXP, SQRT for example, are vector mode (computation done in vector registers) and are used in DO-loops. Others, like ATAN, may be used in vector loops but are not really vector mode. These pseudo-vector functions pass segments of arguments in-register (V1, or V1 and V2), but process them element-by-element in scalar registers. These pseudo-vector (see [2, Appendix B]) routines exist because nobody has rewritten them. They do permit CFT to vectorize the rest of a loop, however. Table II summarizes the intrinsic CFT operations and functions that are vector-mode.

Neither DOUBLE PRECISION nor CHARACTER data computations are done in vector registers. The CRAY-1 has no double precision hardware. Data transfer and I/O for double precision does use vector register memory access, however.

\*Positive difference function DIM gives  $X = \text{DIM}(Y, Z) = Y, Z$  if  $Y > Z$ ,  $X = 0$  otherwise.

TABLE II. CFT Vector Operations and Functions by Data Type

| SINGLE      | COMPLEX    | INTEGER    | DOUBLE | BOOLEAN |
|-------------|------------|------------|--------|---------|
| + -         | * /        | * /        |        | .AND.   |
| /           |            | /          |        | .OR.    |
| logical     |            | logical    |        | .XOR.   |
| SQRT        |            |            |        |         |
| EXP, "      |            |            |        |         |
| COS/SIN     | CSIN/CCOS  |            |        |         |
| ALOG/ALOG10 |            |            |        |         |
| ABS         | CABS       | IABS       | DABS   |         |
| DIM         |            | IDIM*      |        |         |
| INT         |            | AINT       | IDINT  |         |
| AMOD        |            |            |        |         |
| SIGN        |            | ISIGN      | DSIGN  |         |
| RANF        |            |            |        |         |
| MAX1/AMAX1  |            | MAX0/AMAX0 | DMAX1  |         |
| MIN1/AMIN1  |            | MIN0/AMINO | DMIN1  |         |
|             | CONJG      |            |        |         |
| JFIX        |            | FLOAT      | SNGL   |         |
| COMPLX/DBLE | REAL/AIMAG |            |        |         |
| SIGN        |            | ISIGN      | DSIGN  |         |

INTEGER type addition (subtraction) is 64 bit in vector mode, but multiplication is done by floating point hardware. Integer multiplication has only 48 bits of precision.

### 3.1 Increments, Dimensions—How to Index the Arrays

Although many operations on regularly spaced data in memory can be done in vector mode, CFT doesn't recognize them all. Figure 7 shows several examples of array subscripts which CFT recognizes as vector indices. Two concepts are important here: invariants within the loop, and constant increment integers. Loop invariants are quantities unmodified by changes in the DO-variable (I in this case). Constant increment integers (CII) "follow the DO-variable" in that each time I is changed, the CII is incremented by a constant amount. In Figure 7, the variables I1, I3, I4, I5, I6, I8, I9, I10, I11, I12, I15, I17, I19, I20, I22, and X3(I13), IA(I14) are invariants. Variables I2, I7, I16, I18, I21 are constant increment integers (CII).

We have the following rules.

1. The subscripts of a vector in an array may appear in only one dimension of that array. In Figure 7 only I21 varies in Y6.
2. These subscripts must be expressed as functions of the DO-loop variable or of constant increment integers. CII are computed in the loop by

$$\text{CII} = \text{CII} \pm \text{constant increment.}$$

The constant increment can be only an invariant integer or a sum of invariant integers. Expressions for CII cannot contain parentheses, nor any operation but addition (subtraction). In Figure 7, index variables I2, I7, I16, I18, I21 are CII, and I3, I9 + I10, I17, I19, and I22 are the constant increments.

3. Subscripts must be one of the forms

$$\text{inv} \pm I \cdot \text{istep}$$

or

$$\text{inv} \pm \text{CII} \cdot \text{istep}$$

where *inv* is any invariant expression not containing a

```

DO I = 1, N
C (1)      Y1(I, I) = X1(I) + I*(I2*I3)
C (2)      Y2(I4) = X2(IPCN(X) + I)
           I4 = I4 + 2*I5
C (3)      Y3(I6 + I) = ATAN(X3(I))
           Y3(I7 + I) = FCN(X3(I))
C (4)      Y4(I) = X4(I + I*I9)
           Y5(I) = Y4(I8 + I)
C (5)      I10 = INDEX(I)
           Y6(I10) = FLOAT(I)
C
           I CONTINUE

```

FIGURE 8. No Vectors Here.

function reference (*inv* may be null), and *istep* must be a simple integer invariant. Step *istep* can be of either sign (see Y2(I4 + I\*I5) and Y3(I11 - I\*I12) in Figure 7).

**Warning:** At the time of this writing if *istep* = 0, results may be erroneous.

As a final example, Figure 8 illustrates array references that will not vectorize, even as self-contained loops.

The reasons CFT will not vectorize these expressions are:

1. The spacing computations in Y1 and X1 are too complicated for CFT: A diagonal in Y1 (two dimensions, see rule 1), and the spacing (I2\*I3) is not a simple invariant integer variable without parenthesis. Both these subscript expressions are easily changed:
  - a. Set LD1 = LD + 1, where LD is the leading dimension of Y1(LD), then Y1(I\*LD1 - LD) satisfies rule 3, making LD1 a loop invariant. Of course, LD1 must be calculated outside the loop.
  - b. Call I23 = I2\*I3, then X1(I1 + I\*I23) satisfies rule 3.

```

DO I = IL, IU, IS
C
C CONSTANT INCREMENT INTEGER - I2
  Y1(I + I1) = X1(I2)
  I2 = I2 + I3
C REGULAR STEPS (I5, I8), WITH CII (I7)
  Y2(I4 + I*I5) = X2(I6 - I7*I8)
  I7 = I7 + I9 + I10
C REGULAR STEP (I12), INVARIANTS (X3, IA)
  Y3(I11 - I*I12) = X3(I13)*X4(IA(I14) - I)
C RECURSIVE SUMMATION OF ARRAY X6
  Y4(I15) = Y4(I15) + X6(I16)
  I16 = I16 + I17
C RECURSIVELY FORMING PRODUCT OF X6 ELEMENTS
  Y5(I17) = Y5(I17)*X6(I18)
  I18 = I18 + I19
C SETTING ELEMENTS IN ROW OF Y6
  Y6(I20, I21) = X7(I)
  I21 = I21 + I22
C
           I CONTINUE

```

FIGURE 7. Indexing of Vector Loops.

2. The increment  $2 \cdot I5$  is not a sum of simple invariants (see rule 2), and the subscript expression for X2 contains a function reference:
  - a. Changing  $2 \cdot I5$  to  $I5 + I5$  corrects the CII calculation of I4.
  - b. Pulling the expression  $IFNX = IFCN(X)$  out of the loop permits  $X2(IFNX + I)$ , with a valid subscript.
3. Reference to array Y3 represents a subtle dependency case if  $I7 \geq I6$  and the storage overlaps. At compile time, this seems to be a dependency which would yield incorrect results; hence CFT will not attempt to use vector hardware for this calculation. (See Sections 1.2 and 4.5.) The reference to ATAN is really a pseudo-vector function (see Table II). Unless FCN is declared by a VFUNCTIO directive (see Section 3.3), CFT does not recognize it as a vector function.

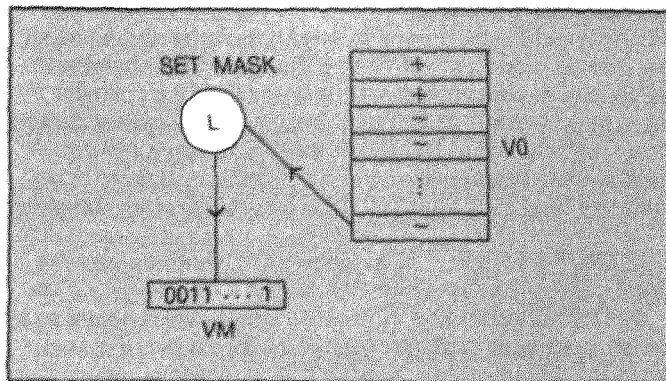


FIGURE 9. Set Mask.

- a. Insertion of an IVDEP directive prior to the loop will permit the vector-mode access to Y3, at the user's risk, if  $I7 \geq I6$ .
  - b. Although ATAN is a pseudo-vector function, it still may be used in a vector loop.
  - c. A CAL-coded (Cray Assembly Language) FCN with call-by-value linkage would have to be provided (see Section 3.3).
4. References Y4 seem to CFT to be a dependency because if I8 is negative, Y5 requires an element  $Y4(I8 + I)$  which is computed at the same step. (See Sections 1.2 and 4.5.) In X4 the subscript is not of the form given in rule 3.
  5. Array Y6 is referenced in a nonlinear way indirectly through the set of pointers INDEX (see Section 4.5).  $FLOAT(I)$  is a direct calculation on the DO-variable I (see Section 2).

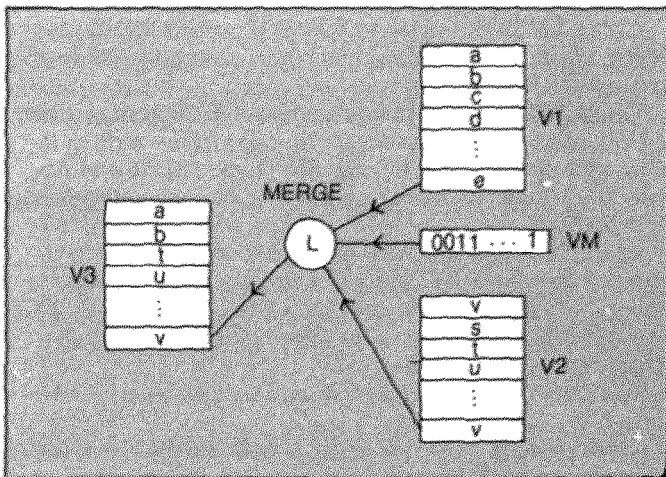


FIGURE 10. Merge Result.

### 3.2 Conditional Statements—IF and GOTO

At this time, CFT (version 1.09) permits no conditional statements IF or GOTO in vector inner loops. Because CRAY-1 is a simple machine, with relatively few instructions, superfluous calculations occur in processing conditionals. For example, if for each I we want

$$A(I) = \text{exp}_1 \quad \text{if } \text{cond}(I) = \text{.false.}$$

$$A(I) = \text{exp}_2 \quad \text{if } \text{cond}(I) = \text{.true.}$$

then both  $\text{exp}_1$  and  $\text{exp}_2$  are evaluated and the resulting vector A is selected from  $\text{exp}_1$  or  $\text{exp}_2$  on the basis of  $\text{cond} = \text{false/true}$ , respectively. This procedure contrasts sharply with sequential machines where  $\text{exp}_2$  may not be evaluated unless  $\text{cond}$  is determined to be true. The hardware operations can be sketched [1, pp. 4-51 and pp. 4-71]:

- V0 ← cond                    set V0 to logical conditions
- VM ← V0, M                set vector mask if true
- V1 ← exp<sub>1</sub>                evaluate exp<sub>1</sub>
- V2 ← exp<sub>2</sub>                evaluate exp<sub>2</sub>
- V3 ← V2!V1 and VM        select result

with the first two operations chaining together. Figures 9 and 10 illustrate the operations of setting the vector mask and merging. An element of V3 is selected from V2 if the vector mask (VM) bit corresponding to that element is set, otherwise an element from V1 is selected.

Macros which generate these instructions are invoked by MAX, MIN (see TABLE II), and CVMG (see Figure 11) functions. The complex vector merge functions summarized in Figure 11 are not portable. Future versions of CFT will proba-

```

DO 1 I = 1, N
C CVMGP selects X1 if TEST ≥ 0, X2 otherwise
  Y1(I) = CVMGP(X1(I), X2(I), TEST(I))
C CVMGM selects X1 if TEST < 0, X2 otherwise
  Y2(I) = CVMGM(X1(I), X2(I), TEST(I))
C CVMGZ likes X1 if TEST = 0, X2 otherwise
  Y3(I) = CVMGZ(X1(I), X2(I), TEST(I))
C CVMGN prefers X1 if TEST = 0, X2 otherwise
  Y4(I) = CVMGN(X1(I), X2(I), TEST(I))
C CVMGT elects X1 if LTEST is truthful
  Y5(I) = CVMGT(X1(I), X2(I), LTEST(I))
C selection for each CVMG is element by element
1 CONTINUE

```

FIGURE 11. CVMG Functions.

bly generate the same instruction sequences for IF statements that perform the same operations.

Despite their nonportability, these functions are attractive because their arguments can be any expression, vector, or scalar.

Many conditionals can be expressed by AMAX1, AMIN1, etc., functions that use the mask/merge hardware instructions. Where applicable, the MAX/MIN conditionals are preferable to the CVMG functions because they are ANSI standard Fortran. An illustration of this can be seen in Floyd's algorithm in Section 4.3.

### 3.3 Compiler Directives

Figure 2 and Figure 8 have examples showing dependencies among array elements inhibiting the generation of vector instructions. A programmer who is wiser than the compiler may wish vector instructions to be generated anyway. Several "comment line" directives exist for CFT to control vector compilation. Beginning with a "C" in column 1, these directives are treated as comment lines by other Fortran compilers.

CDIRS IVDEP tells CFT to ignore the apparent vector dependency in the next inner DO-loop.

CDIRS NOVECTOR = *n* turns off generation of vector code for all loops with iteration counts less than *n* + 1, if *n* is known at compile time. NOVECTOR is a switch and applies to all successive code compiled until this switch is toggled. If *n* is not specified, all vector code is suppressed. A variable *n* is not allowed.

CDIRS VECTOR toggles NOVECTOR and NORECURR-ENCE directives (see below) and causes resumption of vector code generation.

CDIRS VFUNCTION *fname* instructs CFT that there exists a "call-by-value" external vector function *fname* written in CAL with conventional linkage. See [2, pp. (3)1-18 and Appendix F]. Function *fname* is then used like any intrinsic (e.g., SIN, EXP).

CDIRS NORECURRENT = *n* switches off vectorized reduction operations for DO-loops with iteration counts less than *n* + 1, if *n* is known at compile time. NORECURRENT is a switch and applies to all successive compilation until toggled by a VECTOR directive.

CDIRS NEXTSCALAR switches off vectorization for the next DO-loop only.

Compiler directives are not enabled unless the ON = E option is used in the CFT control statement [2, p. (3)1-1]. This option is default at Murray Hill Computer Center. An example of IVDEP usage is illustrated in Section 4.5.2.

## 4. CHOOSING YOUR ALGORITHM

In this section are a number of examples to illustrate the material discussed earlier. Unfortunately, the examples are microscopic and do not indicate global strategies for attacking large problems. A modular approach to larger calculations still demands attention to detail at the subroutine level, however. Some very useful timing information at a global level can be obtained using the FLOWTRACE option of CFT [2, pp. (3)1-5 and pp. (3)1-20, 23]. The total time spent in each subroutine, the percentage of the total time, and some overhead information is computed and printed by FLOWTRACE. The subroutines presented here frequently represent important kernels:

1. Section 4.1. Outer products versus inner products—a linear digital filter.
2. Section 4.2. Inner product formulations, discussed in Section 2.3—a Toeplitz matrix solver.
3. Section 4.3. Conditional statements discussed in Section 3.2, with an example of Floyd's minimum path algorithm.
4. Section 4.4. Maximizing the inner vector length to make the most efficient use of the vector capabilities—a sample Fast Fourier Transform.
5. Section 4.5. Eliminating vector dependencies, noted in Figure 2, and examples 3 and 4 of Figure 8.
  - a. A Gauss-Seidel relaxation step.
  - b. A Red-Black relaxation.
  - c. A multiple tridiagonal solver.
6. Section 4.6. Indirect addressing, noted in example 5 of Figure 8—the problem with sparse matrices.

```

DO 1 J = 1, N-N+1
  F(J) = 0
DO 1 K = 1, N
  F(J) = F(J) + A(K)*D(J+K-1)
1 CONTINUE

```

FIGURE 12. Dot-Product Method.

```

DO 1 J = 1, N-N+1
1  F(J) = 0
DO 2 K = 1, N
DO 2 J = 1, N-N+1
  F(J) = F(J) + A(K)*D(J+K-1)
2 CONTINUE

```

FIGURE 13. Outer-Product Method.

### 4.1 Outer Products—A Linear Digital Filter

A simple nonrecursive digital filter is a convolution of time series data (*d*) with a small number of filter coefficients (*a*). Seismic data, for example, typically has about 2000 points per trace and requires 50 to 100 filter coefficients. Output *f* is

$$f_i = \sum_{k=1}^m a_k d_{i+k-1}$$

where  $1 \leq j \leq n - m$ ,  $1 \leq j + k \leq n$ . Data *d* has *n* + 1 points while the output *f*, containing less information, has *n* - *m* + 1 points. There are two ways to proceed: 1) for each point *j* sum the *k* elements or 2) for each new *k* add all the *j*s from all the previous *ks*. Figure 12 and Figure 13 represent the dot-product method and outer-product method, respectively.

In the above example or that of matrix multiplication [6] or Gaussian elimination [7], outer-product procedures are generally more efficient. On the Murray Hill CRAY-1 (for *N* = 2000, *M* = 100), Figure 13 runs in half the time of Figure 12, precisely for the reasons given in Section 2.3. Both examples compile to vector object code.

If the number of elements to be processed by an outer-product method gives a short vector length while a dot-product procedure gives a much longer one, the differences may be slight or the converse of the above argument may be true. Clearly, a dot-product is not efficiently done as an outer-product of length one. Some procedures are dominantly reduction operations, as the next example shows.

### 4.2 Inner Products, A Symmetric Toeplitz Matrix

An *n* × *n* symmetric Toeplitz matrix *T* with *n* independent elements has  $T_{ij} = T_{j-i+1}$ , for  $1 \leq i, j \leq n$ . Each element of linear array *t* is used to form a pair of symmetric diagonals. The following solution (Figure 14) of the linear system *Tx* = *y* is due to Levinson [8] and requires an operation count  $\propto n^2$ . Although there are algorithms with operation counts of  $O(n \log^2 n)$  [9] and [10] which use discrete Fourier transforms, the present  $O(n^2)$  procedure illustrates the indexing for reduction operations. The method is similar to conjugate gradient, and uses a work space (*C*) of dimension *2n*.

```

SUBROUTINE TOEPLITZ(N,X,Y,C,T)
REAL C(2*N),T(N),X(N),Y(N)

C

X(1) = Y(1)/T(1)
C(1) = T(2)/T(1)

C
C NEXT N-1 RECURSIVE STEPS INCREASE THE VECTOR
  LENGTH

C
DO 1 M = 0,N-2
  IF((M.AND.1).EQ.1)THEN
    CALL CSOLV(M,C(N+1),C(1),T)
    CALL XSOLV(M,X,Y,C(N+1),T)
  ELSE
    CALL CSOLV(M,C(1),C(N+1),T)
    CALL XSOLV(M,X,Y,C(1),T)
  ENDIF
1 CONTINUE
RETURN
END

SUBROUTINE CSOLV(M,C,CC,T)
REAL C(1),CC(1),T(1)
IF(M.LE.0)RETURN
C1N = T(M+2)
C1D = T(1)
DO 1 I = 1,M
  C1N = C1N - CC(I)*T(I+1)
  C1D = C1D - CC(I)*T(M+2-I)
1 CONTINUE
C(1) = C1N/C1D
C1 = C(1)
DO 2 I = 2,M+1
  C(I) = CC(I-1) - C1*CC(M+2-I)
2 CONTINUE
RETURN
END

SUBROUTINE XSOLV(M,X,Y,C,T)
REAL C(1),T(1),X(1),Y(1)
XN = Y(M+2)
XD = T(1)
DO 1 I = 1,M+1
  XN = XN - X(I)*T(M+3-I)
1 CONTINUE
X(M+2) = XN/XD
SX = -X(M+2)
DO 2 I = 1,M+1
  X(I) = X(I) + SX*C(I)
2 CONTINUE
RETURN
END

```

FIGURE 14. Solution to Symmetric Toeplitz System.

In Figure 14 the bulk of computation is in dot-product operations computing local variables C1N, C1D, and XN, XD. In subroutines CSOLV and XSOLV, relabeling variables C(1) as C1 and  $-X(M+2)$  as SX eliminates apparent dependencies. The vector length of each loop increases each iteration  $1 \leq m \leq n$ . Segment overhead (see Section 2.2) is constant, and for short segments becomes appreciable. In this example there is no known way to avoid functional unit overhead for short segments. In the example of an FFT in Section 4.4, an increasing or decreasing vector length can be dealt with effectively, but not in the present case. Nevertheless, this  $O(n^2)$  algorithm is very efficient because every loop is in vector mode. It inverts a 256 dimensional system in 17 milliseconds, but requires 77 milliseconds in scalar mode (i.e., OFF = V option of CFT [2, p. (3)1-1]).

### 4.3 Floyd's Algorithm—Removing the IF Statements

The following shortest path through a network algorithm is due to Floyd [11]. In Figure 15, the inner-loop IF and CVMGM statements are "commented out" to give alternate calculations of  $m_{j,k}$  shown by the "C---" lines. Initially,  $m_{j,k}$  is the length of a

```

SUBROUTINE FLOYD(M,N)
INTEGER M(N,N),S,T
C INF IS EFFECTIVELY INFINITE BUT DOESN'T OVERFLOW
  INF = 10**10
C
DO 1 I = 1,N
  DO 1 J = 1,N
    T = M(J,I)
    IF((T.LT.INF).AND.(I.NE.J))THEN
      DO 2 K = 1,N
        S = T + M(I,K)
C ----- IF(S.LT.M(J,K)) M(J,K) = S
C ----- M(J,K) = CVMGM(S,M(J,K),S-M(J,K))
        M(J,K) = MIN0(S,M(J,K))
      2 CONTINUE
    ENDIF
  1 CONTINUE
RETURN
END

```

FIGURE 15. Minimum Path.

direct link from point  $j$  of a network to point  $k$ . If no direct link exists,  $m_{j,k}$  is assumed to be initially  $\infty \equiv 10^{10}$ . On exit,  $m_{j,k}$  contains the length of the shortest path from  $j$  to  $k$ . In Figure 15 the array  $M$  is of positive integers, but this is not essential—substituting AMIN1 for MIN0 deals with positive reals. It turns out that the introduction of the integer variable  $T$  is necessary. If  $M(J,I)$  were used instead of  $T$ , the compiler would flag the inner loop as a dependency situation;  $M(J,I)$  might be overwritten when  $K=I$  in the  $J$ th column and the new  $M(J,I)$  would be used for  $K>I$ . This will happen as a result of the semantics, not the syntax: namely, the positive array element  $M(J,I)$  will not be overwritten by  $M(J,I) + M(I,I) = 2*M(J,I)$ . The CVMGM function is equivalent to MIN0, and vectorizes. Use of the IF-test replacement in lieu of CVMGM or MIN0 for  $N > 10$  gives a nonvectorizing subroutine that takes five times longer,



#### 4.4 An FFT—Maximizing the Inner Vector Length

In their original paper, Cooley and Tukey [12] noted that their decimated procedure for computing discrete Fourier transforms is a parallel algorithm. The following variant of a power-of-two transform on complex data is from Temperton [13, 14] and apparently is due to T. G. Stockham, Jr. [15]. A similar method is shown by Uhrich [16]. It is ordered, requiring no bit-reversal because of decimation. The price paid for ordering is a work space the same length as the input array. A pipelined computer permits only a bit-reversed algorithm for in-place computation. A careful examination of the signal flow diagram (Figure 17) shows that only simultaneous, as distinct from asynchronous, storage of output at each of the  $\log_2 n$  steps would permit an ordered output to be done in place. A bit-reversed eight point output vector has elements numbered 0, 4, 2, 6, 1, 5, 3, 7—easily seen by writing out the three bit numbers.

Each box in Figure 17 calculates complex numbers  $c$  and  $d$  from inputs  $a$  and  $b$  as in Figure 16. For a box labeled  $k$ , the arithmetic is

$$c = a + b \quad (1)$$

$$d = w^k \cdot (a - b) \quad (2)$$

where  $w$  is the  $n$ th root of unity. In fact, examination of Figure 17 shows there are two possibilities for vector computation of Equations (1) and (2). The first can be seen at step 1: the column of boxes to the left. Reading the elements  $w^k$  from memory as a complex vector, gives

$$C = A + B \quad (3)$$

$$D = W \cdot (A - B) \quad (4)$$

where  $W$  contains  $n/2$  elements  $\{w^k, k = 0, 1, 2, 3$  in the eight point example above]. Vectors  $A$  and  $B$  are vectors of  $a$ 's and  $b$ 's respectively,  $C$  and  $D$  similarly. Vector multiplication in Equation (4) is element-by-element. At step 2, these vectors are now of half-length, with  $W$  containing every other  $w^k$  ( $w^0, w^2$ ) in Figure 17, but Equations (3) and (4) are executed twice. Boxes 1 and 3 are computed first, then boxes 2 and 4, using the same  $W$ . At step 3, the lengths of the vectors are halved again with Equations (3) and (4) executed four times [vector length 1,  $W = \{w^0\}$ , in Figure 17] and so forth if  $n$  is larger.

An alternate method is to read the  $w^k$  elements as complex scalars

$$C = A + B \quad (5)$$

$$D = w^k \cdot (A - B) \quad (6)$$

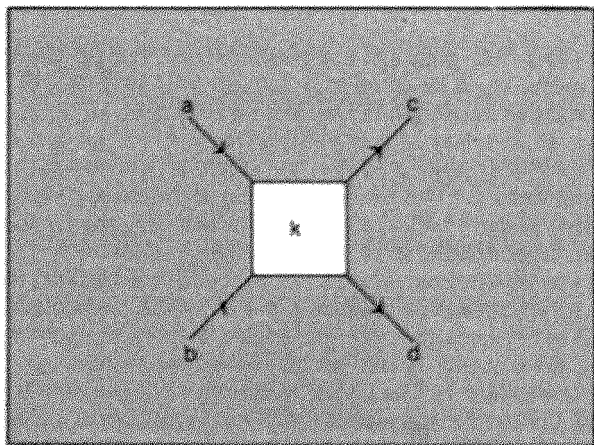


FIGURE 16. Computation Box.

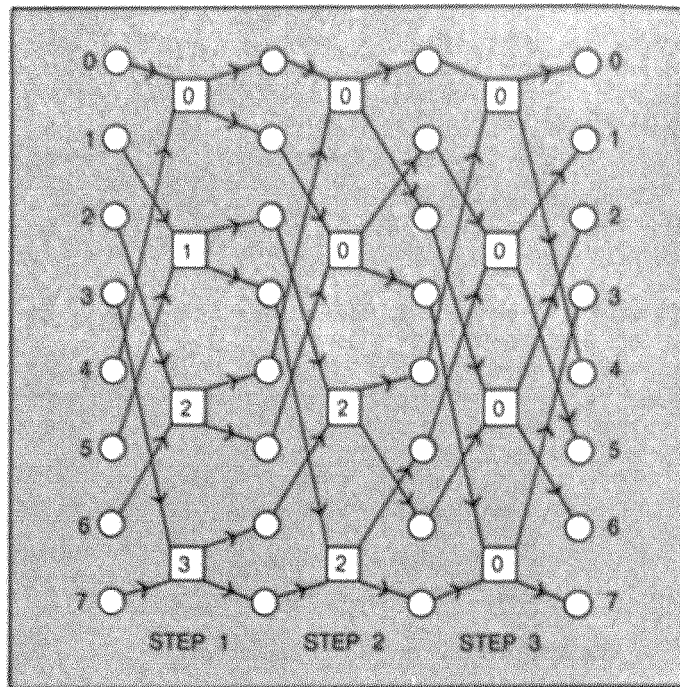


FIGURE 17. Signal Flow Diagram.

which scale the vector  $A - B$ . At the first step, vector Equations (5) and (6) are executed  $n/2$  times (once for each  $k = 0, 1, 2, 3$  in Figure 17); while on the second step the length of the vectors is doubled but executed only  $n/4$  times. Thus, using Equations (3) and (4) the vectors start with length  $n/2$  and shrink by half at each step to length one; while using Equations (5) and (6) the vector length is initially one and doubles at each step to  $n/2$  at the last step. The last step can be done in place. A mixed procedure, using Equations (3) and (4) for the first few steps and (5) and (6) for the last steps, maximizes the inner vector length [17].

The computation in Figure 18 uses Equations (5) and (6) to set

$$y_j = \sum_{k=1}^n \exp\left[\frac{2\pi j}{n} (j-1)(k-1)\right] x_k$$

for  $1 \leq j \leq n$ . Driver CFFT toggles between arrays  $X$  and  $Y$ . Complex vector  $W$  contains the  $n/2$  elements  $\{\exp(2\pi ik/n)\}$  for  $k = 0, 1, 2, \dots, n/2 - 1$ , which must be precomputed. Using the vector SIN and COS functions, the calculation of  $W$  adds about 40 percent to the computation time (for  $n = 1024$ ).

Turning to DO 1 and DO 2 loops insideout in STEP, with appropriate provisions for the increments I2 and I3, computes the transform using Equations (3) and (4). This alternate version of STEP is given in Figure 19. Notice here that  $W$  is read as a vector by steps of I2 in the inner loop. In this mode, when I2 is a multiple of 8, memory bank conflicts start to slow access (see Section 2.2). Despite this technical difficulty, using the alternate form of STEP (Figure 19) for the first four steps of the transform improves the performance, as shown in Figure 20. This plot shows the execution time for a 1024-point complex FFT vs. the last step for which Figure 19 replaces STEP in Figure 18. That is, in Figure 20 IBREAK is the last step [of  $\log_2 n$ ] for which Equations (3) and (4) are used, with Equations (5) and (6) used for the remainder. For transforms longer than 1024, IBREAK = 4 is still used. The 5 millisecond improvement shown in Figure 20 does not change for longer transforms.

To conclude this section, several remarks are appropriate concerning multiple FFTs and general radix transforms. Frequently one is interested in multiple transforms of the same length. For example, solving Poisson's equation on a square might involve calculating 128 independent transforms of length 128. This is easier to vectorize than the single transform case. Indeed, in Figure 21 a variant of STEP computes NT repetitions of step I + 1 for NT transforms of length N. The important thing to notice is that the inner loop has a fixed repetition count, namely NT. The importance of this

concept, that is, of using the inner loop to run **over** the independent cases, cannot be overstated for vector computing.

In Figure 21, the subscripts ranged by L treat the rows of  $X(NT, N)$  as independent. Similarly, to transform the columns independently, the (L, J)s must everywhere be turned into (J, L)s. Padding the leading dimensions of X and Y by one row will avoid memory bank conflicts when the leading dimension of X and Y is a multiple of 8.

Although we have only discussed *radix 2* transforms (N is a power of 2), it turns out that other radix transforms have

```

      SUBROUTINE CF2T(N,X,Y,W)
      COMPLEX X(N),Y(N),W(N/2)
      LOGICAL ITGLE
C   M is LOG2(N), computed using the leading zero count
      M = 63 - LEADZ(N)
      I2 = 1
      ITGLE = .TRUE.
C
      CALL STEP(N,I2,X(1),X(N/2+1),Y(1),Y(I2+1),W)
      DO 1 I = 1,M-2
        I2 = 2*I
        IF(ITGLE)THEN
          CALL STEP(N,I2,Y(1),Y(N/2+1),X(1),X(I2+1),W)
          ITGLE = .FALSE.
        ELSE
          CALL STEP(N,I2,X(1),X(N/2+1),Y(1),Y(I2+1),W)
          ITGLE = .TRUE.
        ENDIF
      1 CONTINUE
C
      I2 = N/2
      IF(ITGLE)THEN
        CALL STEP(N,I2,Y(1),Y(N/2+1),Y(1),Y(I2+1),W)
      ELSE
        CALL STEP(N,I2,X(1),X(N/2+1),Y(1),Y(I2+1),W)
      ENDIF
C
      RETURN
      END

      SUBROUTINE STEP(N,I2,A,B,C,D,W)
      COMPLEX A(1),B(1),C(1),D(1),W(1)
C
      I3 = 2*I2
      MI2 = N/I3
      DO 1 J = 0,MI2-1
        JW = 1 + J*I2
        JA = JW
        JB = JA
        JC = 1 + J*I3
        JD = JC
C
        DO 2 K = 0,I2-1
          C(JC + K) = A(JA + K) + B(JB - K)
          D(JD + K) = W(JW)*(A(JA + K) - B(JB + K))
        2 CONTINUE
C
      1 CONTINUE
      RETURN
      END

```

FIGURE 18. FFT using Eqs. 5 and 6.

```

SUBROUTINE STEP(N,I2,A,B,C,D,W)
COMPLEX APB,AHB,A(1),B(1),C(1),D(1),W(1)
C
I3 = 2*I2
MI2 = N/I3
C
DO 2 K = 1,I2
C
DO 1 J = 0,MI2-1
APB = A(K + J*I2) + B(K + J*I2)
AHB = A(K + J*I2) - B(K + J*I2)
C(K + J*I3) = APB
D(K + J*I3) = W(K + J*I2)*AHB
1 CONTINUE
C
2 CONTINUE
RETURN
END

```

FIGURE 19. Alternate Version of STEP.




FIGURE 20. Execution Time vs. Crossover.

```

SUBROUTINE STEP(NT,N,I2,A,B,C,D,W)
COMPLEX A(NT,1),B(NT,1),C(NT,1),D(NT,1),W(1)
C
I3 = 2*I2
MI2 = N/I3
C
DO 1 J = 0,MI2-1
J2 = J*I2
J3 = 2*J2
J4 = 1 + J2
C
DO 2 K = 1,I2
C
DO 3 L = 1,NT
C(L,J3 + K) = A(L,J2 + K) + B(L,J2 + K)
D(L,J3 + K) = W(J4)*(A(L,J2 + K) - B(L,J2 + K))
3 CONTINUE
C
2 CONTINUE
C
1 CONTINUE
RETURN
END

```

FIGURE 21. STEP for NT Transforms of Length N.

some advantages. In part because of memory bank conflicts, binary radix transforms are less efficient than those of radix 3 or 5 (see [18]), which in some cases have nearly twice the processing rates.

#### 4.5 Getting Around Dependencies

When an updated element of an array is required for the subsequent calculation of an element in the same array, the elements may not be treated independently. In CRAY-1 documentation [2], this situation is regarded as a "dependency." In fact, because of the register architecture, two types of dependencies exist. In Figure 22 both types are illustrated.

In the top example, if  $I_0 > 0$ , then  $A(I)$  requires the previous value  $A(I - I_0)$  to have been reset, as in Figure 2. Hence, there is a sequential ordering, and  $A$  cannot be regarded as a vector with independent elements. At compile time CFT cannot determine whether  $I_0$  is negative and inhibits generation of vector hardware instructions.

Storage of  $B$  in the second part of Figure 22 may overlap if  $I_1 < I_2$ . This kind of dependency is somewhat more subtle, a result of the register architecture of CRAY-1. For example, try setting  $I_1 = 0$ ,  $I_2 = I$ ,  $I_L = 1$ ,  $I_U = 3$ . The storage of a segment  $\{B(2), B(3), B(4)\} = \{2., 2., 2.\}$  over the segment  $\{B(1), B(2), B(3)\} = \{1., 1., 1.\}$  would give  $B = \{1., 2., 2., 2.\}$ , which is not the same as the desired  $B = \{1., 1., 1., 2.\}$ . At compile time, if  $I_1$  and  $I_2$  are not known, CFT flags this as a dependency case. In the following, examples are given that illustrate some ways to sidestep these nonvectorizing dependencies.

**4.5.1 Gauss-Seidel Relaxation.** It is sometimes possible to find directions or subsets of multiply dimensioned arrays in which all the elements along those rays may be treated as independent. For example, all the elements in one column can be regarded as independent of elements in other columns. A simple relaxation step on the interior points of a rectangular grid is shown in Figure 23.

In this figure, computation of the  $J$ th element in row  $I$  depends on the updated value of the  $(J - 1)$  element, inhibiting vectorization. Just drawing a picture of the grid of  $I, J$  elements shows that diagonals depend only on the updated elements of previous (lower) diagonals. Each of these diagonals may be treated as independent, as in Figure 24.

Scanning by diagonals has the disadvantage that the vector length in the inner loop keeps changing. For very small iteration counts (1 or 2), the overhead to fill the functional unit pipelines is appreciable. Nevertheless, Figure 24 runs five times faster than Figure 23 if  $N, M > 100$ .

**4.5.2 A red-black ordering.** A method asymptotically equivalent to the above is a red-black ordering [19]. Figure 25 illustrates a simple red-black relaxation step, which has two additional features: use of the IVDEP directive, and parsing to minimize memory fetches. Within the DO 2 and DO 4 inner loops, the last pair of points in each equation is shared, eliminating two additional fetches. This example runs at 70 million floating point operations/second. A relatively easy modification of Figure 25 for Poisson's equation (not Laplace's equation, as above) performs floating point operations faster than the 80 megaHertz clock. This modification requires a relaxation parameter  $\omega \neq 1$  and a "source" term. Because of the additional computations, the operation rate actually goes up.

**4.5.3 Tridiagonal systems.** In both the forward elimination and back-substitution steps in the solution of a tridiagonal linear system, each element depends on its predecessors. To solve a single tridiagonal linear system by a parallel or vector

```

DO 1 I = IL, IU
C  DEPENDENCE ON A PREVIOUS ELEMENT
   A(I) = A(I - I0)
C  OVERLAPPING STORAGE DEPENDENCY
   B(I1 + I) = 1.
   B(I2 + I) = 2.
1 CONTINUE

```

FIGURE 22. Non-Vectorizing Dependencies.

```

SUBROUTINE GS(U, M, N)
REAL U(M, N)
DO 1 I = 2, M-1
DO 1 J = 2, N-1
   U(I, J) = .25*(U(I-1, J) + U(I+1, J) +
&              U(I, J-1) + U(I, J+1))
1 CONTINUE
RETURN
END

```

FIGURE 23. Gauss-Seidel Relaxation Step.

```

SUBROUTINE VGS(U, M, N)
REAL U(M-1, 1)
DO 1 I = 4, M+N-2
   L = I-M
   DO 1 K = MAX0(2, I-M+1)+1, MIN0(N-1, I-2)+1
      U(L, K) = .25*(U(L-1, K) + U(L+1, K) +
&                 U(L-M, K) + U(L+M, K))
1 CONTINUE
RETURN
END

```

FIGURE 24. Vector Gauss-Seidel Relaxation Step.

algorithm represents a problem of some difficulty. There are several approaches: Buneman's variant of cyclic reduction [20], Stone's recursive doubling [21], and a parallel Cramer's rule method by Swarztrauber [22]. The fastest of these methods on CRAY-1 is cyclic reduction [23]. Unfortunately, each method relies on a recursive doubling of step size in memory, while halving the vector length of each operation, and vice versa. Because of overhead introduced by short vectors, and memory bank conflicts inherent at each stage, even cyclic reduction runs only twice as fast as a simpler scalar method [24] for very large problems. In fact, a simple scalar tridiagonal solver is faster than cyclic reduction for solution vectors of length less than 63. Since this is not a delightful result, let us sidestep the issue and do another problem.

Multiple tridiagonal systems are much more tractable. Block tridiagonal equations, many line-relaxation methods, and three dimensional problems usually need solutions to many totally independent (unrelated) tridiagonal systems. So, now the inner loops can be made to range over each unrelated system in turn.



Once again, it must be emphasized that using inner loops to range over identical operations on independent data or systems is the key to successful vector processing. Indeed, using the above procedure to solve 100 tridiagonal systems of length 100 is 13 times faster than solving 100 one at a time. If each system has a different matrix, Figure 26 is easily modified for this purpose.

#### 4.6 Sparse Matrices—Nonlinear Indexing

Most general purpose sparse matrix solvers [25] keep pointers to nonzero elements. Allocating storage only for nonzero elements and potential fill-in minimizes both memory requirements and the number of null operations. All six [26] forms of Gaussian elimination require a compression and decompression of rows or columns into indexed lists as the elimination proceeds. For example, the most efficient form of elimination [27] has the following reduction step on the working row (Figure 27):

Here, Y is the working row, X is any lower packed row with INDEX an array of pointers to the positions of X in its expanded form. For an arbitrary sparse matrix, INDEX will not point to regularly spaced elements.

This is a difficult problem with no vector hardware solution. At present, only CAL-coded modules can manage to approach vector-mode floating point operation rates. The operations in Figure 28 are available for sparse matrices on CRAY-1 with the 1.11 CFT software release. Respectively, the calling sequences for these modules are as shown in Figure 29. For  $N > 5$ , all these subprograms are more efficient than in-line FORTRAN. For large N these modules give a factor of 3 timing improvement.

```
DO 1 I = IL, IU
  Y(INDEX(I)) = A*X(I) + Y(INDEX(I))
1 CONTINUE
```

FIGURE 27. Sparse:  $= ax + y$  Operations.

```
DO 1 I = 1, N
C SCATTER OPERATION
  Y(INDEX(I)) = X(I)
C GATHER OPERATION
  Y(I) = X(INDEX(I))
C SPAXPY OPERATION
  Y(INDEX(I)) = A*X(I) + Y(INDEX(I))
C SPARSE DOT-PRODUCT
  SPDOT = SPDOT + Y(INDEX(I))*X(I)
1 CONTINUE
```

FIGURE 28. Sparse Operations.

```
CALL SCATTER(N, Y, INDEX, X)
CALL GATHER(N, Y, X, INDEX)
CALL SPAXPY(N, A, X, Y, INDEX)
DOT = SPDOT(N, Y, INDEX, X)
```

FIGURE 29. CAL Sparse Operations.

#### REFERENCES

1. Croy Hordwore Reference Manual. 15 May, 1979. Cray Research Inc. Publ. 2240004. Cray Research Inc., 1440 Northland Dr., Mendota Hts Minn.
2. Croy Fortron (CFT) Reference Manual. May 1980 Cray Research Inc. Publ. SR-0009, Version 1.10, Revision H (ANSI standard).
3. Higbie, L.C. Speeding up FORTRAN (CFT) programs on the CRAY-1 Cray Research Inc. Tech. Note 2240207. 1978.
4. Johnson, P.M. An introduction to vector processing. Comput. Des. (Feb. 1978). 000-000
5. CAL Assembler Reference Manual. Cray Research Inc. Publ. SR-0000, (Cray Assembly Language).
6. Petersen, W. MXM—unit spaced fast matrix multiply. Cray Research Inc. Tech. Note SN-0213, Dec. 1980.
7. Fong, K. and Jordan, T.L. Some linear algebraic algorithms and their performance on CRAY-1. Los Alamos Scientific Laboratory Rep. LA-6774 (Univ. of California Rep. No. UC-32). June 1977.
8. Levinson, N. J. Math. Phys. 25, 4 (Jan. 1947). 261-278.
9. Brent, R.P., Gustavson, F.G., and Yun, D.Y.Y. Fast computation of Padé approximants and the solution of Toeplitz systems of equations IBM Res. Rep. RC 8173 (No. 34952), IBM Research Center, Yorktown Heights, N. Y., Jan. 1980.
10. Bitmead, R.R. and Anderson, B.D.O. Asymptotically fast solution of Toeplitz and related systems of equations. Linear Algebra Appl. 34 (1980). 103-116.
11. Floyd, R.W. Shortest path, Algorithm 97. Commun. ACM 5 (1962), 345.
12. Cooley, J.W., and Tukey, J.W. An algorithm for machine calculation of complex Fourier series. Math. Comput. 19 (April 1965). 297-301.
13. Temperton, C. Mixed-radix fast Fourier transforms without reordering. European Centre for Medium Range Weather Forecasts, Tech. Rep. No. 3, Feb. 1977.
14. Swartztrauber, P.N. Vectorizing the FFTs. In Parallel Computations. G. Rodrigue, Ed., Academic Press, N. Y., 1982.
15. Cochran, W.T., et al. What is Fast Fourier Transform? IEEE Trans. Audio Electroacoustics AU-15, 2 (June 1967). 45-55.
16. Uhrich, M.L. Fast Fourier transforms without sorting. IEEE Trans. Audio Electroacoustics AU-17 (1969), 170-172.
17. Petersen, W. CFFT2—complex fast Fourier transform binary radix subroutine. Cray Research Inc. Tech. Note 2240203, March 1978.
18. Temperton, C. Fast Fourier transforms on CRAY-1. European Centre for Medium Range Weather Forecasts, Tech. Rep. 21, Jan. 1979.
19. Young, D.M. Iterative Solutions of Large Linear Systems. Academic Press, New York, 1971.
20. Buneman, O. A compact noniterative Poisson solver. Stanford Univ. Institute for Plasma Research. Rep. 294, 1969.
21. Stone, H.S. An efficient parallel algorithm for the solution of a tridiagonal system of equations. J. ACM 20 (1973) 27-38.
22. Swartztrauber, P.N. A parallel algorithm for solving general tridiagonal equations. Math. Comput. 33, 145 (Jan. 1979). 185-199.
23. Swartztrauber, P.N. The parallel solution of tridiagonal systems on Cray-1. Infotech State of the Art Report on Supercomputers, Infotech Int. Ltd., Maidenhead, Berkshire, UK, 1979.
24. Penumalli, R. Private communication.
25. Eisenstat, S.C., Schultz, M.H., and Sherman, A.H. Considerations in design of software for sparse matrix computation. In Sparse Matrix Computotron, J.R. Bunch and D.J. Rose, Eds., Academic Press, N.Y., 1976.
26. Dembart, B., and Nevas, K. Sparse triangular factorization on vector computers. In Exploring Applications of Parallel Processing to Power System Analysis, Electric Power Res. Institute Rep. EL-566-SR, Oct. 1977.
27. Dodson, D.S., and Petersen, W. Sparse triangular factorization on CRAY-1, Cray Research Inc. Tech. Note SN-0217, June 1981.

CR Categories and Subject Descriptors: D.3.2 [Programming Languages]: Language Classifications—FORTRAN D.3.3 [Programming Languages]: Language Constructs—control structures: G.1.0. [Numerical Analysis]: General  
General Terms: Algorithms, Design, Languages

Received 6/83; accepted 7/83