


**A Comparative Study of  
Automatic Vectorizing Compilers**

*David Levine*  
*David Callahan*  
*Jack Dongarra*

**CRPC-TR91138**  
**1991**

Center for Research on Parallel Computation  
Rice University  
P.O. Box 1892  
Houston, TX 77251-1892



# A Comparative Study of Automatic Vectorizing Compilers

David Levine \*      David Callahan †      Jack Dongarra ‡

**Abstract.** We compare the capabilities of several commercially available, vectorizing Fortran compilers using a test suite of Fortran loops. We present the results of compiling and executing these loops on a variety of supercomputers, mini-supercomputers, and mainframes.

## 1 Introduction

This paper describes the use of a collection of Fortran loops to test the analysis capabilities of automatic vectorizing compilers. An automatic vectorizing compiler is one that takes code written in a serial language (usually Fortran) and translates it into vector instructions. The vector instructions may be machine specific or in a source form such as the proposed Fortran 90 array extensions or as subroutine calls to a vector library.

Most of the loops in the test suite were written by people involved in the development of vectorizing compilers, although several we wrote ourselves. All of the loops test a compiler for a specific feature. These loops reflect constructs whose vectorization ranges from easy to challenging to extremely difficult. We have collected the results from compiling and executing these loops using commercially available, vectorizing Fortran compilers.

The results reported here expand on our earlier work [3]. In that paper, we focused principally on analyzing each compiler's output listing. For the present study, we ran the loops in both scalar and vector modes. In addition, the set of loops has been expanded.

The remainder of this paper is organized into eight sections. Section 2 describes our classification scheme for the loops used in the test. In Section 3 we describe the structure of the test program. In Section 4 we describe the methodology used to perform the test. Section 5 reports on the number of loops that vectorized according to the compiler's output listing. Section 6 presents two aspects of the speedup results. In Section 7 we discuss our model of optimal vector performance and present the results of comparing the actual performance with the model. Section 8 discusses several aspects of the test. In Section 9 we make some remarks about future work.

---

\*Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439-4801. This work was supported by the Applied Mathematical Sciences subprogram of the Office of Energy Research, U. S. Department of Energy, under Contract W-31-109-Eng-38.

†Tera Computer, 400 North 34th Street, Suite 300, Seattle, WA 98103

‡Computer Science Department, University of Tennessee, Knoxville, TN 37996-1301. This work was supported in part by the Applied Mathematical Sciences Research Program, Office of Energy Research, U. S. Department of Energy, under Contract DE-AC05-84OR21400 and in part by NSF cooperative agreement CCR-8809615.

## 2 Classification of Loops

The objective of the test suite is to test four broad areas of a vectorizing compiler: dependence analysis, vectorization, idiom recognition, and language completeness. All of the loops in this suite are classified into one of these categories.

We define all terms and transformation names but discuss dependence analysis and program transformation only briefly. Recent discussions of these topics can be found in Allen and Kennedy [2], Padua and Wolfe [11], and Wolfe [14]. For a practical exposition of the application of these techniques, see Levesque and Williamson [6].

### 2.1 Dependence Analysis

Dependence analysis comprises two areas: global data-flow analysis and dependence testing. Global data-flow analysis refers to the process of collecting information about array subscripts. Dependence testing refers to the process of testing for memory overlaps between pairs of variables in the context of the global data-flow information.

Dependence analysis is the heart of vectorization, but it can be done with very different levels of sophistication ranging from simple pattern matching to complicated procedures that solve systems of linear equations. Many of the loops in this section test the aggressiveness of the compiler in normalizing subscript expressions into linear form for the purpose of enhanced dependence testing.

1. *Linear Dependence Testing.* Given a pair of array references whose subscripts are linear functions of the loop control variables that enclose the references, decide whether the two references ever access the same memory location. When the references do interact, additional information can be derived to establish the safety of loop restructuring transformations.
2. *Induction Variable Recognition.* Recognize auxiliary induction variables (e.g., variables defined by statements such as  $K=K+1$  inside the loop). Once recognized, occurrences of the induction variable can be replaced with expressions involving loop control variables and loop invariant expressions.
3. *Global Data-Flow Analysis.* Collect global (entire subroutine) data-flow information, such as constant propagation or linear relationships among variables, to improve the precision of dependence testing.
4. *Nonlinear Dependence Testing.* Given a pair of array references whose subscripts are not linear functions, test for the existence of data dependencies and other information.
5. *Interprocedural Data-Flow Analysis.* Use the context of a subroutine in a particular program to improve vectorization. Possibilities include in-line expansion, summary information (e.g., which variables may or must be modified by an external routine), and interprocedural constant propagation.
6. *Control Flow.* Test to see whether certain vectorization hazards exist and whether there are implied dependencies of a statement on statements that control its execution.

7. *Symbolics*. Test to see whether subscripts are linear after certain symbolic information is factored out or whether the results of dependence testing do not, in fact, depend on the value of symbolic variables.

## 2.2 Vectorization

A simple vectorizer would recognize single-statement Fortran DO loops that are equivalent to hardware vector instructions. When this strict syntactic requirement is not satisfied, more sophisticated vectorizers can restructure programs so that it is. Here, program restructuring is divided into two categories: transformations to enhance vectorization and idiom recognition. The first is described here, and the other in the next section.

1. *Statement Reordering*. Reorder statements in a loop body to allow vectorization.
2. *Loop Distribution*. Split a loop into two or more loops to allow partial vectorization or more effective vectorization.
3. *Loop Interchange*. Change the order of loops in a loop nest to allow or improve vectorization. In particular, make a vectorizable outer loop the innermost in the loop nest.
4. *Node Splitting*. Break up a statement within a loop to allow (partial) vectorization.
5. *Scalar and Array Expansion*. Expand a scalar into an array or an array into a higher-dimensional array to allow vectorization and loop distribution.
6. *Scalar Renaming*. Rename instances of a scalar variable. Scalar renaming eliminates some interactions that exist only because of reuse of a temporary variable and allows more effective scalar expansion and loop distribution.
7. *Control Flow*. Convert forward branching in a loop into masked vector operations; recognize loop invariant IF's (loop unswitching).
8. *Crossing Thresholds (Index Set Splitting)*. Allow vectorization by blocking into two sets. For example, vectorize the statement  $A(I) = A(N-I)$  by splitting iterations of the I loop into iterations with I less than  $N/2$  and iterations with I greater than  $N/2$ .
9. *Loop Peeling*. Unroll the first or last iteration of a loop to eliminate anomalies in control flow or attributes of scalar variables.
10. *Diagonals*. Vectorize diagonal accesses (e.g.,  $A(I,I)$ ).
11. *Wavefronts*. Vectorize two-dimensional loops with dependencies in both dimensions by restructuring the loop for diagonal access.

## 2.3 Idiom Recognition

Idiom recognition refers to the identification of particular program forms that have (presumably faster) special implementations.

1. *Reductions.* Compute a scalar value or values from a vector, such as sum reductions, min/max reductions, dot products, and product reductions.
2. *Recurrences.* Identify special first- and second-order recurrences that have logarithmically faster solutions or hardware support.
3. *Search Loops.* Search for the first or last instance of a condition, possibly saving index value(s).
4. *Packing.* Scatter or gather a sparse vector from or into a dense vector under the control of a bit mask or an indirection vector.
5. *Loop Rerolling.* Vectorize loops where the inner loop has been unrolled.

## 2.4 Language Completeness

This section tests how effectively the compilers understand the complete Fortran language. Simple vectorizers might limit analysis to DO loops containing only floating point and integer assignments. More sophisticated compilers will analyze all loops and vectorize wherever possible.

1. *Loop Recognition.* Recognize and vectorize loops formed by backward GO TO's.
2. *Storage Classes and Equivalencing.* Understand the scope of local vs. common storage; correctly handle equivalencing.
3. *Parameters.* Analyze symbolic named constants, and vectorize statements that refer to them.
4. *Nonlogical IF's.* Vectorize loops containing computed GO TO's and arithmetic IF's.
5. *Intrinsic Functions.* Vectorize functions that have elemental (vector) versions such as SIN and COS or known side effects.
6. *Call Statements.* Vectorize statements in loops that contain CALL statements or external function invocations.
7. *Nonlocal GO TO's.* Branches out of loops, RETURN statements or STOP statements inside of loops.
8. *Vector Semantics.* Load before store, and preserve order of stores.
9. *Indirect Addressing.* Vectorize subscripted subscript references (e.g, A(INDEX(I))) as Gather/Scatter.
10. *Statement Functions.* Vectorize statements that refer to Fortran statement functions.

### 3 Test Program Structure

The test program consists of 122 loops that represent different constructs intended to test the analysis capabilities of a vectorizing compiler. Using the classification scheme in Section 2, there are 29 loops in the Dependence Analysis category, 41 loops in the Vectorization category, 24 loops in the Idiom Recognition category, and 28 loops in the Language Completeness category. Also included are 13 additional "control" loops we expect all compilers to be able to vectorize. These allow us to measure the rates of certain basic operations for use with the model discussed in Section 7.

The majority of the test loops operate on one-dimensional arrays; a small number operate on two-dimensional arrays. Most of the loops in the test are fairly short; many are a single statement and others usually no more than several statements. Many of the loops access memory with a stride of one. Each loop is contained in a separate subroutine. A driver routine calls each subroutine with vector lengths of 10, 100, and 1000.

An example loop is shown in Figure 1. Relevant operands are initialized once at the start of the loop. An outer repetition loop is used to increase the granularity of the calculation, thereby avoiding problems with clock resolution. A call to a dummy subroutine is included in each iteration of the repetition loop so that, in cases where the inner loop calculation is invariant with respect to the repetition loop, the compiler is still required to execute each iteration rather than just recognizing that the calculation needs to be done only once.

After execution of the loop is complete, a checksum is computed by using the result array(s). The checksum and the time used are then passed to a check subroutine. The check subroutine verifies the checksum with a precomputed result and prints out the time to execute the loop. The time is calculated by calling a timer at the start of the loop and again at the end of the loop and taking the difference of these times minus the cost of the timing call and the cost of the multiple calls to the dummy subroutine.

### 4 Test Methodology

The test program is distributed in two files: a driver program in one file, and the test loops in the other. The files were distributed to interested vendors, who were asked to compile the loops without making any changes\* using only the compiler options for automatic vectorization. Thus, the use of compiler directives or interactive compilation features to gain additional vectorization was not tested. Vendors were asked to make two separate runs of the test: one using scalar optimizations only, and the other using the same scalar optimizations and, in addition, all automatic vectorization options. Vendors with multiprocessor computers submitted uniprocessor results only. Appendix A contains details of the exact machine configurations and versions of the software used.

The rules require separate compilation of the two files. The rules for compilation of the

---

\*One vendor was allowed to (1) separate a 135-way IF-THEN-ELSEIF-ELSE construct in order to overcome a self-imposed limit, and (2) include the array declarations in a common block in the driver program (only) in order to overcome a self-imposed limit on memory allocation size. Neither modification had any impact on performance.

```

subroutine s111 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
integer ntimes, ld, n, i, nl
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, cs1d
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s111 ')
t1 = second()
do 1 nl = 1,2*ntimes
do 10 i = 2,n,2
    a(i) = a(i-1) + b(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(2*ntimes) )
chksum = cs1d(n,a)
call check (chksum,2*ntimes*(n/2),n,t2,'s111 ')
return
end

```

Figure 1: Example loop

driver file require that no compiler optimizations be used and that the file not be analyzed interprocedurally to gather information useful in optimizing the test loops.

The file containing the loops was compiled twice—once for the scalar run and once for the vector run. For the scalar run, global (scalar) optimizations were used. For the vector run, in addition to the same global optimizations specified in the scalar run, vectorization and — if available — automatic call generation to optimized library routines, function inlining, and interprocedural analysis were used.

All files were compiled to use 64-bit arithmetic. Most runs were made on standalone systems.\* For virtual memory computers, the runs were made with a physical memory and working-set size large enough that any performance degradation from page faults was negligible. In all cases the times reported to us were *user* CPU time.

After compiling and executing the loops, the vendors sent back the compiler's output listing (source echo, diagnostics, and messages) and the output of both the scalar and vector runs. We then examined the compiler's output listings to see which loops had been vectorized, and analyzed the scalar and vector results. In addition to measuring the execution time of the loops, we checked the numerical result in order to verify correctness. However, the check was strictly for correctness of the numerical result; no attempt was made to see whether possibly unsafe transformations had been used.

---

\*The Cray Computer and Hitachi runs were not.



## 5 Number of Loops Vectorized

In this section we discuss the number of loops that were vectorized, as reported by the compiler's output listing. All of the loops in our test are amenable to some degree of vectorization. For some loops, this may only be partial vectorization; for others, vectorization may require the use of optimized library routines or special hardware.

### 5.1 Definition of Vectorization

We define a *statement* as vectorizable if one or more of the *expressions* in the statement involve array references or may be converted to that form. We define three possible results for a compiler attempting to vectorize a loop. A *loop* is *vectorized* if the compiler generates vector instructions for all vectorizable statements in the loop. A loop is *partially vectorized* if the compiler generates vector instructions for some, but not all, vectorizable statements in the loop. No threshold is defined for what percentage of a loop needs to be vectorized to be listed in this category, only that some expression in a statement in the loop is vectorized. A loop is *not vectorized* if the compiler does not generate vector instructions for any vectorizable statements within the loop.

For some loops the Cray Research, FPS Computing, IBM, and NEC compilers generated a runtime IF-THEN-ELSE test which executed either a scalar loop or a vectorized loop. These loops have been scored as either *vectorized* or *not vectorized* according to whether or not vectorized code was actually executed at runtime.

The Cray Computer compiler "conditionally vectorized" certain loops. That is, for loops with ambiguous subscripts, a runtime test was compiled that selected a safe vector length.\* These loops have been scored as either *vectorized* if the safe vector length was greater than one, otherwise *not vectorized*.

For a number of loops, the Fujitsu compiler generated scalar code even though the compiler indicated that partial vector code could be generated. In these cases, the compiler listing contained the message "Partial vectorization overhead is too large," indicating that although partial vectorization was possible, for these loops the compiler considered scalar code more efficient. These loops have been scored as *partially vectorized*.

Our definition of vectorization counts as vectorized those loops that are recognized by the compiler and *automatically* replaced by calls to optimized library routines. In some cases a compiler may generate a call to an optimized library routine rather than explicitly generating vector code. Typical examples are for certain reduction and recurrence loops. Often the library routines use a mix of scalar and vector instructions; while perhaps not as fast as pure vector loops, since the construct itself is not fully parallel, they are usually faster than scalar execution. In all cases where the compiler automatically generated a call to a library routine, we have scored the loop as *vectorized*.

---

\*A safe vector length is one that allows the compiler to execute vector instructions and still produce the correct result. For example, the statement  $A(I)=A(I-7)$  with loop increment one may be executed in vector mode with any vector length less than or equal to 7.

Table 1: Full Vectorization (122 loops)

Computer	V	P	N
CONVEX C210	68.0 (83)	10.7 (13)	21.3 (26)
CCC CRAY-2	60.7 (74)	1.6 (2)	37.7 (46)
CRI CRAY Y-MP	77.9 (95)	8.2 (10)	13.9 (17)
DEC VAX 9000-210	60.7 (74)	3.3 (4)	36.1 (44)
FPS M511EA-2	72.1 (88)	4.9 (6)	23.0 (28)
Fujitsu VP2600/10	71.3 (87)	16.4 (20)	12.3 (15)
Hitachi S-820/80	71.3 (87)	7.4 (9)	21.3 (26)
IBM 3090-600J	77.9 (95)	4.9 (6)	17.2 (21)
NEC SX-X/14	72.1 (88)	5.7 (7)	22.1 (27)
Average	70.2 (85)	7.0 (8)	22.8 (27)

V - vectorized  
P - partially vectorized  
N - not vectorized  
V/P - fully or partially vectorized

Figure 2: Key to symbols for Tables 1-8, 12-13

## 5.2 Results

Tables 1-6 list the results of analyzing the compilers' listings. Each table contains the percentage of loops in each column, followed by the actual number in parentheses. Table 1 summarizes the results for all 122 loops. Table 2 is also a summary of all the loops; here, however, the column V/P counts the loops that were either fully or partially vectorized. Tables 3-6 contain results by category as defined in Section 2.

Table 2: Full and Partial Vectorization (122 loops)

Computer	V/P	N
CONVEX C210	78.7 (96)	21.3 (26)
CCC CRAY-2	62.3 (76)	37.7 (46)
CRI CRAY Y-MP	86.1 (105)	13.9 (17)
DEC VAX 9000-210	63.9 (78)	36.1 (44)
FPS M511EA-2	77.0 (94)	23.0 (28)
Fujitsu VP2600/10	87.7 (107)	12.3 (15)
Hitachi S-820/80	78.7 (96)	21.3 (26)
IBM 3090-600J	82.8 (101)	17.2 (21)
NEC SX-X/14	77.9 (95)	22.1 (27)
Average	77.2 (94)	22.8 (27)

Table 3: Dependence Analysis (29 loops)

Computer	V	P	N
CONVEX C210	65.5 (19)	17.2 (5)	17.2 (5)
CCC CRAY-2	69.0 (20)	0.0 (0)	31.0 (9)
CRI CRAY Y-MP	86.2 (25)	0.0 (0)	13.8 (4)
DEC VAX 9000-210	69.0 (20)	0.0 (0)	31.0 (9)
FPS M511EA-2	82.8 (24)	0.0 (0)	17.2 (5)
Fujitsu VP2600/10	65.5 (19)	24.1 (7)	10.3 (3)
Hitachi S-820/80	55.2 (16)	10.3 (3)	34.5 (10)
IBM 3090-600J	86.2 (25)	0.0 (0)	13.8 (4)
NEC SX-X/14	75.9 (22)	6.9 (2)	17.2 (5)
Average	72.8 (21)	6.5 (1)	20.7 (6)

Table 4: Vectorization (41 loops)

Computer	V	P	N
CONVEX C210	73.2 (30)	14.6 (6)	12.2 (5)
CCC CRAY-2	34.1 (14)	4.9 (2)	61.0 (25)
CRI CRAY Y-MP	56.1 (23)	22.0 (9)	22.0 (9)
DEC VAX 9000-210	58.5 (24)	7.3 (3)	34.1 (14)
FPS M511EA-2	61.0 (25)	14.6 (6)	24.4 (10)
Fujitsu VP2600/10	68.3 (28)	24.4 (10)	7.3 (3)
Hitachi S-820/80	78.0 (32)	9.8 (4)	12.2 (5)
IBM 3090-600J	75.6 (31)	12.2 (5)	12.2 (5)
NEC SX-X/14	65.9 (27)	12.2 (5)	22.0 (9)
Average	63.4 (26)	13.6 (5)	23.0 (9)

Table 5: Idiom Recognition (24 loops)

Computer	V	P	N
CONVEX C210	66.7 (16)	4.2 (1)	29.2 (7)
CCC CRAY-2	70.8 (17)	0.0 (0)	29.2 (7)
CRI CRAY Y-MP	87.5 (21)	4.2 (1)	8.3 (2)
DEC VAX 9000-210	54.2 (13)	4.2 (1)	41.7 (10)
FPS M511EA-2	70.8 (17)	0.0 (0)	29.2 (7)
Fujitsu VP2600/10	87.5 (21)	8.3 (2)	4.2 (1)
Hitachi S-820/80	91.7 (22)	4.2 (1)	4.2 (1)
IBM 3090-600J	58.3 (14)	0.0 (0)	41.7 (10)
NEC SX-X/14	87.5 (21)	0.0 (0)	12.5 (3)
Average	75.0 (18)	2.8 (0)	22.2 (5)

Table 6: Language Completeness (28 loops)

Computer	V	P	N
CONVEX C210	64.3 (18)	3.6 (1)	32.1 (9)
CCC CRAY-2	82.1 (23)	0.0 (0)	17.9 (5)
CRI CRAY Y-MP	92.9 (26)	0.0 (0)	7.1 (2)
DEC VAX 9000-210	60.7 (17)	0.0 (0)	39.3 (11)
FPS M511EA-2	78.6 (22)	0.0 (0)	21.4 (6)
Fujitsu VP2600/10	67.9 (19)	3.6 (1)	28.6 (8)
Hitachi S-820/80	60.7 (17)	3.6 (1)	35.7 (10)
IBM 3090-600J	89.3 (25)	3.6 (1)	7.1 (2)
NEC SX-X/14	64.3 (18)	0.0 (0)	35.7 (10)
Average	73.4 (20)	1.6 (0)	25.0 (7)

### 5.3 Analysis of Results

The average number of loops vectorized was 70%, and vectorized or partially vectorized was 77%. The best results were 78% and 88%, respectively. Of the 122 loops, only two were not vectorized or partially vectorized by any of the compilers; both loops are vectorizable. There is probably no significant difference between compilers within a few percent of each other. Slight differences may be due to different hardware, the availability of special software libraries, the architecture of a machine being better suited to executing scalar or parallel code for certain constructs, or the makeup of the loops used in our test.

From Table 1 we see that the Cray Research and IBM compilers vectorized the most loops. A large number of other compilers are grouped closely together and only a few loops behind these two. Comparing Table 1 to Table 2, we see that counting partially vectorized loops in the totals allows the Fujitsu compiler to vectorize the most loops. It is interesting to note, however, that of the 20 loops we counted as partially vectorized by the Fujitsu compiler, only two actually resulted in (partial) vector code being executed at runtime. For the other 18 loops the Fujitsu compiler made the decision that it would not be cost effective to partially vectorize them. The Convex compiler also did a significant amount of partial vectorization.

Tables 3-6 show that some compilers did particularly well in certain categories. The Cray Research, FPS Computing, and IBM compilers had the best results in the Dependence Analysis category. The Convex, Hitachi, and IBM compilers had the best results in the Vectorization category. The Cray Research, Fujitsu, Hitachi, and NEC compilers had the best results in the Idiom Recognition category. In the Language Completeness category the Cray Research and IBM compilers had the best results. The Vectorization category seemed the most difficult, with approximately 10% fewer loops vectorized overall than for the other sections.

Certain sections seemed fairly easy, with most vendors vectorizing or partially vectorizing almost all of the loops. Using the classification scheme of Section 2 these sections were linear dependence testing, global data-flow analysis, statement reordering, loop distribution, node splitting, scalar renaming, control flow, diagonals, loop rerolling, parameters, intrinsic functions, indirect addressing, and statement functions.

In some sections, while many vendors vectorized or partially vectorized most loops, various

individual vendors did not do particularly well. These sections were induction variable recognition, interprocedural data-flow analysis, symbolics, scalar and array expansion, reductions, search loops, packing, and nonlogical IF'S.

Some sections were difficult for many compilers. Typically, at least half the vendors missed at least some, and sometimes most, of the loops in these sections. These sections were control flow, loop interchange, index set splitting, loop peeling, recurrences, loop recognition, storage classes and equivalencing, and nonlocal GO TO's.

A few sections were particularly difficult, with only one or two compilers doing any vectorization at all. These sections were nonlinear dependence testing, wavefronts, and call statements.

We found that some vendors with approximately equal results did much better in one section than another. Certain induction variable tests, interprocedural data-flow analysis, loop interchange, recurrences, loop recognition, storage classes and equivalence statements, and loops with exits were the sections that showed the greatest variation. We conclude that the compiler vendors have focused their efforts on particular subsets of the features tested by the suite. Possible reasons might include hardware differences or (self-imposed) limits on compilation time, compilation memory use, or the size of the generated code.

Individual results, on a loop-by-loop basis, may be found in Appendix B.

## 6 Speedup

The goal of vectorization is for the vectorized program to execute in less time than the unvectorized program. The metric used is the speedup,  $s_p$ , defined as  $s_p = t_s/t_v$ , where  $t_s$  is the scalar time and  $t_v$  is the vector time. In this section we look at two aspects of speedup. First, does the vector code run slower than the corresponding scalar code? Second, how large a speedup can be gained with vectorization?

### 6.1 Vectorized Loops Revisited

Ideally the speedup from vectorization (or partial vectorization) should be as large as possible. At a minimum, though the vector code should run at least as fast as the scalar code. However, this minimum is not always achieved, particularly at short vector lengths where there may not be enough work in the loop to overcome the vector startup cost.

Tables 7 and 8 revisit the results in Table 1. The number of loops in each of the different categories is again taken from the compiler listing. In Tables 7 and 8 however, we have not counted as vectorized or partially vectorized any loops where  $s_p < .95$ .\* The results in Table 7 are for vector length 100, and the results in Table 8 are for vector length 1000. We have not presented these results for vector length ten since almost all vendors suffer some performance degradation for short vectors.

The results in Tables 7 and 8 are mostly consistent with Table 1. Four of the compilers show

---

\*We use .95 instead of 1 to allow for the possibility of measurement error.

Table 7: Loops Vectorized ( $s_p > .95$ , Vector length = 100, 122 loops)

Computer	V	P	N
CONVEX C210	68.0 (83)	9.0 (11)	23.0 (28)
CCC CRAY-2	60.7 (74)	0.8 (1)	38.5 (47)
CRI CRAY Y-MP	77.9 (95)	7.4 (9)	14.8 (18)
DEC VAX 9000-210	49.2 (60)	2.5 (3)	48.4 (59)
FPS M511EA-2	71.3 (87)	3.3 (4)	25.4 (31)
Fujitsu VP2600/10	68.9 (84)	13.1 (16)	18.0 (22)
Hitachi S-820/80	69.7 (85)	1.6 (2)	28.7 (35)
IBM 3090-600J	71.3 (87)	4.1 (5)	24.6 (30)
NEC SX-X/14	72.1 (88)	2.5 (3)	25.4 (31)
Average	67.7 (82)	4.9 (6)	27.4 (33)

Table 8: Loops Vectorized ( $s_p > .95$ , Vector length = 1000, 122 loops)

Computer	V	P	N
CONVEX C210	68.0 (83)	8.2 (10)	23.8 (29)
CCC CRAY-2	60.7 (74)	0.8 (1)	38.5 (47)
CRI CRAY Y-MP	77.9 (95)	7.4 (9)	14.8 (18)
DEC VAX 9000-210	54.9 (67)	2.5 (3)	42.6 (52)
FPS M511EA-2	71.3 (87)	4.1 (5)	24.6 (30)
Fujitsu VP2600/10	69.7 (85)	15.6 (19)	14.8 (18)
Hitachi S-820/80	70.5 (86)	1.6 (2)	27.9 (34)
IBM 3090-600J	73.0 (89)	4.1 (5)	23.0 (28)
NEC SX-X/14	72.1 (88)	2.5 (3)	25.4 (31)
Average	68.7 (83)	5.2 (6)	26.1 (31)

no degradation on any of the vectorized loops. Three others show a degradation on only one or two loops. Only two compilers show a degradation on any significant number of loops. The results for partial vectorization are also fairly consistent with Table 1, with only one compiler showing any serious number of loops being degraded. There is a large variance in the test suite as to which loops have degraded performance. No particular trend is obvious.

Two compilers also suffered noticeable performance degradations (below 90%) for a significant number of loops (10 or more) that were *not vectorized*. We believe somehow that the attempt to vectorize interfered with the generation of good scalar code. We view this as a performance bug and have advised the vendors. Other than these cases, the vectorizers rarely generated code that was inferior to the scalar code on vector lengths of 100 or more. An exception is the nine loops the CRAY-2 compiler-generated vector code for with a safe vector length of one. These loops, although scored as *not vectorized*, had vector execution times that were frequently twice the scalar execution times.

## 6.2 Aggregate Speedup Results

The speedup that can be achieved on a particular vector computer depends on several factors: the speed of the vector hardware relative to the speed of the scalar hardware, the inherent

Table 9: Aggregate Speedup Results by Section (122 loops)

Section	10		100		1000		All VL	
Data Dependence	1.44	0.90	5.09	1.87	12.25	2.15	6.26	1.42
Vectorization	1.37	0.97	5.11	1.74	21.32	1.91	9.27	1.41
Idiom Recognition	0.95	0.73	3.66	1.73	11.52	2.29	5.38	1.26
Language Completeness	1.55	1.07	4.72	1.90	9.63	2.13	5.30	1.55
All Sections	1.35	0.92	4.73	1.80	14.55	2.08	8.61	1.41

Table 10: Aggregate Speedup Results by Type of Vectorization

Vectorization Level	10		100		1000		All VL	
Full Vectorization	1.54	0.96	6.29	3.11	20.20	4.48	9.34	1.90
Full or Partial Vectorization	1.47	0.93	5.85	2.61	18.56	3.48	8.63	1.72
Partial Vectorization	0.83	0.68	1.41	1.00	2.21	1.07	1.49	0.88

vector parallelism in the code of interest, and the sophistication of the compiler in detecting opportunities to generate code to run on the vector hardware. From the perspective of our test, we would like to measure the speedup achieved just from the compiler's vectorization capabilities. However, speedups are too strongly influenced by architecture and implementation to be meaningful indicators of compiler performance. Therefore, we prefer not to give speedup results for individual vendors which may be misinterpreted as representing compiler performance *only*. Instead, we present speedup statistics using the aggregate results from all vendors.

Table 9 presents a summary of the speedup results of all vendors. The first four rows present results according to the classification scheme in Section 2. Results are given for vector lengths of 10, 100, and 1000 and, in the last column, the sum over all three vector lengths. Each column contains the arithmetic and harmonic means of the speedups for the loops in that section. The results in the last row are summed over all four sections.

Table 10 contains aggregate statistics for three different levels of vectorization. The format of the table is similar to Table 9. The first row contains speedup statistics for the 771 loops scored as fully vectorized. The second row contains speedup statistics for the 848 loops scored as either fully or partially vectorized. The last row contains speedup statistics for the 77 loops that were partially vectorized.

### 6.3 Discussion of Speedup

As might be expected, at the relatively short vector length of 10, the speedups were not very large. This is particularly true of the Idiom Recognition section, where the methods used to vectorize some of the loops are not amenable to the full speedup that can be provided by the hardware. At vector length 100 most speedups were between three and six. At the longest vector length, 1000, the individual speedups were slightly higher for most. Three vendors however, had very large average speedups (29.4, 33.2, and 39.8) over the scalar speed.

The choice of mean clearly affects the results. In the Vectorization section, the arithmetic mean at vector length 1000 is 21.32, while the harmonic mean is only 1.91. These results show that a relatively small number of large speedups can greatly affect the arithmetic mean.

McMahon [10] and Smith [12] discuss the different means.

If we compare the last row of Table 9 with the first two rows in Table 10, we see better speedups at all vector lengths when we consider only the loops fully or partially vectorized. Of course which loops were included, and how many, varies for each vendor.

In several loops in the test suite, not all statements can be vectorized. A compiler can still improve performance by partial vectorization — vectorizing some, but not all, of the statements. As Table 10 shows, the speedups from partial vectorization are significantly less than those from full vectorization. There are several reasons for this result. First, since by definition partial vectorization vectorizes only some of the statements in a loop, others still run at scalar speeds. Second, our definition of partial vectorization classifies as such a loop that uses *any* vector instructions, no matter how much of the loop is executed in scalar mode. Finally, many techniques for partial vectorization introduce extra work, such as extra loads and stores and additional loop overhead, which is not required in the original loop.

Even with these caveats we see from the last row in Table 10 that there is still a benefit to be gained from partial vectorization, but primarily at the longer vector lengths. Even more so than with full vectorization, partial vectorization — at least on the test loops — degrades performance at vector length of ten.

## 7 Percent Vectorization

In this section we focus on the performance of the compiler independent of the computer architecture. We do this by developing a machine-specific model of what optimal vector performance is for each of the loops in our test suite. We then compare the optimal performance predicted by this model with the actual vector execution results to determine the percent of the optimal vector performance actually achieved.

### 7.1 A Model of Compiler Performance

A simple model of vector performance as a function of vector length is given by the formula [8]

$$t = t_o + nt_e, \tag{1}$$

where  $t$  is the time to execute a vector loop of length  $n$ ,  $t_o$  is the vector startup time, and  $t_e$  is the time to execute a particular vector element. Equivalent to (1) is the well-known model of Hockney (see Hockney and Jesshope [5]),

$$t = r_\infty^{-1}(n + n_{1/2}), \tag{2}$$

where  $r_\infty$  is the asymptotic performance rate and  $n_{1/2}$  is the vector length necessary to achieve one half the asymptotic rate. Equations (1) and (2) can be shown to be equivalent if we use the definitions  $r_\infty = t_e^{-1}$  and  $n_{1/2} = t_o/t_e$  [5].

As Lubeck [7] points out, neither equation models the stripmining process used by compilers on register-to-register vector computers. Also, (1) and (2) may not reflect the behavior of cache-



Table 11: Basic Operation Classes

Class	Operation
0	Load
0	Gather (Load indirect)
1	Store
1	Scatter (Store indirect)
2	Arithmetic (Add, Multiply)
2	Reductions

based machines under increasing vector lengths (see, for example [1]). Nevertheless, for the purposes of our model we believe (1) and (2) to be sufficient.

By analogy with  $r_\infty$ , for each loop, we define three rates:  $r_s$  for the optimized scalar code,  $r_v$  for the vector code, and  $r_o$  for optimal vector code for the target machine. These rates are defined in units of the number of iterations per second of the loop. We assume  $r_s < r_o$ , and we expect  $r_s \leq r_v \leq r_o$ , although (as the previous section indicated) it is possible to have  $r_v < r_s$ .

Using the scalar and vector data collected, we can solve (2), for each loop, for  $r_s$  and  $r_v$ , respectively. Since we cannot necessarily assume  $r_v = r_o$ , we must *estimate*  $r_o$ . To do this, we assume that the execution time of a loop is determined by the *basic* operations in the loop. To determine the rate at which basic operations (e.g., addition or load) can be performed, we use the control loops, which we assume *can* be optimally vectorized.

We divide the basic operations into classes. Each class contains operations that utilize a specific functional unit. For example, Table 11 lists the basic operations in each class for a generic computer with separate load and store pipes.\*

The list of which operations belong to which classes varies by vendor, primarily with respect to the memory operations. For example, on a machine with separate load and store pipes, the load and gather operations are in one class (they compete for the load pipe), and the store and scatter operations are in another class (they compete for the store pipe). For machines with only one pipe for all memory accesses the four memory operations are all in the same class. Even though these operations all have their own execution rates, when they compete for the same resources they are in the same class.

To model control flow, we assume an “execute under mask” model in which every operation is assumed to be executed in vector mode, and the results of various control paths are merged together. Alternative strategies are possible, such as using compress and expand to perform arithmetic only where selected, but we found that execute under mask was sufficient for our purposes.

On each computer, and for each loop  $L$ , we estimate its optimal execution rate  $r_o$ , using

---

\*This table could be extended by subdividing classes into special cases. For example, the arithmetic class could be divided into separate addition and multiplication classes. For machines that can execute adds and multiplies concurrently — all machines in this study — these multiple functional units are modeled as simply a higher arithmetic-processing rate. The difference in execution times between computing the elementwise sum of three vectors and the elementwise product was insignificant for all computers. This fact is not surprising, since the rate limiting step for almost all loops in the suite is memory references, and so this distinction would not change our results significantly.

```

maxtime ← 0
foreach c ∈ C occurring in L
    time ← 0
    foreach o ∈ c
        time ← time + No/Ro
    endfor
    maxtime ← max(time, maxtime)
endfor
ro ← 1/maxtime

```

Figure 3: Algorithm for Estimating Optimal Execution Rate

the algorithm shown in Figure 3. Here  $C$  represents the set of classes defined for a particular computer,  $o$  the operations in a class,  $N_o$  the number of instances of  $o$  in  $L$ , and  $R_o$  the rate for operation  $o$  (in units of operations per second) measured with the control loops. The algorithm assumes that operations in different classes execute concurrently while operations in the same class execute sequentially.

This model is based on the notion of a resource limit, similar to the model used to calculate performance bounds in [9, 13]. We assume that for each loop there exists a particular class of operations that use the same function unit and that the time to execute these operations provides a lower bound on the time to execute the loop. The algorithm in Figure 3 calculates that bound, and we use its reciprocal as  $r_o$ .

In addition to measuring the basic vector operation rates, we also measure the basic scalar operation rates. For each loop, we then determine which operations can be executed in vector mode and which must be executed in scalar mode. We then modify the algorithm in Figure 3 to use the appropriate rate (vector or scalar) for  $R_o$  for each operation.

For each loop and each vendor, we have now determined the three execution rates:  $r_o$  using the algorithm given in Figure 3, and  $r_s$  and  $r_v$  using (2). All three rates were computed by using the data for vector lengths of 100 and 1000. We now define *percent vectorization*,  $p_v$ , by the formula

$$p_v = \frac{r_v - r_s}{r_o - r_s}. \quad (3)$$

With this definition, if a loop's vector execution rate is the same as the scalar rate,  $p_v = 0\%$ , and if a loop's vector execution rate is the same as the optimal vector execution rate,  $p_v = 100\%$ . We can now classify a loop as vectorized, partially vectorized, or not vectorized according to the value of  $p_v$ . We do this according to the rule

$$Result = \begin{cases} n & p_v < 10\% \\ p & 10\% \leq p_v < 50\% \\ v & 50\% \leq p_v. \end{cases} \quad (4)$$

## 7.2 Example

In this section we show an example of the computation of  $p_v$  for two computers,  $C_1$  and  $C_2$ . We assume that  $C_1$  has two load pipes and a store pipe and that  $C_2$  has one pipe used for both loads and stores.

The example used is the loop shown in Figure 1. For this loop we have the following profile of basic operations,\*  $N_o$ :

Load	Store	Gather	Scatter	Arithmetic	Reductions
0 2	0 1	0 0	0 0	0 1	0 0

The first number in each pair is the number of scalar operations, and the second is the number of vectorizable operations. In this example, executing the loop requires two vector loads, one vector store, and a vector addition. No scalar operations are required (our model takes into account scalar operations that occur within the loop body, but not scalar operations, such as incrementing the loop control variable or testing for loop termination, that have to do with the loop control itself).

Using the results of the control loops, we have calculated the following basic vector operation rates. The units are in million of operations per second.

Computer	Load	Store	Arithmetic
$C_1$	227	150	269
$C_2$	186	207	286

Using these values and the loop profile above, we can estimate  $r_o$  with the algorithm shown in Figure 3. The result of these calculations is that, for  $C_1$ , the optimal vector execution rate is 114 million iterations per second, and for  $C_2$  it is 64 million iterations per second.

Using the scalar and vector results for vector lengths 100 and 1000, we determined the following results for  $r_s$  and  $r_v$  by solving (2):

Computer	$r_s$	$r_v$
$C_1$	12.3	115.
$C_2$	10.7	19.8

Substituting the appropriate values for  $r_s$ ,  $r_v$ , and  $r_o$ , into (3), we calculated  $p_v = 100\%$  for  $C_1$  and  $p_v = 17\%$  for  $C_2$ . Applying (4), we determined that  $C_1$  fully vectorizes this loop and that  $C_2$  partially vectorizes this loop.

## 7.3 Results

Table 12 is similar to Table 1, except here the number of loops vectorized or partially vectorized has been determined by applying (3) and (4) as opposed to analyzing the compiler's output

---

\*From Appendix C.

Table 12: Full Vectorization according to (3) and (4) (122 loops)

Computer	V	P	N
CONVEX C210	51.6 (63)	14.8 (18)	33.6 (41)
CCC CRAY-2	36.1 (44)	24.6 (30)	39.3 (48)
CRI CRAY Y-MP	54.1 (66)	25.4 (31)	20.5 (25)
DEC VAX 9000-210	45.1 (55)	8.2 (10)	46.7 (57)
FPS M511EA-2	52.5 (64)	18.0 (22)	29.5 (36)
Fujitsu VP2600/10	53.3 (65)	13.1 (16)	33.6 (41)
Hitachi S-820/80	52.5 (64)	15.6 (19)	32.0 (39)
IBM 3090-600J	51.6 (63)	18.9 (23)	29.5 (36)
NEC SX-X/14	46.7 (57)	21.3 (26)	32.0 (39)
Average	49.3 (60)	17.8 (21)	33.0 (40)

listing.

In comparing Table 12 to Table 1 we observe that the results are mostly consistent with Table 1, with a somewhat tighter grouping among vendors with the most loops vectorized. Most compilers vectorized between 20 and 30 loops which did not achieve full vector performance ( $p_v \geq 50\%$ ). Appendix B contains complete results.

Casual inspection of the data indicates that there are a number of loops for which at most one vendor successfully achieved vector performance and all other vendors that vectorized did not. Approximately 23 loops account for most of the differences between the two measures of vectorization. For the most part, these loops are scattered across categories but they include most of the scalar expansion loops, the search loops, the packing loops, and the loops with multiway branching.

Factors other than simple detection of vectorizability are reflected in the computation of vectorization percentages. In particular, traditional optimizations such as common subexpression elimination, register allocation, and instruction scheduling will all influence the quality of the generated code and hence the percentage of vectorization. In this sense, the percentage is more a measure of the overall quality of the compiler generated code.

Optimal code generation and, in particular, instruction scheduling for very simple loops are extremely difficult. For loops with large bodies, heuristic algorithms will usually get within a small number of instructions of what is optimal. When the loop body contains only five or ten instructions, however, being off by a “small” number could cost 25% of achievable performance. Thus, since almost all of the loops in the suite are very simple, the compilers may perform substantially better on “real” codes than is suggested by Table 12.

That the measured execution rates are lower than what might be expected from “vector” code may be due to model limitations. For example, the model treats unit and nonunit stride vector accesses as equal in cost: there was no convincing evidence that nonunit stride was a factor worth adding to the operation classes listed in Table 11. The other major factor not modeled is the presence of a data cache, its size and its organization. This is discussed in Section 8.3.1.

One issue that biases the results presented here is that we use the measured performance on simple loops to calibrate the model. Thus our “optimal rates” may be significantly below “ma-

chine peaks” since those peaks may be achievable only assuming optimal compilation. Further, if the code generation capabilities of one compiler are generally poor compared with another, then its ability to vectorize may appear inflated, since our estimate of optimal execution rate may be too low. This situation can be corrected by replacing the control loops with numbers derived from hand-crafted assembly routines that would provide estimates of “achievable peaks.” We did not have the resources to generate these numbers for each machine.

## 8 Discussion

### 8.1 Validity of the Test Suite

How good is this test suite? The question can be answered in several ways, but we will address three specific areas: coverage, stress, and accuracy.

#### 8.1.1 Coverage

By “coverage” we refer to how well the test suite represents typical, common, or important Fortran programming practices. We would like to assert that high effectiveness on the test suite will correspond to high effectiveness in general. Unfortunately, there is no accepted suite of Fortran programs that can be called representative, and so we have no quantitative way of determining the coverage of our suite. We believe, however, that the method used to select the tests has yielded reasonable coverage. This method consisted of two phases.

In the first phase, a large number of loops were collected from several vendors and interested parties. This gave a diverse set of viewpoints, each with a different machine architecture and hence somewhat different priorities. In some cases the loops represented “real” code from programs that had been benchmarked. The majority, however, were specifically written to test a vectorizing compiler for a particular feature. Independently, the categorization scheme used in Section 2 was developed based on experience and on published literature about vectorization.

In the second phase, the test suite was culled from the collected loops by classifying each loop into one or more categories and then selecting a few representative loops from each category. Our interest was in coverage; and since “representative” is not well defined, we made no attempt to weight some of the subcategories more than others by changing the number of loops. Where we felt that testing a subcategory required a range of situations, we included several loops; in other cases we felt that one or two loops sufficed. There is significant weighting between major categories. For example, the test suite places greater emphasis on basic vectorization (41 loops) than on idiom recognition (24 loops). This weighting was an artifact of the selected categories and was reflected in the original collection of samples. We felt that this weighting was reasonable and made no attempt to adjust it.



a more balanced distribution of tests between “easy” and “difficult” when compared to Table 13. Statistics also support this view. In [3] the average number of loops vectorized was 55%, and vectorized or partially vectorized was 61%. Even if we restrict ourselves to just the eight vendors also participating in this test, the previous results are still only 59% and 64%, respectively. In this test the average number of loops vectorized was 70%, and vectorized or partially vectorized was 77%, an improvement of about 15%

Several factors may be at work. First, compilers have evolved and improved over time. Second, specialized third-party compiler technology is now readily available to interested vendors. Third, for various reasons approximately half the vendors who participated in the previous test did not participate this time. While those who did not participate span the spectrum of previous results, most had results in the lower or middle part of the previous test. While we added new loops to this test (and also deleted a small number), this does not seem to have provided adequate stress. Since one valid use of this test suite is for compiler writers to diagnose system deficiencies, we expect over time that the test will lose its effectiveness to stress compilers.

### 8.1.3 Accuracy

By “accuracy” we refer to how well the test can measure the quality of a vectorizing compiler. Since the difficulty of the tests was determined by the performance of the compilers, it would be circular now to judge the absolute quality of the compilers by their performance on this suite. What about relative performance? It is tempting to distill the results for each compiler into a single number and use that to compare the systems. Such an approach, however, is clearly incorrect, since these compilers cannot be compared in isolation from the machine environment and target application area for which they were designed.

We conclude that the suite represents reasonable coverage, that the stress may no longer be adequate, and that we cannot determine the accuracy of the suite.

## 8.2 Beating the Test

Some of the loops were vectorized in ways that defeated the intent of the test. One example is the use of a runtime test. If the compiler cannot determine at compile time whether a loop is safe to vectorize, because of, say, an unknown parameter value, it must either not vectorize the loop or else generate an alternative code runtime test. At runtime, based on the value of the unknown parameter, the test executes either a scalar or a vector version of the loop, as appropriate. In general, we view runtime testing as a good thing to do. It allows vectorization of loops that would not otherwise be vectorized and allows cost-effectiveness decisions to be deferred until runtime. However, it has a negative side. First, the cost of the test is incurred each time the loop is executed. Second, for large loop nests, it is possible to have a combinatorial explosion in the number of tests generated. All of the loops in our test suite can be determined to be vectorizable at compile time, and thus runtime testing is not necessary. The Cray Research, FPS, IBM, and NEC compilers, however, can generate runtime tests and in a few cases were able to “beat the test” this way.

A technique similar to runtime testing is conditional vectorization, which was used by the

Cray Computer compiler. With conditional vectorization, a safe vector length\* is calculated at runtime. While conditional vectorization is also good for a compiler to be able to do, it also has a negative side. First, there is the overhead involved in calculating the safe vector length at runtime. Second, if the calculated safe vector length is one, it is more efficient to execute a scalar instruction rather than a vector instruction. None of the loops in our test require conditional vectorization. Nevertheless, the Cray Computer compiler conditionally vectorized 20 loops, 11 of which resulted in a safe vector length greater than one.

Another way compilers defeated the intent of the test was by their ability to vectorize recurrences, using either library routines or special hardware. Several of the tests call for the compiler to split up a loop (loop distribution, node splitting) or change the order of a loop nest (loop interchange) in order to vectorize a loop containing an “unvectorizable” recurrence. Several of the compilers – notably those from Fujitsu, Hitachi, and NEC – were able to directly vectorize some of these loops.

We emphasize that “beating the test” is not a bad thing. While there may be more efficient ways to vectorize the loops, the techniques above *are* beneficial.

### 8.3 Caveats

We caution that the results presented here test only one aspect of a compiler and should in no way be used to judge the overall performance of a vectorizing compiler or computer system. The results reflect only a limited spectrum of Fortran constructs. We do not claim these loops are representative of a “real” workload, just that they make an interesting test. Some additional factors are discussed below.

#### 8.3.1 Cache Effects

Two issues may impact machines with data caches. First, to ensure a large enough granularity for timing purposes, we included a repetition loop around the loop of interest. While considered a necessary evil for test purposes, this artificial repetition raises an important question about data locality. The concern is that a cache machine will benefit from the reuse of data loaded into cache on the first trip through the repetition loop and that additional references to main memory will not be necessary.

The second issue concerns the data set size relative to the cache size. A small data set will always fit in the cache. A large data set may not fit in the cache and will cause many performance-degrading cache misses to occur. The paper by Abu-Sufah and Maloney [1] contains a discussion of this issue and its impact on performance. Their uniprocessor performance results on an Alliant FX/8 show that there is only a narrow range of vector lengths for which optimal performance was achieved. Our choice of 10, 100, and 1000 as the vector lengths was somewhat intuitive and was not made with any particular cache size in mind.

---

\*See Section 5.1



```

do i = 2,n
  a(i) = a(i) + b(i)
  b(i) = b(i-1)*b(i-1)*a(i)
  a(i) = a(i) - b(i)
enddo

do vector i = 2,n
  a(i) = a(i) + b(i)
enddo
do i = 2,n
  b(i) = b(i-1)*b(i-1)*a(i)
enddo
do vector i = 2,n
  a(i) = a(i) - b(i)
enddo

```

Figure 4:

### 8.3.2 Loop Granularity

Because of the small granularity of our loops (at most a few statements) the speedups achievable with a certain technique may not be achievable on our particular loops. As an example, vectorizing the loop shown on the left in Figure 4 requires splitting the loop into two vectorizable loops and one scalar loop containing the nonlinear recurrence as shown on the right.

For this transformation to be successful, there needs to be enough work in the loop to justify the two additional loop overheads introduced and the extra loads and stores which are not required in the original loop. For this loop, inspection of the compiler listing showed that eight of the nine compilers had partially vectorized the loop, but only three achieved more than 15% of the estimated optimal performance, and only one achieved more than 50%.

### 8.3.3 Hardware and Software

Some of the loops are really tests of the underlying hardware and may not accurately reflect the ability of the compiler itself. For example, in the statement  $A(I)=B(\text{INDEX}(I))$  a compiler may detect the indirect addressing of array B but not generate vector instructions because the computer does not have hardware support for array references of this form. Other examples are loops containing IF tests that may require mask registers, or recurrences that require special library routines.

Several of the computers tested are multiprocessors whose compilers support the generation of both parallel and vector code. Our test involved strictly uniprocessors and may have penalized vendors who have put considerable effort into parallel execution. On some of these machines, parallel execution may be more efficient than vectorization for certain loops.

Another example where the computer architecture may influence the compiler is on machines that have a data cache. Compilers for such machines may concentrate on loop transformations that improve data locality at the expense of adding "simple" vectorization capabilities.

Several vendors have sophisticated tools to aid the user in vectorization. For example, both Fujitsu and NEC offer vectorization tools that interactively assist the user in vectorizing a program. Another example is an interprocedural analysis compiler from Convex, which analyzes

an entire program at once. While all are very sophisticated tools, their use was against our rules.

## 9 Conclusions and Future Work

Our results indicate that most of the compilers tested are fairly sophisticated, able to use a wide variety of vectorization techniques and transformations. Loops that were considered challenging several years ago, such as indirect addressing or vectorizing loops containing multiple IF tests, now seem routine. While there are still various vectorization challenges left to be met, we are not sure how much they will be addressed in the future. Our perception is that most current compiler work is going into memory hierarchy management, parallel loop generation, highly pipelined scalar processors, and interactive and interprocedural tools. We may well be nearing a plateau as far as how much additional work vendors will put into vectorization techniques alone.

Our test suite continues to evolve from simple inspection of the compiler's output listing to trying to judge the quality of the execution results. To make the test more meaningful, we plan to add the types of "real" loops found in applications. Real loops present combinations of vectorization problems rather than individual challenges. It will then be interesting to compare results on the "simple" loops with those on the real loops.

A copy of the source code used in the test is available from the NETLIB electronic mail facility [4] at Oak Ridge National Laboratory. To receive a copy of the code, send electronic mail to *netlib@ornl.gov*. In the mail message, type *send vectors from benchmark* or *send vectord from benchmark* to get either the REAL or DOUBLE PRECISION versions, respectively.

### Acknowledgments

We thank John Lesvesque, Murray Richman, Steve Wallach, Joel Williamson, and Michael Wolfe for providing many of the loops used in this test. Thanks also to the many people involved in running this test, providing results and constructive feedback on earlier versions.

## References

- [1] W. Abu-Sufah and A. Malony. Vector processing on the Alliant FX/8 multiprocessor. In *Proceedings of the 1986 International Conference on Parallel Processing*, pages 559–566, 1986.
- [2] J. Allen and K. Kennedy. Automatic translation of Fortran programs to vector form. *TOPLAS*, 9(4):491–542, 1987.
- [3] D. Callahan, J. Dongarra, and D. Levine. Vectorizing compilers: A test suite and results. In *Proceedings of Supercomputing '88*, pages 98–105, 1988.
- [4] J. Dongarra and E. Grosse. Distribution of mathematical software via electronic mail. *Communications of the ACM*, 30(5):403–407, July 1987.
- [5] R. Hockney and C. Jesshope. *Parallel Computers: Architecture, Programming and Algorithms*. Adam Hilger, Ltd., Bristol, United Kingdom, 1981.

- [6] J. Levesque and J. Williamson. *A Guidebook to Fortran on Supercomputers*. Academic Press, New York, New York, 1988.
- [7] O. Lubeck. Supercomputer performance: The theory, practice, and results. Technical Report LA-11204-MS, Los Alamos National Laboratory, 1988.
- [8] O. Lubeck, J. Moore, and R. Mendez. A benchmark comparison of three supercomputers: Fujitsu VP-200, Hitachi S810/20, and Cray X-MP/2. *IEEE Computer*, 18(12):10-23, 1985.
- [9] W. Mangione-Smith, S. Abraham, and E. Davidson. A performance comparison of the IBM RS/6000 and the Astronautics ZS-1. *IEEE Computer*, 24(1):39-46, 1991.
- [10] F. McMahon. The Livermore Fortran kernels: A computer test of the numerical range. Technical Report UCRL-53745, Lawrence Livermore National Laboratory, 1986.
- [11] D. Padua and M. Wolfe. Advanced compiler optimizations for supercomputers. *Communications of the ACM*, 29(12):1184-1201, 1986.
- [12] J. Smith. Characterizing computer performance with a single number. *Communications of the ACM*, 31(10):1202-1206, 1988.
- [13] J. Tang and E. Davidson. An evaluation of Cray-1 and Cray X-MP performance on vectorizable Livermore Fortran kernels. In *Proceedings of the 1988 International Conference on Supercomputing*, pages 510-518, St. Malo, France, 1988.
- [14] M. Wolfe. *Optimizing Supercompilers for Supercomputers*. The MIT Press, Cambridge, Massachusetts, 1989.

## Appendix A

Table 15: Hardware and Software used in this test.

Company Computer	Compiler Version Compiler Options	OS Version Main/Cache Memory
CONVEX Computer Corp. CONVEX C210	fc 6.1 -O2 -uo -is	OS 9.0 512MB/None
Cray Computer Corp. CRAY-2	cft77 4.0.1.1 Defaults	UNICOS 6.0 1GB/None
Cray Research, Inc. CRAY Y-MP	CF77 4.0 -Wd"-e78b"	UNICOS 5.1 1GB/None
Digital Equipment Corp. VAXvector 9000-210	FORTRAN V5.5, HPO V1.0 /HPO/VECTOR/BLAS=(INLINE, MAPPED)/ASSUME=NOACC/OPT	VMS 5.4 512MB/128KB
FPS Computing FPS M511EA-2	f77 4.3 -u -O -Oc inl+ -Oc vec+ -Oc pi+	FPX, 4.3.2 256MB/64KB
Fujitsu VP2600/10	Fortran77EX/VP V11L10 VP(2600),OPT(F),INLINE(EXT(S151S)) VMSG(DETAIL)	OSIV/MSP AFII & VPCF V10L10 1GB/None
Hitachi S-820/80	fort77/hap V24-0f sopt,xfunc(xfr), hap(model80,vist),uinline	vos3/as jss4 01-02 512MB/256KB
IBM Corp. IBM 3090-600J	VS FORTRAN 2.4.0, VAST-2 R2 vopt(opton=r8 inline=s151s,s152s) copt(opt(3) vec(rep(xlist)))	MVS/ESA SP3.1.0E JES2 SP3.1.1 256MB/256KB 512MB Extended Memory
NEC Corp. SX-X/14	f77sx 010 -pi *:s151s *:s152s	SUPER-UX R1.11 1GB/64KB

## Appendix B

The tables below contain the results, on a loop-by-loop basis, for the 122 loops in the test suite. For each loop three columns of data are given. The first column is the result according to the compiler listing: vectorized (v), partially vectorized (p), or not vectorized (n), as described in Section 5.

The second and third columns were calculated with the model described in Section 7. The second column was calculated by applying (4) to  $p_v$ . The third column (enclosed in parentheses) is  $p_v$ , the percentage of vectorization, calculated with equation (3). An entry of "(D)" indicates that the vector execution rate was less than the scalar execution rate. Occasionally, the  $p_v$  value will be higher than 100%. Some of these loops — notably s176, s352, and s4116 — seem to have uniformly higher values, while most of the others are scattered throughout the test suite. We have no general explanation for these cases. Two possibilities are limitations of the model or measurement error.\*

Computer	s111	s112	s113	s114	s115
CONVEX C-210	v v (100)	v v (99)	v v (93)	v p (22)	v v (88)
CCC CRAY-2	v p (17)	v v (73)	v v (56)	n n (0)	v <sup>a</sup> v (66)
CRI CRAY Y-MP	v v (103)	v v (79)	v v (100)	v v (50)	v v (107)
DEC VAX 9000-210	v v (97)	v v (181)	v v (147)	v n (5)	v v (83)
FPS M511EA-2	v p (49)	v v (86)	v v (65)	v p (31)	v v (86)
Fujitsu VP2600/10	v p (33)	v v (104)	v v (121)	v p (41)	v v (84)
Hitachi S-820/80	v v (113)	v v (102)	v v (105)	v p (32)	v v (108)
IBM 3090-600J	v v (92)	v v (103)	v v (97)	v n (7)	v v (84)
NEC SX-X/14	v v (74)	v v (75)	v v (69)	v p (43)	v v (61)

Computer	s116	s118	s119	s121	s122
CONVEX C-210	v v (92)	v p (16)	v v (81)	v v (100)	v v (98)
CCC CRAY-2	n n (0)	n n (0)	v <sup>a</sup> v (67)	v v (97)	v v (77)
CRI CRAY Y-MP	v v (55)	v v (84)	v v (81)	v v (73)	v v (76)
DEC VAX 9000-210	v n (D)	v n (3)	v v (59)	v v (60)	v v (110)
FPS M511EA-2	v p (42)	v p (29)	v v (100)	v v (86)	v v (86)
Fujitsu VP2600/10	v p (17)	v p (35)	v v (104)	v v (105)	v v (92)
Hitachi S-820/80	v v (85)	v p (18)	v v (99)	v v (104)	n n (D)
IBM 3090-600J	v v (53)	v n (8)	v v (55)	v v (100)	v v (101)
NEC SX-X/14	v v (189)	v p (37)	v v (87)	v v (84)	v v (75)

<sup>a</sup> These loops were conditionally vectorized. For loops with ambiguous subscripts, a runtime test was compiled which selected a safe vector length.

<sup>b</sup> These loops were executed in scalar mode. The compiler indicated that partial vectorization was possible but that the overhead was too large.

<sup>c</sup> For these loops a runtime IF-THEN-ELSE test was compiled which executed either a scalar loop or a vectorized loop.

\* We estimate slightly less than two decimal digits of significance in the timing information collected. Hence, the percentage of vectorization calculations may have error terms of approximately 10% to 15%.

Computer	s123		s124		s125		s126		s127						
CONVEX C-210	p	n	( 2)	p	n	( 2)	v	v	( 95)	p	n	( 0)	v	v	( 99)
CCC CRAY-2	n	n	( 0)	n	n	( 0)	v	v	( 99)	n <sup>a</sup>	n	(D)	v	v	( 56)
CRI CRAY Y-MP	n	n	( 0)	v	p	( 37)	v	v	(122)	v	v	( 64)	v	v	( 76)
DEC VAX 9000-210	n	n	(D)	n	n	(D)	v	p	( 45)	n	n	( 0)	v	v	(119)
FPS M511EA-2	n	n	( 0)	v	v	( 68)	v	v	(111)	v	p	( 41)	v	v	( 88)
Fujitsu VP2600/10	p <sup>b</sup>	n	( 0)	p <sup>b</sup>	n	( 0)	v	v	(102)	p <sup>b</sup>	n	(D)	p <sup>b</sup>	n	( 0)
Hitachi S-820/80	n	n	(D)	n	n	(D)	v	v	(100)	v	p	( 13)	v	v	( 51)
IBM 3090-600J	n	n	( 0)	v	p	( 46)	v	v	( 58)	v	v	( 55)	v	v	( 74)
NEC SX-X/14	n	n	( 0)	n	n	( 0)	v	v	(103)	v	n	( 5)	n	n	( 0)

Computer	s128		s131		s132		s141		s151						
CONVEX C-210	v	v	(100)	v	v	(100)	v	v	( 91)	n	n	( 0)	n	n	(D)
CCC CRAY-2	v	p	( 35)	v	v	( 98)	v <sup>a</sup>	v	( 66)	n	n	( 0)	v <sup>a</sup>	v	( 68)
CRI CRAY Y-MP	v	v	( 74)	v	v	( 76)	v	v	(110)	n	n	( 0)	v	v	( 78)
DEC VAX 9000-210	v	v	( 53)	v	v	(110)	v	v	(108)	n	n	( 1)	n	n	( 4)
FPS M511EA-2	n	n	( 0)	v	v	( 86)	v	v	( 86)	n	n	( 0)	v	v	( 86)
Fujitsu VP2600/10	p <sup>b</sup>	n	( 0)	v	v	(105)	v	v	(106)	p	n	( 0)	v	v	(106)
Hitachi S-820/80	p	n	(D)	v	v	(104)	v	v	(103)	p	n	(D)	n	n	(D)
IBM 3090-600J	n	n	( 0)	v	v	(101)	v	v	( 95)	n	n	( 1)	v	v	(101)
NEC SX-X/14	p	n	(D)	v <sup>c</sup>	v	( 83)	v <sup>c</sup>	v	( 79)	p	n	( 0)	v	v	( 84)

Computer	s152		s161		s162		s171		s172						
CONVEX C-210	v	v	( 99)	p	n	( 4)	p	n	( 4)	v	v	( 98)	v	v	( 99)
CCC CRAY-2	n	n	( 0)	n	n	( 0)	v <sup>a</sup>	v	( 73)	v	v	( 78)	v	v	( 79)
CRI CRAY Y-MP	v	v	(118)	n	n	( 0)	n <sup>c</sup>	n	(D)	v	v	(104)	v	v	(104)
DEC VAX 9000-210	n	n	(D)	v	v	(127)	v	v	(100)	n	n	( 0)	v	v	(110)
FPS M511EA-2	v	v	( 96)	n	n	( 0)	n <sup>c</sup>	n	( 0)	v	v	( 85)	v	v	( 86)
Fujitsu VP2600/10	n	n	( 0)	v	v	(147)	p <sup>b</sup>	n	( 0)	v	v	(102)	v	v	(105)
Hitachi S-820/80	v	v	(114)	n	n	(D)	p	n	(D)	n	n	(D)	n	n	(D)
IBM 3090-600J	v	v	(100)	v	v	( 78)	n <sup>c</sup>	n	( 0)	v	v	(100)	v	v	(100)
NEC SX-X/14	v	v	( 79)	n	n	( 0)	v <sup>c</sup>	v	( 98)	v <sup>c</sup>	v	( 64)	v	v	( 84)

Computer	s173		s174		s175		s176		s211						
CONVEX C-210	n	n	( 0)	n	n	( 0)	n	n	( 8)	v	v	(200)	v	v	( 98)
CCC CRAY-2	v <sup>a</sup>	v	( 68)	v <sup>a</sup>	v	( 68)	v <sup>a</sup>	v	( 68)	v <sup>a</sup>	v	( 57)	n	n	( 0)
CRI CRAY Y-MP	v <sup>c</sup>	v	( 83)	v	v	( 84)	v	v	( 81)	v	v	(106)	v	v	( 94)
DEC VAX 9000-210	n	p	( 12)	v	v	(118)	n	n	( 1)	v	v	(207)	v	v	(132)
FPS M511EA-2	v <sup>c</sup>	v	( 85)	v	v	( 86)	v	v	( 86)	v	v	(121)	v	v	( 93)
Fujitsu VP2600/10	n	n	( 0)	n	n	( 0)	v	v	(104)	v	v	(187)	v	v	(128)
Hitachi S-820/80	n	n	(D)	n	n	(D)	n	n	(D)	v	v	(175)	v	v	( 91)
IBM 3090-600J	v <sup>c</sup>	v	(102)	v	v	(107)	v	v	(100)	v	v	( 58)	v	v	( 81)
NEC SX-X/14	v <sup>c</sup>	v	( 96)	n	n	( 0)	v	v	( 82)	v	v	(106)	v	v	( 64)

Computer	s212	s221	s222	s231	s232
CONVEX C-210	v v (98)	p p (11)	p n (6)	v v (85)	n n (0)
CCC CRAY-2	n n (0)	n n (0)	n n (0)	n <sup>a</sup> n (D)	n <sup>a</sup> n (D)
CRI CRAY Y-MP	v v (92)	p p (38)	p v (65)	v v (80)	n n (0)
DEC VAX 9000-210	v v (132)	p n (5)	p n (D)	v v (65)	n n (0)
FPS M511EA-2	v v (107)	v v (72)	p p (38)	v v (87)	n n (0)
Fujitsu VP2600/10	v v (107)	v v (62)	p n (0)	v v (132)	n n (0)
Hitachi S-820/80	v v (98)	v v (160)	p n (8)	v v (98)	n n (D)
IBM 3090-600J	v v (100)	p p (16)	p p (47)	v v (65)	n n (2)
NEC SX-X/14	v v (81)	v p (29)	p n (10)	v v (87)	n n (0)

Computer	s233	s234	s235	s241	s242
CONVEX C-210	v p (39)	n n (0)	v v (114)	v v (99)	p n (5)
CCC CRAY-2	p <sup>a</sup> n (D)	n <sup>a</sup> n (D)	n <sup>a</sup> n (D)	n n (0)	n n (0)
CRI CRAY Y-MP	p n (4)	n n (2)	v v (78)	v v (52)	p p (39)
DEC VAX 9000-210	v n (3)	n n (0)	v v (78)	v v (81)	n n (0)
FPS M511EA-2	p n (2)	n n (0)	v v (128)	v v (65)	v v (96)
Fujitsu VP2600/10	p n (0)	n n (0)	v v (197)	v v (107)	v v (64)
Hitachi S-820/80	v p (15)	n n (0)	v v (99)	v v (71)	v v (165)
IBM 3090-600J	p n (D)	v v (66)	v v (89)	v v (67)	p p (25)
NEC SX-X/14	v p (30)	n n (0)	v v (104)	v p (32)	v p (30)

Computer	s243	s244	s251	s252	s253
CONVEX C-210	v v (98)	v v (81)	v v (95)	v p (50)	v v (98)
CCC CRAY-2	n n (0)	n n (0)	v v (94)	n n (0)	v p (22)
CRI CRAY Y-MP	v v (67)	p n (4)	v v (79)	p n (8)	v p (24)
DEC VAX 9000-210	v p (48)	v v (51)	v v (117)	n n (1)	v v (71)
FPS M511EA-2	v v (66)	p n (D)	v v (82)	p n (0)	v v (62)
Fujitsu VP2600/10	v v (104)	v v (105)	v v (153)	p <sup>b</sup> n (0)	v v (84)
Hitachi S-820/80	v v (87)	p n (D)	v v (93)	v p (41)	v v (92)
IBM 3090-600J	v v (52)	v p (29)	v v (74)	v p (12)	v v (80)
NEC SX-X/14	v v (54)	n n (0)	v v (95)	p n (D)	v p (35)

Computer	s254	s255	s256	s257	s258
CONVEX C-210	v p (42)	v p (24)	n n (D)	v n (1)	p n (8)
CCC CRAY-2	n n (0)	n n (0)	n n (0)	v p (13)	n n (0)
CRI CRAY Y-MP	n n (0)	n n (0)	n n (0)	v p (37)	p n (3)
DEC VAX 9000-210	n n (0)	n n (1)	p n (0)	n n (0)	n n (D)
FPS M511EA-2	n n (0)	n n (0)	n n (0)	v p (11)	p n (4)
Fujitsu VP2600/10	p <sup>b</sup> n (0)	p <sup>b</sup> n (0)	v n (5)	v p (16)	p <sup>b</sup> n (0)
Hitachi S-820/80	v p (36)	v p (25)	v n (5)	v p (14)	n n (D)
IBM 3090-600J	v p (27)	v p (12)	n n (0)	v n (D)	p n (9)
NEC SX-X/14	n n (0)	n n (0)	v n (3)	n n (0)	n n (0)

Computer	s261	s271	s272	s273	s274
CONVEX C-210	v v (98)	v v (100)	v v (99)	v v (98)	v v (95)
CCC CRAY-2	n n (0)	v p (18)	v p (31)	v v (68)	v p (36)
CRI CRAY Y-MP	v v (118)	v p (23)	v p (42)	v v (84)	v p (34)
DEC VAX 9000-210	v v (166)	v v (71)	v v (129)	v v (60)	v p (42)
FPS M511EA-2	v v (93)	v v (58)	v v (61)	v v (112)	v v (74)
Fujitsu VP2600/10	v v (96)	v v (101)	v v (132)	v v (113)	v v (121)
Hitachi S-820/80	v v (96)	v v (108)	v v (93)	v v (108)	v v (91)
IBM 3090-600J	v v (99)	v v (70)	v v (81)	v v (61)	v v (67)
NEC SX-X/14	p n (D)	v v (82)	v v (72)	v v (54)	v p (45)

Computer	s275	s276	s277	s278	s279
CONVEX C-210	p n (0)	v p (49)	v v (83)	v v (96)	v v (61)
CCC CRAY-2	n <sup>a</sup> n (D)	v p (21)	n n (D)	v p (42)	v p (39)
CRI CRAY Y-MP	n n (0)	v p (22)	n n (0)	v p (30)	v p (30)
DEC VAX 9000-210	n n (0)	v n (1)	v n (D)	v p (31)	v n (D)
FPS M511EA-2	n n (0)	v p (18)	n n (0)	v v (54)	v p (39)
Fujitsu VP2600/10	p <sup>b</sup> n (0)	v v (59)	v v (105)	v v (90)	v v (60)
Hitachi S-820/80	p p (13)	v p (47)	n n (D)	v v (76)	v v (62)
IBM 3090-600J	n n (D)	v p (19)	v n (D)	v p (35)	v n (1)
NEC SX-X/14	p n (6)	v p (30)	n n (0)	v v (57)	v p (34)

Computer	s2710	s2711	s2712	s281	s291
CONVEX C-210	v p (35)	v v (100)	v v (98)	p n (D)	v p (39)
CCC CRAY-2	v p (30)	v p (19)	v p (17)	n n (0)	n n (0)
CRI CRAY Y-MP	v p (23)	v p (20)	v p (19)	v p (25)	p n (5)
DEC VAX 9000-210	v n (D)	v v (70)	v v (53)	n n (0)	n n (0)
FPS M511EA-2	v p (27)	v v (100)	v p (49)	v p (22)	p n (1)
Fujitsu VP2600/10	v p (17)	v v (102)	v v (101)	p <sup>b</sup> n (0)	p <sup>b</sup> n (0)
Hitachi S-820/80	v v (84)	v v (108)	v v (103)	p n (D)	v v (108)
IBM 3090-600J	v n (D)	v v (100)	v v (72)	v p (27)	v n (8)
NEC SX-X/14	v p (31)	v v (82)	v v (86)	p n (D)	v v (63)

Computer	s292	s293	s2101	s2102	s2111
CONVEX C-210	v p (35)	n n (2)	v n (1)	v v (84)	n n (D)
CCC CRAY-2	n n (0)	n <sup>a</sup> n (D)	v v (81)	p n (5)	n <sup>a</sup> n (D)
CRI CRAY Y-MP	n n (0)	v v (100)	v v (102)	p v (79)	n n (0)
DEC VAX 9000-210	n n (2)	n n (0)	v n (D)	v v (77)	n n (0)
FPS M511EA-2	n n (0)	v p (23)	v v (100)	n <sup>c</sup> n (0)	n n (0)
Fujitsu VP2600/10	p <sup>b</sup> n (0)	n n (0)	v v (75)	v p (33)	v n (3)
Hitachi S-820/80	v v (90)	n n (D)	v v (72)	v p (21)	v p (11)
IBM 3090-600J	v p (17)	v v (73)	v n (D)	n n (0)	n p (17)
NEC SX-X/14	n n (0)	v v (96)	v v (92)	v v (140)	v n (3)



Computer	s311	s312	s313	s314	s315
CONVEX C-210	v v (100)	v v (52)	v v (73)	v v (91)	v v (88)
CCC CRAY-2	v v (99)	v v (98)	v v (73)	v p (28)	n n (0)
CRI CRAY Y-MP	v v (100)	v v (100)	v v (163)	v p (38)	v p (27)
DEC VAX 9000-210	v v (104)	v v (103)	v v (98)	v v (67)	v v (72)
FPS M511EA-2	v v (100)	v p (44)	v v (107)	v v (100)	v p (45)
Fujitsu VP2600/10	v v (98)	v n (0)	v v (179)	v v (99)	v v (100)
Hitachi S-820/80	v v (100)	v n (0)	v v (117)	v v (93)	v v (92)
IBM 3090-600J	v v (100)	n n (0)	v v (97)	v n (5)	v n (5)
NEC SX-X/14	v v (102)	v v (70)	v v (124)	v p (40)	v p (25)

Computer	s316	s317	s318	s319	s3110
CONVEX C-210	v v (90)	v p (43)	n n (0)	v v (98)	n n (0)
CCC CRAY-2	v p (28)	v v (75)	n n (0)	v v (95)	n n (0)
CRI CRAY Y-MP	v p (38)	v v (84)	v p (42)	v v (67)	p n (D)
DEC VAX 9000-210	v v (71)	v v (469)	v p (45)	p p (21)	n n (D)
FPS M511EA-2	v v (100)	n n (0)	v p (47)	v v (58)	v p (23)
Fujitsu VP2600/10	v v (99)	v n (D)	v v (132)	p <sup>b</sup> n (0)	v v (99)
Hitachi S-820/80	v v (89)	v n (D)	v v (88)	v v (78)	p n (D)
IBM 3090-600J	v n (0)	n n (3)	n n (0)	n n (0)	v n (6)
NEC SX-X/14	v p (40)	v v (123)	v p (26)	v v (54)	v p (25)

Computer	s3111	s3112	s3113	s321	s322
CONVEX C-210	v v (85)	n n (0)	n n (0)	n n (D)	n n (6)
CCC CRAY-2	v p (21)	n n (0)	v p (35)	n n (0)	n n (0)
CRI CRAY Y-MP	v p (41)	n n (0)	v v (68)	v p (17)	v p (45)
DEC VAX 9000-210	v v (80)	n n (4)	v v (62)	n n (3)	n n (0)
FPS M511EA-2	v v (78)	n n (0)	v v (110)	v v (104)	v p (31)
Fujitsu VP2600/10	v v (125)	v n (1)	v v (130)	v v (183)	n n (4)
Hitachi S-820/80	v v (75)	v n (6)	v v (90)	v v (184)	n n (D)
IBM 3090-600J	v p (47)	n n (4)	v n (8)	n n (8)	n p (10)
NEC SX-X/14	v p (36)	n n (0)	v p (35)	v v (113)	n n (0)

Computer	s323	s331	s332	s341	s342
CONVEX C-210	p n (2)	v p (49)	n n (0)	v p (20)	v p (21)
CCC CRAY-2	n n (0)	v p (40)	v p (20)	v p (17)	v p (29)
CRI CRAY Y-MP	n n (0)	v p (46)	v p (29)	v p (27)	v p (19)
DEC VAX 9000-210	n n (5)	n n (D)	n n (D)	n n (D)	n n (D)
FPS M511EA-2	n n (0)	v n (6)	n n (0)	n n (0)	n n (0)
Fujitsu VP2600/10	p <sup>b</sup> p (10)	v p (20)	v p (44)	v p (44)	v p (30)
Hitachi S-820/80	v v (326)	v p (32)	v p (43)	v v (81)	v v (65)
IBM 3090-600J	n n (7)	n n (0)	v v (70)	v v (53)	v p (12)
NEC SX-X/14	v v (56)	v p (11)	n n (0)	v p (16)	v p (16)

Computer	s343	s351	s352	s353	s411
CONVEX C-210	v n (1)	v v (101)	v v (155)	v p (22)	n n (0)
CCC CRAY-2	v p (12)	v v (91)	v v (117)	v p (21)	v p (35)
CRI CRAY Y-MP	v p (19)	v v (107)	v v (144)	v p (31)	v v (100)
DEC VAX 9000-210	n n (0)	v n (0)	v v (860)	v p (11)	n n (D)
FPS M511EA-2	n n (0)	v v (85)	v v (121)	v p (29)	v v (100)
Fujitsu VP2600/10	v p (20)	v v (72)	v v (113)	v p (25)	n n (0)
Hitachi S-820/80	v p (19)	v v (92)	v v (222)	v p (26)	n n (0)
IBM 3090-600J	n n (0)	v v (99)	v v (186)	v p (14)	v v (100)
NEC SX-X/14	v p (13)	v v (287)	v v (175)	v p (47)	n n (0)

Computer	s412	s413	s414	s415	s421
CONVEX C-210	n n (0)	n n (1)	p n (D)	n n (0)	v v (99)
CCC CRAY-2	v p (25)	v p (46)	n <sup>a</sup> n (D)	n n (0)	v v (98)
CRI CRAY Y-MP	v v (100)	v v (76)	v v (79)	n n (0)	v v (81)
DEC VAX 9000-210	n n (D)	n n (D)	n n (0)	n n (D)	v v (112)
FPS M511EA-2	v v (100)	v v (74)	v v (100)	n n (D)	v v (86)
Fujitsu VP2600/10	n n (0)	n n (0)	p <sup>b</sup> n (0)	n n (0)	v v (105)
Hitachi S-820/80	n n (0)	n n (0)	v p (11)	n n (0)	v v (103)
IBM 3090-600J	v v (100)	v v (67)	v v (68)	n n (0)	v v (96)
NEC SX-X/14	n n (0)	n n (0)	v n (5)	n n (0)	v v (77)

Computer	s422	s423	s424	s431	s432
CONVEX C-210	v v (100)	v v (96)	n n (D)	v v (100)	v v (100)
CCC CRAY-2	v v (97)	v v (96)	v v (76)	v v (97)	v <sup>a</sup> v (68)
CRI CRAY Y-MP	v v (80)	v v (79)	v v (104)	v v (104)	v <sup>c</sup> v (80)
DEC VAX 9000-210	v v (111)	v v (113)	n n (7)	v v (136)	v v (109)
FPS M511EA-2	v v (86)	n n (0)	n n (0)	v v (86)	v <sup>c</sup> v (86)
Fujitsu VP2600/10	n n (0)	n n (0)	n n (0)	v v (104)	v v (105)
Hitachi S-820/80	n n (D)	n n (D)	n n (D)	v v (96)	n n (D)
IBM 3090-600J	v v (99)	p p (18)	n p (13)	v v (101)	v <sup>c</sup> v (101)
NEC SX-X/14	v v (81)	v v (82)	n n (0)	v v (83)	v <sup>c</sup> v (82)

Computer	s441	s442	s443	s451	s452
CONVEX C-210	v p (36)	n n (1)	v v (67)	v v (53)	v v (72)
CCC CRAY-2	n n (0)	n n (0)	v p (30)	v p (43)	v v (76)
CRI CRAY Y-MP	v n (10)	v n (4)	v v (56)	v p (30)	v v (67)
DEC VAX 9000-210	n n (1)	n n (7)	v p (16)	v v (64)	v v (56)
FPS M511EA-2	v n (3)	v n (D)	v v (96)	v v (91)	v p (44)
Fujitsu VP2600/10	v p (36)	n n (0)	v v (61)	v v (71)	v v (75)
Hitachi S-820/80	v p (41)	p n (D)	v v (92)	v p (42)	v v (74)
IBM 3090-600J	v n (D)	v n (D)	v p (48)	v n (3)	v p (40)
NEC SX-X/14	v p (21)	n n (0)	v p (42)	v p (33)	v v (63)

Computer	s453	s471	s481	s482	s491
CONVEX C-210	v v (61)	n n (0)	n n (1)	n n (1)	v v (118)
CCC CRAY-2	v p (45)	n n (0)	v v (56)	v v (56)	v v (123)
CRI CRAY Y-MP	v p (41)	n n (0)	v v (59)	v v (70)	v v (106)
DEC VAX 9000-210	n n (0)	v n (D)	n n (D)	n n (D)	v v (149)
FPS M511EA-2	v p (36)	n n (2)	n n (0)	n n (0)	v v (131)
Fujitsu VP2600/10	v n (2)	v v (58)	v v (97)	v n (0)	v v (66)
Hitachi S-820/80	v n (9)	n n (0)	v v (82)	n n (0)	v v (122)
IBM 3090-600J	v p (18)	v v (115)	v v (90)	v v (68)	v v (130)
NEC SX-X/14	n n (0)	n n (0)	n n (0)	n n (0)	v v (71)

Computer	s4112	s4113	s4114	s4115	s4116
CONVEX C-210	v v (99)	v v (125)	v v (119)	v v (79)	v v (112)
CCC CRAY-2	v v (97)	v v (107)	v v (115)	v v (86)	v v (185)
CRI CRAY Y-MP	v v (110)	v v (81)	v v (157)	v v (126)	v v (190)
DEC VAX 9000-210	v v (74)	v v (85)	v v (124)	v v (61)	v p (41)
FPS M511EA-2	v v (112)	v v (138)	v v (109)	v v (128)	v v (150)
Fujitsu VP2600/10	v v (109)	v v (68)	v v (151)	v v (131)	v v (175)
Hitachi S-820/80	v v (105)	v v (117)	v v (89)	v v (65)	v v (103)
IBM 3090-600J	v v (77)	v v (140)	v v (123)	v v (73)	v v (132)
NEC SX-X/14	v v (108)	v v (63)	v v (101)	v v (101)	v v (181)

Computer	s4117	s4121
CONVEX C-210	v p (39)	v v (100)
CCC CRAY-2	v p (41)	v v (100)
CRI CRAY Y-MP	v p (36)	v v (100)
DEC VAX 9000-210	v v (67)	v v (138)
FPS M511EA-2	v p (37)	v v (101)
Fujitsu VP2600/10	v p (38)	v v (99)
Hitachi S-820/80	v v (63)	v v (100)
IBM 3090-600J	v p (30)	v v (100)
NEC SX-X/14	v n (9)	v v (100)

## Appendix C

The tables below contain the operation counts for each loop in the test suite.

Loop	Load	Store	Gather	Scatter	Arithmetic	Reductions
s111	0 2	0 1	0 0	0 0	0 1	0 0
s112	0 2	0 1	0 0	0 0	0 1	0 0
s113	0 1	0 1	0 0	0 0	0 1	0 0
s114	0 2	0 1	0 0	0 0	0 1	0 0
s115	0 2	0 1	0 0	0 0	0 2	0 0
s116	0 5	0 5	0 0	0 0	0 5	0 0
s118	0 2	0 0	0 0	0 0	0 2	0 0
s119	0 2	0 1	0 0	0 0	0 1	0 0
s121	0 2	0 1	0 0	0 0	0 1	0 0
s122	0 2	0 1	0 0	0 0	0 1	0 0
s123	0 4	0 1	0 0	0 0	0 5	0 0
s124	0 4	0 1	0 0	0 0	0 5	0 0
s125	0 3	0 1	0 0	0 0	0 2	0 0
s126	0 3	0 1	0 0	0 0	0 2	0 0
s127	0 4	0 2	0 0	0 0	0 4	0 0
s128	0 3	0 2	0 0	0 0	0 2	0 0
s131	0 2	0 1	0 0	0 0	0 1	0 0
s132	0 2	0 1	0 0	0 0	0 2	0 0
s141	0 1	0 1	0 1	0 1	0 1	0 0
s151	0 2	0 1	0 0	0 0	0 1	0 0
s152	0 4	0 2	0 0	0 0	0 3	0 0
s161	0 6	0 2	0 0	0 0	0 7	0 0
s162	0 3	0 1	0 0	0 0	0 2	0 0
s171	0 2	0 1	0 0	0 0	0 1	0 0
s172	0 2	0 1	0 0	0 0	0 1	0 0
s173	0 2	0 1	0 0	0 0	0 1	0 0
s174	0 2	0 1	0 0	0 0	0 1	0 0
s175	0 2	0 1	0 0	0 0	0 1	0 0
s176	0 2	0 0	0 0	0 0	0 2	0 0
s211	0 5	0 2	0 0	0 0	0 4	0 0
s212	0 5	0 2	0 0	0 0	0 3	0 0
s221	0 4	0 2	0 0	0 0	1 3	0 0

Loop	Load	Store	Gather	Scatter	Arithmetic	Reductions
s222	2 6	1 2	0 0	0 0	2 4	0 0
s231	0 2	0 1	0 0	0 0	0 1	0 0
s232	0 2	0 1	0 0	0 0	0 2	0 0
s233	0 4	0 2	0 0	0 0	0 2	0 0
s234	0 3	0 1	0 0	0 0	0 2	0 0
s235	0 3	0 1	0 0	0 0	0 2	0 0
s241	0 4	0 2	0 0	0 0	0 4	0 0
s242	0 3	0 1	0 0	0 0	1 3	0 0
s243	0 5	0 2	0 0	0 0	0 6	0 0
s244	0 4	0 2	0 0	0 0	0 5	0 0
s251	0 3	0 1	0 0	0 0	0 3	0 0
s252	0 2	0 1	0 0	0 0	0 2	0 0
s253	0 4	0 2	0 0	0 0	0 6	0 0
s254	0 1	0 1	0 0	0 0	0 2	0 0
s255	0 1	0 1	0 0	0 0	0 3	0 0
s256	0 2	0 1	0 0	0 0	0 2	0 0
s257	0 2	0 1	0 0	0 0	0 2	0 0
s258	0 4	0 2	0 0	0 0	0 7	0 0
s261	0 5	0 2	0 0	0 0	0 3	0 0
s271	0 3	0 1	0 0	0 0	0 4	0 0
s272	0 5	0 2	0 0	0 0	0 7	0 0
s273	0 5	0 3	0 0	0 0	0 8	0 0
s274	0 5	0 2	0 0	0 0	0 4	0 0
s275	0 4	0 1	0 0	0 0	0 3	0 0
s276	0 3	0 1	0 0	0 0	0 2	0 0
s277	0 5	0 2	0 0	0 0	0 8	0 0
s278	0 5	0 3	0 0	0 0	0 1	0 0
s279	0 5	0 3	0 0	0 0	0 1	0 0
s2710	0 5	0 4	0 0	0 0	0 2	0 0
s2711	0 3	0 1	0 0	0 0	0 2	0 0
s2712	0 3	0 1	0 0	0 0	0 4	0 0
s281	0 2	0 1	0 0	0 0	0 2	0 0
s291	0 2	0 1	0 0	0 0	0 2	0 0
s292	0 3	0 1	0 0	0 0	0 2	0 0
s293	0 0	0 1	0 0	0 0	0 0	0 0
s2101	0 3	0 1	0 0	0 0	0 2	0 0
s2102	0 0	0 1	0 0	0 0	0 0	0 0
s2111	0 2	0 1	0 0	0 0	0 1	0 0

Loop	Load	Store	Gather	Scatter	Arithmetic	Reductions
s311	0 1	0 0	0 0	0 0	0 0	0 1
s312	0 1	0 0	0 0	0 0	0 0	0 1
s313	0 2	0 0	0 0	0 0	0 1	0 1
s314	0 1	0 0	0 0	0 0	0 0	0 1
s315	0 1	0 0	0 0	0 0	0 0	0 1
s316	0 1	0 0	0 0	0 0	0 0	0 1
s317	0 0	0 0	0 0	0 0	0 0	0 1
s318	0 1	0 0	0 0	0 0	0 1	0 1
s319	0 3	0 2	0 0	0 0	0 2	0 2
s3110	0 1	0 0	0 0	0 0	0 0	0 1
s3111	0 1	0 0	0 0	0 0	0 1	0 1
s3112	0 1	0 1	0 0	0 0	0 0	0 1
s3113	0 1	0 0	0 0	0 0	0 1	0 1
s321	0 2	1 0	0 0	0 0	2 0	0 0
s322	0 3	1 0	0 0	0 0	4 0	0 0
s323	0 3	0 2	0 0	0 0	2 2	0 0
s331	0 1	0 0	0 0	0 0	0 0	0 1
s332	0 1	0 0	0 0	0 0	0 1	0 0
s341	0 1	0 1	0 0	0 0	0 1	0 0
s342	0 1	0 1	0 0	0 0	0 1	0 0
s343	0 2	0 1	0 0	0 0	0 1	0 0
s351	0 10	0 5	0 0	0 0	0 1	0 0
s352	0 10	0 5	0 0	0 0	0 9	0 1
s353	0 2	0 1	0 1	0 0	0 1	0 0
s411	0 3	0 1	0 0	0 0	0 2	0 0
s412	0 3	0 1	0 0	0 0	0 2	0 0
s413	0 4	0 2	0 0	0 0	0 4	0 0
s414	0 3	0 1	0 0	0 0	0 2	0 0
s415	0 3	0 1	0 0	0 0	0 2	0 0
s421	0 2	0 1	0 0	0 0	0 1	0 0
s422	0 2	0 1	0 0	0 0	0 1	0 0
s423	0 2	0 1	0 0	0 0	0 1	0 0
s424	0 2	0 1	0 0	0 0	0 1	0 0
s431	0 2	0 1	0 0	0 0	0 1	0 0
s432	0 2	0 1	0 0	0 0	0 1	0 0
s441	0 3	0 1	0 0	0 0	0 1	0 0
s442	0 5	0 1	0 0	0 0	0 1	0 0
s443	0 3	0 1	0 0	0 0	0 7	0 0

Loop	Load	Store	Gather	Scatter	Arithmetic	Reductions
s451	0 28	0 1	0 0	0 0	0 2	0 0
s452	0 2	0 1	0 0	0 0	0 2	0 0
s453	0 1	0 1	0 0	0 0	0 1	0 1
s471	0 3	0 1	0 0	0 0	0 2	0 0
s481	0 4	0 1	0 0	0 0	0 3	0 0
s482	0 4	0 1	0 0	0 0	0 3	0 0
s491	0 4	0 0	0 0	0 1	0 2	0 0
s4112	0 1	0 1	0 1	0 0	0 2	0 0
s4113	0 1	0 0	0 1	0 1	0 1	0 0
s4114	0 3	0 1	0 1	0 0	0 3	0 0
s4115	0 1	0 0	0 1	0 0	0 1	0 1
s4116	0 0	0 0	0 2	0 0	0 1	0 1
s4117	0 2	0 1	0 1	0 0	0 3	0 0
s4121	0 3	0 1	0 0	0 0	0 2	0 0
va	0 1	0 1	0 0	0 0	0 0	0 0
vpv	0 2	0 1	0 0	0 0	0 1	0 0
vtv	0 2	0 1	0 0	0 0	0 1	0 0
vpvts	0 2	0 1	0 0	0 0	0 2	0 0
vpvtv	0 3	0 1	0 0	0 0	0 2	0 0
vpvpv	0 3	0 1	0 0	0 0	0 2	0 0
vtvtv	0 3	0 1	0 0	0 0	0 2	0 0
vbor	0 6	0 1	0 0	0 0	0 59	0 0
vif	0 2	0 1	0 0	0 0	0 2	0 0
vag	0 0	0 1	0 1	0 0	0 0	0 0
vas	0 2	0 0	0 0	0 1	0 0	0 0
vsumr	0 1	0 0	0 0	0 0	0 0	0 1
vdotr	0 2	0 0	0 0	0 0	0 1	0 1

# Appendix D

```

c$1.1  subroutine s111 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c   linear dependence testing
c   no dependence - vectorizable
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s111 ')
t1 = second()
do 1 n1 = 1,2*ntimes
do 10 i = 2,n,2
  a(i) = a(i-1) + b(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(2*ntimes) )
chksum = csld(n,a)
call check (chksum,2*ntimes*(n/2),n,t2,'s111 ')
return
end

c$1.1  subroutine s112 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c   linear dependence testing
c   loop reversal
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s112 ')
t1 = second()
do 1 n1 = 1,ntimes
do 10 i = n-1,1,-1
  a(i+1) = a(i) + b(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a)
call check (chksum,ntimes*(n-1),n,t2,'s112 ')
return
end

c$1.1  subroutine s113 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c   linear dependence testing
c   a(i)=a(i) but no actual dependence cycle
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s113 ')
t1 = second()
do 1 n1 = 1,ntimes
do 10 i = 2,n
  a(i) = a(1) + b(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a)

```

```

call check (chksum,ntimes*(n-1),n,t2,'s113 ')
return
end

c$1.1  subroutine s114 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c   linear dependence testing
c   transpose vectorization
c
integer ntimes, ld, n, i, n1, j
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, cs2d

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s114 ')
t1 = second()
do 1 n1 = 1,2*(ntimes/n)
do 10 j = 1,n
do 20 i = 1,j-1
  aa(i,j) = aa(j,i) + bb(i,j)
20 continue
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(2*(ntimes/n)) )
chksum = cs2d(n,aa)
call check (chksum,2*(ntimes/n)*((n*n-n)/2),n,t2,'s114 ')
return
end

c$1.1  subroutine s115 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c   linear dependence testing
c   triangular saxpy loop
c
integer ntimes, ld, n, i, n1, j
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s115 ')
t1 = second()
do 1 n1 = 1,2*(ntimes/n)
do 10 j = 1,n
do 20 i = j+1, n
  a(i) = a(i) - aa(i,j) * a(j)
20 continue
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(2*(ntimes/n)) )
chksum = csld(n,a)
call check (chksum,2*(ntimes/n)*((n*n-n)/2),n,t2,'s115 ')
return
end

c$1.1  subroutine s116 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c   linear dependence testing
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s116 ')
t1 = second()
do 1 n1 = 1,5*ntimes
do 10 i = 1,n-5,5

```





```

      a(j) = c(i)+ d(i) * e(i)
    endif
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - cttime - ( dttime * float(2*ntimes) )
chksum = csld(n,a)
call check (chksum,2*ntimes*(n/2),n,t2,'s123 ')
return
end
c#1.2
subroutine s124 (ntimes,ld,n,ctime,dttime,a,b,c,d,e,aa,bb,cc)
c
c Induction variable recognition
c Induction variable under both sides of if (same value)
c
integer ntimes, ld, n, i, nl, j
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dttime, csld

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s124 ')
t1 = second()
do 1 nl = 1,2*ntimes
  j = 0
  do 10 i = 1,n/2
    if(b(i).gt. 0.) then
      a(j) = b(i) + d(i) * e(i)
    else
      j = j + 1
      a(j) = c(i) + d(i) * e(i)
    endif
  end
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - cttime - ( dttime * float(2*ntimes) )
chksum = csld(n,a)
return
end
c#1.2
subroutine s125 (ntimes,ld,n,ctime,dttime,a,b,c,d,e,aa,bb,cc)
c
c Induction variable recognition
c Induction variable in two loops; collapsing possible
c
integer ntimes, ld, n, i, nl, j, k, nn
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
parameter(nn=1000)
common /cdata/ array(nn*nn)

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s125 ')
t1 = second()
do 1 nl = 1,ntimes/n
  k = 0
  do 20 i = 1,n
    k = k + 1
    array(k) = aa(i,j) + bb(i,j) * cc(i,j)
  end
20 continue
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - cttime - ( dttime * float(2*ntimes) )
chksum = csld(n,a)
return
end
c#1.2
subroutine s126 (ntimes,ld,n,ctime,dttime,a,b,c,d,e,aa,bb,cc)
c
c Induction variable recognition
c Induction variable in two loops; recurrence in inner loop
c
integer ntimes, ld, n, i, nl, j, k, nn
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dttime, cs2d, array
parameter(nn=1000)
common /cdata/ array(nn*nn)

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s126 ')
t1 = second()
do 1 nl = 1,ntimes/n
  k = 1
  do 10 i = 1,n
    do 20 j = 2,n
      bb(i,j) = bb(i,j-1) + array(k) * cc(i,j)
      k = k + 1
    end
  10 continue
  call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
  1 continue
  t2 = second() - t1 - cttime - ( dttime * float(ntimes/n) )
  chksum = cs2d(n,bb)
  call check (chksum,(ntimes/n)*(n-1),n,t2,'s126 ')
  return
end
c#1.2
subroutine s127 (ntimes,ld,n,ctime,dttime,a,b,c,d,e,aa,bb,cc)
c
c Induction variable recognition
c Induction variable with multiple increments
c
integer ntimes, ld, n, i, nl, j
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dttime, csid

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s127 ')
t1 = second()
do 1 nl = 1,2*ntimes
  j = 0
  do 10 i = 1,n/2
    j = j + 1
    a(j) = b(i) + c(i) * d(i)
    j = j + 1
    a(j) = b(i) + d(i) * e(i)
  end
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - cttime - ( dttime * float(2*ntimes) )
chksum = csld(n,a)
call check (chksum,2*ntimes*(n/2),n,t2,'s127 ')
return
end
c#1.2
subroutine s128 (ntimes,ld,n,ctime,dttime,a,b,c,d,e,aa,bb,cc)
c
c Induction variables

```

```

c
integer ntimes, ld, n, i, ni, j, k
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s128 ')
t1 = second()
do 1 ni = 1,2*ntimes
  j = 0
  do 10 i = 1, n/2
    k = j + 1
    a(i) = b(k) - d(i)
    j = k + 1
    b(k) = a(i) + c(k)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(2*ntimes) )
chksum = csld(n,a) + csld(n,b)
call check (chksum,2*ntimes*(n/2),n,t2,'s128 ')
return
end

c#1.3
subroutine s131 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c Global data flow analysis
c forward substitution
c
integer ntimes, ld, n, i, ni, m
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld

m = 1
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s131 ')
if(a(1).gt.0)then
  a(1) = b(1)
endif
t1 = second()
do 1 ni = 1,ntimes
  do 10 i = 1,n-1
    a(i) = a(i+m) + b(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a)
call check (chksum,ntimes*(n-1),n,t2,'s131 ')
return
end

c#1.3
subroutine s132 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c Global data flow analysis
c loop with multiple dimension ambiguous subscripts
c
integer ntimes, ld, n, i, ni, j, k, m
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, cs2d

m = 1
j = m
k = m+1
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s132 ')
t1 = second()
do 1 ni = 1,ntimes
  do 10 i = 1,2,n
    aa(i,j) = aa(i-1,k) + b(i) * c(k)
    continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = cs2d(n,aa)
call check (chksum,ntimes*n-1,n,t2,'s132 ')
return
end

c#1.4
subroutine s141 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c nonlinear dependence testing
c walk a row in a symmetric packed array
c element a(i,j) for (j>1) stored in location j*(j-1)/2+1
c
integer ntimes, ld, n, i, ni, j, k, nn
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld, array
parameter(nn=1000)
common /cdata/ array(nn*nn)

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s141 ')
t1 = second()
do 1 ni = 1,ntimes/n
  do 10 i = 1,n
    k = i*(i-1)/2+1
    do 20 j = i,n
      array(k) = array(k) + bb(i,j)
      k = k + j
20 continue
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes/n) )
chksum = csld(n,array)
call check (chksum,(ntimes/n)*n*n,n,t2,'s141 ')
return
end

c#1.5
subroutine s151 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c Interprocedural data flow analysis
c passing parameter information into a subroutine
c
integer ntimes, ld, n, ni
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s151 ')
t1 = second()
do 1 ni = 1,ntimes
  call s151s(a,b,n,1)
  call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a)
call check (chksum,ntimes*(n-1),n,t2,'s151 ')
return
end

subroutine s151s(a,b,n,m)
integer j, n, m
real a(n), b(n)
do 10 i = 1,n-1
  a(i) = a(i+m) + b(i)
10 continue
end

```





```

real t1, t2, second, chksum, cttime, dttime, csld
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s212 ')
t1 = second()
do 1 n1 = 1, ntimes
do 10 i=1, n-1
a(i) = a(i) * c(i)
b(i) = b(i) + a(i+1) * d(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - cttime - ( dttime * float(ntimes) )
chksum = csld(n,a) + csld(n,b)
call check (chksum, ntimes*(n-1), n, t2, 's212 ')
return
end
c#2.2
subroutine s221 (ntimes,ld,n,cttime,dttime,a,b,c,d,e,aa,bb,cc)
c
c loop distribution
c loop that is partially recursive
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, cttime, dttime, csld
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s221 ')
t1 = second()
do 1 n1 = 1, ntimes
do 10 i = 2, n
a(i) = a(i) + c(i) * d(i)
b(i) = b(i-1) + a(i) + d(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - cttime - ( dttime * float(ntimes) )
chksum = csld(n,a) + csld(n,b)
call check (chksum, ntimes*(n-1), n, t2, 's221 ')
return
end
c#2.2
subroutine s222 (ntimes,ld,n,cttime,dttime,a,b,c,d,e,aa,bb,cc)
c
c loop distribution
c partial loop vectorization, recurrence in middle
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, cttime, dttime, csld
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s222 ')
t1 = second()
do 1 n1 = 1, ntimes
do 10 i = 2, n
a(i) = a(i) + b(i) * c(i)
b(i) = b(i-1) * b(i-1) * a(i)
a(i) = a(i) - b(i) * c(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - cttime - ( dttime * float(ntimes) )
chksum = csld(n,a) + csld(n,b)
call check (chksum, ntimes*(n-1), n, t2, 's222 ')
return
end
c#2.3
subroutine s231 (ntimes,ld,n,cttime,dttime,a,b,c,d,e,aa,bb,cc)
c
c loop interchange
c loop with multiple dimension recursion
c
integer ntimes, ld, n, i, n1, j
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, cttime, dttime, cs2d
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s231 ')
t1 = second()
do 1 n1 = 1, ntimes/n
do 10 i=1, n
do 20 j=2, n
aa(i,j) = aa(i,j-1) + bb(i,j)
20 continue
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - cttime - ( dttime * float(ntimes/n) )
chksum = cs2d(n,aa)
call check (chksum, ntimes/n*(n-1), n, t2, 's231 ')
return
end
c#2.3
subroutine s232 (ntimes,ld,n,cttime,dttime,a,b,c,d,e,aa,bb,cc)
c
c loop interchange
c interchanging of triangular loops
c
integer ntimes, ld, n, i, n1, j
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, cttime, dttime, cs2d
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s232 ')
t1 = second()
do 1 n1 = 1, 2*(ntimes/n)
do 10 j = 2, n
do 20 i = 2, j
aa(i,j) = aa(i-1,j)*aa(i-1,j)+bb(i,j)
20 continue
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - cttime - ( dttime * float(2*(ntimes/n)) )
chksum = cs2d(n,aa)
call check (chksum, 2*(ntimes/n)*(n*(n-1)/2), n, t2, 's232 ')
return
end
c#2.3
subroutine s233 (ntimes,ld,n,cttime,dttime,a,b,c,d,e,aa,bb,cc)
c
c loop interchange
c interchanging with one of two inner loops
c
integer ntimes, ld, n, i, n1, j
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, cttime, dttime, cs2d
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s233 ')
t1 = second()
do 1 n1 = 1, ntimes/n
do 10 i = 2, n
do 20 j = 2, n
aa(i,j) = aa(i,j-1) + cc(i,j)
20 continue
10 continue
end

```

```

do 30 j = 2,n
  bb(i,j) = bb(i-1,j) + cc(i,j)
continue
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes/n) )
chksum = cs2d(n,aa) + cs2d(n,bb)
call check (chksum,(ntimes/n)*(n-1)*(2*n-2),n,t2,'s233 ')
return
end

c#2.3
subroutine s234 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c loop interchange
c if loop to do loop, interchanging with if loop necessary
c
integer ntimes, ld, n, i, n1, j
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, cs2d

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s234 ')
do 1 n1 = 1,ntimes/n
  i = 1
11 if(i.gt.n) goto 10
  j = 2
  if(j.gt.n) goto 20
  aa(i,j) = aa(i,j-1) + bb(i,j-1) * cc(i,j-1)
  j = j + 1
  goto 21
20 i = i + 1
  goto 11
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes/n) )
chksum = cs2d(n,aa)
call check (chksum,(ntimes/n)*n*(n-1),n,t2,'s234 ')
return
end

c#2.3
subroutine s235 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c loop interchanging
c imperfectly nested loops
c
integer ntimes, ld, n, i, n1, j
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld, cs2d

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s235 ')
t1 = second()
do 1 n1 = 1,ntimes/n
  do 10 i = 1,n
    a(i) = a(i) + b(i) * c(i)
    do 20 j = 2,n
      aa(i,j) = aa(i,j-1) + bb(i,j) * a(i)
    continue
  10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes/n) )
chksum = cs2d(n,aa) + csld(n,a)
call check (chksum,(ntimes/n)*n*(n-1),n,t2,'s235 ')
return
end

do 30 j = 2,n
  bb(i,j) = bb(i-1,j) + cc(i,j)
continue
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes/n) )
chksum = cs2d(n,aa) + cs2d(n,bb)
call check (chksum,(ntimes/n)*(n-1)*(2*n-2),n,t2,'s233 ')
return
end

c#2.4
subroutine s241 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c node splitting
c preloading necessary to allow vectorization
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s241 ')
t1 = second()
do 1 n1 = 1,ntimes
  do 10 i = 1,n-1
    a(i) = b(i) * c(i) * d(i)
    b(i) = a(i) * a(i+1) * d(i)
  10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a) + csld(n,b)
call check (chksum,ntimes*(n-1),n,t2,'s241 ')
return
end

c#2.4
subroutine s242 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc,sl,s2)
c
c node splitting
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld, sl, s2

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s242 ')
t1 = second()
do 1 n1 = 1,ntimes
  do 10 i = 2,n
    a(i) = a(i-1) + sl + s2 + b(i) + c(i) + d(i)
  10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a)
call check (chksum,ntimes*(n-1),n,t2,'s242 ')
return
end

c#2.4
subroutine s243 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c node splitting
c false dependence cycle breaking
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s243 ')
t1 = second()
do 1 n1 = 1,ntimes
  do 10 i = 1,n-1
    a(i) = b(i) + c(i) * d(i)
    b(i) = a(i) + d(i) * c(i)
    a(i) = b(i) + a(i+1) * d(i)
  10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
end

```

```

1 continue
t2 = second() - t1 - cttime - ( dtime * float(ntimes) )
chksum = csld(n,a) + csld(n,b)
call check (chksum,ntimes*(n-1),n,t2,'s243 ')
return
end

ct#2.4
subroutine s244 (ntimes,ld,n,cttime,dtime,a,b,c,d,e,aa,bb,cc)
c
c node splitting
c false dependence cycle breaking
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, cttime, dtime, csld

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s244 ')
t1 = second()
do 10 i = 1,n-1
  a(i) = b(i) + c(i) * d(i)
  b(i) = c(i) + b(i)
  a(i+1) = b(i) + a(i+1) * d(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - cttime - ( dtime * float(ntimes) )
chksum = csld(n,a) + csld(n,b)
call check (chksum,ntimes*(n-1),n,t2,'s244 ')
return
end

ct#2.5
subroutine s251 (ntimes,ld,n,cttime,dtime,a,b,c,d,e,aa,bb,cc)
c
c scalar and array expansion
c scalar expansion
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, cttime, dtime, csld, s

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s251 ')
t1 = second()
do 10 i = 1,n
  s = b(i) + c(i) * d(i)
  a(i) = s + s
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - cttime - ( dtime * float(ntimes) )
chksum = csld(n,a) + csld(n,b)
call check (chksum,ntimes*(n-1),n,t2,'s251 ')
return
end

ct#2.5
subroutine s252 (ntimes,ld,n,cttime,dtime,a,b,c,d,e,aa,bb,cc)
c
c scalar and array expansion
c loop with ambiguous scalar temporary
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, cttime, dtime, csld, s, t

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s252 ')

```

```

t1 = second()
do 10 i = 1,ntimes
  t = 0.
  do 10 i=1,n
    s = b(i) * c(i)
    a(i) = s + t
    t = s
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - cttime - ( dtime * float(ntimes) )
chksum = csld(n,a)
return
end

ct#2.5
subroutine s253 (ntimes,ld,n,cttime,dtime,a,b,c,d,e,aa,bb,cc)
c
c scalar and array expansion
c scalar expansion, assigned under if
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, cttime, dtime, csld, s

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s253 ')
t1 = second()
do 10 i = 1,ntimes
  do 10 i = 1,n
    if(a(i) .gt. b(i)) then
      s = a(i) - b(i) * d(i)
      c(i) = c(i) + s
      a(i) = s
    endif
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - cttime - ( dtime * float(ntimes) )
chksum = csld(n,a) + csld(n,c)
call check (chksum,ntimes*n,n,t2,'s253 ')
return
end

ct#2.5
subroutine s254 (ntimes,ld,n,cttime,dtime,a,b,c,d,e,aa,bb,cc)
c
c scalar and array expansion
c carry around variable
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, cttime, dtime, csld, x

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s254 ')
t1 = second()
do 10 i = 1,ntimes
  x = b(n)
  do 10 i = 1,n
    a(i) = (b(i) + x) * .5
    x = b(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - cttime - ( dtime * float(ntimes) )
chksum = csld(n,a)
call check (chksum,ntimes*n,n,t2,'s254 ')

```



```

end
c%2.5 subroutine s255 (ntimes,ld,n,ctime,dttime,a,b,c,d,e,aa,bb,cc)
c scalar and array expansion
c carry around variables, 2 levels
c
c integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dttime, csld, x, y
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s255 ')
t1 = second()
do 10 i = 1, ntimes
  x = b(n)
  y = b(n-1)
do 10 i = 1, n
  a(i) = (b(i) + x + y) * .333
  y = x
  x = b(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a)
call check (chksum,ntimes*n,t2,'s255 ')
return
end
c%2.5 subroutine s256 (ntimes,ld,n,ctime,dttime,a,b,c,d,e,aa,bb,cc)
c scalar and array expansion
c array expansion
c
c integer ntimes, ld, n, i, n1, j
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dttime, csld, cs2d
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s256 ')
t1 = second()
do 1 n1 = 1, ntimes/n
do 10 i = 1, n
do 20 j = 2, n
  a(j) = aa(i,j) - a(j-1)
  aa(i,j) = a(j) + bb(i,j)
20 continue
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes/n) )
chksum = csld(n,a) + cs2d(n,aa)
call check (chksum,ntimes/n)*(n-1),n,t2,'s256 ')
return
end
c%2.5 subroutine s257 (ntimes,ld,n,ctime,dttime,a,b,c,d,e,aa,bb,cc)
c scalar and array expansion
c array expansion
c
c integer ntimes, ld, n, i, n1, j
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dttime, csld, cs2d
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s257 ')
t1 = second()
do 1 n1 = 1, ntimes/n
do 20 j = 1, ntimes
  t = a(i) + b(i)
  a(i) = t + c(i-1)
  t = c(i) * d(i)
  c(i) = t
20 continue
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a) + csld(n,c)
call check (chksum,ntimes*(n-1),n,t2,'s261 ')
return
end
do 1 n1 = 1, ntimes/n
do 10 i = 2, n
do 20 j = 1, n
  a(i) = aa(i,j) - a(i-1)
  aa(i,j) = a(i) + bb(i,j)
20 continue
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes/n) )
chksum = csld(n,a) + cs2d(n,aa)
call check (chksum,ntimes/n)*(n-1),n,t2,'s257 ')
return
end
c%2.5 subroutine s258 (ntimes,ld,n,ctime,dttime,a,b,c,d,e,aa,bb,cc)
c scalar and array expansion
c wrap-around scalar under an if
c
c integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dttime, csld, s
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s258 ')
t1 = second()
do 1 n1 = 1, ntimes
  s = 0.
do 10 i = 1, n
  if (a(i) .gt. 0.) then
    s = d(i) * d(i)
  endif
  b(i) = s * c(i) + d(i)
  e(i) = (s + 1.) * aa(i,1)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,b) + csld(n,e)
call check (chksum,ntimes*n,t2,'s258 ')
return
end
c%2.6 subroutine s261 (ntimes,ld,n,ctime,dttime,a,b,c,d,e,aa,bb,cc)
c scalar renaming
c
c integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dttime, csld, t
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s261 ')
t1 = second()
do 1 n1 = 1, ntimes
do 10 i = 2, n
  t = a(i) + b(i)
  a(i) = t + c(i-1)
  t = c(i) * d(i)
  c(i) = t
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a) + csld(n,c)
call check (chksum,ntimes*(n-1),n,t2,'s261 ')
return
end

```

```

c42.7
end
subroutine s271 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c control flow
c loop with singularity handling
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s271 ')
t1 = second()
do 10 i = 1, n
  if (b(i) .gt. 0.) a(i) = a(i) + b(i) * c(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a)
call check (chksum,ntimes*n,t2,'s271 ')
return
end
c42.7
subroutine s272 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc,t)
c
c control flow
c loop with independent conditional
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld, t
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s272 ')
t1 = second()
do 10 i = 1, n
  if (e(i) .ge. t) then
    a(i) = a(i) + c(i) * d(i)
    b(i) = b(i) + c(i) * c(i)
  endif
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a)
call check (chksum,ntimes*n,t2,'s272 ')
return
end
c42.7
subroutine s273 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c control flow
c simple loop with dependent conditional
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s273 ')
t1 = second()
do 10 i = 1, ntimes
  a(i) = a(i) + d(i) * e(i)
  if (a(i) .lt. 0.) b(i) = b(i) + d(i) * e(i)
10 continue
end
c42.7
subroutine s274 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c control flow
c complex loop with dependent conditional
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s274 ')
t1 = second()
do 10 i = 1, n
  a(i) = c(i) + e(i) * d(i)
  if (a(i) .gt. 0.) then
    b(i) = a(i) + b(i)
  else
    a(i) = d(i) * e(i)
  endif
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a) + csld(n,b)
call check (chksum,ntimes*n,t2,'s274 ')
return
end
c42.7
subroutine s275 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c control flow
c if around inner loop, interchanging needed
integer ntimes, ld, n, i, n1, j
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, cs2d
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s275 ')
t1 = second()
do 10 i = 1, ntimes/n
  do 10 i = 2, n
    if (aa(i,1) .gt. 0.) then
      do 20 j = 2, n
        aa(i,j) = aa(i,j-1) + bb(i,j) * cc(i,j)
      20 continue
    endif
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
t2 = second() - t1 - ctime - ( dtime * float(ntimes/n) )
chksum = cs2d(n,aa)
call check (chksum, (ntimes/n) * (n-1) * (n-1), n, t2, 's275 ')
return
end
c42.7

```

```

c
c control flow
c if test using loop index
c
integer ntimes, ld, n, i, n1, mid
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s276 ')
t1 = second()
mid = n/2
do 10 i = 1, ntimes
do 10 i = 1, n
if ( i.lt. mid ) then
a(i) = a(i) + b(i) * c(i)
else
a(i) = a(i) + b(i) * d(i)
endif
endif
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a)
call check (chksum,ntimes*n,t2,'s276 ')
return
end

c%2.7
subroutine s277 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c control flow
c test for dependences arising from guard variable computation.
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s277 ')
t1 = second()
do 10 i = 1, ntimes
do 10 i = 1, n-1
if (a(i) .ge. 0.) go to 20
if (b(i) .ge. 0.) go to 30
a(i) = a(i) + c(i) * d(i)
30 continue
b(i+1) = c(i) + d(i) * e(i)
20 continue
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a) + csld(n,b)
call check (chksum,ntimes*(n-1),n,t2,'s277 ')
return
end

c%2.7
subroutine s278 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c control flow
c if/goto to block if-then-else
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s278 ')
t1 = second()

```

```

do 1 n1 = 1, ntimes
do 10 i = 1, n
if (a(i) .gt. 0.) goto 20
b(i) = -b(i) + d(i) * e(i)
20 continue
30 continue
10 continue
a(i) = b(i) + c(i) * d(i)
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a) + csld(n,b) + csld(n,c)
call check (chksum,ntimes*n,t2,'s278 ')
return
end

c%2.7
subroutine s279 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c control flow
c vector if/gotos
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s279 ')
t1 = second()
do 1 n1 = 1, ntimes
do 10 i = 1, n
if (a(i) .gt. 0.) goto 20
b(i) = -b(i) + d(i) * d(i)
if (b(i) .le. a(i)) goto 30
c(i) = c(i) + d(i) * e(i)
30 continue
20 continue
c(i) = -c(i) + e(i) * e(i)
30 continue
a(i) = b(i) + c(i) * d(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a) + csld(n,b) + csld(n,c)
call check (chksum,ntimes*n,t2,'s279 ')
return
end

c%2.7
subroutine s2710 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc,x)
c
c control flow
c scalar and vector ifs
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld, x

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s2710')
t1 = second()
do 1 n1 = 1, ntimes
do 10 i = 1, n
if (a(i) .gt. b(i)) then
a(i) = a(i) + b(i) * d(i)
if (n .gt. 10) then
c(i) = c(i) + d(i) * d(i)

```

```

else
  c(i) = 1.0 + d(i) * e(i)
endif
else
  b(i) = a(i) + e(i) * e(i)
  if(x.gt. 0.)then
    c(i) = a(i) + d(i) * d(i)
  else
    c(i) = c(i) + e(i) * e(i)
  endif
endif
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a) + csld(n,b) + csld(n,c)
call check (chksum,ntimes*n,n,t2,'s2710')
return
end
c%2.7
subroutine s2711(ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c control flow
c semantic if removal
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s2711')
t1 = second()
do 10 i = 1,n
  if(b(i) .ne. 0.) a(i) = a(i) + b(i) * c(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a)
call check (chksum,ntimes*n,n,t2,'s2711')
return
end
c%2.8
subroutine s2712(ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c control flow
c if to elemental min
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s2712')
do 10 i = 1,ntimes
  do 10 i = 1,n
    if(a(i) .gt. b(i)) a(i) = a(i) + b(i) * c(i)
  10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a)
call check (chksum,ntimes*n,n,t2,'s2712')
return
end
c%2.9
subroutine s2713(ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c crossing thresholds
c index set splitting
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld, x
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s281 ')
t1 = second()
do 10 i = 1,ntimes
  x = a(n-i+1) + b(i) * c(i)
  a(i) = x - 1.0
  b(i) = x
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a) + csld(n,b)
call check (chksum,ntimes*n,n,t2,'s281 ')
return
end
c%2.9
subroutine s291 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c loop peeling
c wrap around variable, 1 level
c
integer ntimes, ld, n, i, n1, im1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s291 ')
t1 = second()
do 10 i = 1,ntimes
  im1 = n
  do 10 i = 1,n
    a(i) = (b(i) + b(im1)) * .5
    im1 = i
  10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a)
call check (chksum,ntimes*n,n,t2,'s291 ')
return
end
c%2.9
subroutine s292 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c loop peeling
c wrap around variable, 2 levels
c
integer ntimes, ld, n, i, n1, im1, im2
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s292 ')
t1 = second()
do 10 i = 1,ntimes
  im1 = n
  im2 = n-1
  do 10 i = 1,n

```

```

    im2 = im1
    im1 = 1
10 continue
    call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
    t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
    chksum = csld(n,a)
    call check (chksum,ntimes*n,n,t2,'s292 ' )
    return
end
c%2.9
subroutine s293 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c loop peeling
c a(i)-a(i) with actual dependence cycle, loop is vectorizable
c
c integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s293 ' )
t1 = second()
do 10 i = 1, ntimes
    do 10 i = 1, n
        a(i) = a(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a)
call check (chksum,ntimes*n,n,t2,'s293 ' )
return
end
c%2.10
subroutine s2101(ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c diagonals
c main diagonal calculation
c
c integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, cs2d

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s2101' )
t1 = second()
do 10 i = 1, ntimes
    do 10 i = 1, n
        aa(i,1) = aa(i,1) + bb(i,1) * cc(i,1)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = cs2d(n,aa)
call check (chksum,ntimes*n,n,t2,'s2101' )
return
end
c%2.12
subroutine s2102(ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c diagonals
c identity matrix, best results vectorize both inner and outer loops
c
c integer ntimes, ld, n, i, n1, j
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, cs2d

```

```

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s2102' )
t1 = second()
do 10 i = 1, ntimes/n
    do 10 i = 1, n
        do 20 j = 1, n
            aa(i,j) = 0.
20 continue
        aa(i,1) = 1.
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes/n) )
chksum = cs2d(n,aa)
call check (chksum, (ntimes/n)*n,n,t2,'s2102' )
return
end
c%2.11
subroutine s2111 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c wavefronts
c
c integer ntimes, ld, n, i, n1, j
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, cs2d

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s2111' )
t1 = second()
do 10 i = 1, (ntimes/n)
    do 10 j = 2, n
        do 20 i = 2, n
            aa(i,j) = aa(i-1,j) + aa(i,j-1)
20 continue
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes/n) )
chksum = cs2d(n,aa)
if (chksum .eq. 0.) chksum = 3.0
call check (chksum, (ntimes/n)*(n-1),n,t2,'s2111' )
return
end
c*****
c IDIOM RECOGNITION
c*****
c%3.1
subroutine s311 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c reductions
c sum reduction
c
c integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, sum

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s311 ' )
t1 = second()
do 10 i = 1, ntimes
    sum = 0.
    do 10 i = 1, n
        sum = sum + a(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,sum)
1 continue

```

```

t2 = second() - t1 - cttime - ( dtime * float(ntimes) )
chksum = sum
call check (chksum,ntimes*n,n,t2,'s311 ')
return
end
c#3.1
subroutine s312 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c reductions
c product reduction
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, prod
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s312 ')
t1 = second()
do 1 n1 = 1,ntimes
prod = 1.
do 10 i = 1,n
prod = prod * a(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,prod)
1 continue
t2 = second() - t1 - cttime - ( dtime * float(ntimes) )
chksum = prod
call check (chksum,ntimes*n,n,t2,'s312 ')
return
end
c#3.1
subroutine s313 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c reductions
c dot product
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, dot
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s313 ')
t1 = second()
do 1 n1 = 1,ntimes
dot = 0.
do 10 i = 1,n
dot = dot + a(i) * b(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,dot)
1 continue
t2 = second() - t1 - cttime - ( dtime * float(ntimes) )
chksum = dot
call check (chksum,ntimes*n,n,t2,'s313 ')
return
end
c#3.1
subroutine s314 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c reductions
c if to max reduction
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, x
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s314 ')
t1 = second()
do 1 n1 = 1,ntimes
x = a(i)
if (a(i) .lt. x) x = a(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,x)
1 continue
t2 = second() - t1 - cttime - ( dtime * float(ntimes) )
chksum = x
call check (chksum,ntimes*n,n,t2,'s314 ')
return
end
c#3.1
subroutine s315 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c reductions
c if to max with index reduction, 1 dimension
c
integer ntimes, ld, n, i, n1, index
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, x
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s315 ')
t1 = second()
do 1 n1 = 1,ntimes
x = a(i)
index = 1
do 10 i = 2,n
if(a(i) .gt. x)then
x = a(i)
index = i
endif
10 continue
chksum = x+float(index)
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,chksum)
1 continue
t2 = second() - t1 - cttime - ( dtime * float(ntimes) )
chksum = x + float(index)
call check (chksum,ntimes*n,n,t2,'s315 ')
return
end
c#3.1
subroutine s316 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c reductions
c if to min reduction
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, x
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s316 ')
t1 = second()
do 1 n1 = 1,ntimes
x = a(i)
do 10 i = 2,n
if (a(i) .lt. x) x = a(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,x)
1 continue
t2 = second() - t1 - cttime - ( dtime * float(ntimes) )
chksum = x
call check (chksum,ntimes*n,n,t2,'s316 ')
return
end
c#3.1

```



```

1      call dummy(ld,n,a,b,c,d,e,aa,bb,cc,sum)
      continue
      t2 = second() - t1 - cttime - ( dtime * float(ntimes) )
      chksum = sum
      call check (chksum,ntimes*n,t2,'s3111')
      return
      end
c#3.1
c      subroutine s3112(ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c      reductions
c      sum reduction saving running sums
c
c      integer ntimes, ld, n, i, n1
      real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
      real t1, t2, second, chksum, ctime, dtime, csld, sum
      call init(ld,n,a,b,c,d,e,aa,bb,cc,'s3112')
      t1 = second()
      do 1 n1 = 1,ntimes
      sum = 0.
      do 10 i = 1,n
      sum = sum + a(i)
      b(i) = sum
      10 continue
      call dummy(ld,n,a,b,c,d,e,aa,bb,cc,sum)
      1 continue
      t2 = second() - t1 - cttime - ( dtime * float(ntimes) )
      chksum = csld(n,b) + sum
      call check (chksum,ntimes*n,t2,'s3112')
      return
      end
c#3.1
c      subroutine s3113(ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c      reductions
c      maximum of absolute value
c
c      integer ntimes, ld, n, i, n1
      real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
      real t1, t2, second, chksum, ctime, dtime, max
      call init(ld,n,a,b,c,d,e,aa,bb,cc,'s3113')
      t1 = second()
      do 1 n1 = 1,ntimes
      max = abs(a(1))
      do 10 i = 2,n
      if(abs(a(i)) -gt. max) max = abs(a(i))
      10 continue
      call dummy(ld,n,a,b,c,d,e,aa,bb,cc,max)
      1 continue
      t2 = second() - t1 - cttime - ( dtime * float(ntimes) )
      chksum = max
      call check (chksum,ntimes*(n-1),n,t2,'s3113')
      return
      end
c#3.2
c      subroutine s321 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c      recursences
c      first order linear recurrence
c
c      integer ntimes, ld, n, i, n1
      real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
      real t1, t2, second, chksum, ctime, dtime, csld
      call init(ld,n,a,b,c,d,e,aa,bb,cc,'s321 ')
      t1 = second()
      do 1 n1 = 1,ntimes
      do 10 i = 2,n
      a(i) = a(i) + a(i-1) * b(i)
      call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
      10 continue
      t2 = second() - t1 - cttime - ( dtime * float(ntimes) )
      chksum = csld(n,a)
      call check (chksum,ntimes*(n-1),n,t2,'s321 ')
      return
      end
c#3.2
c      subroutine s322 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c      recursences
c      second order linear recurrence
c
c      integer ntimes, ld, n, i, n1
      real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
      real t1, t2, second, chksum, ctime, dtime, csld
      call init(ld,n,a,b,c,d,e,aa,bb,cc,'s322 ')
      t1 = second()
      do 1 n1 = 1,ntimes
      do 10 i = 3,n
      a(i) = a(i) + a(i-1) * b(i) + a(i-2) * c(i)
      call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
      10 continue
      t2 = second() - t1 - cttime - ( dtime * float(ntimes) )
      chksum = csld(n,a)
      call check (chksum,ntimes*(n-2),n,t2,'s322 ')
      return
      end
c#3.2
c      subroutine s323 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c      recursences
c      coupled recurrence
c
c      integer ntimes, ld, n, i, n1
      real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
      real t1, t2, second, chksum, ctime, dtime, csid
      call init(ld,n,a,b,c,d,e,aa,bb,cc,'s323 ')
      t1 = second()
      do 1 n1 = 1,ntimes
      do 10 i = 2,n
      a(i) = b(i-1) + c(i) * d(i)
      b(i) = a(i) + c(i) * e(i)
      10 continue
      call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
      1 continue
      t2 = second() - t1 - cttime - ( dtime * float(ntimes) )
      chksum = csld(n,a) + csld(n,b)
      call check (chksum,ntimes*(n-1),n,t2,'s323 ')
      return
      end
c#3.3
c      subroutine s331 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c      search loops
c      if to last-1
c

```



```

integer ntimes, ld, n, i, n1, j
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s331 ')
t1 = second()
do 1 n1 = 1, ntimes
  j = -1
  do 10 i = 1, n
    if(a(i) .lt. 0) j = 1
  10 continue
  chksum = float(j)
  call dummy(ld,n,a,b,c,d,e,aa,bb,cc,chksum)
  1 continue
  t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
  chksum = float(j)
  call check (chksum, ntimes*n, n, t2, 's331 ')
  return
end

c#3.3
subroutine s332 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc,t)
c
c search loops
c first value greater than threshold
c
integer ntimes, ld, n, i, n1, index
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, t, value

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s332 ')
t1 = second()
do 1 n1 = 1, ntimes
  index = -1
  value = -1
  do 10 i = 1, n
    if ( a(i) .gt. t ) then
      index = i
      value = a(i)
      goto 20
    endif
  10 continue
  20 continue
  chksum = value + float(index)
  call dummy(ld,n,a,b,c,d,e,aa,bb,cc,chksum)
  1 continue
  t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
  chksum = value + float(index)
  call check (chksum, ntimes*n, n, t2, 's332 ')
  return
end

c#3.4
subroutine s341 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c packing
c pack positive values
c
integer ntimes, ld, n, i, n1, j
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csid

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s341 ')
t1 = second()
do 1 n1 = 1, ntimes
  j = 0
  do 10 i = 1, n
    if(b(i) .gt. 0.)then
      j = j + 1
    endif
  10 continue
  call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
  t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
  chksum = csid(n,a)
  call check (chksum, ntimes*n, n, t2, 's341 ')
  return
end

c#3.4
subroutine s342 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c packing
c unpadding
c
integer ntimes, ld, n, i, n1, j
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csid

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s342 ')
t1 = second()
do 1 n1 = 1, ntimes
  j = 0
  do 10 i = 1, n
    if(a(i) .gt. 0.)then
      j = j + 1
      a(i) = b(j)
    endif
  10 continue
  call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
  t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
  chksum = csid(n,a)
  call check (chksum, ntimes*n, n, t2, 's342 ')
  return
end

c#3.4
subroutine s343 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c packing
c pack 2-d array into one dimension
c
integer ntimes, ld, n, i, n1, j, k, nn
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld, array
parameter(nn=1000)
common /cdata/ array(nn*nn)

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s343 ')
t1 = second()
do 1 n1 = 1, ntimes/n
  k = 0
  do 10 i = 1, n
    do 20 j = 1, n
      if (bb(i,j) .gt. 0) then
        k = k + 1
        array(k) = aa(i,j)
      endif
    20 continue
  10 continue
  call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
  t2 = second() - t1 - ctime - ( dtime * float(ntimes/n) )
  chksum = csld(n,array)
end

```

```

call check (chksum, (ntimes/n)*n,n,t2,'s343 ')
return
end
c#3.5
subroutine s351 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c loop rerolling
c unrolled saxpy
c
integer ntimes, ld, n, l, nl
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, alpha, csld

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s351 ')
t1 = second()
alpha = c(l)
do 1 nl = 1,5*ntimes
do 10 i = 1,n,5
a(i) = a(i) + alpha * b(i)
a(i+1) = a(i+1) + alpha * b(i+1)
a(i+2) = a(i+2) + alpha * b(i+2)
a(i+3) = a(i+3) + alpha * b(i+3)
a(i+4) = a(i+4) + alpha * b(i+4)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,l.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(5*ntimes) )
chksum = csld(n,a)
call check (chksum,5*ntimes*(n/5),n,t2,'s353 ')
return
end
c
c*****
c LANGUAGE COMPLETENESS
c
c*****
c#4.1
subroutine s411 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c loop recognition
c if loop to do loop, zero trip
c
integer ntimes, ld, n, i, nl
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s411 ')
t1 = second()
do 1 nl = 1,ntimes
i = 0
10 continue
i = i + 1
if (i.gt.n) goto 20
a(i) = a(i) + b(i) * c(i)
goto 10
20 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,l.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a)
call check (chksum,ntimes*n,t2,'s411 ')
return
end
c#4.1
subroutine s412 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc,inc)
c
c loop recognition
c if loop with variable increment
c
integer ntimes, ld, n, i, nl, inc
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s412 ')
t1 = second()
do 1 nl = 1,ntimes
i = 0
10 continue

```

```

i = i + inc
if (i .gt. n) goto 20
a(i) = a(i) + b(i) * c(i)
goto 10
20 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ct ime - ( dt ime * float(ntimes) )
chksum = csld(n,a)
call check (chksum,ntimes*n,t2,'s412 ')
return
end
c$4.1 subroutine s413 (ntimes,ld,n,ct ime,dt ime,a,b,c,d,e,aa,bb,cc)
c
c loop recognition
c if loop to do loop, code on both sides of increment
c
integer ntimes, ld, n, i, nl
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ct ime, dt ime, csld
call init (ld,n,a,b,c,d,e,aa,bb,cc,'s413 ')
t1 = second()
do 1 nl = 1,ntimes
i = 1
10 continue
if (i .ge. n) goto 20
b(i) = b(i) + d(i) * e(i)
i = i + 1
a(i) = c(i) + d(i) * e(i)
goto 10
20 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
c
c storage classes and equivalencing
c
c equivalence- no overlap
c
integer ntimes, ld, n, i, nl, nn
parameter(nn=1000)
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ct ime, dt ime, csld
real x(nn),y(nn)
equivalence (x(1),y(1))
call setld(n,x,0,0,1)
call setld(n,y,1,0,1)
call init (ld,n,a,b,c,d,e,aa,bb,cc,'s421 ')
t1 = second()
do 1 nl = 1,ntimes
do 10 i = 1,n-1
x(i) = y(i+1) + a(i)
10 continue
call dummy(ld,n,x,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ct ime - ( dt ime * float(ntimes) )
chksum = csld(n,x)
call check (chksum,ntimes*(n-1),n,t2,'s421 ')
return
end
c$4.2 subroutine s422 (ntimes,ld,n,ct ime,dt ime,a,b,c,d,e,aa,bb,cc)
c
c storage classes and equivalencing
c
c common and equivalence statement
c
c anti-dependence, threshold of 4
c
integer ntimes, ld, n, i, nl, nn, vl
parameter(nn=1000,vl=64)

```

```

real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
common /cdata/ array(1000000)
real x(nn), array
equivalence (x(1),array(5))
real t1, t2, second, chksum, ctime, dtime, csld

call setid(n,x,0.0,1)
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s422 ')
t1 = second()
do 10 i = 1, ntimes
do 10 i = 1, n
x(i) = array(i*8) + a(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,x)
call check (chksum,ntimes*(n-8),n,t2,'s422 ')
return
end
c44.2
subroutine s423 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c storage classes and equivalencing
c common and equivalenced variables - with anti-dependence
c
integer ntimes, ld, n, i, n1, nn, v1
parameter(nn=1000,v1=64)
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real x(nn), array
common /cdata/ array(1000000)
equivalence (array(v1),x(1))
real t1, t2, second, chksum, ctime, dtime, csld

call setid(n,x,1.0,1)
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s423 ')
t1 = second()
do 10 i = 1, ntimes
do 10 i = 1, n-1
array(i+1) = x(i) + a(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,array)
call check (chksum,ntimes*(n-1),n,t2,'s423 ')
return
end
c44.2
subroutine s424 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c storage classes and equivalencing
c common and equivalenced variables - overlap
c vectorizable in strips of 64 or less
c
integer ntimes, ld, n, i, n1, nn, v1
parameter(nn=1000,v1=64)
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real x(nn), array
common /cdata/ array(1000000)
equivalence (array(v1),x(1))

call setid(n,x,0.0,1)
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s424 ')
t1 = second()

```

```

do 10 i = 1, ntimes
do 10 i = 1, n-1
x(i+1) = array(i) + a(i)
10 continue
call dummy(ld,n,x,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,x)
call check (chksum,ntimes*(n-1),n,t2,'s424 ')
return
end
c44.3
subroutine s431 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c parameters
c parameter statement
c
integer ntimes, ld, n, i, n1, k, k1, k2
parameter(k1=1, k2=2, k=2*k1-k2)
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s431 ')
t1 = second()
do 10 i = 1, ntimes
do 10 i = 1, n
a(i) = a(i+k) + b(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a)
call check (chksum,ntimes*n,n,t2,'s431 ')
return
end
c44.3
subroutine s432 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c parameters
c data statement
c
integer ntimes, ld, n, i, n1, k, k1, k2
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld
data k1,k2 /1,2/

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s432 ')
k=2*k1-k2
t1 = second()
do 10 i = 1, ntimes
do 10 i = 1, n
a(i) = a(i+k) + b(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a)
call check (chksum,ntimes*n,n,t2,'s432 ')
return
end
c44.4
subroutine s441 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c non-logical if's
c arithmetic if
c

```

```

integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csid

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s441 ')
t1 = second()
do 1 n1 = 1, n
do 10 i = 1, n
if (d(i)) 20,30,40
a(i) = a(i) + b(i) * c(i)
goto 50
a(i) = a(i) + b(i) * b(i)
goto 50
a(i) = a(i) + c(i) * c(i)
50 continue
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csid(n,a)
call check (chksum,ntimes*n,t2,'s441 ')
return
end

c#4.4 subroutine s442 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc,indx)
c non-logical if's
c computed goto
c
integer ntimes, ld, n, i, n1, indx(n)
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csid

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s442 ')
t1 = second()
do 1 n1 = 1, ntimes
do 10 i = 1, n
goto (15,20,30,40) indx(i)
a(i) = a(i) + b(i) * b(i)
goto 50
a(i) = a(i) + c(i) * c(i)
goto 50
a(i) = a(i) + d(i) * d(i)
goto 50
a(i) = a(i) + e(i) * e(i)
50 continue
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csid(n,a)
call check (chksum,ntimes*n,t2,'s442 ')
return
end

c#4.4 subroutine s443 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c non-logical if's
c arithmetic if
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csid

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s443 ')
t1 = second()
do 1 n1 = 1, ntimes
if (d(i)) 20,30,40
a(i) = a(i) + b(i) * c(i)
goto 50
a(i) = a(i) + b(i) * b(i)
goto 50
a(i) = a(i) + c(i) * c(i)
50 continue
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csid(n,a)
call check (chksum,ntimes*n,t2,'s443 ')
return
end

c#4.5 subroutine s451 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c intrinsic functions
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csid

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s451 ')
t1 = second()
do 1 n1 = 1, ntimes
do 10 i = 1, n
a(i) = sin(b(i)) + cos(c(i))
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csid(n,a)
call check (chksum,ntimes*n,t2,'s451 ')
return
end

c#4.5 subroutine s452 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c intrinsic functions
c seq function
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csid

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s452 ')
t1 = second()
do 1 n1 = 1, ntimes
do 10 i = 1, n
a(i) = b(i) + c(i) * float(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csid(n,a)
call check (chksum,ntimes*n,t2,'s452 ')
return
end

c#4.5 subroutine s453 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c intrinsic functions
c seq function

```

```

c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld, s
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s453 ')
t1 = second()
do 1 n1 = 1, ntimes
s = 0.
do 10 i = 1, n
s = s + 2.
a(i) = s * b(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a)
call check (chksum, ntimes*n, t2, 's453 ')
return
end
c44.7
subroutine s471 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc,c471s)
c
c call statements
integer ntimes, ld, n, i, n1, nn, m
parameter(nn=1000)
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld, c471s, x(nn)
m = n
call setld(n,x,0,0,1)
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s471 ')
t1 = second()
do 1 n1 = 1, ntimes
x(i) = b(i) + d(i) * d(i)
call s471s
b(i) = c(i) + d(i) * e(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) ) -
+ ( ( n*ntimes) * c471s )
chksum = csld(n,x) + csld(n,b)
call check (chksum, ntimes*n, t2, 's471 ')
return
end
c44.8
subroutine s481 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c non-local goto's
c stop statement
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s481 ')
t1 = second()
do 1 n1 = 1, ntimes
do 10 i = 1, n
if (d(i) .lt. 0.) stop 'stop 1'
a(i) = a(i) + b(i) * c(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
c
c44.9
call setld(n,x,0,0,1)
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s471 ')
t1 = second()
do 1 n1 = 1, ntimes
x(i) = b(i) + d(i) * d(i)
call s471s
b(i) = c(i) + d(i) * e(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) ) -
+ ( ( n*ntimes) * c471s )
chksum = csld(n,x) + csld(n,b)
call check (chksum, ntimes*n, t2, 's471 ')
return
end
c44.8
subroutine s482 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c non-local goto's
c other loop exit with code before exit
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s482 ')
t1 = second()
do 1 n1 = 1, ntimes
a(i) = a(i) + b(i) * c(i)
if (c(i) .gt. b(i)) goto 20
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a)
call check (chksum, ntimes*n, t2, 's482 ')
20 continue
return
end
c44.9
subroutine s491 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc,ip)
c
c vector semantics
c indirect addressing on lhs, store in sequence
c
integer ntimes, ld, n, i, n1, ip(n)
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s491 ')
t1 = second()
do 1 n1 = 1, ntimes
do 10 i = 1, n
a(ip(i)) = b(i) + c(i) * d(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a)
call check (chksum, ntimes*n, t2, 's491 ')
return
end
c44.11
subroutine s4112 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc,ip,s)
c
c indirect addressing
c sparse saxpy
c
integer ntimes, ld, n, i, n1, ip(n)
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld, s
call init(ld,n,a,b,c,d,e,aa,bb,cc,'s4112')
t1 = second()

```



```

c*4.12
subroutine s4121 (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c statement functions
c elementwise multiplication
c
integer ntimes, ld, n, i, nl
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld, f, x, y
f(x,y) = x*y

call init(ld,n,a,b,c,d,e,aa,bb,cc,'s4121')
t1 = second()
do 10 i = 1, ntimes
do 10 i = 1, n
a(i) = a(i) + f(b(i),c(i))
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a)
call check (chksum,ntimes*n,t2,'s4121')
return
end

c*5.1
subroutine va (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c control loops
c vector assignment
c
integer ntimes, ld, n, i, nl
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld

call init(ld,n,a,b,c,d,e,aa,bb,cc,'va ')
t1 = second()
do 10 i = 1, ntimes
do 10 i = 1, n
a(i) = b(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a)
call check (chksum,ntimes*n,t2,'va ')
return
end

c*5.1
subroutine vif (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c control loops
c vector if
c
integer ntimes, ld, n, i, nl
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld

call init(ld,n,a,b,c,d,e,aa,bb,cc,'vif ')
t1 = second()
do 10 i = 1, ntimes
do 10 i = 1, n
if (b(i) .gt. 0.) then
a(i) = b(i)
endif
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a)
call check (chksum,ntimes*n,t2,'vif ')
return
end

c*5.1
subroutine vpv (ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c
c control loops
c vector plus vector
c
integer ntimes, ld, n, i, nl
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld

call init(ld,n,a,b,c,d,e,aa,bb,cc,'vpv ')
t1 = second()
do 10 i = 1, ntimes

```



```

do 10 i = 1,n
  a(i) = a(i) + b(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a)
call check (chksum,ntimes*n,n,t2,'vpv ' )
return
end

c$5.1 subroutine vtv(ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c control loops
c vector times vector
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld

call init(ld,n,a,b,c,d,e,aa,bb,cc,'vtv ' )
t1 = second()
do 1 n1 = 1,ntimes
do 10 i = 1,n
  a(i) = a(i) + b(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
chksum = csld(n,a)
call check (chksum,ntimes*n,n,t2,'vtv ' )
return
end

c$5.1 subroutine vpvv(ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c control loops
c vector plus vector vector
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld

call init(ld,n,a,b,c,d,e,aa,bb,cc,'vpvv' )
t1 = second()
do 1 n1 = 1,ntimes
do 10 i = 1,n
  a(i) = a(i) + b(i) + c(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a)
call check (chksum,ntimes*n,n,t2,'vpvv' )
return
end

c$5.1 subroutine vvtv(ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc)
c control loops
c vector times vector times vector
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld

call init(ld,n,a,b,c,d,e,aa,bb,cc,'vvtv' )
t1 = second()
do 1 n1 = 1,ntimes
do 10 i = 1,n
  a(i) = a(i) + b(i) * c(i)
10 continue
call dummy(ld,n,a,b,c,d,e,aa,bb,cc,1.)
1 continue
t2 = second() - t1 - ctime - ( dtime * float(ntimes) )
chksum = csld(n,a)
call check (chksum,ntimes*n,n,t2,'vvtv' )
return
end

c$5.1 subroutine vpvts(ntimes,ld,n,ctime,dtime,a,b,c,d,e,aa,bb,cc,s)
c control loops
c vector plus vector times scalar
c
integer ntimes, ld, n, i, n1
real a(n), b(n), c(n), d(n), e(n), aa(ld,n), bb(ld,n), cc(ld,n)
real t1, t2, second, chksum, ctime, dtime, csld, s

```

