

A Collection of Equation-Solving  
Codes for the CRAY-1

D. A. Calahan

W. G. Ames

E. J. Sesek

S. Arya

Systems Engineering Laboratory  
University of Michigan  
Ann Arbor, Michigan 48109

August 1, 1979

. SEL Report #133

Sponsored Jointly by  
Directorate of Mathematical and Information Sciences  
Air Force Office of Scientific Research, and  
Air Force Flight Dynamics Laboratory  
Under Grant 75-2812

Engn  
UMR1076

## Abstract

This report contains user documentation and timing results for a collection of assembly language equation-solving codes for the CRAY-1.

## Acknowledgement

The development of these codes was made possible by the development of a CRAY-1 simulator by D. A. Orbits. The programming assistance of S. Arya is also acknowledged.

## Revision Notice

The discussion on tridiagonal matrix solution was revised on 11/8/79.

## Program Availability

These programs are available on a user-supplied 9 track tape. Specify desired density, labeling, and blocking.

## TABLE OF CONTENTS

	PAGE
I. Introduction	1
A. Implicitness	1
B. Vector length	1
C. Data flow	2
D. Report Summary	3
II. Codes for the Solution of a Small, Dense System	4
A. Solution of a single full unsymmetric system of equations	4
B. Solution of a single banded unsymmetric system of equations	6
III. Codes for the Solution of Simultaneous Systems	8
A. Introduction	8
B. Solution of full unsymmetric simultaneous systems of equations	10
C. Solution of banded unsymmetric simultaneous systems of equations	12
D. Solution of block tridiagonal unsymmetric simultaneous systems of equations	14
IV. Code for the Solution of a Single Tridiagonal System	16
References	18

## I. Introduction

### A. Implicitness

The most common application envisioned in the design of vector processors has been the solution of partial differential equations. It is now clear that many application-oriented researchers are planning to use this vastly increased computational capability to make solution algorithms more implicit [1]. For example, an algorithm implicit by line (i.e., in one dimension) would become implicit by strips (1+n dimensions); or, more variables may be coupled in a multi-variable problem. It is usually found that the larger the problem and/or the more implicit the algorithm, the greater fraction of total formulation and solution time is devoted to the latter. The coding of the equation-solver then becomes a critical issue.

### B. Vector Length

The vector length is the most obvious and general concern in vector processing. The length results either from simultaneous operations on a number of similar systems or from the density of structure (coupling of variables and grid nodes) within a single system.

In the solution of partial differential equations, the frequency of variable updating determines the number of grid node equations that can be formulated simultaneously. A point-Jacobi iteration would allow simultaneous formulation and updating of all variables. An ADI method would allow simultaneous updating of variables along a line. The coupling between grid points and between lines of grid points assumed in the equation solution determines the number of systems of equations which can be solved simultaneously. Thus, if grid points are assumed coupled in only one of two dimensions, then one can expect a system of tridiagonal or block tridiagonal matrices which can be solved simultaneously. This in turn yields a vector length equal to the number of uncoupled grid lines. On the other hand, an ADI method necessitates a solution of a single tridiagonal system.

The most obviously vectorizable algorithm would therefore be one in which variables are simultaneously updated and minimal coupling is assumed between grid nodes and/or variables at a grid point.

Iterative methods based on such schemes have notoriously poor convergence. However, if simultaneity--and hence vector length--is reduced to increase convergence rate, one must exploit single system density to achieve long vectors. This density is usually manifested by relatively small block sizes, bandwidths, or profiles in the matrix structure.

### C. Data Flow

With attention only to the vector length and with the use of a high level language such as Fortran, it is not uncommon to obtain only 20%-50% of the optimum CRAY-1 performance in the equation solution. The data flow must also be considered to achieve high performance, as illustrated by the following instances associated with equation-solving codes.

(a) With short vector chained operations, the CRAY-1 protocol results in large bubbles in the arithmetic pipelines [2]. This occurs in the processing of small dense blocks or narrow bandwidths.

(b) Vectors of any length on which little computation is performed can create excessive data flow between memory hierarchies. This situation prevails in the above-mentioned simultaneous solution of equations, and results from the inherent decoupling of such systems.

(c) General equation-solving codes--ones which can accomodate arbitrary problem size parameters--may suffer from excessive data flow visa vis a special code written for small problem sizes which can maintain critical data in the cache memory. For example, a simultaneous block tridiagonal solver specialized to a fixed small block size can yield much higher execution rates than a general block tridiagonal solver (to be demonstrated).

The goal of high performance in spite of short vectors and apparent data flow bottlenecks appears to suggest the need for a plethora of specialized and highly-tuned codes to fill the same functions as a single code executing on a scalar processor. However, a mitigating effect is the computational dominance of the equation formulation over the equation solution as the coupling shrinks. This observation is based on (a) the number of entries in a matrix being of complexity  $O(n^r)$ , and (b) the triangular factorization operation count being  $O(n^{r+\epsilon})$ , where  $\epsilon > 0$  and  $n$  represents the number of vari-

ables coupled to each other. Thus, more inefficiency can be tolerated in the equation solver with small blocks, bands, and profiles.

#### D. Report Summary

This report provides user documentation for three classes of codes either expected to be of general utility or else resulting from on-going specialized algorithm research for the CRAY-1. All were developed with aid of a CRAY-1 timing simulator [3]. Many of the accumulation kernels on which the high performance of the codes depend are described in [4].

(a) Small dense systems. Highly-tuned accumulation kernels yield codes which achieve high execution rates with small full and banded systems.

(b) Simultaneous systems. Simultaneous full, banded, and block tridiagonal equation solvers have been developed around kernels which reduce the previously-mentioned memory traffic. A variety of block tridiagonal solvers are included representing the utility of such codes.

(c) Single system, single variable, odd-even tridiagonal solver. The use of a simulator was deemed essential to develop this challenging code which involves high memory traffic.

A fourth code, which can solve general single system sparse problems ranging from block tridiagonal systems, general full and banded matrices (which must be partitioned into 64 x 64 blocks for the CRAY-1), and arbitrary-sparsity finite element problems, is also being prepared [5].

## II. Codes for the solution of small, dense systems

### A. Solution of a single full unsymmetric system of equations (Calahan)

#### Description.

A system of equations  $\underline{A} \underline{X} = \underline{B}$  is solved for  $\underline{X}$ , where  $\underline{A}$  is an  $n \times n$  real matrix and  $\underline{X}$  and  $\underline{B}$  are  $n \times m$  real matrices. Double column accumulation [4] is used to achieve full cache utilization during both the triangular factorization of  $\underline{A}$  and the forward and back substitution of  $\underline{B}$ . When  $m = 1$ , an alternate substitution routine is provided.

Subroutine call to triangularly factor A.

```
CALL FULFAC (N, A, NDIMA, IERR)
```

where N is the dimension of  $\underline{A}$

A is the array representing  $\underline{A}$

NDIMA is the row dimension of array A

IERR contains a return code:

IERR = 0 implies N = 0

IERR > 0 is normal exit

IERR < 0 implies zero-valued pivot; position given  
by |IERR|.

Subroutine call to forward and back substitute

```
CALL FULSOL (N, A, NDIMA, B, M, NDIMB)
```

where A contains the factored matrix

B is the array representing  $\underline{B}$

M is the number of columns of  $\underline{B}$

NDIMB is the row dimension of array B

#### Restrictions.

(a)  $N \leq 64$ .

(b) no pivoting

(c) Fetches (but not stores) from main memory will occur in FULFAC and FULSOL from the (n+1)st and (m+1)st columns of A and B; this space should contain data, not instructions.



When  $M = 1$ , a special call\*

CALL FULSOL1 (N, A, NDIMA, D, M, DNIMB)

should be used for best efficiency. The A array is not altered in either FULSOL OR FULSOL1.

Performance (simulated)

Matrix size	Factorization	Substitution
4	6.5/.47	4.5/.49
8	23/1.1	12/.77
16	58/3.6	27/1.5 (52/1.5)
32	95/18	44/3.7
64	122/113	60/11

[Execution rate (MFLOPS)]/[time (kilo clocks)]  
for solution of a full system of equations.  
Result in parentheses for substitution of two  
columns of B ( $M = 2$ ).

---

\*FULSOL1 will solve systems where  $M > 1$  and may be useful when  $M$  is odd and the extra column fetch of FULSOL is undesirable.

B. Solution of a single banded unsymmetric system of equations  
(Calahan)\*

Description

A system of equations  $\underline{A} \underline{x} = \underline{b}$  is solved for  $\underline{x}$ , where  $\underline{A}$  is a banded real matrix and  $\underline{x}$  and  $\underline{b}$  are  $n \times 1$  real vectors. The matrix  $\underline{A}$  is stored in compressed form.

Subroutine call to triangularly factor  $\underline{A}$ .

CALL BANFAC (N, NB, A, NDIM)

where N is the dimension of  $\underline{A}$

NB is the half bandwidth (i.e.,  $2*NB+1$  is the full bandwidth.)

A is an array containing the elements of  $\underline{A}$

NDIM is the row dimension of A.

Subroutine call to forward and back substitute.

CALL BANSOL (N, NB, A, NDIM, B)

where B contains the elements of  $\underline{b}$  on entry and  $\underline{x}$  on exit.

Comments

$\underline{A}$  is stored in packed form so that the  $(i,j)$  position of  $\underline{A}$  is stored in the  $(i - j + NB) + (j - 1)*NB$  address of A.

---

\*The factorization algorithm is a recoded version of a band solver written by T. Jordan of LASL [6][7]. The substitution code is identical to Jordan's.

Half Bandwidth	Factorization	Substitution
2	3.8/14	6.2/8.1
4	9.6/19	11/8.1
8	23/28	20/8.1
16	52/45	36/8.1
32	88/210	65/18
64	117/1260	93/49

[Execution rate (MFLOPS)]/[time (kiloclocks)] for solution of a banded system of equations. Sixty four equations were solved except for half bandwidths of 32 and 64, where 128 and 256 equations were solved, respectively.

III. Codes for the Solution of Simultaneous Systems

A. Introduction

Let  $\hat{\underline{A}} \hat{\underline{x}} = \hat{\underline{b}}$  represent the simultaneous system of equations described by the block-diagonal matrix

$$\begin{bmatrix} \underline{A}^{(1)} & & & & & & \\ & \underline{A}^{(2)} & & & & & \\ & & \underline{A}^{(3)} & & & & \\ & & & \ddots & & & \\ & & & & \underline{A}^{(m)} & & \\ & & & & & & \end{bmatrix} \begin{bmatrix} \underline{x}^{(1)} \\ \underline{x}^{(2)} \\ \underline{x}^{(3)} \\ \vdots \\ \underline{x}^{(m)} \end{bmatrix} = \begin{bmatrix} \underline{b}^{(1)} \\ \underline{b}^{(1)} \\ \underline{b}^{(3)} \\ \vdots \\ \underline{b}^{(m)} \end{bmatrix} \quad (1)$$

where the  $\underline{A}^{(k)}$  are identically-structured  $n \times n$  real unsymmetric matrices containing  $a_{ij}^{(k)}$  and the  $\underline{x}^{(k)}$  and  $\underline{b}^{(k)}$  are  $n \times 1$  real matrices containing  $x_i^{(k)}$  and  $b_i^{(k)}$ , respectively.

Solutions of such systems on a vector machine have a number of common characteristics.

1. Vectors are defined across the systems; e.g., the  $(i,j)$  positions of all  $\underline{A}^{(k)}$ ,  $k = 1, 2, \dots, m$ , constitute a single vector.

2. Two storage array maps are common.

Map I.  $\hat{\underline{A}}$ ,  $\hat{\underline{x}}$ , and  $\hat{\underline{b}}$  stored by column, then by row, then by system, i.e.,

$a_{ij}^{(k)}$  :  $i = 1, 2, \dots, n$ ;  $j = 1, 2, \dots, n$ ;  $k = 1, 2, \dots, m$

$x_i^{(k)}$  :  $i = 1, 2, \dots, n$ ;  $k = 1, 2, \dots, m$

$b_i^{(k)}$  :  $i = 1, 2, \dots, n$ ;  $k = 1, 2, \dots, m$

Map II.  $\hat{\underline{A}}$ ,  $\hat{\underline{x}}$ , and  $\hat{\underline{b}}$  stored by system, then by column, then by row.

$a_{ij}^{(k)}$  :  $k = 1, 2, \dots, m$ ;  $i = 1, 2, \dots, n$ ;  $j = 1, 2, \dots, n$

$x_i^{(k)}$  :  $k = 1, 2, \dots, m$ ;  $i = 1, 2, \dots, n$

$b_i^{(k)}$  :  $k = 1, 2, \dots, m$ ;  $i = 1, 2, \dots, n$

Map I may suffer from bank conflicts and loss of critical chaining when the systems are a multiple of 8 address locations apart.

A number of simultaneous system solvers are described in the report.

1. Simultaneous full systems.
2. Simultaneous banded systems.
3. Simultaneous 3 x 3 block tridiagonal systems.
4. Simultaneous 5 x 5 block tridiagonal systems.

B. Solution of full unsymmetric simultaneous systems of equations  
(Ames/Calahan)

Description

m simultaneous systems of the form of (1) are solved when the  $\underline{A}^{(k)}$  are full matrices.

Subroutine call to triangularly factor  $\hat{A}$ .

CALL FUSFAC (A, N, M, IA)

where A is the array containing the full  $\underline{A}^{(k)}$  matrices of (1)

N is the dimension of the  $\underline{A}^{(k)}$

M is the number (m) of  $\underline{A}^{(k)}$  matrices

IA is the address displacement in array A between  $a_{ij}^{(k)}$   
and  $a_{ij}^{(k+1)}$

Subroutine call to forward and back substitute

CALL FUSSOL (B, IB)

where B is the array containing the right hand sides  $\underline{b}$  and  
the solutions  $\underline{x}$ ,

IB is the address displacement in array B between  $b_i^{(k)}$  and  
 $b_i^{(k+1)}$

Restrictions:

- (a)  $M \leq 64$
- (b)  $IA \geq N^2$ , i.e., Map I is used.
- (c) no err monitoring of pivot reciprocation.

Performance (simulated)

Equations per system (N)	Number of Systems (M)				
	4	8	16	32	64
2	3.1/.52	5.8/.55	11/.61	16/.80	21/1.2
4	5.7/2.1	11/2.2	19/2.6	28/3.5	35/5.6
8	9.6/11	18/11	32/13	45/18	53/31

(a) Factorization

Equations per system (N)	Number of Systems (M)				
	4	8	16	32	64
2	3.1/.62	6.1/.63	11/.68	19/.80	26/1.2
4	6.7/1.3	13/1.4	23/1.6	35/2.0	44/3.2
8	11/3.4	21/3.6	37/4.1	53/5.8	62/9.8

(b) Substitution

[Execution rate (MFLOPS)]/[time (kiloclocks)]  
of simultaneous equation solver.

C. Solution of banded unsymmetric systems of equations (Ames/  
Calahan)

Description:

m simultaneous systems of the form of (1) are solved when  
the  $\underline{A}^{(k)}$  are banded. The  $\underline{A}^{(k)}$  are stored in compressed form.

Subroutine call to triangularly factor  $\hat{A}$  and forward and back  
substitute

CALL BANSIM (A, B, N, NB, M)

where A is an  $M*(2*NB+1)*N$  array of elements of the banded  
 $\underline{A}^{(k)}$  matrices of (1)

B is an  $M*N$  array of elements of the right hand side  $\underline{b}^{(k)}$   
and the solution  $\underline{x}^{(k)}$

N is dimension of the matrices  $\underline{A}^{(k)}$

NB is the half bandwidth (i.e.,  $2*NB+1$  is the full bandwidth)

M is the number of systems.

Comments

$\hat{A}$ , a band matrix, is stored in packed form so that the (i,j)  
position of  $\underline{A}^{(k)}$  is stored in the  $k + (i - j + NB)*M +$   
 $(j - 1)*M*(2*NB+1)$  address of A.

$\underline{b}$  is stored so that the ith position of  $\underline{b}^{(k)}$  is stored in the  
 $k + (i - 1)*M$  address of B



Performance (simulated)

Half- bandwidth (NB)	Number of systems (M)					
	2	4	8	16	32	64
2	2.1/47.3	4.1/47.8	7.9/49.3	14.5/53.9	23./68.0	31.3/100.
4	3.1/80.7	6.2/81.5	12./84.0	21.6/93.6	32.9/123.	43.0/188.
8	4.6/156.	9.2/157.	17.9/161.	31.1/186.	45.1/256.	56.5/409.

[Execution rates (MFLOPS)]/[time (kiloclocks)]  
for factorization, forward, and back substitution  
of simultaneous banded systems. Each system has  
32 equations.



Comments

Storage Map II is used for all  $\hat{A}_{ij}$ ,  $\hat{x}_i$ , and  $\hat{b}_i$ . For block size  $\ell$ , the storage is given below for  $n$  diagonal blocks.

Array A ( $m \times \ell \times \ell n$ )

$$[\hat{A}_{11} \hat{A}_{22} \cdots \hat{A}_{nn}]$$

Array B ( $m \times \ell \times \ell(n-1)$ )

$$[\hat{A}_{12} \hat{A}_{23} \cdots \hat{A}_{n-1,n}]$$

Array C ( $m \times \ell \times \ell(n-1)$ )

$$[\hat{A}_{21} \hat{A}_{32} \cdots \hat{A}_{n,n-1}]$$

Array D ( $m \times \ell n$ )

$$[\hat{b}_1 \hat{b}_2 \cdots \hat{b}_n] \text{ on entry to code}$$

$$[\hat{x}_1 \hat{x}_2 \cdots \hat{x}_n] \text{ on exit from code}$$

Performance (simulated).

Execution rates (MFLOPS)

Number of Systems	3x3 FACTOR.	3x3 SUBS.	5x5 FACTOR.	5x5 SUBS.
8	28.8	27.6	40.8	40.1
16	46.6	46.1	54.7	53.6
32	60.3	58.7	65.5	64.1
64	67.9	66.3	72.0	70.0

MFLOP rates for solution of block tridiagonal systems for two block sizes, as a function of the number of simultaneous systems (= vector length).

IV. Code for the Solution of a Single Tridiagonal System (Arya/Sesek/Calahan)

Description

A single tridiagonal system is completely solved by cyclic (odd-even) reduction, extended to include any number of equations.

Subroutine call to triangularly factor A.

CALL TRIFAC (IPT, A, B, C, N)

where IPT contains the premature termination factor,

A contains the diagonal stripe,  $A(i), i=1, \dots, N$

B contains the super-diagonal stripe,  $B(i), i=1, \dots, N-1, B(N)=0$

C contains the sub-diagonal stripe,  $C(i), i=2, \dots, N, C(1)=0$

N is the number of diagonal elements.

Subroutine call to forward and back substitute

CALL TRISUB (IPT, A, B, C, R, N)

where R contains the right hand side on entry and the solution on exit.

Performance (simulated) (large IPT)

Matrix Size (N)	Factorization	Forward and Back Substitution
15	8.92/1.28	4.75/1.89
16	8.87/1.71	4.42/2.60
31	14.5/1.73	7.93/2.53
32	13.6/2.18	7.06/3.29
63	22.4/2.37	12.7/3.37
64	20.5/2.84	11.2/4.15
127	32.0/3.40	19.1/4.63
128	27.8/4.17	16.0/5.83
255	40.9/5.42	26.0/6.95
256	35.4/6.49	21.7/8.58
511	47.3/9.45	31.8/11.5
512	42.2/10.8	27.3/13.6
1023	51.0/17.6	35.5/20.7
1024	47.0/19.4	31.8/23.4
2047	53.0/34.1	37.7/39.1
2048	50.3/36.1	35.0/42.3
4095	54.0/66.9	38.9/75.9
4096	52.4/69.2	37.2/79.6

Timings of a cyclic reduction of a single tridiagonal system; results given as [execution rate(MFLOPS)]/[time(kiloclocks)]

Premature Termination:

Users not familiar with the cyclic reduction technique or premature termination should set IPT to a large value (e.g., 1000). Users familiar with these subjects can use IPT as follows:

a)  $IPT \leq 0$  -

Only the reciprocation of A will be done in the factorization routine and R will be multiplied to the reciprocate of A in the forward and back substitution.

b)  $0 < IPT < \lfloor \log_2^N \rfloor$

For the first IPT cycles normal cyclic solution will be carried out, but after the first IPT cycles only the reciprocation of A and its multiplication by R will be carried out.

c)  $IPT \geq \lfloor \log_2^N \rfloor$

No premature termination. Complete cyclic-reduction algorithm will be carried out.

The following table depicts the speed up of the algorithm when IPT is used.

N	IPT	FAC (MFLOPS/KCLOCKS)
63	large	22.4/2.37
63	2	33.4/1.28
2047	large	53.0/34.1
2047	6	54.7/32.6
2047	3	57.9/27.7
2047	2	61.3/22.8

Restrictions:

- (1) Single precision (10 digit) reciprocation of pivots.
- (2) Arrays B and C should have a dimension of 2 times the dimension of A.

NOTE: Simulation has been performed assuming a 16-bank CRAY-1 configuration.

Comments:

Odd-even reduction requires  $\approx 2.7$  times the floating point computation of a scalar solution (important when evaluating the MFLOPS rates above), but yields, for larger matrices, nearly all 64-length vector operations.

This code has been optimized for the large matrix case. The execution is then easily shown to be memory-path bound. The algorithm was therefore chosen to make a minimum number of memory accesses per loop, utilizing shifting instead to align operand vectors in the vector registers. The resultant coding was then optimized to achieve full memory-path utilization. Simulation shows that 93,89,78% utilization is achieved for the factorization, forward, and back substitutions, respectively, as  $N \rightarrow \infty$ . Based on the nature of the algorithm, even-sized matrices take longer to execute than odd-sized matrices. Therefore, the best ( $N = 2^r - 1$ ) and worst ( $N = 2^r$ ,  $r$  integer) execution timings are given above.

Further algorithm discussion is given in [4] and [8].

## References

- [1]. Boley, D., B. L. Buzbee, and S. V. Parter, "On Block Relaxation Techniques," TR #318, Computer Sciences Dept., University of Wisconsin, Madison, June, 1978.
- [2]. Orbits, D. A., and D. A. Calahan, "A CRAY-1 Simulator and Its Application to the Development of High Performance Codes," Proc. 1978 LASL Workshop on Vector and Parallel Processors, Los Alamos Scientific Laboratories, pp. 42-56.
- [3]. Orbits, D. A., "A CRAY-1 Timing Simulator," SEL Report No. 118, Systems Engineering Laboratory, University of Michigan, September, 1978.
- [4]. Ames, W. G., et al, "Sparse Matrix and Other High Performance Algorithms for the CRAY-1," Report #124, Systems Engineering Laboratory, University of Michigan, January 24, 1979.
- [5]. Calahan, D. A., "Vectorized Sparse Equation Solution on the CRAY-1," Proc. 1979 International Conf. on Parallel Computing, Wayne State University, August, 1979.
- [6]. Fong, K., and T. L. Jordan, "Some Linear Algebraic Algorithms and Their Performance on the CRAY-1," Report LA-6774, Los Alamos Scientific Laboratory, June, 1977.
- [7]. Calahan, D. A., "Performance of Linear Algebra Codes on the CRAY-1," Fifth SPE Symposium on Reservoir Simulation, Denver, February, 1979.
- [8]. Lambiotte, J. J., and R. G. Voigt, "The Solution of Tridiagonal Linear Systems on the CDC STAR-100 Computer," ACM Trans. Math. Soft., vol 1, no 4, December, 1975, pp. 308-329.

## A NEW TRIDIAGONAL SOLVER

MATRIX SIZE	TOTAL TIME (clocks)	CLOCKS/EQUATION
5	440	88
15	781	52.1
20	921	46.0
50	1956	39.1
101	3748	37.1
200	7752	38.8
401	15867	39.6
800	31312	39.1
1601	62411	39.0

Timing includes subroutine call overhead, approximately 40 clocks/call

Times to factorize and forward & back substitute a single tridiagonal system using an algorithm by Jordan (LASL) programmed by Forrest Brown (Nuc. Eng., UofM). Measured timings on UCS CRAY-1 (6/28/80).

- D. A. Calahan



UNIVERSITY OF MICHIGAN



3 9015 02652 7534