A CRAY RESEARCH, INC. PUBLICATION

CRAY CHANNELS

Volume 4, Number 3

ANNOUNCEMENT!
CRAY-1/M price/performance breakthrough

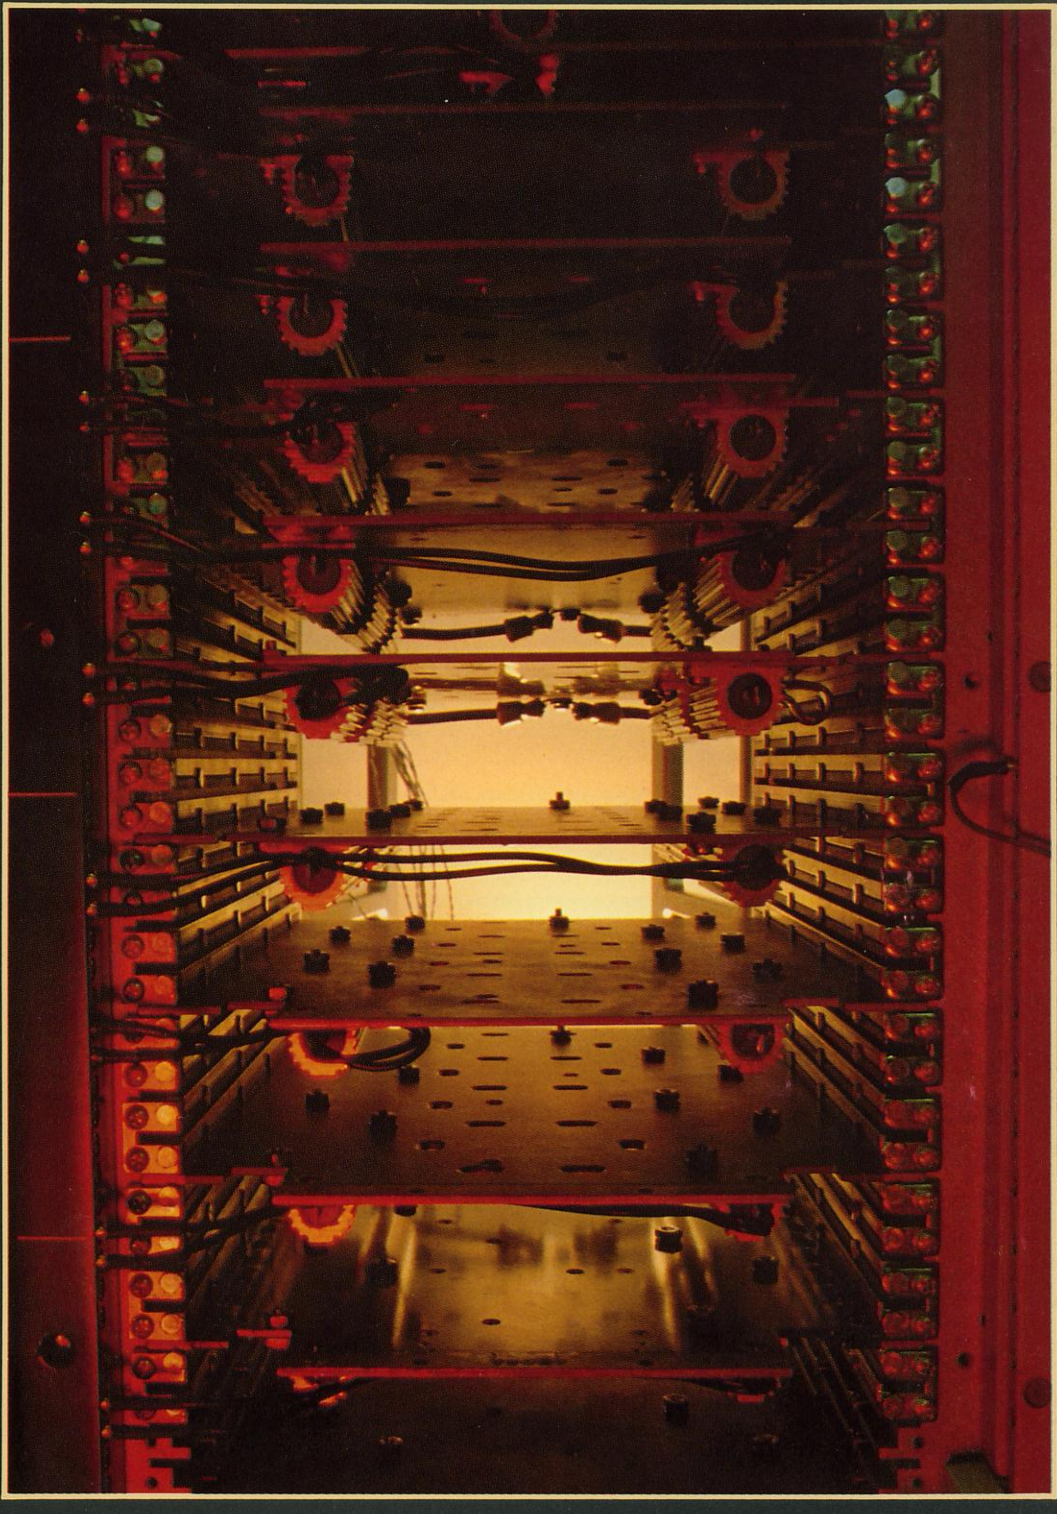

FEATURE ARTICLES:

The solution of sparse linear equations on the CRAY-1

Lawrence Livermore: 30 years of technical excellence

Cray Research's new system test equipment focuses on quality

REGULAR COLUMNS:

Corporate register

Applications in depth

User news

From the editor's desk

Fittingly, Cray Research has passed many milestones during its tenth anniversary year. The announcement of the CRAY X-MP in April allowed the company to offer customers a new level of supercomputer performance. The shipment of our fiftieth computer system in November was a truly satisfying accomplishment. But the year's climax occurred back in September, when Cray Research announced a supercomputing breakthrough. Our new CRAY-1 M Series of Computer Systems clearly sets a new industry standard for supercomputer price/performance. You can read about this exciting new product in a feature article beginning on page 2.

Also in this issue, we offer you an article on the solution of sparse linear equations on the CRAY-1. The author of this article, Dr. Iain S. Duff of AERE Harwell in the United Kingdom, has authored numerous papers on numerical linear algebra and mathematical software.

In another feature article, we discuss Cray Research's newest resources aimed at providing customers with products of the highest quality and reliability. And there's a look at Lawrence Livermore National Laboratory as it celebrates an anniversary of its own—its thirtieth.

Looking ahead to 1983, we see a year of change and further growth for Cray Research. The year promises more on the CRAY-2 and further enhancements to existing products. As ever, we're pleased that through **CRAY CHANNELS**, you'll be moving along with us into the year to come!

—T.M.B.

About the cover

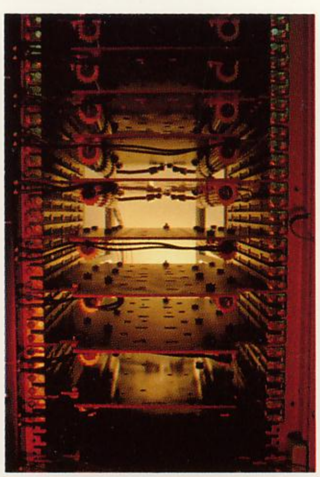

Heat-load modules moderate cooling in certain CRAY systems

Heat dissipated by the dense concentration of components in a CRAY computer must be handled using innovative cooling techniques. Cray Research's answer has been to run "cold bars" through each system. These bars carry a cold liquid refrigerant kept at a temperature of about 50°F, which is ten or more degrees cooler than the computer room environment.

Certain CRAY-1/M mainframes, I/O Processors, and the Solid-state Storage Device have "growing room"—that is, they hold one or more empty columns. When cold bars operate in a column that does not house heat-producing modules, conditions are ideal for the formation of condensation. To eliminate this possibility, Cray Research has developed special heat-load modules to even the heat throughout the system. Heat-load modules are spaced as shown in a system's unpopulated columns.

CRAY CHANNELS

A Cray Research, Inc. publication Volume 4, Number 3

CONTENTS

FEATURE ARTICLES

2

REGULAR COLUMNS

CREDITS

Editor	Tina M. Bonetti
Assistant editor	Carol A. Smith
Art/Design	Chad Jewett
	Kathleen Van Kirk
	Jeanne M. Jacobowski

CRAY CHANNELS is a quarterly publication of Marketing Publications, Cray Research, Inc., P.O. Box 154, Minneapolis, Minnesota 55440. **CRAY CHANNELS** is intended for users of Cray Research computer systems and others interested in the company and its products.

Requests for copies of **CRAY CHANNELS** may be addressed to the Marketing Publications Department of Cray Research at the above address. If you have a feature article or news item to submit for publication, please contact the editor of **CRAY CHANNELS** either by mail or by telephone at (612) 333-5889.

A N N O U N C I N G

The CRAY-1 M Series of Computer Systems

Now, Cray Research offers you more computational performance per supercomputer dollar than ever before. Announcing the new CRAY-1 M Series of computers, an evolutionary family of systems that sets a new standard for supercomputing.

Introduction of the CRAY-1 computer system in 1976 firmly established Cray Research as leader in the supercomputer market. The new CRAY-1 M Series carries on this tradition. Cray Chairman John A. Rollwagen said, "We're very excited about this new development. With the M Series, Cray Research is changing the price/performance curve for supercomputing—a move that ensures our continued leadership in the supercomputer market."

1982 has been Cray Research's year for rewriting the price/performance standards for supercomputing. Earlier this year, Cray introduced a new performance leader, the CRAY X-MP Series of Computer Systems. Now, acknowledging that leadership consists of price as well as performance, Cray Research announces the CRAY-1 M Series, offering CRAY-class power at traditional mainframe prices.

The CRAY-1/M is an extremely well-balanced system, offering equally strong scalar and vector performance. Its large Central Memory and powerful I/O structure make it ideal for a diverse range of applications including computer graphics, structural analysis, oil reservoir modeling, seismic analysis, and nuclear safety.

At the foundation of the CRAY-1 M Series is the field-proven design of the CRAY-1/S. However, whereas the S Series had a bipolar memory, the M Series Central Processing Unit features MOS memory, a technology currently used in Cray's Solid-state Storage Device and Buffer Memory.

The CRAY-1/M is available in three basic models with Central Memory sizes of one, two, or four million words. Central Memory is arranged in eight banks for the one-million word M/1200 and 16 banks for the two- and four-million word M/2200 and M/4200. The CRAY-1/M's Central Processing Unit features the same functional units and operating registers found on the CRAY-1/S, providing full software compatibility between the two products.

Cray Research's I/O Subsystem, which is an integral part of the CRAY-1/M, also contributes to the new system's outstanding performance. Two I/O Subsystem Processors and one million words of Buffer Memory are standard on all models of the M Series. But the CRAY-1/M has been designed to meet your future computing needs as well as your current performance requirements. One or two more I/O Processors may be added to the I/O Subsystem for supporting additional mass storage or magnetic tape devices. The I/O Subsystem's Buffer Memory may be upgraded to four or eight million words, accommodating more and larger I/O buffer areas and allowing memory-resident datasets. An optional Solid-state Storage Device can provide extremely high-performance data access to meet special demands, such as the handling of large datasets generated and manipulated repetitively by user programs. The M Series models are field upgradable from the smallest to the largest system.

The first CRAY-1/M is now available for benchmarking and demonstration at company facilities in Chippewa Falls, Wisconsin. CRAY-1/M delivery to customer sites is scheduled to begin in mid-1983.□

Features of the CRAY-1 M Series Systems

- *Thirteen segmented functional units that operate in parallel to support vector and scalar processing of fixed and floating point arithmetic as well as Boolean and related operations*
- *Four instruction buffers, each of which holds 64 consecutive 16-bit instruction parcels*
- *MOS Central Memory with either 1M 64-bit words arranged in 8 banks (M/1200), or 2M (M/2200) or 4M words (M/4200) arranged in 16 banks (SECDED)*
- *Input/output channel configuration featuring one or two 100-Mbyte/sec channels for transferring data between the IOS and Central Memory and four 6-Mbyte/sec channels for I/O control*
- *One standard and two optional front-end interfaces*
- *1M, 4M, or 8M 64-bit words of I/O Buffer Memory*
- *Two to 48 DD-29 Disk Storage Units*
- *One to 16 Block Multiplexer Channels, which can support user-supplied on-line magnetic tape units*
- *Optional Solid-state Storage Device (SSD) with 8M, 16M, or 32M 64-bit words of memory arranged in 16, 32, or 64 banks, respectively, providing exceptionally fast-access storage*
- *An I/O Subsystem composed of two standard and up to two optional high-speed I/O Processors*
- *A Peripheral Expander providing maintenance functions*
- *Power and cooling equipment*

The solution of sparse linear equations on the CRAY-1

Iain S. Duff, AERE Harwell, United Kingdom

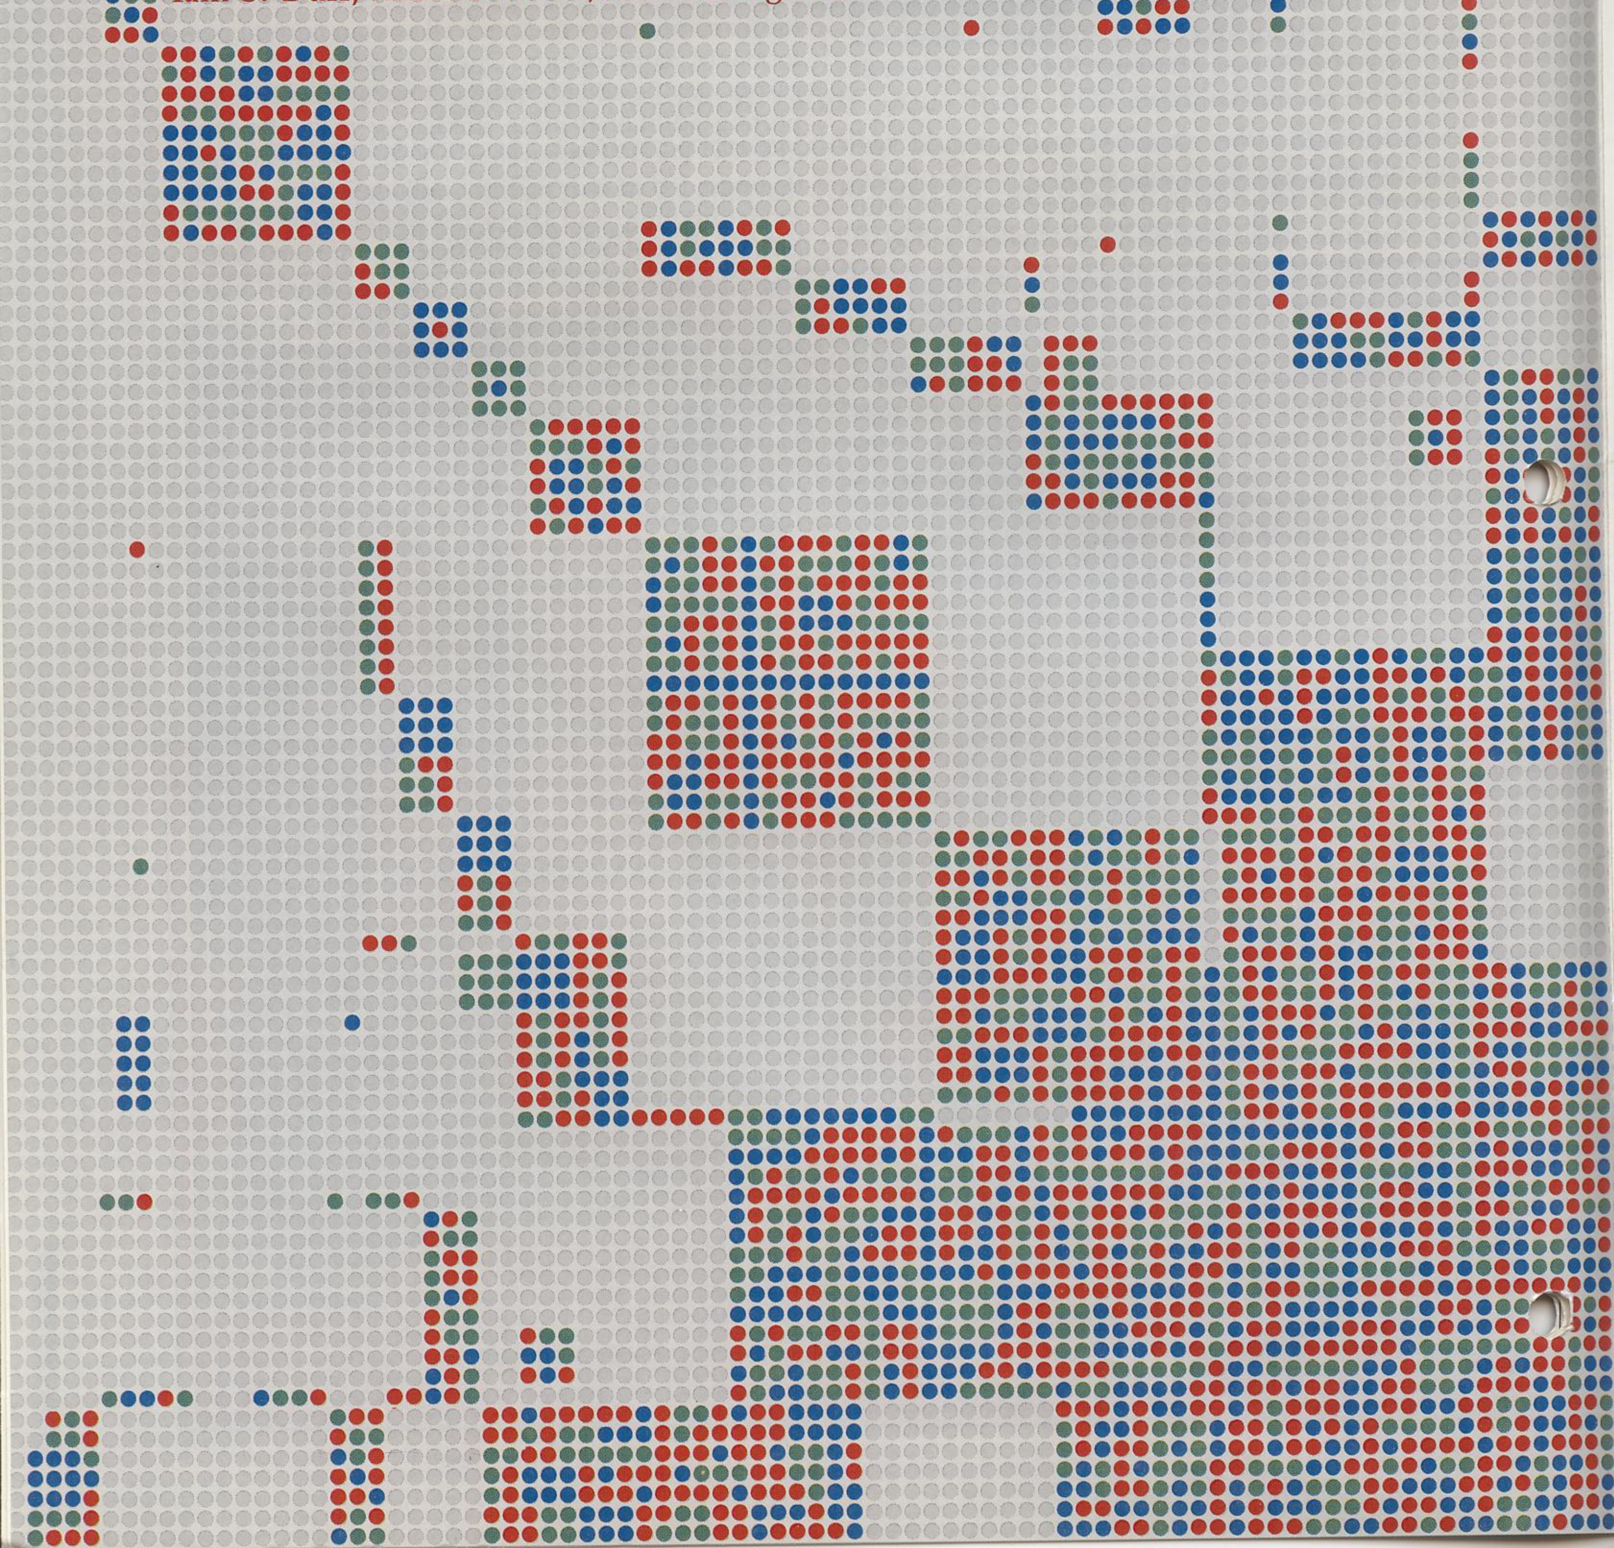

The solution to a large sparse set of linear equations lies at the kernel of many problems in computational science. In some applications it is possible to take advantage of special features or regularities in the matrix structure. This article is, however, concerned with techniques for the direct solution of sparse linear equations applicable to matrices of any sparsity pattern. We at Harwell are interested in solving general systems, largely because the subroutines developed for the Harwell Subroutine Library are used in a very wide range of applications, many of which are not known when the codes are initially designed.

The CRAY-1 computer has had a significant impact on large scale scientific computing, particularly in computations involving full matrices. Rates in excess of 130 million floating point operations per second (MFLOPS) have been recorded for the solution of linear equations and even higher rates for the multiplication of matrices. When the matrix is sparse however, the innermost loops generally involve at least one level of indirect addressing, thus inhibiting vectorization on the CRAY-1. This article discusses efficient implementation of indirect addressing and techniques for solving general sparse equations which avoid indirect addressing in the inner loops.

The vectorization of sparse codes

The principal problem with the vectorization of a general purpose code for sparse matrices lies in the use of indirect addressing in the innermost loop. Typically this is of the form

```
      DO 10 JJ = J1,J2
      J = ICN(JJ)
      W(J) = W(J) + AMULT*A(JJ)                (1)
   10 CONTINUE
```

and is often referred to as a sparse SAXPY.

One approach to enhancing the performance of a sparse code on a vector processor is to tackle loop (1) directly. This can be done by using a GATHER operation of the form

```
   I = J1,J2                                   (2)
   B(I) = W(ICN(I))
```

followed by a vectorizable inner loop

```
      DO 10 JJ = J1,J2
      B(JJ) = B(JJ) + AMULT*A(JJ)              (3)
   10 CONTINUE
```

followed by a SCATTER of the form

```
   W(ICN(I)) = B(I), I = J1,J2                 (4)
```

Since assembly-coded routines or microcode for implementing GATHER and SCATTER are likely to remain a standard feature of vector processors, this approach has the merit of comparative portability. Of course, for optimization on the CRAY-1, it may be more beneficial to code all of (1) in CRAY Assembly Language (CAL), as in CRAYPACK in the Boeing Computer Services Library. Table 1 shows the relative times and asymptotic rates for (1) on the CRAY-1, using the CRAY-1 FORTRAN Compiler (CFT Version 1.09) and the sparse SAXPY from CRAYPACK and using GATHER/SCATTER. The GATHER was performed within a vectorizable loop of the form

```
   DO 10 JJ = J1,J2
   B(JJ) = GATHR(LOCW + ICN(JJ)) + AMULT*A(JJ)
10 CONTINUE                                    (5)
```

where LOCW is the starting address of array W and GATHR, written at ECMWF, is a vector function. B was then SCATTERed to W using the CRAY SCILIB routine SCATTER.

Table 1
Vectorization of sparse SAXPY on the CRAY-1

	TIME FOR VECTOR LENGTH OF 10 (Micro-seconds)	ASYMPTOTIC SPEED (Vector length n) (Micro-seconds)	ASYMPTOTIC MFLOP RATE
CFT	6.1	.55n	3.6
CFT and GATHR/SCATTER	8.5	.29n	7.0
CRAYPACK	3.6	.16n	12.4

The CRAYPACK routine has been used to replace the inner loop of a general code for sparse symmetric equations, SPARSPAK, giving an overall rate of about twice the scalar speed. The degradation from the value in Table 1 is caused by the considerable amount of non-vectorized fixed point overhead in the code.

Gaussian elimination can also be implemented using inner (scalar or dot) products as in compact elimination. Conventional folklore suggests that the scalar product

```
      DO 10 I = 1,N
      S = S + X(I)*Y(I)                        (6)
   10 CONTINUE
```

often termed SDOT, is not easily vectorizable but, on the CRAY-1, a recursive vector sum instruction enables the CAL-coded CRAY SCILIB SDOT to attain an asymptotic rate of 74 MFLOPS. Although most sparse codes do not use an inner product formulation, a few of the SPARSPAK subroutines do, and these loops have been replaced by CAL-coded sparse

SDOTs in CRAYPACK. The asymptotic rate of the CRAYPACK sparse SDOT is 16.3 MFLOPS, and an overall rate of nearly 6 MFLOPS has been recorded for the appropriate modified SPARSPAK routines.

The effort in coding sparse SAXPYs and sparse SDOTs is very valuable and can significantly improve the performance of existing sparse codes on the CRAY-1 with little subsequent effort. However, algorithm redesign that increases performance efficiency is preferable. There are two main reasons for this. The first is that greater machine independence can be achieved by this approach. Second, although achieving an MFLOP rate on the CRAY-1 of over 16 for an inner loop employing indirect addressing is no mean performance, far higher rates can be attained if indirect addressing can be avoided altogether.

When redesigning an algorithm for improved vectorization, it is important that one does not merely increase the MFLOPS at a cost of performing more floating point operations. Harwell wishes to develop software which not only is applicable to a wide range of problems but also will perform well on non-vectorizing computers. The techniques discussed later all make use of the fact that code employing direct addressing in the solution of full linear systems can be vectorized easily. For example, rates in excess of 130 MFLOPS for the solution of sets of equations on the CRAY-1 have been reported. Indeed, experience with running a FORTRAN-coded linear equation solver on the CRAY-1 at Harwell indicates that a rate of 30 MFLOPS is possible even with code that was not designed for vectorization. The remainder of this article discusses three codes for the solution of sparse sets of linear equations which, for most or all of the inner loop operations, replace the sparse SAXPY of equation (1) with a direct SAXPY of the form

```
      DO 10 J = J1,J2
      W(J) = W(J) + AMULT*A(J)          (7)
   10 CONTINUE
```

While the maximum rate for this loop on the CRAY-1 is a 50 MFLOPS, by combining sequences of these, rates in excess of 100 MFLOPS can be attained.

Hybrid full and sparse codes on the CRAY-1

Even if the original matrix is quite sparse, the non-zeros created by the Gaussian elimination operation

$$a_{ij} := a_{ij} - a_{ik}\,(a_{kk})^{-1}\,a_{kj} \qquad (8)$$

where a_{ik}, a_{kk}, a_{kj} are all non-zero and a_{ij} was originally zero, cause the remaining active or reduced matrix to become increasingly dense. This is illustrated in Figure 1 where a sparse matrix is shown before and after Gaussian elimination with the pivots chosen down the diagonal in order. In the final stages of the decomposition, the reduced matrix becomes completely full, so code for Gaussian elimination on full matrices can be employed without any loss of efficiency. Even on a scalar machine this is advantageous, because not only does the removal of indirect addressing speed up the inner loop, but the complicated data manipulation present in a general sparse code can be avoided. The gains are much greater on a vector processor since the part of the elimination when the reduced matrix is treated as full can be performed at the asymptotic rate of codes using direct addressing (130 MFLOPS or more on the CRAY-1).

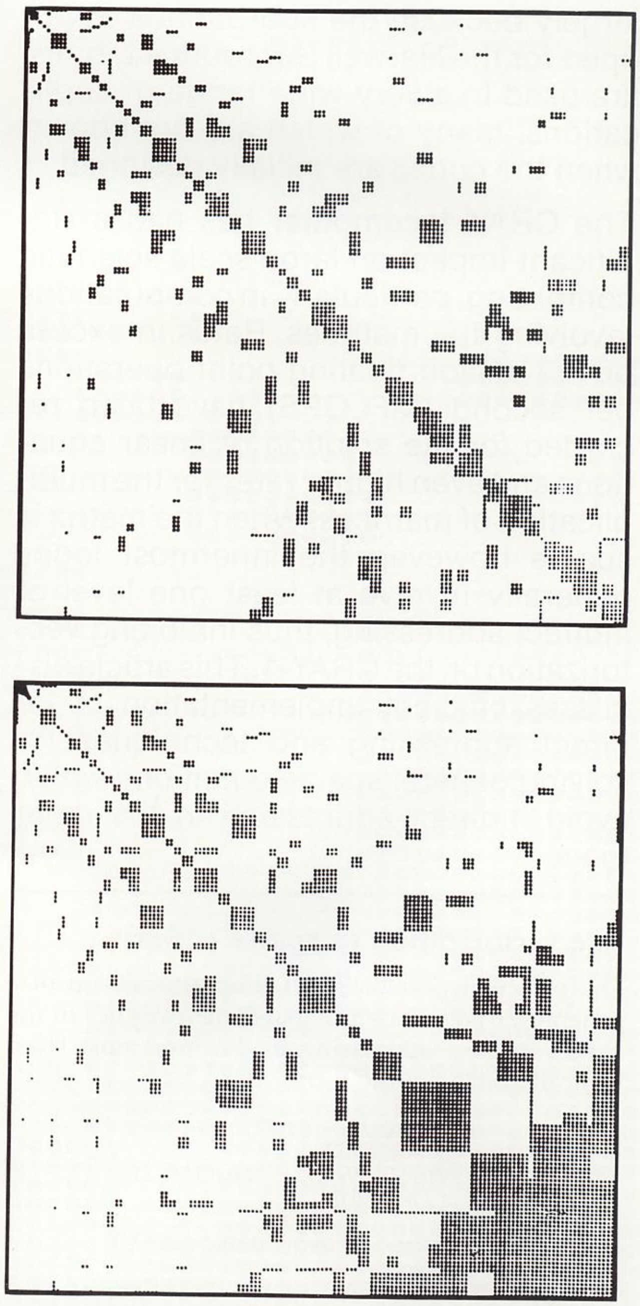

Figure 1. Pattern of a sparse matrix before and after Gaussian elimination with pivots chosen down the diagonal in order.

Table 2
Time (in seconds) for runs of MA28 on the IBM 3033 and the CRAY-1

The density of reduced matrix at which the switch to full code is made is shown in the left hand columns

IBM 3033

Order	147	1176	292	541
Non-zeros	2441	18552	2208	4285
	Time for factorization			
No switch	.724	4.143	.459	1.207
1.0	.611	4.281	.454	1.180
0.8	.501	2.772	.445	1.160
0.6	.466	2.872	.416	1.137
0.4	.467	3.373	.428	1.175

CRAY-1

Order	147	1176	292	541
Non-zeros	2441	18552	2208	4285
	Time for factorization			
No switch	.427	2.585	.276	.749
1.0	.365	2.610	.278	.739
0.8	.277	1.541	.250	.723
0.6	.207	1.429	.228	.696
0.4	.153	1.055	.201	.657
0.2	.114	.806	.169	.587
0.05	.132	.806	.252	.810

Harwell has recently incorporated a switch to full code in Harwell Subroutine MA28 for the solution of unsymmetric sparse linear equations. The results of running this code on some standard sparse test problems are presented in Table 2. While the switch to full code gives gains of up to 30% on the IBM 3033, the savings on the CRAY-1 can be up to 70%. Also, it pays to switch to full code at a much lower density of the reduced matrix. The differences between the results in Table 2 are not unexpected. Indeed, if the full matrix code were written in CAL rather than FORTRAN, the gains when using the CRAY-1 would be even greater and the fastest times would occur at even lower densities of the reduced matrix. The threshold at which it pays to use sparse code shifts considerably when using the CRAY-1.

Naturally, switching to full code before the reduced matrix is completely full will normally require more storage for the entries in the factors because some zeros will be held explicitly. At first glance, this may appear to be the penalty that must be paid for gains due to vectorization. However, in addition to avoiding indirect addressing when using full code, it is not necessary to store integer indexing information on the position of the non-zeros. When this saving is taken into account, less storage is required until the switch-over density becomes very low. Indeed, in most cases, the switch-over density can be set for the fastest times with minimal extra storage required.

These results have been obtained using a prototype code and are likely to be more favorable to switching when further refinements have been made.

The use of the Harwell frontal code MA32

A common algorithm that avoids indirect addressing in the inner loop is the variable band or profile scheme. This is a generalization of a band matrix approach that permits the first non-zero in each row of the matrix to be a varying distance from the diagonal. All entries between the first non-zero and the diagonal are stored and operated upon; with this it is possible to use direct addressing in the inner loop. Matrices of a suitable form for profile elimination arise with some simple orderings of discretizations of partial differential equations. The SPARSPAK package has an option for implementing this ordering. With CAL-coded inner loops overall MFLOP rates of nearly 7 for the matrix factorization and over 9 for the subsequent solution of equations are obtained. The full asymptotic rates of direct SAXPY and direct SDOT are not achieved because of non-vectorizable data manipulation in the code and the short length of the vectors processed in the inner loops. It is possible that a complete redesign of the algorithm could yield higher rates, but we prefer to concentrate on the frontal approach which can be viewed as a generalization of the variable band technique.

The frontal method as implemented in the MA32 package in the Harwell Subroutine Library can be used to solve any set of unsymmetric linear equations. The approach in terms of finite element problems is described below. In a finite element problem the matrix A is a sum

$$A = \sum_{\ell} B^{\ell} \qquad (9)$$

where each B^{ℓ} has non-zeros in only a few rows and columns and corresponds to contributions to the matrix from finite element ℓ. It is normal to hold B^{ℓ} in packed form as a small full matrix together with an indexing vector to identify where the non-zeros belong in A. The basic "assembly" operation when forming A can be written as

$$a_{ij} := a_{ij} + b_{ij} \qquad (10)$$

If we examine the basic step in Gaussian elimination,

$$a_{ij} := a_{ij} - a_{ik}\,(a_{kk})^{-1}\,a_{kj} \qquad (11)$$

it is evident that it can be performed before all assemblies (10) are completed, provided only that the terms in the triple product in (11) have been fully summed (that is, have had all sums of the form (10) completed). In a frontal code the elements are assembled one by one and each variable can be eliminated whenever its row and column is fully summed, that is after its last occurrence in a B^{ℓ}. This permits all intermediate working to be performed in a full matrix whose size increases when a variable appears for the first time and decreases when one is eliminated. The full matrix in which all arithmetic is performed is called the frontal matrix. For symmetric positive definite matrices, all variables may be eliminated as soon as they are fully summed. For more general systems, some form of pivoting is required to ensure numerical stability. For non-element problems, the rows (equations) are "assembled" one at a time and a variable becomes full summed whenever there are no further equations in which it appears. This entry is the generalization of variable band methods discussed earlier. For the purpose of this study, all pivoting and elimination operations are performed within a full submatrix so direct addressing can be used in the inner loop.

The frontal code MA32 has recently been used on the CRAY-1 by the Theoretical Physics Division at Harwell in the solution of problems addressing buoyancy driven flow in a square cavity using finite elements. Some results are shown in Table 3 where the elements used for the runs in the first two columns were five-node rectangular elements with three variables at each corner and one at the centroid. The element for the run in column three was a rectangle having nodes at the corners, midpoints and centroid with three variables defined at each node. The results have been augmented by runs on two artificial problems. The first (column 4) is a grid of nine-node elements similar to those used in the runs in column 3 but with five variables at each node. The second (column 5) is a finite difference problem arising from the five-point discretization of the Laplacian operator on a rectangular grid. In every case the elements or equations were assembled in a pagewise ordering along the side of shorter dimension. In all the runs in Table 3, the innermost loop of the frontal code was a direct SAXPY coded in FORTRAN and vectorized by the CFT compiler (version 1.09).

Table 3
Performance of frontal code on the CRAY-1

Dimensions of grid of elements or equations	28x28	36x36	20x32	16x16	64x64
Order (degrees of freedom)	3023	5039	9480	5445	4096
Total time in seconds for factorization of matrix and solution of equations	3.83	8.63	13.43	12.58	3.55
Time in innermost loop (in seconds)	1.70	4.29	6.96	7.63	1.49
Number of operations in inner loop (in millions)	44.2	120.9	198.6	238.4	32.6
Inner loop MFLOPS	26	28.2	28.5	31.2	21.9
Total MFLOPS	11.5	14.0	14.9	19.4	9.2

The two artificial cases have been run on the IBM 3033 at Harwell and the times for the element and the equation input were 133.3 and 26.5 seconds respectively. The overall increase in speed due to the vectorizing capability of the CRAY-1 was 10.6 and 7.5 with the inner loop running over 15 and 11 times faster on the element and equation problem respectively.

Tuning the frontal code for the CRAY-1

As discussed previously, the innermost loop of the frontal code is a direct SAXPY. Table 4 shows the MFLOP rate for a range of implementations of this basic loop on the CRAY-1. In this table, the relative costs of start up times are reflected in the performance at different vector lengths.

Table 4
MFLOP rates for different implementations of direct SAXPY

Vector length	FORTRAN CFT version 1.09	Cray SCILIB	Boeing CRAYPACK	Mostyn Lewis' CAL code
20	17.7	11.3	16.7	22.5
30	21.2	15.0	21.4	28.6
50	25.4	20.4	27.8	35.7
100	29.7	28.0	35.7	39.1
150	31.5	31.9	39.5	40.2
200	32.4	34.3	41.7	40.7
300	33.4	37.1	44.4	42.8
Asymptotic	35.7	44.4	50.0	48.0

Since the SAXPY operation

$$\underline{x} := \underline{x} + a\underline{y} \qquad (12)$$

requires two memory loads and one store for each pair of floating point operations, the optimal rate is limited by memory accesses and is bounded by 50 MFLOPS when the multiplication and addition are chained with one of the loads. Thus, very little further improvement can be obtained over the figures in columns 2-4 of Table 4 which fall a long way short of the full potential of the CRAY-1.

However, higher rates can be obtained by observing that, at each elimination stage in the frontal code, a sequence of direct SAXPYs corresponding to eliminations by the pivot row on all other rows of the frontal matrix is performed. The sequence can be represented by several SAXPYs of the form (12) where the vector $\underline{x}$ and the scalar a change but $\underline{y}$ (which represents the pivot row) remains constant. Thus, if we can keep $\underline{y}$ in the vector registers throughout the sequence of SAXPYs we need only one load and one store for each two floating point operations, yielding a maximum asymptotic rate of nearly 80 MFLOPS. The CAL code in column 5 of Table 4 implements this idea and although further optimization is possible, it is clearly faster than the maximum rate for a single SAXPY.

It is possible to do even better. In any realistic problem, there are several fully-summed rows and columns in the frontal matrix at each stage. If just two elimination steps can be combined, then our inner loop can run at well over 100 MFLOPS. This is done by loading the two pivotal rows into the vector registers and chaining the load of the first non-pivotal row with the floating point operations from the first pivotal row. The SAXPY operations from the second pivotal row are then chained and the second non-pivotal row is simultaneously loaded from memory. The operations of the two pivot rows on this second non-pivotal row are then overlapped with the store of the first non-pivotal row and the load of the third non-pivotal row respectively. We continue in this way always keeping both arithmetic pipes busy, and so achieving optimal performance. A prototype CAL code implementing this runs at over 90 MFLOPS and it is hoped to improve this substantially.

Although the prospect of a fairly general sparse code whose inner loop runs at over 100 MFLOPS may raise one's pulse rate considerably, a small caveat is in order. As shown in Table 3, the ratios of inner loop time to total time for the 16 x 16 element example on the IBM 3033 and the CRAY-1 are .88 and .61 respectively, reflecting the fact that the inner loop time is reduced much more than the rest of the code through vectorization. Thus, by speeding up the inner loop even more, this ratio will decrease further until the computation time in the inner loop ceases to be dominant. This is a common phenomenon of vectorization. Indeed, if we can get the inner-loop running at 130 MFLOPS, the overall MFLOP rate for the five problems in Table 3 is only increased from 11,14,15,19 and 9 to 18,23,25,35 and 14 respectively. Of course, this is on the assumption that the other now dominant parts of the code cannot be further vectorized. However, it is believed that they can be, and this will be the subject of further investigation.

The performance of the Harwell multifrontal code, MA27

The frontal code will become less efficient if variables remain in the front for many assemblies without becoming fully summed. However, the frontal matrix can be viewed at any stage as one of the B^ℓ of equation (9), albeit in general of larger size than the original elements. If the frontal matrix is stored, and the computation continues by assembling other elements (including previously stored frontal matrices) and eliminating fully-summed variables, the overhead of retaining variables in the front and performing operations on them at each step can be reduced. This class of method is called multifrontal since several previous frontal matrices might be held in store during the computation. Code MA27 in the Harwell Subroutine Library implements this method efficiently using a simple stack to hold any stored frontal matrices. Furthermore, MA27 obtains a numerically stable solution for sparse symmetric systems whose coefficient matrix is not definite.

Since this multifrontal method also uses only direct addressing in the inner loop, it should perform well on vector processors. Its performance is illustrated in Table 5 where MA27 is compared on standard sparse test matrices with a sparse code, YSMP, for general systems whose coefficient matrix is symmetric and definite. Although the YSMP code is comparable to the more general MA27 code on the IBM 3033, the MA27 code is significantly faster on the CRAY-1, where the inner loops of MA27 vectorize.

Table 5
Times in seconds
on the IBM 3033 and CRAY-1 for MA27 and YSMP

Order	**1561**		**1005**		**900**	
Non-zeros	**6121**		**4813**		**4322**	
	IBM	**CRAY**	**IBM**	**CRAY**	**IBM**	**CRAY**
Factorization times						
MA27	1.10	.255	.897	.193	.575	.144
YSMP	1.04	.420	1.059	.410	.578	.237

Conclusions

It is possible to design algorithms for the direct solution of sparse linear equations that use direct addressing in their innermost loops and will perform well on vector processors. The algorithms discussed also perform well on scalar machines and make no particular demands on the structure of the coefficient matrix. The performance of our algorithms on the CRAY-1 have been illustrated and an indication given of how special coding can be used to obtain very high computational speeds.□

Further References

The complete version and references for this article are found in: Duff, Iain S., **"The Solution of Sparse Linear Equations on the CRAY-1," Harwell Report CCS 125,** presented at the Science, Engineering and the CRAY-1 Conference, April 5-7, 1982. Copies are available through Cray Research, Inc., Applications Department, 1440 Northland Drive, Mendota Heights, Minnesota 55120.

About the Author

Iain S. Duff received his B.S.c. degree from the University of Glasgow in 1969, the Diploma in Advanced Mathematics and D. Phil degree in Mathematics from the University of Oxford in 1970 and 1972 respectively. He then spent a year in the United States as a postdoctoral research worker on a Harkness Fellowship, visiting Stanford University and the State University of New York at Stony Brook. From 1973 until 1975, he was a lecturer in the Computing Laboratory at the University of Newcastle upon Tyne. Since 1975, he has been working in a research and advisory capacity, assisting in maintaining and developing the Harwell Subroutine Library. His main interests are in numerical linear algebra and mathematical software, with strong emphasis on sparse matrices. Dr. Duff has authored numerous papers on the same and is an Associate Fellow of the Institute of Mathematics and its Applications.

Lawrence Livermore: 30 years of technical excellence

Lawrence Livermore National Laboratory (LLNL) celebrated its 30th anniversary on September 2, 1982. The Laboratory, located about an hour's drive from San Francisco is an example of founder E.O. Lawrence's modern large-scale team research program at its best. For 30 years many of our country's most important research programs have been undertaken by Livermore scientists. In 1978 Lawrence Livermore installed its first CRAY-1 computer. Since then, five additional CRAY systems have been installed in two Livermore-based computer centers. The availability of supercomputers has had a significant impact on the manner and type of research that is conducted at LLNL. We thought we'd take this opportunity to tell you a little about what that research has been all about. At the same time we'll look at why the Laboratory was established and where it is going as it enters its fourth decade.

Research needs and Laboratory are established

Since its founding in 1952, the Lab's primary mission has been research and development for national security. Livermore has also been committed to bringing the technologies born out of those research efforts into practical application. The Lab has been successful in extending its nuclear science and engineering technologies to many other key application areas. Among the most prominent non-military programs are the magnetic and laser fusion energy programs, non-nuclear energy development, biomedical and environmental science, pure physics, chemistry and engineering.

The single most important event triggering the establishment of LLNL in 1952 was the Soviet's unexpected detonation of their first atomic bomb in 1949. Ernest O. Lawrence, 1939 Nobel laureate in physics and founder of the University of California Radiation Laboratory at Berkeley, had been concerned for some time after World War II that a larger effort was required for the United States' nuclear development. Edward Teller, one of the world's most brilliant physicists and a key participant in the Manhattan Project, was also a proponent of accelerating development efforts in light of the Soviet Union's unexpected advances.

Three years after the Soviets' first bomb blast, Lawrence and Teller were successful in obtaining the Atomic Energy Commission's approval for a new laboratory. The site was an abandoned naval base in Livermore, California—only an hour's drive from Lawrence's laboratory on the Berkeley campus. Thus, the University of California Radiation Laboratory at Livermore, later to become the Lawrence Livermore National Laboratory, was established. The Laboratory is still operated for the U.S. government by the University of California. LLNL's charter committed it to addressing immediate and long-term national security concerns by focusing large-scale research efforts on thermonuclear weapons design, diagnostic weapons experiments and basic physics research. At the same time LLNL's steering committee, of which Teller was a member, decided that a major secondary goal for the Lab was the development of useful power from controlled fusion energy.

Early endeavors

The talented young cadre of scientists at the Laboratory began their research with great resolve. However, the results of some of their early efforts stand in sharp contrast to their later successes. The first Livermore-designed shot was fired during the Upshot-Knothole Operation at the Nevada Test Site in early 1953. With the device poised atop a 300 foot tower, the only indications that the shot had gone off were a small spark of light on the horizon and a swirl of dust. Though mangled, even the tower had been left standing. "It was really embarrassing," Wally Decker, an engineer on the project recalls. "We were all very disappointed, especially since we knew Los Alamos (LLNL's older sister laboratory) was watching our every move." Shortly thereafter, in 1953, the Soviets announced the detonation of their first thermonuclear bomb. Research efforts redoubled with that hard reality. In the three decades since, Livermore scientists' successes have included the provision of many of the country's major defense systems.

The necessity of computers realized early on

The advances made in these and other areas of research at LLNL have largely been dependent on the availability of sophisticated computing power. Recognizing this correlation at the outset, the Laboratory installed its first computer, the Univac 1, in early 1953. Since then, Lawrence Livermore has continually upgraded its computing resources as its research projects have grown. Today, the test ground for nuclear reactions has moved from the outdoors into computer systems. Simulation computations that once required 100 hours to solve now typically take five hours and provide better results and finer resolution. However, 60 hours is still not unusual for a major simulation effort. Fortunately, computers have come a long way since the Univac 1—Livermore computer experts have calculated that one of their CRAY-1s has 22,200 times the power of their original Univac 1.

Major fusion energy research programs

It is often assumed that Livermore's great computational capacity is tied only to defense systems development, but these resources have also been put to non-defense oriented uses. The most important ongoing non-defense project began in 1952 at the Laboratory's founding. The quest to generate useful power from controlled nuclear fusion reactions has been a significant and expanding part of Livermore's work for 30 years. The motivation has been constant: to respond to an anticipated national need—the inevitable exhaustion of fossil fuel resources. By tapping the virtually inexhaustible and very low cost fuel deuterium, the kind of heavy hydrogen that is found in the ocean, scientists expect to reshape the way we generate energy. Fusion may quite likely become man's primary energy source in the future. Conceptually simple but technically demanding, fusion is a process in which heated and confined atomic nuclei collide and join to release tremendous amounts of energy. To be practical, the process must release more fusion energy than is required to confine and heat the fuel. (The same type of fusing process in the sun produces its light and heat.) Two separate strategies, magnetic and laser research efforts, are underway in pursuit of the practical manufacture of fusion energy at LLNL.

Magnetic fusion energy research

One strategy contains the very hot fusion plasma in a very strong magnetic field. Intensive efforts in magnetic fusion research are being conducted by the Magnetic Fusion Energy program. Magnetic fusion relies on heating thermonuclear fuel, a mixture of deuterium and tritium, to about 100,000,000° C and containing it in the magnetic field long enough for efficient fuel burnup. Once this is achieved, the energy released must be converted to a useful form (electricity or process heat), the combusted fuel must be replaced and the "ashes", which are helium, must be removed.

Magnetic fusion research aims at a practical solution to these various requirements using specially configured strong magnetic fields produced by powerful magnetic coils surrounding the fusion combustion chamber. The magnetic field acts to confine and isolate the heated fusion fuel from contact with the vacuum chamber in which it is contained. The quality of confinement, ability to reach combustion temperature and conversion of the energy released to useful form are key challenges to developing efficient magnetic fusion energy.

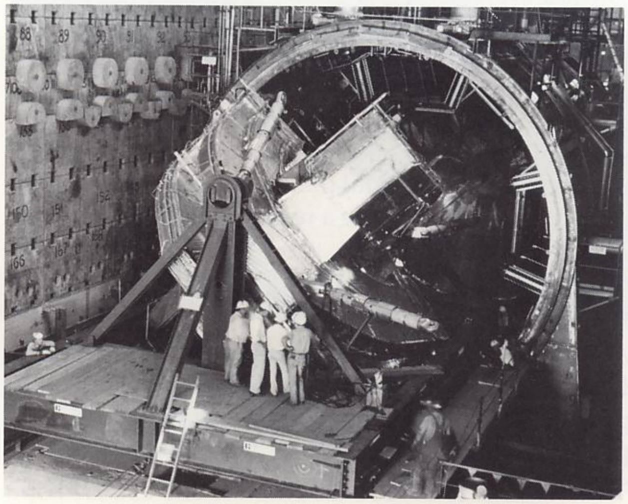

The 375-ton superconducting magnet pair is readied for fusion experiment at LLNL. When operating, this magnet will produce a field up to 150,000 times that of the earth. The force produced by the magnets repelling each other is up to 22 million pounds—the weight of 30 jumbo jets.

The extreme complexity of this problem led to the establishment of the National Magnetic Fusion Energy Computer Center (NMFECC) at LLNL in 1974. Two CRAY-1 systems are the major workhorse computers at the NMFECC. The center provides large-scale computational support to the National Magnetic Fusion Energy community of researchers in national laboratories, universities and industry throughout the United States. Complex time-consuming codes modeling plasma in one and two-dimensions are effective predictors of the fusion material's behavior. Three-dimensional plasma simulation codes also provide accurate models by carrying many calculations in discrete steps.

The laser fusion research program

A second strategy in fusion energy development, called Inertial Confinement Fusion, uses very short impulses of intense laser light. In this process, the pulses of a powerful laser are focused on small targets, some the size of a salt grain. A jolt of laser energy causes the target to implode, compressing the target's fuel and producing fusion reactions. Laser fusion is dependent on the inertia of the fuel to maintain a highly compressed state long enough to react and produce significant energy.

Encouraging computation results led to the establishment of LLNL's Laser Fusion Program in the early 1970's. Shiva, a $25 million multi-beam laser system, was commissioned in 1977 for an intensive and successful schedule of target experiments. A wealth of physics data came out of these experiments before it was decommissioned late in 1981. Subsequent lasers named Novette and Nova are part of continued laser fusion development efforts.

Juxtaposed with the great rewards for achieving fusion are the comparably great difficulties of attaining it, as three decades of worldwide research attest. However, Lawrence Livermore's fusion energy researchers have opened a major path to this goal. Today it is no longer a question of whether fusion will work, but how soon and by what means. Scientists at Livermore are hopeful that by the close of this century fusion energy will be a reality on a practical scale.

Biomedical and environmental research yields important findings

Another major program that was initiated in the early 1960's is the Biomedical and Environmental Research program. Born out of the concern that many health problems, current and future, have their origins in slowly developing injurious effects of the technological environment, the program was set up to study the elusive links between pollutants and human health. Many of the findings coming out of this program affect the direction other projects take. "We have much more to learn about human health and the complex workings of the environment before we can pinpoint the costs or risks of energy alternatives (for example)," says Mort Mendelsohn, program director. The program has made significant contributions to developments in cell biology, biochemistry and cell kinetics. Today, Lawrence Livermore is known as the world's foremost center for analytical cytology.

LLNL computing, yesterday and today

Lawrence Livermore has never been shy about voicing its desire for faster and larger computers. It has been a motivating force in the development of each computer generation's fastest systems and a major user of the same. In addition, LLNL has been instrumental in developing many computing concepts that are now taken for granted. LLNL was represented on a team that designed FORTRAN. Livermore scientists were among the first to use batch processing systems and developed one of the earliest time sharing systems.

The advent of interactive time sharing at the Laboratory was closely associated with another development—networking. The combination of the two yielded a unique computing environment that has been duplicated at only one other place on anywhere near the same scale, at LLNL's sister laboratory in Los Alamos. Over 15 years Livermore has built a flexible web of interacting computers in networks that are more general and powerful than most networks in operation today. At the heart of one of these networks are four CRAY-1 computer systems. These systems are major contributors to a majority of the diverse research efforts at LLNL. The Lab's broad experience in large scale computing for scientific applications was invaluable when LLNL was chosen as the central site for the National Magnetic Fusion Energy Computer Center, which operates the additional two CRAYs that are installed at Livermore.

According to LLNL computer experts, speed is the most important component of scientific large-scale computing. George Michael, leader of the Computation Department's research group commented, "The way to increase speed is with a judicious mixture of improvements in computer architecture, language and algorithms." The Computation Department plans to continue breaking ground in computer architecture, language, and algorithms. However, John Ranelletti, head of the department, notes, "The department is not in the business of hardware development. Software and mathematical algorithms are what we're all about." According to Livermore computer scientists, parallel processors that simultaneously do parts of a single job are one of the coming innovations in supercomputers. The CRAY X-MP multiprocessing system with two identical Central Processing Units that enable multiprocessor jobs, is an example of the type of computing capability that will be required in years to come.

Lawrence Livermore continues to address the country's largest technological challenges

Lawrence Livermore National Laboratory has pursued the largest scientific and technological challenges facing the United States for three decades. From the onset, LLNL has approached its research with rigor and competency, augmented with the most advanced computing tools of the day. The results of this research help shape the way we live today and address our needs in the future. The Laboratory is recognized the world over for the progress it has made in challenging scientific frontiers. However, in the final analysis, without the commitment of LLNL's people, this work would not have been accomplished. "Without occasionally overextending yourself," Edward Teller once said in recalling the Lab's early disappointments, "you will get nowhere." Those words have been taken to heart by those at Lawrence Livermore who work toward ensuring our national security, a quality environment and energy independence. □

Acknowledgements

Special thanks are extended to all those at LLNL who graciously gave their time in the development of this article, especially Dr. Gus Dorough and Michael Ross.

Cray Research's new system test equipment focuses on quality

One measure of a company's commitment to product quality is the value of resources devoted to quality control. Cray Research's past record and current activities in this area are strong indicators of an ongoing dedication. This article describes several of Cray's newest resources aimed at providing customers with products of the highest quality and reliability.

The complexity of a supercomputer calls for weeks of testing and analysis, from the individual component level to testing of completed systems. This thorough testing and analysis can point out general areas of system weakness or failure, but it does not directly answer the questions of *why* a computer system does not work and *how* to keep it from happening again.

Now, however, technicians at Cray Research's Chippewa Falls facilities can run tests to provide the valuable answers to these important questions rapidly and accurately. The company's newly-established inhouse system test facility includes a JEOL Scanning Electron Microscope (SEM) and a Tracor Northern X-ray spectrum analyzer. The SEM and X-ray spectrum analyzer units together offer an extremely wide magnification range (10 to 180,000 times actual size) and analysis of materials composition.

Dr. Charles Fuller, Sr. Chemical Engineer for Cray Research, has been responsible for setting up the facility and establishing a reliability enhancement group. Prior to purchasing the SEM, he said, some of the more complicated field and assembly problems were analyzed by sources outside of the company. Now, says Fuller, "We have the ability to do a lot of that analysis ourselves. An added bonus is that frequently the engineers involved can be there with us, watching while we do an analysis. Not only does this give instant turnaround for results, but it also provides a special kind of interaction that would be almost impossible with an outside lab." Testing services are available to Cray engineers, analysts and other Cray employees requiring analysis of component or mechanical problems.

What the SEM shows

Images created by the SEM differ from those seen through conventional microscopes. An SEM image is generated when a tiny electron beam (100 Ångstroms wide) is rastered across the sample and then displayed on a terminal. By contrast, conventional microscopes use a lightbulb source of illumination and a series of powerful lenses for magnification.

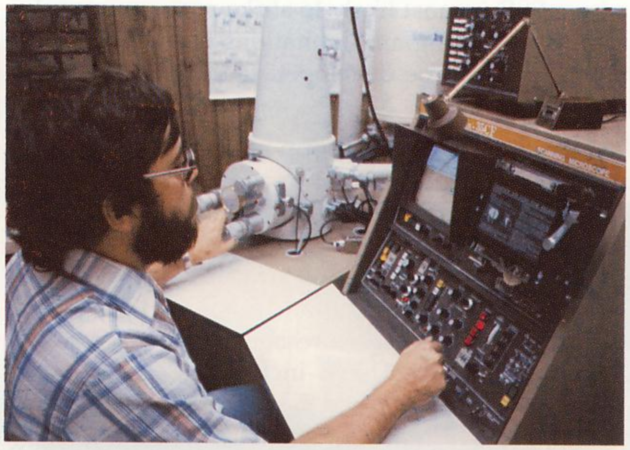

Dr. Charles Fuller reading results of the X-ray spectrum analyzer. The Scanning Electron Microscope is the tall cylindrical tool in the background.

A magnified SEM picture (known as a secondary electron image) has several advantages over conventional images. These include better depth of field and provision of a measurement scale. The improved depth of field makes it possible to have a three-dimensional picture of the sample. Additionally, the viewer can zoom in for a better look at a targeted region, which is often difficult or impossible to do on a standard optical microscope. Images from the SEM are very realistic and of high resolution, making minute, otherwise invisible segments of the sample available for viewing.

Some major SEM uses

The SEM proved helpful recently in studying a problem with a wire bonder, the tool used for bonding wires within the chips. A malfunctioning bonding tool can destroy thousands of chips and waste valuable employee hours, so early detection is critical. In this particular instance, the bonding wedge on one of the wire bonders had gone through several visual and optical checks without detection of flaws. Explains Kevin Stanton, Sr. Electronics Technician, "Our preliminary tests had indicated that the bonding wedge was functional, yet the employees doing the bonding work knew that something was wrong." When technicians viewed the tool under the SEM, they detected a sharp fracture in the bonding wedge. Once this diagnosis was made, the broken part was replaced and work resumed.

In development efforts, the SEM can assist in identifying possible circuit design improvements. Maximizing circuit speeds while keeping the incidence of electrostatic discharge (ESD) to a minimum is an important part of Cray Research's quality goal. Therefore, recognition and redesign of weak circuit areas where ESD is most likely to strike proves very valuable. By verifying any limitations in a design, new procedures can be developed to keep the problem from recurring.

The product reliability group in Chippewa Falls has requested use of the SEM to analyze the connector alloys used in construction of CRAY systems. For this purpose, a 5400-power magnification was sufficient to provide pictures and an X-ray spectrum of good quality and high resolution. Remarked Pete Peloquin, head of the product reliability group, "The SEM is a tool that will have increasingly greater use throughout the company."

Compositional analysis is also possible with the new equipment. Heavier elements, such as mercury or lead, give off more electrons than lighter elements like aluminum and carbon. By looking back at scattered electron readings, the knowledgeable viewer can easily see whether the components are consistent with what the physical composition of the sample should be. In other words, by looking at the compositional reading of a sample, a technician can determine whether an unwelcome contaminant is creating failure.

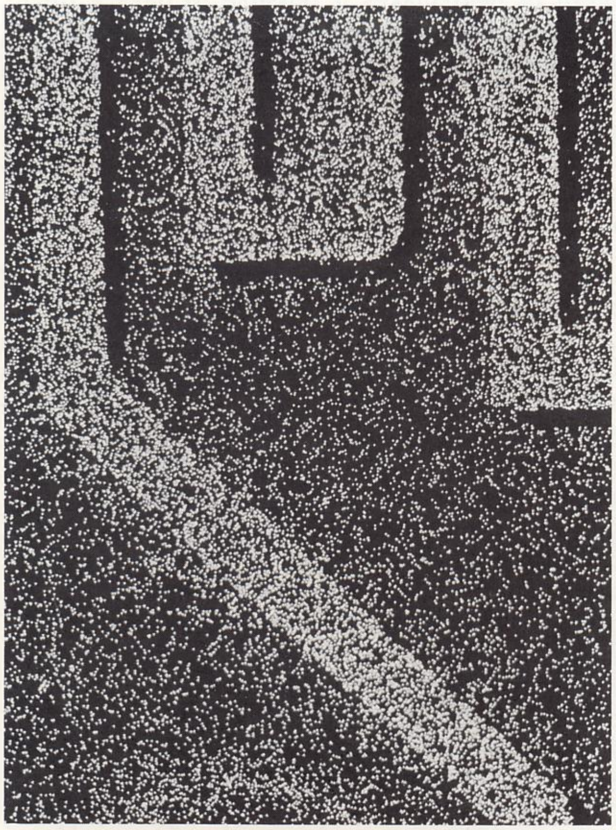

Example of an X-ray dot map.

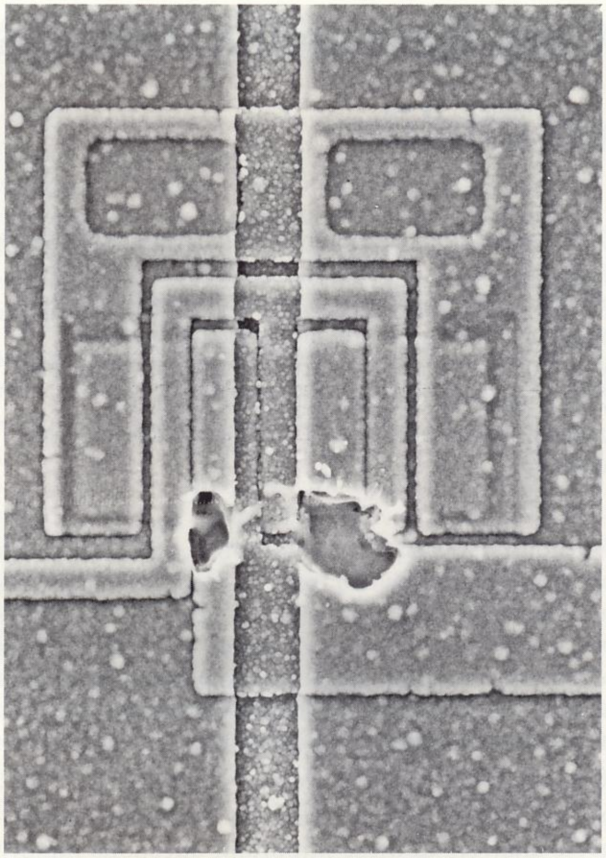

SEM image of an integrated circuit transistor that was "zapped" with electrostatic discharge. The square transistor structure seen in this picture is about 60 microns on a side (about 0.0025 inches).

Other equipment uses

In addition to magnification and physical composition, the X-ray spectrum analyzer attached to the SEM can be used to tell which chemical elements are present in a sample. For example, it might be important to know whether a particular part had been plated with tin, and the analyzer can determine this.

The SEM also can be used to run a pulse through the pathway of an integrated circuit, thereby showing how the circuit is functioning. In this process, the voltage path of the circuit is illuminated and the brightness is monitored on the screen. If a circuit has a short or an open cut, the illumination will show a bias, which is a differential between expected brightness and registered brightness. With this tool, bad chips can be identified.

An emphasis on quality

To ensure high quality and reliability in all of its products, Cray Research must continue to develop methods for measurement and detection of potential problems. The addition of the SEM is one such building block in the company's quality control activities. □

CORPORATE REGISTER

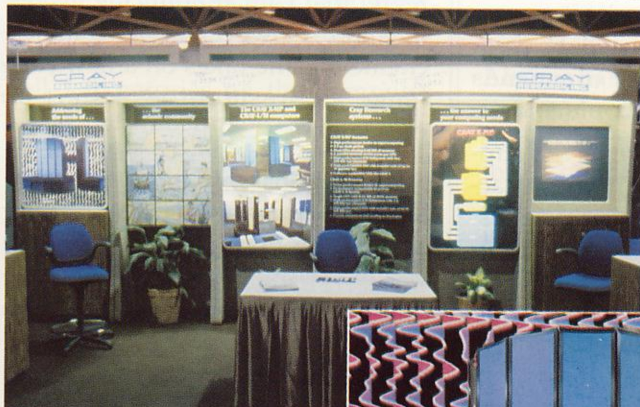

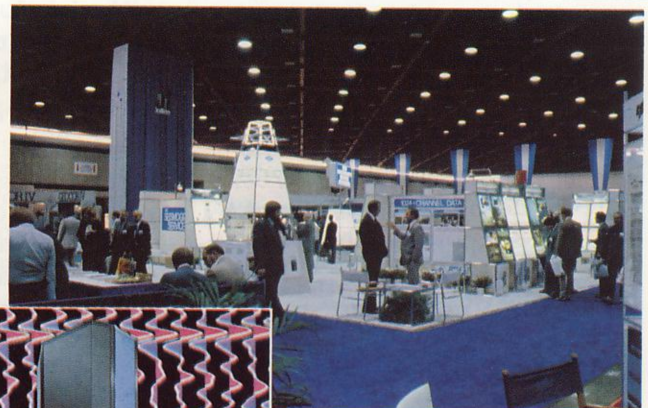

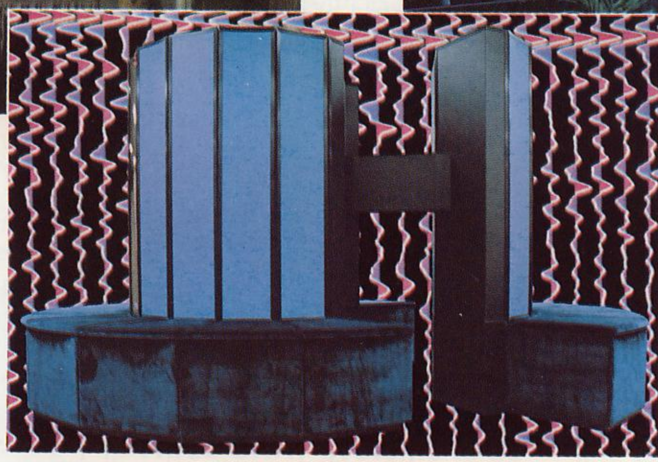

Cray Research exhibited at the 1982 SEG conference.

Cray attends SEG conference

The Cray Research exhibit received a warm reception at this year's Annual International Meeting of the Society of Exploration Geophysicists (SEG). Over 12,000 attendees of the conference held in Dallas in mid-October, made this year's meeting one of the most successful ever. Guided by the theme, "Geophysics—New Dimensions for the 80's", a packed program of "technical rejuvenation" sessions provided a broad sampling of many disciplines which integrate to form exploration geophysics today. Cray Research was glad to be a part of this dynamic 52nd Annual Meeting of the SEG.

London Museum exhibits CRAY technology

During 1982, the famed London Museum sponsored a special exhibition of computer technology as part of its Year of Information Technology celebration. Fittingly, among the technologies featured in the Computing and Data Processing collection at the Science Museum was the CRAY-1. Cray Research was asked to contribute to the exhibit because the original CRAY-1 design "broke the mold" in many ways. The company's contributions to the exhibition included samples of its unique CPU modules, memory modules and refrigeration cold bars. After the exhibition closed at the museum, the CRAY components were placed in the permanent museum archives. London Museum officials told a Cray Research representative, "The London Museum is pleased to be able to add parts of the CRAY-1 to the collection as to document the technology of the fastest processing available in the seventies."

VAX/VMS Station Software Service available

Now, users of VAX computer systems can retain the user-friendly qualities of their VMS environment yet have immediate access to the power and performance capability of the CRAY-1 and CRAY X-MP computer systems. Cray Research introduces the VAX/VMS Station Software Service, providing a new level of distributed processing capability to CRAY system users.

The VAX/VMS Station Software Service will be released in three versions. The initial version of the service, which provides facilities for CRAY job input and output and dataset movement between the VAX and CRAY systems, will be available for installation at customer sites by year-end. Versions 2 and 3, which will provide interactive facilities and DECNET support, will be available in the third quarter of 1983 and the first quarter of 1984, respectively.

The VAX/VMS Station Software Service enables Digital Equipment Corporation VAX systems running under control of Version 3 of the Virtual Memory System (VMS) Operating System to be linked with a CRAY computer system running under control of the Cray Operating System (COS), Version 1.11 or later. The CRAY/VAX Station can be configured with either the Network Systems Corp.'s HYPERCHANNEL hardware (Version 1) or a Cray front-end interface (Version 2 and beyond). Following are more details on the content of each version.

In Version 1, a Digital Control Language-compatible command allows submission of user job files to the CRAY. When running on the CRAY, a job can stage datasets to and from the VAX using the COS ACQUIRE and DISPOSE control statements. Datasets may be staged to VMS disks, magnetic tape, a printer, or a VMS batch job queue. The COS reference to a dataset to be transferred between the VAX and the CRAY system may specify character mode or transparent mode. For character mode, the Station supports VMS sequential and relative file organizations, and the records within the organization can be fixed or variable. For transparent mode, the Station uses the VMS block I/O file access method.

Version 2 will provide interactive facilities, permitting a CRAY program development environment and providing an applications interface to allow for distributed application. Also in Version 2, additional user commands and displays will be provided, including status commands and job manipulation commands.

Version 3 will include DECNET support, enabling DECNET users to access a subset of the Station facilities. Also in this version, new CRAY operator commands and displays will allow full control of the CRAY system by the VMS operator.

The VAX/VMS Station software is simple to install and adapts easily to meet site-specific requirements. VMS command procedures enable the system manager to generate the Station, and tools and utilities allow the manager to change the configuration parameters that control day-to-day operation of the Station. Further tailoring of the VAX/VMS Station Software Service to meet specific installation needs is simplified by the provision of installation exit points at strategic positions in the station code. The components of the Station are modular, enabling multiple users to access Station functions. Only standard facilities and interfaces within the VMS architecture are used.

Currently, Version 1 of the VAX/VMS Station Software Service is installed and available for demonstration at Cray Research's Mendota Heights facility. For more information about this exciting new service, contact your nearest Cray Research sales office.

Organizational changes announced

Late in September, the board of directors announced that Peter Appleton Jones, executive vice president of Cray Research, resigned to accept the presidency of APPLITEK, a new computer networking company in Boston. Chairman John A. Rollwagen said, "In the five years Peter Appleton Jones has been with us, he has made very significant contributions to our success. It is with real regret that I am announcing his resignation. However, I share his excitement as he takes on a new challenge as president of a start-up company."

Rollwagen has assumed Appleton Jones's essential marketing responsibilities for now. He stated that the vacated position will be filled eventually. "But in the meantime," he said, "I am looking forward to becoming more deeply involved once again in our marketing activities."

Rollwagen also announced a number of organizational changes to reflect the realignment of marketing responsibilities and to acknowledge the changing management relationships at Cray Research. Bruce Kasson, vice president of U.S. sales, and Michael Dickey, vice president of international marketing, will now report directly to Rollwagen. "I will be relying heavily on Bruce and Mike to provide continuing strong support for our customers and products." said Rollwagen.

Margaret A. Loftus, vice president of software development, will now report to executive vice president Lester T. Davis, thereby consolidating the company's diverse development efforts. Loftus has assumed responsibility for the Mendota Heights Data Center, which supports marketing and technical development efforts for the company.

In a separate move, John Carlson was promoted to executive vice president of the company. This move was made in recognition of Carlson's contributions to the company's success as senior vice president of finance and chief financial officer.

CRAY BLITZ exhibits skill in playing chess

This year at the annual chess tournament sponsored by the Association for Computing Machinery (ACM) held in Dallas late in October, CRAY BLITZ, the Cray computer chess program came closer than ever before to unseating BELLE of Bell Laboratories as the computer world chess champion. In non-elimination play, CRAY BLITZ and BELLE drew the match in the final round. Robert Hyatt, who has been entering CRAY BLITZ in tournament play since 1980, is very excited with this year's results. He said, "This year I worked with Harry Nelson of Lawrence Livermore National Laboratory—and it shows. By revectorizing the code, we were able to get a lot more speed out of the system. At the rate we're going, we should be able to take the title in 1983."

This chess board exemplifies possible positioning in this year's championship play.

Cray's 50th machine is delivered to a welcome home

On November 5 Cray Research, Inc. shipped the 50th CRAY-1 computer to the recipient of the first CRAY-1 in 1976. Los Alamos National Laboratory (LANL) will add this system to its existing computer network that currently includes four CRAY-1 computers. Robert Ewald, leader of the Los Alamos Computing Division said that the CRAY has become a standard for current high-performance computing carried out at Los Alamos. He went on to relate that, "Los Alamos has played a pioneering role in the development of electronic computation. Today, the laboratory's scientific computing facilities are among the most powerful in the world." Cray Research is proud to be an ongoing contributor to LANL's computing resources.

Lester Davis, Executive Vice President and one of the original founders of Cray Research, said, "The engineers, software staff and production workers take genuine pride in being the first to design, build and deliver this large number of computer systems. The systems we have built are performing to customer specifications and satisfaction as evidenced by repeat orders."

Three CRAY systems to be installed in Germany

During October 1982, Cray Research accepted contracts for the installation of three systems at different sites in West Germany. The systems will be installed in 1983 subject to export license approval. One of those systems, a CRAY-1 S/1000 will be installed at Deutsche Forschungs und Versuchanstalt Fuer Luft und Raumfahrt (DFVLR) located near Munich. DFVLR is a governmental center for aerospace and aircraft research and will use the CRAY-1 for combustion and aerodynamics design studies.

The University of Stuttgart will receive a CRAY-1 for use in nuclear safety, plasma physics and fluid dynamics research. A third system will be installed at Kernforschungsanlage (KFA) in Juelich, West Germany. The system will be a two-million word CRAY X-MP/22 and will be engaged in basic nuclear research, reactor safety and nuclear fusion for KFA.

Cray Research chairman John Rollwagen said, "We are delighted with the strength of our international business this year. At the beginning of 1982 there were nine CRAY-1 systems installed outside of the United States. We expect that by the end of 1983 the international installed base will have doubled."

CFT and $FTLIB 1.10 software released

Cray Research, Inc. released the CFT 1.10 FORTRAN Compiler and the accompanying $FTLIB Library on September 1. The software is compatible with the CRAY Operating System (COS) versions 1.10 and 1.11. Performance of CFT 1.10 is approximately equivalent to that of CFT 1.09.

Significant features in this release include:

- Compatability with the ANSI '78 FORTRAN Standard, including:
  - CHARACTER variable type,
  - CHARACTER operators and functions,
  - Generic intrinsic functions,
  - DIRECT ACCESS I/O,
  - List-directed I/O, and
  - New I/O control list options and FORMAT specifications.
- New compiler directives:
  - NEXTSCALAR,
  - BLOCK, and
  - NORECURRENCE.
- New compiler options:
  - OPT to optimize constant increment integers, and
  - AIDS to control listing of certain error messages.

Customers may order using the same processes as for all other standard Cray products.

APPLICATIONS IN DEPTH

DRC converted for use on CRAY

Design Rule Checker (DRC), a computer-aided design program for electronic circuits was recently converted for use on CRAY systems. The program identifies design rule violations in integrated circuits. The user inputs a circuit description and a test request file to specify the design rules to check and instructions to generate other useful mask layers. The program then verifies whether the design rules are met on the input layers and on the generated layers. DRC has been in use since 1975 and is generally accepted as an industry standard. For further information about DRC on CRAY systems, contact the NCA Corporation, 388 Oakmead Parkway, Sunnyvale, CA 94086, telephone (408) 245-7990.

Updated Directory of Software available

In November Cray Research released revision B of the Cray Research Directory of Software. This new release replaces revision A issued in May 1982. The new edition is expanded by about 30% as new applications software tools continue to become available on the CRAY from a variety of sources. Most of the software is available from third party sources while a few packages are distributed by Cray Research. To obtain a copy of the new Directory of Software, contact any Cray Research regional sales office or Cray Research Inc., Applications Department, 1440 Northland Drive, Mendota Heights, MN 55120, telephone (612) 452-6650.

ACSL converted for use on CRAY

Recently, conversion of ACSL to the CRAY was completed. ACSL is a Continuous System Simulation Language (CSSL) which follows the standards established by Simulation Council's CSSL Technical Committee in 1967. It is a language based on FORTRAN, designed to help the engineer or scientist mathematically model and analyze the behavior of a continuous system described by time dependent non-linear differential equations or transfer functions. Typical application areas are control system design, aerospace simulation, fluid flow and heat transfer analysis. For additional information about ACSL, contact Mitchell & Gauthier Associates, Inc., P.O. Box 685, Concord, MA 01742, tel. (617) 369-5115.

Graphics packages become available on the CRAY

Over the past few months a number of graphics packages have been converted for use on the CRAY. Below is a brief summary of a few of those packages and their features:

- QUIKRAY/QUIKSHOT—A 3-D interactive seismic modeling system that performs 3-D model input and editing, synthetic seismic trace generation, raypath, seismic section and time slice display. The package is a complete and independent tool for synthesis of seismic exploration and interpretation of geologic structures. For more information about QUIKRAY/QUIKSHOT, contact Sierra Geophysics, 15446 Bell-Red Road, Redmond, WA 98052, telephone (206) 881-8833.
- TEMPLATE—A high-level general-purpose computer graphics software support system designed for dynamic and static applications in 2- and 3-D environments. This product is available from MEGATEK Corporation, 3931 Sorrento Valley Road, San Diego, CA 92121, telephone (714) 455-5590.
- UNIRAS—A set of high-level raster graphic programs geared for the seismic, business and landsat applications. Detailed information about this comprehensive package can be obtained in the U.S.A. and Canada by contacting American Software Contractors, Inc., 48 Cummings Park, Woburn, MA 01801, telephone (617) 933-6102. In Europe, contact European Software Contractor ApS, 385 A, Lyngbyvej, DK-2820 Gentoffe, Denmark, tel. +45 1 65 74 12.

Finite element package available for use on the CRAY

STARDYNE, a finite element, static and dynamic structural analysis program which can handle up to 24,000 degrees of freedom for any single model has recently been converted for use on CRAY systems. It is user-oriented, containing automatic node and element generation features that reduce the effort required to generate input. Additional information about this package can be obtained by contacting System Development Corporation, Suite 3001, 3600 Sepulveda Boulevard, El Segundo, CA 90245, telephone (312) 615-1188.

USER NEWS

Mature convective storms like this one were the focus of CCOPE investigations.

CCOPE project at NCAR deemed a success

Imagine for a moment that you are in the high plains of Montana enjoying a beautiful blue-sky summer day. And then imagine that far on the horizon, bulging white clouds begin to congregate ominously. A short time later the sky is filled with a towering mass of thunderheads, piercing lightning, winds and torrential rains. Chances are that your peaceful afternoon has been ruined as you race for cover. However, for a group of 125 meteorologists stationed in the same neighborhood, this is just the kind of weather they want. With the bubbling clouds on the horizon, a group of pilots scrambles to their planes for a flight that most of us hope never to experience. Flying directly into the storm, they undertake their latest dogfight with the elements.

During the summer of 1981, the National Center for Atmospheric Research (NCAR—recipient of the second CRAY-1 computer) and the Bureau of Reclamation co-sponsored the Cooperative Convective Precipitation Experiment (CCOPE). It was the largest collaborative field program ever undertaken to study convective storms. CCOPE involved scientists and technicians from the United States, Canada, Italy, China, France and Great Britain. Participating U.S. government agencies included the National Aeronautics and Space Administration and the National Oceanic and Atmospheric Administration. The experiment was headquartered at Miles City, Montana.

The goal of CCOPE was to gather data that will help scientists predict, limit, and perhaps someday modify thunderstorm activity. Although almost 2,000 storms buffet the earth at any given time, the mechanisms that drive them remain very elusive. It is ironic these phenomena affecting social, economic and military events, still mystify cloud physicists. For instance, if rainfall could be controlled effectively by cloud-seeding, it has been estimated that the Dakotas could receive an additional 1-2 inches of rainfall during the growing season and increase annual crop yields by about $100 million.

In large part, the Montana-based project evolved out of the National Hail Research Experiment dating back to 1975. At that time, scientists hoped to learn enough to use weather modification technology to reduce the destructive effects of hailstorms. While they obtained a detailed record of the hailstorm process, they found that their efforts to control the storm had no measurable effect. The scientists concluded that they needed a clearer understanding of the actual thunderstorm process before they could be controlled. Hence, CCOPE was born.

As storms evolved during that summer of 1981, up to eight of the 14 available heavily instrumented weather research aircraft would take off and head into the storm. A Schweizer 2-32 Sailplane already in the air would be detached from its tow plane. It would attempt to enter the cloud and catch its updraft. Drifting into the growing storm, the Schweizer 2-32 would measure everything from liquid water content, to vertical air motion, precipitation-particle-size spectra and electric fields. If the storm continued to mature, a twin jet Sabreliner would fly into and probe the anvil of windblown ice crystals at the top of the storm. Likewise, the powerful T-28 aircraft, armored with quarter inch aluminum paneling and a hail-resistant windshield, would await instructions to fly directly into the heart of the storm against driving hail to record activity of its violent phase. The T-28 was the only aircraft able to penetrate the core. It would often encounter violent updrafts and return heavily iced. NASA's Convair 990 research aircraft would wait in the wings for orders to fly through the storm's gust front, the area of the high speed outwash of potentially destructive winds generated as the storm's downdraft across the ground. Leaving no portion of the storm unobserved, the aircraft would fly around, through and under the storm system on pre-assigned routes while radar monitored the storm activity.

Since CCOPE was completed in August of 1981, the collected data has undergone extensive processing. Carl Mohr, an NCAR computer scientist who participated in the experiment, said that the CRAY-1 has been heavily involved in processing the data collected by the aircraft, radar and ground instruments. He explained that, "Our first order of business was to process the data into common formats. The CRAY has been plotting aircraft data as a function of time in addition to putting radar data such as Doppler velocities and reflectivity through a Cartesian interpolation process. This involves extensive computation on a couple million pieces of data at a time. Analysis of this information that takes 20 seconds on the CRAY would take minutes on our other computers." While much of the data has been readied for analysis, the long process of synthesizing and generalizing the data—all the real learning—will not be completed for a number of years. Scientists are excited with the results of the Montana project because it provides information that allows them for the first time to check theories out against actual data.

CRAY-1 earns its keep at the University of Minnesota

The University of Minnesota, one of the country's leading academic institutions committed to research, installed a CRAY-1 system late in 1981. Increasing its compute power by an order of magnitude, the University sees the CRAY-1 as a way to meet the computing demands of research in the 1980's. This system is the first CRAY placed in an academic environment. Other universities that have since ordered and/or installed CRAY computers, are the University of London and the University of Stuttgart.

The University of Minnesota's decision to install the CRAY is part of a multi-tiered approach to computing wherein the CRAY is now at the apex of the institution's computing resources. The University Computer Center is committed to providing cost/effective computing at each level of technological sophistication, ranging from the microcomputer level up through supercomputers. Peter C. Patton, Director of the University Computer Center explained, "In years past, each individual academic department was encouraged to have its own computing resources. That is changing here at the University of Minnesota. We are committed to maintaining a dedicated research center that offers superior computer resources at a lower cost—and to a diverse group of users."

Cost effectiveness was a key consideration in the decision to install the CRAY because the University Computer Center operates on a cost recovery basis. "Because of the CRAY's ability to execute complex problems very quickly, jobs can be run in a fraction of the time it would take on any of our other systems," Dr. Patton said. "The result is that users end up paying less for the same amount of computation because the CRAY handles it so well." He went on to say that, "In academia this is very important. Researchers competing for scarce resource dollars from the granting institutions need to show the efficiency with which the funds are being spent. The CRAY gives us an edge in helping those people capture that money."

Users within the university include those involved in aeronautical, biomedical, geological and soil sciences research. The Chemical Engineering and Materials Science department accounts for the largest ongoing use of the CRAY-1 within the university. The many external users are as diverse as the Minnesota Pollution Control Agency, the Federal Reserve Bank of Minneapolis, U.S. Geological Survey and the Minnesota Zoo. Most of these groups were able to transfer programs running on other equipment to the CRAY when it became available. Peter Patton says, "We wanted a system that could be used by the many different types of current and future users we expect to have. A major reason that we selected the CRAY was because of its relative accessibility to different groups. Minimal effort has been expended by many of our users in converting programs to the CRAY—and yet we are seeing runtimes improve by a factor of 20 or more on the average. Its uptime of over 98% also helps that situation. The CRAY is a good, cost effective, general purpose machine that will fulfill our needs through the 1980's."

CORPORATE ADDRESSES

Regional Sales Offices

Central Regional Sales Office
5330 Manhattan Circle, Suite F
Boulder, CO 80303
Tel: (303) 499-3055

Eastern Regional Sales Office
11120 New Hampshire Avenue
Silver Spring, MD 20904
Tel: (301) 681-9514

Petroleum Regional Sales Office
5858 Westheimer, Suite 500
Houston, TX 77057
Tel: (713) 975-8998

Western Regional Sales Office
Sunset Office Plaza
1874 Holmes Street
Livermore, CA 94550
Tel: (415) 447-0201

Sales Offices

Albuquerque, NM

Atlanta, GA

Boston, MA

Chicago, IL

Dallas, TX

Laurel, MD

Los Angeles, CA

Mountain View, CA

Pittsburgh, PA

Seattle, WA

International Subsidiaries

Cray Research France, S.A.
21 Avenue Victor Hugo
75016 Paris
France
Tel: 01-501-5270

Cray Research (UK) Limited
Seymour House, The Courtyard
Denmark Street, Wokingham
Berkshire, England
Tel: 0734-791180
TLX: 848841

Cray Research GmbH
Wartburgplatz 7
8000 Munich 40
West Germany
Tel: 089-363076
TLX: 8415213211

Cray Research Japan, Limited
New Aoyama Bldg, West 1661
1-1 Minami-Aoyama 1-chome
Minato-Ku, Tokyo 107
Japan
Tel: 03-403-3471
TLX: T8128893

CRAY RESEARCH, INC.

Corporate Headquarters
Cray Research, Inc.
608 Second Avenue South
Minneapolis, Minnesota 55402
Tel: (612) 333-5889