

2

~~UNCLASSIFIED~~  
SECURITY CLASSIFICATION OF THIS PAGE

DOCUMENTATION PAGE

Form Approved  
OMB No. 0704-0188

AD-A204 321

1b. RESTRICTIVE MARKINGS  
DTIC FILE COPY

3. DISTRIBUTION / AVAILABILITY OF REPORT  
Approved for public release;  
distribution unlimited

2b. DECLASSIFICATION / DOWNGRADING

4. PERFORMING ORGANIZATION REPORT NUMBER(S)

5. MONITORING ORGANIZATION REPORT NUMBER(S)  
AFOSR-TR-89-0011

6a. NAME OF PERFORMING ORGANIZATION  
University of Michigan  
Dept. of EE & CS

6b. OFFICE SYMBOL  
(if applicable)

7a. NAME OF MONITORING ORGANIZATION  
AFOSR/NM

6c. ADDRESS (City, State, and ZIP Code)  
Ann Arbor, Michigan

7b. ADDRESS (City, State, and ZIP Code)  
Building 410  
Bolling, AFB DC 20332-6448

8a. NAME OF FUNDING / SPONSORING ORGANIZATION  
AFOSR

8b. OFFICE SYMBOL  
(if applicable)  
NM

9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER  
AFOSR-84-0096

8c. ADDRESS (City, State, and ZIP Code)  
Building 410  
Bolling, AFB DC 20332-6448

10. SOURCE OF FUNDING NUMBERS  
PROGRAM ELEMENT NO. 61102F  
PROJECT NO. 2304  
TASK NO. AR  
WORK UNIT ACCESSION NO.

11. TITLE (Include Security Classification)  
Sparse Elimination on Vector Multiprocessors

12. PERSONAL AUTHOR(S)  
D.A. Calahan

13a. TYPE OF REPORT  
FINAL

13b. TIME COVERED  
FROM 5/1/84 TO 4/30/88

14. DATE OF REPORT (Year, Month, Day)  
2 May 88

15. PAGE COUNT  
9

15. SUPPLEMENTARY NOTATION

17. COSATI CODES		
FIELD	GROUP	SUB-GROUP

18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)

19. ABSTRACT (Continue on reverse if necessary and identify by block number)  
The research of this grant spanned a number of topical areas in its four years duration  
(1) Blocked parallel solution of dense and sparse systems. Closely-related to the original proposal, this research involved a study of the relationship between task granularity and block partitioning size in the solution of linear algebra problems. The rationale for this blocking was the restricted effective memory bandwidth of the shared-memory CRAY-2 due to memory conflicts.  
→ page 2

DTIC  
SELECTED  
S 13 FEB 1989 D  
E

20. DISTRIBUTION / AVAILABILITY OF ABSTRACT  
 UNCLASSIFIED/UNLIMITED  SAME AS RPT.  DTIC USERS

21. ABSTRACT SECURITY CLASSIFICATION  
UNCLASSIFIED

22a. NAME OF RESPONSIBLE INDIVIDUAL  
DR. ABRAHAM WAKSMAN

22b. TELEPHONE (Include Area Code)  
(202) 767-5026

22c. OFFICE SYMBOL  
NM

AFOSR-TR- 89-0011

Final Technical Report  
for Period 5-1-84 to 4-30-88  
Grant ~~AF~~-AFOSR-84-0096

**SPARSE ELIMINATION ON VECTOR MULTIPROCESSORS**

D. A. Calahan  
Principal Investigator  
Department of Electrical Engineering & Computer Science  
University of Michigan  
Ann Arbor, MI

May 2, 1988



Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By _____	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

89 2 10 108

## I. RESEARCH SUMMARY

The research of this grant spanned a number of topical areas in its four years duration.

(1) Blocked parallel solution of dense and sparse systems. Closely-related to the original proposal, this research involved a study of the relationship between task granularity and block partitioning size in the solution of linear algebra problems. The rationale for this blocking was the restricted effective memory bandwidth of the shared-memory CRAY-2 due to memory conflicts. This prompted the

(2) development of conflict-resistant algorithms, i.e., a preliminary study of properties of algorithms which were insensitive to memory conflicts. This study led to an extended study of

(3) memory conflict modeling of shared-memory multiprocessors. The final result was development of unique "black-box" models of the CRAY-2 memory system based on dedicated machine measurements. In the realization that the limited parallelism of the CRAY-2 was restrictive for future algorithm studies, a new effort precursing future research cooperative with WPAFB personnel was initiated. (on → JES)

(4) distributed-memory (massively-parallel) solution of CFD problems. This effort is being continued under a new AFOSR grant.

The remainder of the report is chronological and taken largely from previous interim progress reports.

## II. 1984-85 PROGRESS REPORT

Access to (proprietary) preliminary design information for the CRAY-2, together with early access to the MFECC and NAS CRAY-2's, indicated a number of areas for related future research. These included the following.

(a) Parallelization at the vector instruction level. This research in small-grain parallelization introduced the concept of "microtasking" [1][5][10], a term later adopted by Cray Research and applied to general small-grain task control from Fortran.

(b) Design of conflict-resistant algorithms. It became clear from this brief study that by blocking the solution of simultaneous linear equations, the performance degradation associated with memory conflicts present in the CRAY X-MP and dominant in the CRAY-2 could be mitigated [11][12].

More background on these topics is included in the next section.

### III. 1985-86 PROGRESS REPORT

#### A. Introduction

The above studies continued on the following topics.

#### B. Equation-solving on the CRAY-2

The CRAY-2 combines the algorithmic requirements of (a) vector inner-loop syntax and high-level taskability, similar to the X-MP, and (b) partitioning to exploit a local memory (LM) that is necessitated by a massive but slow common memory (CM). Although this appears to be the most demanding of any supercomputer architecture, in fact all of these can be satisfied in equation solution by block partitioning of the matrix.

In [3] and [16], the performances of several uniprocessor CRAY-2 implementations were compared. The above block partitioning based on a matrix-matrix multiply (M\*M) kernel was shown to produce a speedup greater than 6:1 over conventionally-coded Gauss elimination, and nearly 2:1 over the best previous algorithms [Dongarra] based on matrix-vector multiplies. To date, this was the most dramatic evidence on a supercomputer architecture of the value of basing linear algebra algorithms on an M\*M kernel (also termed a third-level BLAS). A heretofore problem with block partitioning when pivoting is required was solved by using a two-level algorithm, which preserved the asymptotic performance of the M\*M kernel at the higher level and permitted pivot searches and row exchanges at the lower level.

The same block partitioning permitted parallel solution when multitasking became available on the CRAY-2 in May-June 1986.

#### C. Task Granularity Studies

A simulation study of up to a 16-processor X-MP was completed [1]. The effect of choosing different task sizes in solving a full set of equations was studied. This was the largest number of processors studied until the availability of the (slow, non-vector) message-passing machines.

#### D. Memory Conflict Studies

The memory conflict problem is the result of attempting to service a multiprocessor architecture with a critical uniprocessor resource - main or common memory. To the extent that conflicts can be avoided, a major rationale for having local memory - and the algorithmic complexity it introduces - can be avoided. A modest conflict problem was observed for the X-MP in the range of 5-10% performance degradation; with early Cray-2 memory systems, this degradation approached 60%.

The common memory conflict problem is largely unappreciated in the academic community; thus, it is possible to make a contribution without an extensive architectural design background. It had previously been shown that a critical memory organization feature of the CRAY X-MP-2 permitted worst case steady-state delays of 25% in vector accesses of CM. Based on this and similar internal studies, Cray Research re-designed the memory system of the X-MP-4. Unfortunately, the new design aggravated the *startup* access delay; this feature was documented by this author in [3], which has led to a new design for X-MP memories.

#### IV. 1986-87 PROGRESS REPORT

##### A. Introduction

Besides continuation on the above topics, the procurement of a small hypercube under grant auspices permitted research on distributed-memory algorithms.

##### B. Equation-solving on the CRAY-2

The CRAY-2 combines the algorithmic requirements of (a) vector inner-loop syntax and high-level taskability, similar to the X-MP, and (b) partitioning to exploit a local memory (LM) that is necessitated by a massive but slow common memory (CM). Although this appears to be the most demanding of any CRAY-class supercomputer architecture, in fact all of these can be satisfied in equation solution by block partitioning of the matrix.

In [7] the performances of several uniprocessor CRAY-2 implementations were compared.

It was expected that the same block partitioning would be used for parallel solution when multitasking became available on the CRAY-2 at NAS in the summer of 1986. Unfortunately, the current version of the UNI-COS operating system at NAS has a latency (bug) beyond normal task startup that will not consistently permit concurrent operation with tasks less than 1 million clocks. This algorithm effort was put on hold, and the associated grant resources were directed to procurement of a \$19400 hypercube prototyping system in a proposal revision of December, 1986.

##### C. Memory Conflict Studies

In a jointly-sponsored effort with the Research Institute for Advanced Computer Science (RIACS) using the NAS CRAY-2, experimental studies have shown an extraordinary performance degradation in the presence of moderate scalar activity in a shared memory. This is interpreted as a dependence on both the amount and the regularity of memory accesses. Current efforts involve development of empirical mathematical models of this

degradation in algorithmic terms, such as the rate and the average vector length of accesses of users sharing the memory. Ultimately, this could evolve a standard test and an associated design figure of merit for shared-memory systems.

## V. 1987-88 PROGRESS REPORT

### A. Introduction

In this final year of the grant,

(a) research was initiated on distributed-memory algorithms involving both linear algebra and CFD problems, and

(b) conflict modeling was accelerated due to regular access to dedicated CRAY-2s at the Air Force Weapons Laboratory and NASA Ames Research Center.

### B. Distributed-memory CFD algorithms

This study, together with a small startup contract with Dr. Joseph Shang at AFFDL (WPAFB), began the effort of examining partitioning strategies for parallel implementation of explicit and implicit aerodynamic CFD algorithm. Although no publication was forthcoming, the groundwork was laid for a new AFOSR research effort to begin in May, 1988. Included in this study will be an emphasis on realistic (non-ideal) problem structures, where the designer and an automatic partitioning algorithm must work cooperatively to balance computational workload across a massively-parallel architecture.

### C. Two-parameter memory conflict modeling

Previous anecdotal studies on memory-load-related degradations associated with the CRAY-2 [4] were incorporated in a two-parameter "black-box" model, based on extensive dedicated-time measurements on the NAS CRAY-2s at NASA Ames Research Center. Among other features, this model quantifies the above-mentioned severe effects of scalars on memory performance [8][9].

VI. COUPLING ACTIVITIES (with national laboratories)

A. Los Alamos National Laboratory; consultant.

FY 84-85: Study of conflict problems on a many-processor X-MP.

FY 85-86: Study of linear algebra algorithms on INTEL hypercube.

B. NASA Ames Research Center.

FY 83-86: University research on conversion of Computational Chemistry codes to the CRAY-2.

FY 85-88: Consultant to RIACS, developing math library software for the CRAY-2.

C. Air Force Flight Dynamics Laboratory.

FY 84-85: Visiting scientist, studying vectorization of structural transient optimization codes and vector multiprocessor tasking of structural analysis algorithms.

FY 87: University research, studying massively-parallel algorithms for CFD.

D. Lawrence Livermore National Laboratory (MFECC); consultant

FY 85: Study of conflict problems on the CRAY-2.

E. San Diego Supercomputer Center; consultant

FY 88: Study of graphics-based algorithms for parallel partitioning of whole-body aerodynamic simulations.

## VII. GRANT-SUPPORTED PUBLICATIONS & ACTIVITIES

### Journal Publications

- [1] Calahan, D. A., "Task Granularity Studies on a Many-processor CRAY X-MP," *Parallel Computing*, vol. 2, October, 1985, pp109-118.
- [2] Arya, Siamak, "An Optimal Instruction-Scheduling Model for a Class of Vector Processors," *IEEETC*, vol. C-34, no. 11, November, 1985, pp981-995.
- [3] Calahan, D. A., "An Analysis of Vector Startup Access Delays," *IEEE Transactions on Computers* (accepted as correspondence)

### Edited Books

- [4] Calahan, D. A., and D. H. Bailey, "Measurement & Analysis of Memory Conflicts on Vector Multiprocessors," chapter in *Performance Analysis of Supercomputers*, Ed. by J. L. Martin, Elsevier Science Publishers B.V. (North-Holland), 1988.

### Conference Proceedings

- [5] Calahan, D. A., "Influence of Task Granularity on Vector Multiprocessor Performance," *Proc. 1985 Intl. Conf. on Parallel Processing*, Bellaire, MI, August, 1985, pp. 278-84.
- [6] Calahan, D. A., "An Analysis and Simulation of the CRAY X-MP Memory System," *Proc. First Intl. Conf. on Supercomputing Systems*, Tarpon Springs, FL, December, 1985, pp568-574.
- [7] Calahan, D. A., "Block-Oriented Local-Memory-Based Linear Equation Solution on the CRAY-2: Uniprocessor Algorithms," *1986 International Conference on Parallel Processing*, St. Charles Ill., August 20, 1986, pp 375-378.
- [8] Calahan, D. A., "Characterization of Memory Conflict Loading on the CRAY-2," to appear in *Proc. 1988 Intl. Conf. on Parallel Processing*, St. Charles, Ill., August, 1988.
- [9] Calahan, D. A., "Performance Evaluation of Static and Dynamic Memory Systems on the CRAY-2," to appear in *Proc. 1988 ACM Conference on Supercomputing*, St. Malo, France, July, 1988.

## Reports

[10] Calahan, D. A. , "Tasking Studies in Solving a Linear Algebra Problem on a CRAY-class Multiprocessor," Report SARL #5, Department of Electrical Engineering and Computer Science, University of Michigan, December, 1983.

[11] Calahan, D. A., "Conflict Sensitivity of Algorithms," Report SARL #7, Dept. of Electrical Engineering and Computer Science, University of Michigan, April 15, 1985.

[12] Calahan, D. A., "Memory Conflict Simulation of a Many-Processor CRAY Architecture. Part I: A CRAY X-MP Study," Report SARL #6, Dept. of Electrical Engineering and Computer Science, University of Michigan, April 15, 1985.

[13] Calahan, D. A., "Block-oriented, Local-memory Linear Equation Solution on the CRAY-2. Part I: Uniprocessor Algorithms," Report SARL#9, Department of Electrical Engineering & Computer Science, University of Michigan, December 15, 1985.

## Presentations to Air Force Laboratories

[14] Calahan, D. A., "CRAY-2 Performance Evaluation," seminar at AFWL, February 19, 1987.

[15] Calahan, D. A., "Evaluation of the CRAY-2 Memory System," seminar at AFWL, April 27, 1988.

## Workshops and Other Presentations

[16] Calahan, D. A., "Equation-solving on the CRAY-2," Proc. 2nd Intl. Conf. on Scientific and Parallel Computing, Norfolk, November, 1985.

[17] Calahan, D. A., "Performance Evaluation of the CRAY-2," 2nd Air Force Supercomputer Conference, Eglin AFB, November 4, 1986.

[18] Calahan, D. A., "CRAY-2 Performance Evaluation," Workshop on Performance Evaluation of Parallel Computers and Programs, Argonne National Laboratory, November 18, 1986.

## Related Presentations (not grant-sponsored)

[19] Calahan, D. A., "CRAY-2 Performance Evaluation," seminar at San Diego Supercomputer Center, La Jolla, February 24, 1987.

[20] Calahan, D. A., "Evaluation of the CRAY-2 Memory System," seminar at Los Alamos National Laboratory. April 28, 1988.

[21] Calahan, D. A., Lecturer, NSF-sponsored Institute on Supercomputers at San Diego Supercomputer Center, August, 1986, and August, 1987.