



How Moderate-Sized RISC-Based SMPs Can Outperform Much Larger Distributed Memory MPPs

by D. M. Pressel, Walter B. Sturek,
J. Sahu, and K. R. Heavey

ARL-TR-2062

October 1999

Approved for public release; distribution is unlimited.

DTIC QUALITY INSPECTED 4

19991029 093

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.

Army Research Laboratory

Aberdeen Proving Ground, MD 21005-5067

ARL-TR-2062

October 1999

How Moderate-Sized RISC-Based SMPs Can Outperform Much Larger Distributed Memory MPPs

D. M. Pressel, Walter B. Sturek, J. Sahu, and K. R. Heavey
Corporate Information and Computing Directorate, ARL

Approved for public release; distribution is unlimited.

Abstract

Historically, comparisons between computer systems were based primarily on theoretical peak performance. Today, comparisons based on delivered levels of performance are frequently used. This, of course, raises a whole host of questions concerning methodology. From the standpoint of the user, delivered performance frequently refers to how fast a job runs. However, is it reasonable to base this measurement on running the same algorithm on all of the computers? When comparing some combination of mainframes and vector supercomputers, the answer is probably *yes*. The same holds true when comparing the performance of large distributed memory **MIMD MPPs**. However, when comparing the algorithms of choice used on these two classes of platforms, one frequently finds that the algorithms are quite different. Furthermore, the amount of work (the number of **FLOPS**) associated with each algorithm can also be quite different.

While troubling, this dichotomy has been largely unavoidable. This implies that for an MPP to have the same level of delivered performance as the mainframes and the vector supercomputers, it must have a significantly greater level of performance when measured in terms of **FLOPS**. Recent advances involving moderate-sized **RISC-based SMPs** have allowed us to solve this problem. The net result is that for some problems a 128 processor Origin 2000 can deliver levels of performance that might require the use of a 500+ processor MPP using more traditional approaches.

Acknowledgments

This work was made possible through a grant of computer time by the Department of Defense (DOD) High Performance Computing Modernization Program. Additionally, it was funded as part of the Common High Performance Computing Software Support Initiative administered by the DOD High Performance Computing Modernization Program.

INTENTIONALLY LEFT BLANK.

Table of Contents

	<u>Page</u>
Acknowledgments	iii
List of Figures	vii
List of Tables	vii
1. Introduction	1
2. Delivered Performance	4
3. Loop-Level Parallelism	6
4. Speedup	8
5. Conclusions	14
6. References	15
Glossary	17
Distribution List	19
Report Documentation Page	23

INTENTIONALLY LEFT BLANK.

List of Figures

<u>Figure</u>		<u>Page</u>
1.	Predicted Speedup From the Parallelization of a Problem With a Fixed Problem Size.....	7
2.	Predicted Speedup for a Loop With 15 Units of Parallelism	12
3.	Performance Results for the One-Million-Grid-Point Data Set	13
4.	Performance Results for the 59-Million-Grid-Point Data Set.....	14

List of Tables

<u>Table</u>		<u>Page</u>
1.	The Number of Processors Required to Achieve a Specified Level of Delivered Performance Using Traditional Techniques.....	9
2.	The Speedup Achieved When Using the F3D Code Parallelized Using Loop-Level Parallelism on SGI Origin 2000's.....	9
3.	The Number of Processors That One Would Normally Use When Using an MPP and Traditional Techniques to Process the Test Cases.....	11
4.	Predicted Speedup for a Loop With 15 Units of Parallelism	12

INTENTIONALLY LEFT BLANK.

1. Introduction

For a given job, one can define the *Delivered Performance* such that

$$\textit{Delivered Performance} = \textit{Theoretical Peak Performance} * \textit{Total Efficiency},$$

where

$$\textit{Total Efficiency} = \textit{Algorithmic Efficiency} * \textit{Serial Efficiency} * \textit{Parallel Efficiency}.$$

Traditionally, many researchers using parallel computers have ignored the question of Algorithmic Efficiency and/or Serial Efficiency, preferring to stress Parallel Efficiency. A few people have even gone so far as to assume that all jobs on all machines have similar levels of efficiency, and therefore, all one needs to know is the Theoretical Peak Performance for the machines in question.

A direct consequence of this attitude has been the reaction by many users of vector computers who point out that while the parallel computers may be delivering higher levels of **FLOPS**, the vector computers will frequently have better wall clock time. Even when one takes into account the relative costs of the machines, an important consideration in throughput oriented environments, the vector machines will frequently fare much better than the raw numbers might indicate. Based on these observations, it is clear that one needs to consider Algorithmic Efficiency and Serial Efficiency as well as Parallel Efficiency when evaluating projects that use parallel computers.

When one compares the architectures of Cray vector computers (e.g., the C90), traditional **MPPs** (e.g., the Cray T3E or the IBM SP), and **RISC***-based **SMPs** (e.g., the SGI Origin 2000 or the SUN HPC 10000), one finds significant differences in the design principles on which these systems are based. The Cray vector computers have vector processors, a low latency very high bandwidth

* Definitions for boldface text can be found in the glossary.

memory system, and make very few assumptions about data locality or data reuse. In fact, the very nature of their design tends to discourage attempts to tune for data locality or reuse.

Traditional MPPs have an intermediate level of memory latency and bandwidth. Most of these systems now use RISC processors with moderate-sized caches and some design features that facilitate streaming data into and out of the processor. Experiments reported by the NAS group at NASA Ames and an analysis performed by David O'Neal and John Urbanic at the Pittsburgh Supercomputing Center indicate that the memory system limitations on these systems result in a lower level of efficiency than with comparable codes running on the Cray vector machines (Saini 1996, 1997; O'Neal and Urbanic 1997).

RISC-based SMPs tend to have longer memory latencies and somewhat lower memory bandwidth than the MPPs. In general, they also have no special features designed to facilitate the streaming of data into and out of the processor. On the other hand, they are usually equipped with at least 1 MB of cache per processor. Codes that have been tuned to take advantage of this cache can in many cases reduce the rate of cache misses, that miss all the way back to main memory, to less than 1%. As a result, for some codes, it is possible to achieve serial levels of performance that actually exceed the performance achieved on MPPs (both in terms of absolute single processor performance and in terms of the percentage of peak). However, the serial performance for codes that were only tuned to run on an MPP may fall short of what is normally seen on an MPP (Sahu et al. 1997, 1998).

From this discussion, it should be clear that different classes of machines are likely to deliver different percentages of peak performance. Furthermore, the delivered level of performance is likely to strongly depend on the quality of the tuning (this includes the vector machines, since the production of good vectorizable code requires the mastery of multiple complementary techniques). Finally, the ability to deliver well-tuned code will frequently depend on the design of the hardware itself.

For the rest of the discussion, it will be assumed that the codes are well tuned for serial performance and that the serial performance achievable with the MIPS R10K, HP PA-8000, Alpha 21164 (as configured for the T3E), and the IBM P2SC are all comparable. This means that in terms of efficiency, the MIPS R10K and the IBM P2SC have a significant lead over the other two chips (this agrees with published results as well as information that the author has received in private briefings). It is also assumed that the achievable serial efficiency of the MIPS R10K and the IBM P2SC approaches and, in some cases, matches that seen on Cray vector machines.*

The remainder of this report will focus on parallel efficiency and algorithmic efficiency, with most of the emphasis being on the latter. The assumption here is that in many cases one can produce parallel efficiencies that are close to 100%.†

This leaves the question, What is the algorithmic efficiency? While it is hard to identify what one means by an algorithmic efficiency of 100%, it is generally easy to define the relative algorithmic efficiency of two approaches by comparing the number of floating point operations required to achieve a solution.‡

The important point here is to recognize that many parallelized programs have a significantly lower algorithmic efficiency than do the programs normally run on serial/vector processors. Evidence will be given that it is, in some cases, now possible to avoid the need to use less efficient algorithms. The downside is that the approaches that lead to this conclusion are, in many cases, not highly scalable. In some cases, they may work on a range of SMPs and MPPs, while in other cases, they will only work on SMPs. However, even with these limitations, it is possible that these

* To a first approximation, the preceding assumptions seem to be nearly independent of the processor speeds in question. This probably means that each of these designs is strongly limited by the performance of the memory system and/or other system components.

† In situations where this is not likely to be the case, some discussion will be made of what the predicted behavior is.

‡ Ideally, this comparison should be made on a single machine, even if the different algorithms are not normally all used on the same machine. This way, one is measuring differences in the algorithms, not in the compilers. However, even here, it is important that if operations are normally performed N times, where N is the number of processors, then if one is comparing using three different algorithms, one to be run on a vector machine with 16 processors, one on an SMP with 100 processors, and one on an MPP with 500 processors, then the operation counts should reflect the intended usage.

approaches will allow the performance of a 64-128 processor job to equal or exceed the performance of a job using a 512 processor MPP written using traditional techniques.

2. Delivered Performance

If one asks most users of computer equipment what job characteristics they consider to be important, the replies seem to center on two themes:

- (1) The accuracy of the results.*
- (2) The time to completion.†

While, in general, one assumes that these two themes are independent, in reality, that has rarely been the case. In a resource-constrained environment, it is easy to see how the sophistication of the calculations that one can reasonably hope to carry out will be limited. What may be less obvious is that there are other ways in which the search for a rapid time to completion can adversely affect the accuracy of the results.

A common problem, which comes from the field of CFD, is that, frequently, the most efficient serial or vectorizable algorithm uses what is known as an Implicit approach to solving the Navier-Stokes equations. Unfortunately, such an approach has generally resisted attempts to parallelize it. An example of this problem can be found in the F3D code from NASA Ames Research Center, which uses a *block tridiagonal solver* based on Gaussian Elimination. Due to dependencies in this part of the code, even though this solver handles three-dimensional problems, each of the sweeps through the data can only be parallelized in a single direction. In the past, three methods have been used to get around this problem:

* In general, this is a relative concept and refers to the accuracy conforming to the expectations for a particular run.

† This may either apply to the time required to run a single job or the time required to run a complete set of jobs. In the former case, one is usually limited to using some or all of a single computer. In the latter case, one is usually limited to using some or all of the resources at one or a small number of computer centers.

- (1) In theory, one can replace Gaussian Elimination with a more parallelizable algorithm such as cyclic reduction. Unfortunately, this approach can itself result in two major problems:
- It increases the operation count by a factor of $\text{LOG}_2(N)$, where N is the number of processors being used. Clearly, this has the effect of decreasing the *algorithmic efficiency*.*
 - Since this algorithm requires the use of a large number of relatively small messages, it was much better suited for use on **SIMD** (Flynn 1972) machines than for use with today's **MIMD** (Flynn 1972) machines.
- (2) One can use an entirely different algorithm, such as one of the Explicit algorithms, which are known to be highly parallelizable. Unfortunately, the use of these algorithms will, in general, substantially increase the operation count required to obtain a solution. In other words, once again, the *algorithmic efficiency* suffers.
- (3) Alternatively, one can use **domain decomposition** as the basis for parallelization. Unfortunately, this approach can severely compromise the convergence behavior of the algorithm. A number of approaches have been suggested to deal with this problem, but in all cases, the *algorithmic efficiency* will to some degree suffer[†] (Wang and Tafti 1997; Singh, Uthup, and Ravishanker, 1998[?]).

The key point here is that:

$$\text{Delivered Performance} = \text{Theoretical Peak Performance} * \text{Total Efficiency},$$

where

$$\text{Total Efficiency} = \text{Algorithmic Efficiency} * \text{Serial Efficiency} * \text{Parallel Efficiency}.$$

* Using 512 processors will increase the operation count for this part of the solver by a factor of 9.

† Many of these approaches will also limit the available parallelism and/or adversely affect the parallel efficiency.

Therefore, any changes that result in a decrease in the *Algorithmic Efficiency* will directly affect the *Delivered Performance*, even though the performance as measured by MFLOPS might be quite high.* Using this unit, *Delivered Performance* is inversely proportional to the *Time to Completion* for a job (assuming that the processors are dedicated to this job).

Figure 1 shows an example of this for a fixed-sized problem when one attempts to scale to large numbers of processors. Two things that are important to note here are:

- (1) For sufficiently large numbers of processors, the combined effects of Amdahl's Law and the costs of interprocessor communication will limit the maximum achievable level of performance. Therefore, for all but the largest problem sizes, and given enough processors, the parallel efficiency may be far less than 100%.
- (2) The effect of going to less efficient algorithms in an attempt to improve the parallelizability of the code can virtually eliminate the perceived benefits of having a highly parallelizable code.

If one applies Gustafson's (1988) concept of *scaled speedup*, one can overcome some if not all of the limiting effects attributable to Amdahl's Law and interprocessor communication. However, this concept will have little impact on the loss of algorithmic efficiency. Therefore, the basic premise behind Figure 1 (and this report) remains intact.

3. Loop-Level Parallelism

It turns out that there is an alternative way in which one can parallelize Implicit CFD codes, which does not result in a reduction of their *Algorithmic Efficiency*. This approach is based on parallelizing the individual loops and is therefore referred to as Loop-Level Parallelism. Of course,

* The units for *Delivered Performance* are Useful MFLOPS.




Figure 1. Predicted Speedup From the Parallelization of a Problem With a Fixed Problem Size.

if this method is so great, then one might wonder why it was not the method of choice all along. The following are some of the reasons for this.

Loop-Level Parallelism in general is based on the same parallelism used to produce vector code. Therefore if the program is to run in parallel on a vector computer such as the Cray C90, it will be difficult to produce a code that exhibits both good vector performance and good parallel performance at the same time.

While in theory it is possible to implement Loop-Level Parallelism using some form of message-passing code, the result can be a huge number of calls to the message-passing library (either to implement matrix transpose operations and/or to manually implement some form of coherency protocol). By comparison, when Loop-Level Parallelism is implemented on a shared memory system, it is not uncommon to leave the loops in the boundary condition routines unparallelized (in general, these loops may represent 80% or more of the loops in the program, but less than 1% of the

total work). This makes it both painful to implement Loop-Level Parallelism using message-passing code and, in general, results in code that is very inefficient.

Traditionally, there have been two types of shared memory platforms. The first type is based on a small number of vector processors. This tends to make the system very expensive, while limiting one's ability to show good speedup. As a result, many codes run on vector processors were never parallelized. The second type of system was based on inexpensive mass-market microprocessors. Unfortunately, until recently, the aggregate peak speed of systems based on this design was generally much less than the peak speed of one processor on a state-of-the-art vector machine from Cray Research.

Therefore, until recently, none of the machines commonly used for High Performance Computing were well suited for use with Loop-Level Parallelism. It was not until the advent of the SGI Power Challenge that one could make a clear case for investigating this approach. Even then, enough people equated Loop-Level Parallelism with Automatic Parallelization (a concept that doesn't work very well) that they failed to properly appreciate the potential for this approach (Theys, Braun, and Siegel 1998). In fact, even now there are only a few systems (e.g., the SGI Origin 2000) for which a compelling case can be made (in some cases, the bottleneck is the hardware, while in other cases, limitations in the operating system and/or the compilers are at fault).

Table 1 shows the potential benefit of using Loop-Level Parallelism in conjunction with a well-designed shared memory system. Table 2 shows the actual speedup that was achieved for different problem sizes when using Loop-Level Parallelism with an SGI Origin 2000 to run the F3D code for a common test case.

4. Speedup

Up until now, this discussion has assumed that one can easily achieve linear speedup. In reality, this is frequently not the case. Therefore, let us consider what is likely to be the case when using

Table 1. The Number of Processors Required to Achieve a Specified Level of Delivered Performance Using Traditional Techniques

Speedup Relative to One Processor	Minimum No. of Processors Required When Using	
	Domain Decomposition	Cyclic Reduction
16	64	108
32	181	256
48	333	418
64	512	589
80	716	767

Table 2. The Speedup Achieved When Using the F3D Code Parallelized Using Loop-Level Parallelism on SGI Origin 2000's

No. of Processors Used	Grid Size (Millions of Grid Points)	Speedup Relative to One Processor
23	1.00	16.8
23	3.00	16.3
21	6.00	16.1
20	12.0	16.4
20	23.8	17.0
21	59.4	18.1
21	124.0	17.6
85	1.0	32.4
81	3.0	32.2
55	6.0	32.1
46	12.0	33.0
45	23.8	32.1
41	59.4	33.1
41	124.0	32.8
114	12.0	48.1
87	23.8	49.8
61	59.4	48.1
61	124.0	48.6
117	59.4	66.9
88	124.0	65.7
116	124.0	81.4

Note: Except for the largest test case, runs using fewer than 64 processors were run on either 32 or 64 processor Origin 2000's. Due to the memory requirements of the largest test case, all runs were made on a 128 processor Origin 2000. For all of the remaining cases, runs were made on a preproduction 128 processor Origin 2000.

both the traditional approaches to parallelization and loop level parallelism. Based on the numbers in Table 1, it is clear that when using traditional approaches, one will likely need a large number of processors. However, for fixed-size problems, Amdahl's Law predicts that there is enough serial code remaining that one will asymptotically approach a maximum level of performance when using large numbers of processors. The traditional counter argument has been to use the concept of *scaled speedup* (Gustafson 1988). With this concept, the available parallelism and the available work are assumed to scale linearly with the problem size. Therefore, as the problem size gets bigger, one can use additional processors while keeping the run time constant. This concept also assumes that the amount of work associated with the serial code grows very slowly, if at all, and can therefore be ignored.

A common rule of thumb when parallelizing programs on distributed memory MPPs is that one should use the smallest number of processors possible, with the amount of memory per processor usually being the limiting factor. Most modern MPPs are now equipped with between 64 MB and 1 GB of memory per processor, with somewhere around 10–20 MB of memory per processor reserved for use by the operating system. Based on these numbers, Table 3 shows how many processors one would normally expect to use for the test cases mentioned in Table 2.

There is no guarantee that one will actually get good scalability all the way to the upper bounds listed in Table 3. Rather, the upper bound is based on the impossibility of running the job if there is not enough memory. However, the rule of thumb indicates that it is questionable if one will see linear speedup when using even larger numbers of processors.

When comparing Tables 1 and 3, it becomes apparent that there are some problems. The smaller test cases are unlikely to produce speedups much in excess of a factor of 16. While in theory the larger problem sizes will fare better, there is a second, less obvious problem. Very few of the currently installed MPPs are configured with 512 or more processors. Therefore, in many cases, one will find it difficult, if not impossible, to use enough processors to get speedups of 64 or greater.

Table 3. The Number of Processors That One Would Normally Use When Using an MPP and Traditional Techniques to Process the Test Cases

Grid Size (Millions of Grid Points)	Recommended No. of Processors
1.0	1-10
3.0	2-30
6.0	3-60
12.0	6-120
23.8	12-240
59.4	30-600
124.0	62-1,240

Turning our attention to programs parallelized using Loop-Level Parallelism, the following question comes up: What kinds of speedup is one likely to see from these programs? The answer here is a bit complicated. In general, the available parallelism will be a function of the smallest of the grid dimensions. Therefore, the available parallelism will, at best, scale as the cube root of the size of each zone. A direct result of this is that it no longer makes sense to talk about scaled speedup. Instead, one is back at the problem of obtaining speedup for a fixed problem size.

The second problem is that when using Loop-Level Parallelism, the available parallelism is frequently within an order of magnitude of the number of processors being used. Since there are an integer number of iterations in a loop, the predicted speedup is no longer linear, but rather is a stairstep.

Figure 2 and Table 4 show an example of this. This also means that, for smaller problems, one may run out of parallelism in some/or all of the loops prior to using all of the processors in the machine.

An additional complication with Loop-Level Parallelism is that since many of the loops will be doing very little work, the overhead associated with parallelizing them may be so great as to result in Parallel Slowdown! This situation is especially common in the boundary condition routines.




Figure 2. Predicted Speedup for a Loop With 15 Units of Parallelism.

Table 4. Predicted Speedup for a Loop With 15 Units of Parallelism

No. of Processors	Maximum Units of Parallelism Assigned to a Single Processor	Predicted Speedup
1	15	1.000
2	8	1.875
3	5	3.000
4	4	3.750
5-7	3	5.000
8-14	2	7.500
15	1	15.000

While in theory this problem should be less severe when dealing with larger problem sizes, the reality of the situation is that the code will normally be tuned for the smaller problem sizes. While in many cases it should be possible to reduce the amount of CPU time spent on serial code to 1% or less of the total CPU time, this is enough for Amdahl's Law to be a problem when using more

than about 50 processors. The combination of the stairstepping with Amdahl's law explains why the smaller test cases show limited speedup in Table 2.*

Figures 3 and 4 show all of these effects in a real problem. Figure 3 is for a relatively small problem (less than 500 MB of memory), while Figure 4 is for a relatively large problem (over 20 GB of memory). Our calculations indicate that the primary reason for the difference between the predicted and measured levels of performance in these curves is Amdahl's law.†




Figure 3. Performance Results for the One-Million-Grid-Point Data Set.

* A third reason for the limited speedup is that the average memory latency on the 128 processor Origin 2000 is slightly longer than on the 32 and 64 processor systems. This has the effect of decreasing the serial efficiency from 30–40% to about 25–30% on the 128 processor system.

† The predicted curve is based on the assumption that one can achieve the same percentage of peak performance for a single processor job on both the Cray C90 and on RISC-based machines such as the SGI Origin 2000. This does not mean that one will achieve this without any work. Rather, it is assumed that a significant effort at tuning the code was made on both platforms. Taking into account only the available level of parallelism (the stairstepping effect), this expected level of performance is then extrapolated out for multiprocessor runs. As such, the predicted level of performance will, in general, equal or exceed the observed level of performance and serve as an excellent reference point for determining how well the system is performing.

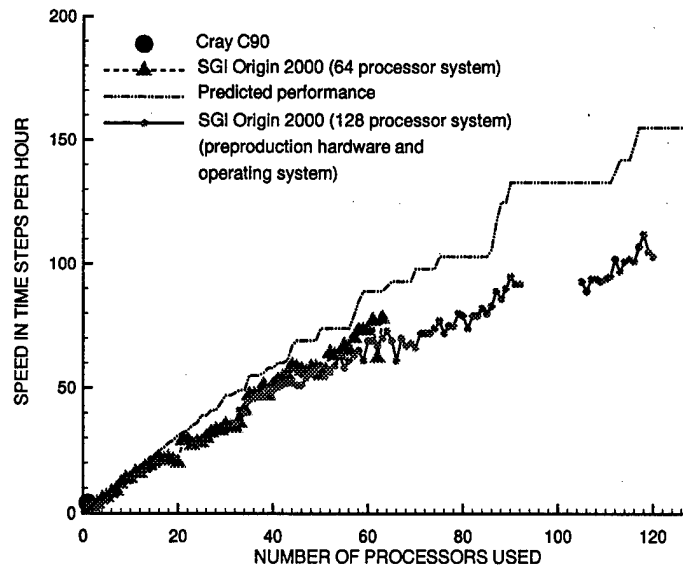


Figure 4. Performance Results for the 59-Million-Grid-Point Data Set.

5. Conclusions

The combination of Loop-Level Parallelism and RISC-based SMPs has been shown to be a promising approach to parallelizing a class of highly efficient algorithms that had previously resisted attempts at parallelization. Additionally, evidence has been presented that demonstrates that, in general, the resulting code achieves a much higher level of delivered performance than traditional techniques might be expected to deliver. While it is not practical to look at all possible approaches in detail and to determine what their effect is on the Total Efficiency in all cases, it seems likely that the benefits of using our approach are real.

An additional consideration is the availability of the hardware. SGI and SUN have both been quite successful at selling moderate-sized RISC-based SMPs. While IBM and Cray (a subsidiary of SGI) have sold a significant number of MPPs, very few of them had 512 or more processors.

Therefore, even when in theory the performance of a large MPP using traditional methods should exceed our results, it is far from certain that one will actually be able to obtain access to enough processors in a single machine at one time.

6. References

- Flynn, M. J. "Some Computer Organizations and Their Effectiveness." *IEEE Transactions Computers*, C-21, pp. 948-60, 1972.
- Gustafson, J. L. "Reevaluating Amdahl's Law." *Communications of the ACM*, vol. 31, no. 5, pp. 532-533, The Association for Computing Machinery, Inc., May 1988.
- O'Neal, D., and J. Urbanic. "On Performance and Efficiency: Cray Architectures." Parallel Applications Group Pittsburgh Supercomputing Center, electronically published at <http://www.psc.edu/~oneal/eff/eff.html>, August 1997.
- Sahu, J., D. M. Pressel, K. R. Heavey, and C. J. Nietubicz. "Parallel Application of a Navier-Stokes Solver for Projectile Aerodynamics." *Parallel Computational Fluid Dynamics, Recent Developments and Advances Using Parallel Computers. Proceedings of the Parallel CFD'97 Conference*, Manchester, UK, 19-21 May 1997, edited by D. R. Emerson, J. Periaux, A. Ecer, N. Satofuka, and P. Fox, Amsterdam: Elsevier, 1998.
- Sahu, J., D. M. Pressel, K. R. Heavey, and C. J. Nietubicz. "Parallel Application of a Navier-Stokes Solver for Projectile Aerodynamics." *Proceedings of the 1998 Army Science Conference*, to be published.
- Saini, S. (ed.). "NAS Parallel Benchmarks, NPB 1 Data." Electronically published at <http://Science.nas.nasa.gov/Software/NPB/NPB1Results/index.html>, 17 November 1996.
- Saini, S. (ed.). "NAS Parallel Benchmarks, NPB 2 Data." Electronically published at <http://Science.nas.nasa.gov/Software/NPB/NPB2Results/index.html>, 17 November 1997.
- Singh, K. P., Biju Uthup, and Laxmi Ravishanker. "Parallelization of Euler and N-S Code on 32 Node Parallel Super Computer PACE+." Presented at the ADA/DRDO-DERA Workshop on CFD, 1998(?).
- Theys, M. D., T. D. Braun, and H. J. Siegel. "Widespread Acceptance of General-Purpose, Large-Scale Parallel Machines: Fact, Future, or Fantasy?" *IEEE Concurrency Parallel, Distributed and Mobile Computing*, IEEE Computer Society, January-March 1998.
- Wang, G., and D. K. Tafti. "Performance Enhancement on Microprocessors With Hierarchical Memory Systems for Solving Large Sparse Linear Systems." *The International Journal of Supercomputing Applications*, February 1997.

INTENTIONALLY LEFT BLANK.

Glossary

CFD	Computational Fluid Dynamics
Domain decomposition	The process of splitting a small number of zones (some of which are assumed to be large) into a moderate to large number of zones (generally all of which are fairly small in size)
FLOPS	Floating-Point Operations Per Second
MFLOPS	Million Floating-Point Operations Per Second
MIMD	Multiple Instruction Multiple Data - a class of parallel computers as defined in Flynn's taxonomy
MPP	Massively Parallel Processor
RISC	Reduced Instruction Set Computer
SIMD	Single Instruction Multiple Data - a class of parallel computers as defined in Flynn's taxonomy
SMP	Symmetric Multiprocessor - a term normally only applied to shared memory systems using hardware memory coherency protocols

INTENTIONALLY LEFT BLANK.

<u>NO. OF COPIES</u>	<u>ORGANIZATION</u>
2	DEFENSE TECHNICAL INFORMATION CENTER DTIC DDA 8725 JOHN J KINGMAN RD STE 0944 FT BELVOIR VA 22060-6218
1	HQDA DAMO FDQ D SCHMIDT 400 ARMY PENTAGON WASHINGTON DC 20310-0460
1	OSD OUSD(A&T)/ODDDR&E(R) R J TREW THE PENTAGON WASHINGTON DC 20301-7100
1	DPTY CG FOR RDA US ARMY MATERIEL CMD AMCRDA 5001 EISENHOWER AVE ALEXANDRIA VA 22333-0001
1	INST FOR ADVNCD TCHNLGY THE UNIV OF TEXAS AT AUSTIN PO BOX 202797 AUSTIN TX 78720-2797
1	DARPA B KASPAR 3701 N FAIRFAX DR ARLINGTON VA 22203-1714
1	NAVAL SURFACE WARFARE CTR CODE B07 J PENNELLA 17320 DAHLGREN RD BLDG 1470 RM 1101 DAHLGREN VA 22448-5100
1	US MILITARY ACADEMY MATH SCI CTR OF EXCELLENCE DEPT OF MATHEMATICAL SCI MADN MATH THAYER HALL WEST POINT NY 10996-1786

<u>NO. OF COPIES</u>	<u>ORGANIZATION</u>
1	DIRECTOR US ARMY RESEARCH LAB AMSRL DD J J ROCCHIO 2800 POWDER MILL RD ADELPHI MD 20783-1197
1	DIRECTOR US ARMY RESEARCH LAB AMSRL CS AS (RECORDS MGMT) 2800 POWDER MILL RD ADELPHI MD 20783-1145
3	DIRECTOR US ARMY RESEARCH LAB AMSRL CI LL 2800 POWDER MILL RD ADELPHI MD 20783-1145
	<u>ABERDEEN PROVING GROUND</u>
4	DIR USARL AMSRL CI LP (BLDG 305)

<u>NO. OF</u> <u>COPIES</u>	<u>ORGANIZATION</u>	<u>NO. OF</u> <u>COPIES</u>	<u>ORGANIZATION</u>
1	PM CHSSI JOHN GROSH SUITE 650 1110 N GLEBE ROAD ARLINGTON VA 22201	1	NAVAL RSRCH LAB GEORGE HEBURN RSRCH OCEANOGRAPHER CNMOC BLDG 1020 RM 178 STENNIS SPACE CENTER MS 39529
1	RICE UNIVERSITY MCHNCL ENGRNG AND MTRLS SCI MAREK BEHR MS 321 6100 MAIN STREET HOUSTIN TX 77005	1	US AIR FORCE WRIGHT LAB WL FIM JOSEPH J S SHANG 2645 FIFTH STREET STE 6 WPAFB OH 45433-7912
1	COMMANDER CODE C2892 CLINT HOUSH 1 ADMINISTRATION CIR CHINA LAKE CA 93555	1	USAF PHILIPS LAB OLAC PL RKFE CPT SCOTT G WIERSCHKE 10 EAST SATURN BLVD EDWARDS AFB CA 93524-7680
2	WL FIMC STEPHEN SCHERR BILL STRANG BLDG 450 2645 FIFTH ST SUITE 7 WPAFB OH 45433-7913	1	USAE WATERWAYS EXPERIMENT STATION CEWES HV C JEFFREY P HOLLAND 3909 HALLS FERRY ROAD VICKSBURG MS 39180-6199
1	NSWC A B WARDLAW CODE B44 SILVER SPRING MD 20903-5640	1	US ARMY CECOM RD&E CTR AMSEL RD C2 BARRY S PERLMAN FT MONMOUTH NJ 07703
1	NAVAL RSRCH LAB CODE 6400 JAY BORIS 4555 OVERLOOK AVE SW WASHINGTON DC 20375-5344	1	SPAWARSYSCEN (D4402) ROBERT A WASILAUSKY BLDG 33 RM 0071A 53560 HULL ST SAN DIEGO CA 92152-5001
1	NAVAL RSRCH LAB CODE 6410 RAVI RAMAMURTI WASHINGTON DC 20375-5344	1	US AIR FORCE RESEARCH LAB INFORMATION DIRECTORATE RICHARD W LINDERMAN 26 ELECTRONIC PARKWAY ROME NY 13441-4514
1	ARMY AEROFLIGHT DYNAMICS DIRECTORATE ROBERT MEAKIN MS 258 1 MOFFETT FIELD CA 94035-1000	1	US AIR FORCE RESEARCH LAB PROPULSION DIRECTORATE LESLIE PERKINS 5 POLLUX DR EDWARDS AFB CA 93524-7048
1	NAVAL RSRCH LAB CODE 7320 J W MCCAFFREY JR HEAD OCEAN DYNAMICS AND PREDICTION BRANCH STENNIS SPACE CENTER MS 39529	1	AIR FORCE RESEARCH LAB/DEHE ROBERT PETERKIN 3550 ABERDEEN AVE SE KIRTLAND AFB NM 87117-5776

NO. OF
COPIES ORGANIZATION

1 SPACE & NAVAL WARFARE SYS CTR
CODE D7305 KEITH BROMLEY
BLDG 606 RM 325
53140 SYSTEMS ST
SAN DIEGO CA 92152-5001

1 UNVRSTY OF MINNESOTA
DEPT OF ASTRONOMY
PROF P WOODWARD
356 PHYSICS BLDG
116 CHURCH STREET SE
MINNEAPOLIS MN 55455

1 RICE UNIVERSITY
MCHNCL ENGRNG AND MTRLS SCI
TAYFUN TEZDUYAR DEPT CHR MN
MS 321 6100 MAIN ST
HOUSTON TX 77005

1 DIRECTOR
ARMY HIGH PERFORMANCE
COMPUTING RSRCH CTR
BARBARA BRYAN
1200 WASHINGTON AVE
SOUTH MINNEAPOLIS MN 55415

1 DIRECTOR
ARMY HIGH PERFORMANCE
COMPUTING RSRCH CTR
GRAHAM V CANDLER
1200 WASHINGTON AVE
SOUTH MINNEAPOLIS MN 55415

1 NAVAL CMND CONTROL AND
OCEAN SURVEILLANCE CTR
L PARNELL HPC CRDNTR & DIR
NCCOSC RDTE DIV D3603
49590 LASSING ROAD
SAN DIEGO CA 92152-6148

NO. OF
COPIES ORGANIZATION

ABERDEEN PROVING GROUND

15 DIR USARL
AMSRL CI
N RADHAKRISHNAN
AMSRL CI H
C NIETUBICZ
AMSRL CI HA
W STUREK
A MARK
R NAMBURU
AMSRL CI HC
D PRESSEL
D HISLEY
C ZOLTANI
A PRESSLEY
T KENDALL
P DYKSTRA
AMSRL WM BC
H EDGE
J SAHU
K HEAVEY
P WEINACHT

INTENTIONALLY LEFT BLANK.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.			
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE October 1999	3. REPORT TYPE AND DATES COVERED Final, Jan 96 - Dec 96	
4. TITLE AND SUBTITLE How Moderate-Sized RISC-Based SMPs Can Outperform Much Larger Distributed Memory MPPs		5. FUNDING NUMBERS 611102H4800	
6. AUTHOR(S) D. M. Pressel, Walter B. Sturek, J. Sahu, and K. R. Heavey			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) U.S. Army Research Laboratory ATTN: AMSRL-CI-HC Aberdeen Proving Ground, MD 21005-5067		8. PERFORMING ORGANIZATION REPORT NUMBER ARL-TR-2062	
9. SPONSORING/MONITORING AGENCY NAMES(S) AND ADDRESS(ES)		10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES			
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.		12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) Historically, comparisons between computer systems were based primarily on theoretical peak performance. Today, comparisons based on delivered levels of performance are frequently used. This, of course, raises a whole host of questions concerning methodology. From the standpoint of the user, delivered performance frequently refers to how fast a job runs. However, is it reasonable to base this measurement on running the same algorithm on all of the computers? When comparing some combination of mainframes and vector supercomputers, the answer is probably yes. The same holds true when comparing the performance of large distributed memory MIMD MPPs. However, when comparing the algorithms of choice used on these two classes of platforms, one frequently finds that the algorithms are quite different. Furthermore, the amount of work (the number of FLOPS) associated with each algorithm can also be quite different. While troubling, this dichotomy has been largely unavoidable. This implies that for an MPP to have the same level of delivered performance as the mainframes and the vector supercomputers, it must have a significantly greater level of performance when measured in terms of FLOPS. Recent advances involving moderate-sized RISC-based SMPs have allowed us to solve this problem. The net result is that for some problems a 128 processor Origin 2000 can deliver levels of performance that might require the use of a 500 processor MPP using more traditional approaches.			
14. SUBJECT TERMS supercomputer, high performance computing, parallel processor, computational fluid dynamics		15. NUMBER OF PAGES 25	
		16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL

INTENTIONALLY LEFT BLANK.

USER EVALUATION SHEET/CHANGE OF ADDRESS

This Laboratory undertakes a continuing effort to improve the quality of the reports it publishes. Your comments/answers to the items/questions below will aid us in our efforts.

1. ARL Report Number/Author ARL-TR-2062 (Pressel) Date of Report October 1999

2. Date Report Received _____

3. Does this report satisfy a need? (Comment on purpose, related project, or other area of interest for which the report will be used.) _____

4. Specifically, how is the report being used? (Information source, design data, procedure, source of ideas, etc.) _____

5. Has the information in this report led to any quantitative savings as far as man-hours or dollars saved, operating costs avoided, or efficiencies achieved, etc? If so, please elaborate. _____

6. General Comments. What do you think should be changed to improve future reports? (Indicate changes to organization, technical content, format, etc.) _____

CURRENT
ADDRESS

Organization

Name

E-mail Name

Street or P.O. Box No.

City, State, Zip Code

7. If indicating a Change of Address or Address Correction, please provide the Current or Correct address above and the Old or Incorrect address below.

OLD
ADDRESS

Organization

Name

Street or P.O. Box No.

City, State, Zip Code

(Remove this sheet, fold as indicated, tape closed, and mail.)
(DO NOT STAPLE)