

# Parallel Conjugate Gradient: Effects of Ordering Strategies, Programming Paradigms, and Architectural Platforms

L. Oliker, X. Li  
NERSC  
Mail Stop 50F  
Lawrence Berkeley Natl. Lab.  
Berkeley, CA 94720

G. Heber  
Cornell Theory Ctr.  
638 Rhodes Hall  
Cornell University  
Ithaca, NY 14853

R. Biswas  
Computer Sciences Corp.  
Mail Stop T27A-1  
NASA Ames Res. Ctr.  
Moffett Field, CA 94035

## Abstract

The Conjugate Gradient (CG) algorithm is perhaps the best-known iterative technique to solve sparse linear systems that are symmetric and positive definite. A sparse matrix-vector multiply (SPMV) usually accounts for most of the floating-point operations within a CG iteration. In this paper, we investigate the effects of various ordering and partitioning strategies on the performance of parallel CG and SPMV using different programming paradigms and architectures. Results show that for this class of applications, ordering significantly improves overall performance, that cache reuse may be more important than reducing communication, and that it is possible to achieve message passing performance using shared-memory constructs through careful data ordering and distribution. However, a multithreaded implementation of CG on the Tera MTA does not require special ordering or partitioning to obtain high efficiency and scalability.

## 1 Introduction

The ability of computers to solve hitherto intractable problems and simulate complex processes using mathematical models makes them an indispensable part of modern science and engineering. Computer simulations of large-scale realistic applications usually require solving a set of non-linear partial differential equations (PDEs) over a finite region. Structured grids are the most natural way to discretize such a computational domain; however, complicated domains must often be divided into multiple structured grids to be completely discretized, requiring a great deal of human intervention. Unstructured meshes, by contrast, can be generated automatically for applications with complex geometries or those with dynamically moving boundaries (but at the cost of higher

storage requirements to explicitly store the connectivity information for every point in the mesh). They also facilitate dynamic grid adaptation to efficiently solve problems with evolving physical features such as shock waves, vortices, detonations, shear layers, and crack propagation.

The process of obtaining numerical solutions to the governing PDEs requires solving large sparse linear systems or eigen systems defined over the unstructured meshes that model the underlying physical objects. The Conjugate Gradient (CG) algorithm is perhaps the best-known iterative technique to solve sparse linear systems that are symmetric and positive definite. Within each iteration of CG, the sparse matrix vector multiply (SPMV) is usually the most expensive operation.

On uniprocessor machines, numerical solutions of such complex, real-life problems can be extremely time consuming, a fact driving the development of increasingly powerful parallel multiprocessor supercomputers. The unstructured, dynamic nature of many systems worth simulating, however, makes their efficient parallel implementation a daunting task. This is primarily due to the load imbalance created by the dynamically changing nonuniform grids and the irregular data access patterns. These cause significant communication at runtime, leaving many processors idle and adversely affecting the total execution time.

Furthermore, modern computer architectures, based on deep memory hierarchies, show acceptable performance only if users care about the proper distribution and placement of their data [2, 8]. Single-processor performance crucially depends on the exploitation of locality, and parallel performance degrades significantly if inadequate partitioning of data causes excessive communication and/or data migration. The traditional approach would be to use a sophisticated partitioning algorithm, and then to post-

process the resulting partitions with an enumeration strategy for enhanced locality. Although, in that sense, optimizations for partitioning and locality may be treated as separate problems, real applications tend to show a rather intricate interplay of both.

In this paper, we investigate the effects of various ordering and partitioning strategies on the performance of CG and SPMV using different programming paradigms and architectures. In particular, we use the reverse Cuthill-McKee [3] and the self-avoiding walks [5] ordering strategies, and the METIS [7] partitioner. We examine parallel implementations of CG using MPI, shared-memory compiler directives, and multithreading, on three state-of-the-art parallel supercomputers: a Cray T3E, an SGI Origin2000, and a Tera MTA. Results show that ordering significantly improves overall performance, that cache reuse can be more important than reducing communication, and that it is possible to achieve message passing performance using shared-memory constructs through careful data ordering and distribution. However, the multithreaded implementation does not require any special ordering or partitioning to obtain high efficiency and scalability.

## 2 Partitioning and Linearization

Space-filling curves have been demonstrated to be an elegant and unified linearization approach for certain problems in N-body and FEM simulations, mesh partitioning, and other graph-related areas [4, 9, 10, 11, 13]. The linearization of a higher-dimensional spatial structure, i.e. its mapping onto a one-dimensional structure, is exploited in two ways: First, the locality preserving nature of the construction fits elegantly into a given memory hierarchy, and second, the partitioning of a contiguous linear object is trivial. For our experiments, we pursued both strategies with some modifications. In the following, we briefly describe the two classes of enumeration techniques and the general-purpose graph partitioner which were used.

### 2.1 Cuthill-McKee Algorithms (CM)

The particular enumeration of the vertices in a FEM discretization controls, to a large extent, the sparseness pattern of the resulting stiffness matrix. The bandwidth, or profile, of the matrix has a significant impact on the efficiency of linear systems and eigensolvers. Cuthill and McKee [3] suggested a simple algorithm based on ideas from graph theory. Starting from a vertex of minimal degree, levels

of increasing distance from that vertex are first constructed. The enumeration is then performed level-by-level with increasing vertex degree (within each level). Several variations of this method have been suggested, the most popular being reverse Cuthill-McKee (RCM) where the level construction is restarted from a vertex of minimal degree in the final level. In many cases, it has been shown that RCM improves the profile of the resulting matrix. The class of CM algorithms are fairly straightforward to implement and largely benefit by operating on a pure graph structure, i.e. the underlying graph is not necessarily derived from a triangular mesh.

### 2.2 Self-Avoiding Walks (SAW)

These were proposed recently [5] as a mesh-based (as opposed to geometry-based) technique with similar application areas as space-filling curves. A SAW over a triangular mesh is an enumeration of the triangles such that two consecutive triangles (in the SAW) share an edge or a vertex, i.e. there are no jumps in the SAW. It can be shown that walks with more specialized properties exist over arbitrary unstructured meshes, and that there is an algorithm for their construction whose complexity is linear in the number of triangles in the mesh. Furthermore, SAWs are amenable to hierarchical coarsening and refinement, i.e. they have to be rebuilt only in regions where mesh adaptation occurs, and can therefore be easily parallelized. SAW, unlike CM, is not a technique designed specifically for vertex enumeration; thus, it cannot operate on the bare graph structure of a triangular mesh. This implies a higher construction cost for SAWs, but several different vertex enumerations can be derived from a given SAW.

### 2.3 Graph Partitioning (e.g. METIS)

Some excellent parallel graph partitioning algorithms have been developed and implemented in the last decade that are extremely fast while giving good load balance quality and low edge cuts. Perhaps the most popular is METIS [7] that belongs to the class of multilevel partitioners. METIS reduces the size of the graph by collapsing vertices and edges using a heavy edge matching scheme, applies a greedy graph growing algorithm for partitioning the coarsest graph, and then uncoarsens it back using a combination of boundary greedy and Kernighan-Lin refinement to construct a partitioning for the original graph. Partitioners strive to balance the computational workload among processors while reducing interprocessor communication.

Improving cache performance is not a typical objective of most partitioning algorithms.

### 3 Sparse Matrix-Vector Multiplication and Conjugate Gradient

Sparse matrix-vector multiplication (SPMV) is one of the most heavily-used kernels in large-scale numerical simulations. To perform a SPMV,  $y \leftarrow Ax$ , we assume that the nonzeros of matrix  $A$  are stored in the Compressed Row Storage (CRS) format [1]. The dense vector  $x$  is stored sequentially in memory with unit stride. Various numberings of the mesh elements/vertices result in different nonzero patterns of  $A$ , which in turn cause different access patterns for the entries of  $x$ . Moreover, on a distributed-memory machine, they imply different amounts of communication.

The Conjugate Gradient (CG) algorithm is the oldest and best-known Krylov subspace method used to solve the linear system  $Ax = b$ . The method starts from an initial guess  $x_0$  of the vector  $x$ . It then successively generates approximate solutions in the Krylov subspace, and search directions used in updating the approximate solution and residual. The algorithm [12] is outlined in Figure 1. Each iteration of CG involves one SPMV for  $Ap_j$ , three vector updates (AXPY) for  $x_{j+1}$ ,  $r_{j+1}$ , and  $p_{j+1}$ , and three inner products (DOT) for the update scalars  $\alpha_j$  and  $\beta_j$  which make the generated sequences satisfy certain orthogonality conditions. For a symmetric and positive definite linear system, these conditions imply that the distance between the approximate solution and the true solution is minimized.

```

Compute  $r_0 = p_0 = b - Ax_0$  for an initial guess  $x_0$ 
for  $j = 0, 1, \dots$ , until convergence
     $\alpha_j = (r_j, r_j) / (Ap_j, p_j)$ 
     $x_{j+1} = x_j + \alpha_j p_j$ 
     $r_{j+1} = r_j - \alpha_j Ap_j$ 
     $\beta_j = (r_{j+1}, r_{j+1}) / (r_j, r_j)$ 
     $p_{j+1} = r_{j+1} + \beta_j p_j$ 
endfor

```

Figure 1: The Conjugate Gradient algorithm.

Suppose the matrix  $A$  is of order  $n$  and has  $nnz$  nonzeros. Then, one SPMV involves  $O(nnz)$  floating-point operations, while AXPY and DOT involve only  $O(n)$  floating-point operations. Thus, for many practical matrices, SPMV dominates the other two oper-

ations. This is demonstrated by the results given in Section 4. Note that both AXPY and DOT are insensitive to mesh orderings.

### 4 Experimental Results

Our experimental test mesh consists of a two-dimensional Delaunay triangulation, generated by the Triangle [14] software package. The mesh is shaped like the letter “A”, and contains 661,054 vertices and 1,313,099 triangles. The underlying matrix was assembled by assigning a random value in  $(0, 1)$  to each  $(i, j)$  entry corresponding to the vertex pair  $(v_i, v_j)$ , where  $1 \leq \text{distance}(v_i, v_j) \leq 3$ . All other off-diagonal entries were set to zero. This simulates a stencil computation where each vertex needs to communicate with its neighbors that are no more than three edge lengths away. The matrix is symmetric with its diagonal entries set to 40, which makes it diagonally dominant (and hence positive definite). This ensures that the CG algorithm converges successfully. The final sparse matrix  $A$  has approximately 39 entries per row and a total of 25,753,034 nonzeros. This sparsity is representative of matrices obtained from discretizing PDEs on three-dimensional meshes; however, the connectivity pattern will be different for three-dimensional problems. The CG algorithm converges in 13 iterations, with the unit vector as the right-hand side  $b$  and the zero vector as the initial guess for  $x$ . For our test matrix, the SPMV computation accounts for approximately 87% of the total number of floating-point operations within each CG iteration.

#### 4.1 Distributed-Memory Implementation

In our experiments, we use the parallel SPMV and CG routines in Aztec [6], implemented using MPI. The matrix  $A$  is partitioned into blocks of rows, with each block assigned to one processor. The associated components of vectors  $x$  and  $b$  are distributed accordingly. Communication may be needed to transfer some components of  $x$ . For example, in  $y \leftarrow Ax$ , if  $y_i$  is updated on processor  $p_1$ ,  $A_{ij} \neq 0$ , and  $x_j$  is owned by processor  $p_2$ , then  $p_2$  must send  $x_j$  to  $p_1$ . In general, a processor may need more than one  $x$ -component from another processor. It is thus more efficient to combine several  $x$ -components into one message so that each processor sends no more than one message to another processor. This type of optimization can be performed in a pre-processing phase. The other two operations, AXPY and DOT in the CG algorithm, are easily parallelized — AXPY requires only local computations,

whereas DOT requires a local sum followed by a global sum reduction.

Three routines within Aztec are of particular interest to us: `AZ_transform`, which initializes the data structures and the communication schedule for SPMV, `AZ_matvec_mult`, which performs the matrix-vector multiply, and `AZ_cg`, which solves a linear system using the CG algorithm. In Table 1, we report the runtimes of the `AZ_matvec_mult` and `AZ_cg` routines on the Cray T3E at NERSC. Each T3E node consists of a 450 MHz DEC Alpha processor (900 Mflops peak theoretical floating-point speed), 96 KB secondary cache, and is interconnected to other nodes through a 3D torus. It was not possible to run our test problem on less than 8 processors of the T3E due to memory constraints.

P	ORIG		METIS	
	SPMV	CG	SPMV	CG
8	0.562	8.652	0.476	7.662
16	0.325	5.093	0.268	2.909
32	0.199	3.167	0.087	1.468
64	0.119	1.929	0.056	0.961
P	RCM		SAW	
	SPMV	CG	SPMV	CG
8	0.381	6.185	0.171	2.916
16	0.193	3.198	0.086	1.491
32	0.095	1.662	0.044	0.795
64	0.045	0.882	0.028	0.462

Table 1: Runtimes (in seconds) of `AZ_matvec_mult` (SPMV) and `AZ_cg` (CG) using different orderings on the Cray T3E.

For the key kernel routine `AZ_matvec_mult`, SAW is always about twice as fast as RCM. In turn, RCM is about 1.5 times faster than METIS on 16 or fewer processors, and about the same on 32 or more processors. Note that when using 32 or more processors, METIS is twice as fast as ORIG (the natural ordering from Triangle). For `AZ_cg`, SAW is again about twice as fast as RCM. However, we do not see a clear advantage of RCM over METIS for this routine. Both RCM and METIS are twice as fast as ORIG on large number of processors. Finally, METIS, RCM, and SAW, all demonstrate excellent scalability (more than 75% efficiency) up to the 64 processors that were used for these experiments, but ORIG seems less scalable (only about 56% efficiency). As expected, there is a strong correlation between the performance of CG and the underlying SPMV for all test cases.

Table 2 shows the pre-processing times spent in `AZ_transform`. The times for METIS, RCM, and SAW are comparable, and are usually an order of magnitude larger than the corresponding times for `AZ_matvec_mult`. The `AZ_transform` times show some scalability up to 32 processors. However, for ORIG, the times are two to three orders of magnitude larger, and show very little scalability. Clearly, the ORIG ordering is too inefficient and unacceptable on distributed-memory machines.

P	ORIG	METIS	RCM	SAW
8	504.2	2.829	2.370	2.023
16	547.9	1.455	1.330	1.157
32	333.7	0.840	0.864	0.804
64	150.0	0.422	0.776	0.537

Table 2: Runtimes (in seconds) of `AZ_transform` using different orderings on the Cray T3E.

To better understand the various partitioning and ordering algorithms, we have built a simple performance model to predict the parallel runtime of `AZ_matvec_mult`. First, using the T3E's hardware performance monitor, we collected the average number of cache misses per processor. This is reported in Table 3, and shows that SAW has the fewest number of cache misses. In comparison, RCM, METIS, and ORIG have between two and three times that number. Second, we gathered statistics on the average communication volume and the maximum number of messages per processor, both of which are also shown in Table 3. Notice that METIS transfers the least

P	Avg. Cache Misses ( $\times 10^6$ )			
	ORIG	METIS	RCM	SAW
8	3.684	3.034	3.749	2.004
16	2.007	1.330	1.905	0.971
32	1.060	0.658	1.017	0.507
64	0.601	0.358	0.515	0.290
P	Avg. Comm. Vol. (in Mbytes) (Max. # Msgs)			
	ORIG	METIS	RCM	SAW
8	3.228 (7)	0.011 (3)	0.031 (2)	0.049 (6)
16	2.364 (15)	0.011 (4)	0.032 (2)	0.036 (9)
32	1.492 (31)	0.009 (5)	0.032 (2)	0.030 (11)
64	0.828 (63)	0.008 (6)	0.032 (2)	0.023 (16)

Table 3: Locality and communication statistics for `AZ_matvec_mult`.

amount of data, whereas RCM has the fewest number of messages.

In our model, we estimate the total parallel runtime  $T$  as

$$T = T_f + T_m + T_c ,$$

where,  $T_f$ ,  $T_m$ , and  $T_c$  are the estimated per-processor times to perform floating-point operations, to service the cache misses, and to communicate the  $x$  vector. Given that a floating-point operation requires 1/900 microseconds and that each cache miss latency is 0.08 microseconds (both from product documentation), and assuming that the MPI bandwidth and latency are 50 MB/second and 10 microseconds (both from measurement), respectively, we can estimate the total runtime based on the information in Table 3.

P	ORIG		METIS	
	$T$ (deviation)	$T_m/T$	$T$ (deviation)	$T_m/T$
8	0.367 (-35%)	0.80	0.250 (-47%)	0.97
16	0.212 (-35%)	0.76	0.110 (-58%)	0.96
32	0.117 (-41%)	0.72	0.055 (-37%)	0.96
64	0.067 (-44%)	0.72	0.030 (-46%)	0.96
P	RCM		SAW	
	$T$ (deviation)	$T_m/T$	$T$ (deviation)	$T_m/T$
8	0.308 (-19%)	0.97	0.169 (-1%)	0.95
16	0.157 (-19%)	0.97	0.082 (-5%)	0.94
32	0.084 (-12%)	0.97	0.043 (-2%)	0.94
64	0.043 (-5%)	0.96	0.025 (-12%)	0.93

Table 4: Predicted runtimes (in seconds) for `AZ_matvec_mult` (SPMV) on the T3E. The fraction of the total time spent servicing cache misses is also shown. In the column of total time  $T$ , the percentage deviation from the measured time is given in parenthesis.

Table 4 shows the predicted total time  $T$  and the ratio  $T_m/T$ .  $T_f$  is comparatively negligible (consistently less than 5% of  $T$ ) for all ordering strategies and processor sets.  $T_c$  is 18–27% of  $T$  for ORIG, but less than 3% of  $T$  for METIS, RCM, and SAW. In parenthesis, we also give the percentage deviation of  $T$  from the measured experimental runtime (that are reported in Table 1). The maximum deviation from the measured runtimes is  $-58\%$ , which gives us a fair degree of confidence in our model. The results in Table 4 clearly indicate that servicing the cache misses is extremely expensive and requires more than 93% of the total time for METIS, RCM, and SAW, and 72–80% for ORIG (which has relatively more communication). Although SAW and RCM both incur more communi-

cation than METIS (in terms of the average message volume as shown in Table 3), their total runtimes are significantly less. This illustrates that for our combination of applications and architectures, improving cache reuse can be more important than reducing interprocessor communication.

## 4.2 Shared-Memory Implementation

The shared-memory version of CG was implemented on the SGI Origin2000, which is a SMP cluster of nodes each containing two 250 MHz MIPS R10000 processors and local memory. The hardware makes all memory equally accessible from a software standpoint, by sending memory requests through routers located on the nodes. Access time to memory is nonuniform, depending on how far away the memory lies from the processor. The topology of the interconnection network is a hypercube, bounding the maximum number of memory hops to a logarithmic function of the number of processors. Each processor also has a relatively large 4 MB secondary cache, where only it can fetch and store data. If a processor refers to data that is not in cache, there is a delay while a copy of the data is fetched from memory. When a processor modifies a word of data, all other copies of the cache line containing that word are invalidated.

This version of the parallel CG code was written using SGI’s native pragma directives, which create IRIX threads. A rewrite to OpenMP would require minimal programming effort but has not been done at this time. Each processor is assigned an equal number of rows in the matrix. The parallel SPMV and AXPY routines do not require explicit synchronizations, since they do not contain concurrent writes. Global reduction operations are required for DOT and the convergence tests. Two basic implementation approaches described below were taken.

The SHMEM strategy naively assumes that the Origin2000 is a flat shared-memory machine. Arrays are not explicitly distributed among the processors, and nonlocal data requests are handled by the cache coherent hardware. Alternatively, the CC-NUMA strategy addresses the underlying distributed-memory nature of the machine by performing an initial data distribution. Sections of the sparse matrix are appropriately mapped onto the memories of their corresponding processors using the default “first touch” data distribution policy of the Origin2000. The computational kernels of both the SHMEM and CC-NUMA implementations are identical, and simpler to implement than the MPI version. Table 5 shows the SPMV and CG runtimes using both approaches with

the ORIG, RCM, and SAW orderings of the mesh. We also present the runtime of CG using an MPI implementation on the Origin2000 with the SAW ordering, as a basis for comparison.

P	SHMEM			MPI
	ORIG	RCM	SAW	
1	46.911	37.183	36.791	
2	28.055	21.867	21.772	
4	30.637	25.350	24.751	
8	16.836	14.431	14.121	
16	16.348	15.516	15.548	
32	16.653	15.350	15.423	
64	10.809	7.782	8.450	

  

P	CC-NUMA			MPI
	ORIG	RCM	SAW	SAW
1	46.911	37.183	36.791	
2	27.053	21.454	21.229	23.145
4	17.608	10.651	10.593	7.880
8	9.824	5.575	5.516	3.815
16	6.205	2.845	2.872	1.926
32	3.584	1.548	1.514	1.075
64	2.365	0.885	0.848	0.905

Table 5: Runtimes (in seconds) of CG for different orderings running in SHMEM and CC-NUMA modes on the SGI Origin2000. The CG runtimes for an MPI implementation on the Origin2000 with the SAW ordering is also given for comparison.

Notice that the CC-NUMA implementation shows significant performance gains over SHMEM. This is expected since the Origin2000 is a distributed-memory system, and therefore should be treated as such. As the number of processors increases, the runtime difference between the two approaches becomes more dramatic, achieving an order of magnitude improvement when using more than 16 processors. Proper data distribution becomes increasingly important for larger numbers of processors since the corresponding communication overhead grows nonuniformly. Within the CC-NUMA approach, the RCM and SAW ordering schemes dramatically reduce the runtimes compared to ORIG, indicating that an intelligent ordering algorithm is necessary to achieve good performance and scalability on distributed shared-memory systems. There is little difference in parallel performance between RCM and SAW because both ordering techniques reduce the number of secondary cache misses and the non-local memory references of the processors. Recall however that on the T3E, SAW was about twice

as fast as RCM. This discrepancy in performance is probably due to the larger cache size of the Origin2000 that reduces the beneficial effects of smart ordering.

The last two columns of Table 5 compare the CC-NUMA and MPI implementations of CG on the Origin2000 using the SAW ordering. Notice that the runtimes are very similar, even though the programming methodologies of these two approaches are quite different. These results indicate that for this class of applications, it is possible to achieve message passing performance using shared-memory constructs, through careful data ordering and distribution.

### 4.3 Multithreaded Implementation

The Tera MTA is a supercomputer recently installed at the San Diego Supercomputing Center (SDSC). The MTA has a radically different architecture than current high-performance computer systems. Each 255 MHz processor has support for 128 hardware streams, where each stream includes a program counter and a set of 32 registers. One program thread can be assigned to each stream. The processor switches among the active streams at every clock tick, while executing a pipelined instruction.

The uniform shared memory of the MTA is flat, and physically distributed across hundreds of banks that are connected through a 3D toroidal network to the processors. All memory addresses are hashed by the hardware so that apparently adjacent words are actually distributed across different memory banks. Because of the hashing scheme, it is impossible for the programmer to control data placement. This enhances programmability compared to standard cache-based multiprocessor systems. Rather than using data caches to hide latency, the MTA processors use multithreading to tolerate latency. If a thread is waiting for its memory reference to complete, the processor executes instructions from other threads. Performance thus depends on having a large number of concurrent computation threads.

Lightweight synchronization among threads is provided by the memory itself. Each word of physical memory contains a full-empty bit, which enables fast synchronization via load and store instructions without operating system intervention. Synchronization among threads may stall one of the threads, but not the processor on which the threads are running, since each processor may run many threads. Explicit load balancing across loops is also not required since the dynamic scheduling of work to threads provides the ability of keeping the processors saturated, even if different iterations require varying amounts of time to

complete. Once a code has been written in the multithreaded model, no additional work is required to run it on multiple processors, since there is no difference between uni- and multiprocessor parallelism.

The multithreaded implementation of CG is trivial, requiring only MTA compiler directives. Since the data structures are dynamically allocated pointers, special pragma assertions were used to indicate that there are no loop-carried dependencies. The compiler was thus able to automatically parallelize the appropriate loop segments. Load balancing is implicitly handled by the operating system which dynamically assigns rows to threads. The reduction operations for DOT and the convergence test were handled automatically as well. Otherwise, special synchronization constructs were not required since there are no other possible race conditions in the multithreaded CG. It is important to highlight that no special ordering was necessary to achieve parallel performance.

P	ORIG		SAW
	SPMV	CG	CG
1	0.378	9.86	9.74
2	0.189	5.02	5.01
4	0.095	2.53	2.64
8	0.051	1.35	1.36

Table 6: Runtimes (in seconds) for the original and SAW orderings on the Tera MTA.

Results using 60 streams per processor are presented in Table 6. Both CG and the underlying SPMV achieve high scalability of over 90% using the ORIG ordering. This indicates that there is enough thread and instruction level parallelism in CG to tolerate the relatively high overhead of memory access. There is a slight drop in performance between four and eight processors. As we increase the number of processors, the number of active threads increases proportionately while the runtimes become very small. As a result, a greater percentage of the overall time is spent on thread management, causing a decrease in efficiency. Notice that the SAW ordering does not significantly change the performance of CG on this cacheless architecture. Thus, the programming and runtime overheads associated with partitioning/linearization schemes are not required on this platform. We look forward to continuing our experiments as more processors become available on the MTA.

## 5 Summary and Future Work

In this paper, we examined three different parallel implementations of the Conjugate Gradient (CG) algorithm using three leading programming paradigms and architectures. The MPI version of the code on the T3E uses the Aztec [6] library, where we compared the parallel performance of ordering the sparse matrix using reverse Cuthill-McKee (RCM) [3], self-avoiding walk (SAW) [5], and the METIS partitioner [7]. Results showed that all three schemes greatly improve the parallel performance of CG compared to the naive natural ordering. In addition, we demonstrated that traditional graph partitioners, which focus on minimizing edge cuts, are not necessarily the best tools for partitioning sparse matrices on multiprocessor systems. Using RCM or SAW as an ordering/partitioning strategy results in a faster CG than METIS, due to better cache reuse. A performance model was also presented which predicts the expected sparse matrix-vector multiply (SPMV) runtime as a function of both cache misses and communication overhead. Within each CG iteration, the SPMV is usually the most expensive operation.

A shared-memory implementation of CG on the Origin2000 showed that ordering algorithms dramatically improve parallel performance. This is because the Origin2000 is a distributed-memory architecture, so proper data distribution is required even when programming in shared-memory mode. A direct comparison with an MPI implementation indicated that it is possible to achieve message passing performance using shared-memory constructs for this class of applications through careful data ordering and distribution. Finally, results of a multithreaded implementation of CG on the Tera MTA indicated that special ordering and/or partitioning schemes are not required on the MTA to obtain high efficiency and scalability.

Realistic scientific applications require preconditioning algorithms, such as ILU, in order for CG to reach convergence. We plan to investigate the effects of preconditioning on the performance of CG, for various ordering strategies and architectural platforms. We also plan to port the distributed-memory implementation of CG onto the newly installed RS/6000 SP machine at NERSC. This system consists of 256 two-CPU SMP nodes, and is the first commercial implementation of the POWER3 microprocessor. In addition, we will examine the effects of partitioning the sparse matrix using METIS, and subsequently performing RCM or SAW orderings on each subdomain. Combining both schemes should minimize interprocessor communication and significantly improve data lo-

cality. Future research will focus on evaluating the effectiveness of the parallel Jacobi-Davidson eigensolver, when various orderings are applied to the underlying sparse matrix. A multithreaded version of the Jacobi-Davidson algorithm will be implemented on the MTA. We also intend to extend the SAW algorithm to three dimensions and modify it to efficiently handle adaptively refined meshes in a parallel environment.

## Acknowledgements

The work of the first two authors was supported by Director, Office of Computational and Technology Research, Division of Mathematical, Information, and Computational Sciences of the U.S. Department of Energy under contract number DE-AC03-76SF00098. The work of the third author was partially supported by the National Science Foundation under grant numbers NSF-CISE-9726388 and NSF-MIPS-9707125 while the author was at the University of Delaware. The work of the fourth author was supported by National Aeronautics and Space Administration under contract number NAS 2-14303 with MRJ Technology Solutions.

## References

- [1] R. Barrett, M. Berry, T. F. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine, and H. van der Vorst, *Templates for the Solution of Linear Systems: Building Blocks for the Iterative Methods*, SIAM, 1994.
- [2] D. A. Burgess and M. B. Giles, "Renumbering Unstructured Grids to Improve the Performance of Codes on Hierarchical Memory Machines," *Advances in Engineering Software*, Vol. 28, pp. 189–201, 1997.
- [3] E. Cuthill and J. McKee, "Reducing the Bandwidth of Sparse Symmetric Matrices," *Proc. ACM National Conference*, pp. 157–192, 1969.
- [4] M. Griebel and G. Zumbusch, "Hash-Storage Techniques for Adaptive Multilevel Solvers and their Domain Decomposition Parallelization," *AMS Contemporary Mathematics Series*, Vol. 218, pp. 279–286, 1998.
- [5] G. Heber, R. Biswas, and G. R. Gao, "Self-Avoiding Walks over Adaptive Unstructured Grids," *Concurrency: Practice and Experience*, Vol. 12, pp. 85–109, 2000.
- [6] S. A. Hutchinson, L. V. Prevoost, J. N. Shadid, and R. S. Tuminaro, *Aztec User's Guide, Version 2.0 Beta*, Sandia National Laboratories, SAND95-1559, 1998.
- [7] G. Karypis and V. Kumar, "A Fast and High Quality Multilevel Scheme for Partitioning Irregular Graphs," *SIAM J. Scientific and Statistical Computing*, Vol. 20, pp. 359–392, 1998.
- [8] R. Löhner, "Renumbering Strategies for Unstructured-Grid Solvers Operating on Shared-Memory, Cache-Based Parallel Machines," *Computer Methods in Applied Mechanics and Engineering*, Vol. 163, pp. 95–109, 1998.
- [9] C.-W. Ou, S. Ranka, and G. Fox, "Fast and Parallel Mapping Algorithms for Irregular Problems," *J. of Supercomputing*, Vol. 10, pp. 119–140, 1995.
- [10] M. Parashar and J. C. Browne, "On Partitioning Dynamic Adaptive Grid Hierarchies," *Proc. 29th Hawaii International Conference on System Sciences*, pp. 604–613, 1996.
- [11] J. R. Pilkington and S. B. Baden, "Dynamic Partitioning of Non-Uniform Structured Workloads with Space-Filling Curves," *IEEE Trans. on Parallel and Distributed Systems*, Vol. 7, pp. 288–300, 1996.
- [12] Y. Saad, *Iterative Methods for Sparse Linear Systems*, PWS Publishing Company, 1996.
- [13] J. Salmon and M. S. Warren, "Parallel, Out-of-Core Methods for Fast Evaluation of Long-Range Interactions," *Proc. 8th SIAM Conference on Parallel Processing for Scientific Computing*, 1997.
- [14] J. R. Shewchuk, "Triangle: Engineering a 2D Quality Mesh Generator and Delaunay Triangulator," *Applied Computational Geometry: Towards Geometric Engineering*, LNCS Vol. 1148, pp. 203–222, 1996.