OAK RIDGE
NATIONAL LABORATORY

MANAGED BY UT-BATTELLE
FOR THE DEPARTMENT OF ENERGY

ORNL/TM-2003/67

Cray X1 Evaluation

March 2003

Prepared by:

A. S. Bland[1]
J. J. Dongarra[1,2]
J. B. Drake[1]
T. H. Dunigan, Jr.[1]
T. H. Dunning, Jr.[1,2]
A. Geist[1]
B. Gorda[3]
W. D. Gropp[4]
R. J. Harrison[1]
R. Kendall[5]
D. Keyes[6]
J. A. Nichols[1]
L. Oliker[3]
H. Simon[3]
R. Stevens[4]
J. B. White, III[1]
P. H. Worley[1]
T. Zacharia[1]

[1]Oak Ridge National Laboratory
[2]University of Tennessee
[3]Lawrence Berkeley National Laboratory
[4]Argonne National Laboratory
[5]Ames Laboratory
[6]Old Dominion University

UT-BATTELLE
ORNL-27 (2-03)

DOCUMENT AVAILABILITY

Reports produced after January 1, 1996, are generally available free via the U.S. Department of Energy (DOE) Information Bridge.

Web site http:///www.osti.gov/bridge

Reports produced before January 1, 1996, may be purchased by members of the public from the following source.

National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
Telephone: 703-605-6000 (1-800-553-6847)
TDD: 703-487-4639
Fax: 703-605-6900
E-mail: info@ntis.fedworld.gov
Web site: http://www.ntis.gov/support/ordernowabout.htm

Reports are available to DOE employees, DOE contractors, Energy Technology Data Exchange (ETDE) representatives, and International Nuclear Information System (INIS) representatives from the following source.

Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, Tennessee 37831
Telephone: 865-576-8401
Fax: 865-576-5728
E-mail: reports@adonis.osti.gov

This report has prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

ORNL-TM-2003/67

CRAY X1 EVALUATION

A. S. Bland
J. J. Dongarra
J. B. Drake
T. H. Dunigan, Jr.
T. H. Dunning, Jr.
A. Geist
B. Gorda
W. D. Gropp
R. J. Harrison
R. Kendall
D. Keyes
J. A. Nichols
L. Oliker
H. Simon
R. Stevens
J. B. White, III
P. H. Worley
T. Zacharia

Date Published: March 2003

CCS Center for Computational Sciences

Prepared by
OAK RIDGE NATIONAL LABORATORY
P.O. Box 2008
Oak Ridge, Tennessee 37831-6285
managed by
UT-Battelle, LLC
for the
U.S. DEPARTMENT OF ENERGY
under contract DE-AC05-00OR22725

CONTENTS

EXECUTIVE SUMMARY v

1. IMPLEMENTATION AND EVALUATION PLAN FOR THE CRAY X1 1

2. EVALUATION OVERVIEW 3

3. CRAY X1 DESCRIPTION 5

4. EVALUATION COMPONENTS 7

   4.1 MICROBENCHMARKS 7

   4.2 KERNELS AND PERFORMANCE OPTIMIZATION 8

   4.3 APPLICATION AND EVALUATION BENCHMARKING 9

   4.4 CLIMATE SCIENCE: CAM, CLM, AND POP 10

   4.5 FUSION SCIENCE: AORSA, GYRO, GTC, AND M3D 11

   4.6 MATERIALS SCIENCE: LSMS AND PWSCF 12

   4.7 CHEMISTRY: NWCHEM AND GAMESS 12

   4.8 ASTROPHYSICS: CACTUS AND MADCAP 14

   4.9 DELIVERABLES 14

   4.10 I/O AND SYSTEM SOFTWARE EVALUATION AND BENCHMARKING 14

   4.11 SCALABILITY EVALUATION 15

5. RELATION TO DOE SCIENCE PROJECTS 17

   5.1 ACKNOWLEDGEMENTS 17

6. REFERENCES 19

APPENDIX A. WORKSHOPS 21

APPENDIX B. USER FACILITIES 27

EXECUTIVE SUMMARY

This document outlines the plan for evaluating a computer system at the Oak Ridge National Laboratory (ORNL) that has the potential to match or exceed the performance objectives of the U.S. Department of Energy (DOE) science missions and programmatic goals. The Center for Computational Sciences (CCS) at ORNL proposes to partner with Cray, our sister laboratories, and university partners to evaluate and drive the development of the Cray X1 and Black Widow systems to provide systems to significantly enhance high-end computing. The Cray X1 and future generations of this product line promise to substantially increase the fraction of peak performance realized by many of the DOE Office of Science research applications. The efforts described here, in close collaboration with Cray, Inc., will reassert U.S. leadership in high-performance computing for science in strategic areas important to DOE and the Nation, and will greatly enhance high-end computing in the U.S.

1. IMPLEMENTATION AND EVALUATION PLAN FOR THE CRAY X1

This document outlines the plan for evaluating a computer system at the Oak Ridge National Laboratory (ORNL) that has the potential to match or exceed the performance objectives of the U.S. Department of Energy (DOE) science missions and programmatic goals. The Center for Computational Sciences (CCS) at ORNL proposes to partner with Cray, our sister laboratories, and university partners to evaluate and drive the development of the Cray X1 and Black Widow systems to provide systems to significantly enhance high-end computing (Ref. 1). The Cray X1 and future generations of this product line promise to substantially increase the fraction of peak performance realized by many of the DOE Office of Science research applications. The efforts described here, in close collaboration with Cray, Inc., will reassert U.S. leadership in high-performance computing for science in strategic areas important to DOE and the Nation, and will greatly enhance high-end computing in the U.S.

The needs of scientific applications motivate this evaluation. The mainstream of parallel supercomputing in the U.S. has for some time been dominated by clusters of commodity shared-memory processors (SMPs). This technology path provides excellent price/performance when measured from the theoretical peak speed, but many applications of strategic importance (e.g., climate, combustion) sustain only a small fraction of this speed. Indeed, over the past decade, the sustained fraction of peak performance achieved by these applications has been falling steadily. Many factors account for this, but foremost are the increasing disparity between the speed of processors and memory and the design of processors for commodity applications rather than high-end scientific simulations. Illustrative of the alternative path we are evaluating is the Japanese Earth Simulator, which recently attracted attention (including several Gordon Bell prizes at the Supercomputing Conference (SC) 2002, Baltimore, Maryland, November 2002) by sustaining a large fraction of its 40 TFLOP/s peak performance on several scientific applications. This performance advantage was realized by employing relatively few (5,120) powerful (8 GFLOP/s) vector processors connected to high-bandwidth memory and coupled by a high-performance network.

The X1 is the first of Cray's new scalable vector systems that offers high-speed custom vector processors, high memory bandwidth, and an exceptionally high-bandwidth, low-latency interconnect linking the nodes. Cray is the only U.S. manufacturer of multi-terascale systems with these characteristics. The efficiency of the processors in the Cray X1 is anticipated to be comparable to the efficiency of the NEC SX-6 processors in the Earth Simulator on many Office of Science research applications. A significant feature of the Cray X1 is that it combines the processor performance of traditional vector systems with the scalability of modern microprocessor-based architectures. The X1 is the first vector supercomputer designed to scale to thousands of processors with a single system image.

As the latest entry in the competitive, high-end computing marketplace, the Cray X1 must be evaluated for both processor and system performance and for production computing readiness. Our evaluation approach focuses on scientific applications and the computational needs of the science and engineering research communities. Among the goals of this project are establishing and maintaining a public record of benchmarks relevant to DOE missions, and establishing an open discussion in the High Performance Computing (HPC) community aimed at improving supercomputer hardware and software performance for scientific and engineering applications.

Cray's development path scales to nearly 1 PFLOP/s in the latter half of this decade, while significantly improving the price-performance of the technology used. Cray has eagerly solicited our participation in a long-term collaboration focused on making their product even better suited for DOE applications and helping prioritize their activities and investments so as to meet our performance expectation and delivery schedule. This partnership is an exciting opportunity for science to be driving high-performance

computing technology once more, rather than merely consuming commodity products and accepting commodity performance.

The primary tasks of this evaluation are to:

- evaluate benchmark and application performance and compare with systems from other HPC vendors,
- determine the most effective approaches for using the Cray X1,
- evaluate system and system administration software reliability and performance,
- predict scalability, both in terms of problem size and number of processors, and
- collaborate with Cray to incorporate this information and experience into their next generation designs.

While the performance of individual application kernels may be predicted with detailed performance models and limited benchmarking, the performance of full applications and the suitability of this system as a production scientific computing resource can only be determined through extensive experiments conducted on a resource large enough to run datasets representative of the target applications. It is not possible to model fully the interactions of computation, communication, input/output (I/O) and other system services, load balancing, job scheduling, networking, resource allocation, compilers, math libraries, and other system software, along with application software and terabytes of data in memory and on disk. This fact is recognized by the "Applications Performance Matrix" (http://www.krellinst.org/matrix), which focuses upon the actual performance of full applications rather than a few simple kernels.

To improve the performance of both applications and systems, to predict performance on larger or future systems, and to draw concrete conclusions from system and application level benchmarks, a fundamental understanding of the machine and application kernels is also essential. In essence, this evaluation plan will result in a much more detailed "Matrix" for selected applications on the Cray X1 along with detailed interpretation of the raw information, thereby providing a firm basis for future decisions in high-performance scientific computing.

The goal of the evaluation is to predict the expected sustained performance of Cray systems running at tens of TFLOP/s on large-scale simulations, using actual applications codes from climate, fusion, materials science, and chemistry. The codes initially selected for evaluation were identified in discussions with application scientists at workshops held at ORNL in the fall of 2002. Additional workshops with wider application-community participation were held in February and March of 2003 and will be used to build consensus on the codes and evaluation plan. The organization of the evaluation and specific sub-goals are described in more detail below.

2. EVALUATION OVERVIEW

The evaluation will be hierarchical, staged, and open. In the hierarchical approach employed in the investigation of the X1 architecture, the low-level functionality of the system will be examined first. Results will then be used to guide and understand the evaluation using kernels and compact or full application codes. This approach is important because the Cray X1 embodies a number of novel architectural features that make it difficult to predict the most efficient coding styles and programming paradigms. Standard benchmarks will be used when appropriate, and to compare with evaluations of other systems, but the emphasis will be on application-relevant studies, as described later.

A hierarchical approach will also be employed in the examination of performance issues. The first step will be to establish functional correctness. A system incorporating as many novel architectural features as the X1 will be difficult for Cray to test adequately in the accelerated system roll out being proposed. As such, it is crucial to work with Cray to establish correctness. The second step will be to establish performance correctness. Performance expectations are the motivation behind the design and the DOE evaluation of the X1. While many of the performance details are unknowable at this time, gross performance, based on hardware and software specifications, can be used to predict performance in a coarse sense. Verification of these performance expectations is important to identify and correct hardware and software implementation errors. The third step is to determine "best practice" for the X1, that is, what coding styles and programming paradigms achieve the best performance.

Finally, a hierarchical approach will be applied to the evaluation of system scalability. As larger systems are installed, variants of earlier experiments will be repeated that examine performance in the new, larger setting. It will be especially important to re-evaluate functional and performance correctness as the system size increases, as these have been problematic in previous large systems from other vendors.

Performance activities will be staged to produce relevant results throughout the duration of the evaluation. For example, subsystem performance will need to be measured as soon as a system arrives, and measured again following a significant upgrade or system expansion. Fortunately, these experiments can be conducted relatively quickly. In contrast, the porting, tuning, and performance analysis of complex applications, such as those involving unstructured or dynamic data structures, will take much longer to complete, and must be started as soon as possible to have an impact on the evaluation. Allocation of system time and manpower to the different aspects of the evaluation will be staged.

The performance evaluation will also be open. One of the strengths of performance evaluations performed within the Office of Science is the ability to fully disclose the test codes and the results. To ensure the correctness of the approach and interpretation of results, evaluation methodologies and results will be described in public forums and on an evaluation web site for public comment. Moreover, the relevant evaluation communities will be engaged to determine the state-of-the-art evaluation tools and methodologies.

3. CRAY X1 DESCRIPTION

The Cray X1 is an attempt to incorporate the best aspects of previous Cray vector systems and massively-parallel-processing (MPP) systems into one design. Like the Cray T90, the X1 has high memory bandwidth, which is key to realizing a high percentage of theoretical peak performance. Like the Cray T3E, the design has a high-bandwidth, low-latency, scalable interconnect, and scalable system software. And, like the Cray SV1, the X1 design leverages commodity CMOS technology and incorporates non-traditional vector concepts, like vector caches and multi-streaming processors.

The Cray X1 multi-streaming processor (MSP) is capable of 12.8 GFLOP/s and has multiple vector pipes designed to increase vector throughput without increasing the inner-loop vector length. Unlike the SX-6, the X1 keeps the relatively short vector length of 64 elements, thus making it easier for algorithms to efficiently use the vector pipes. The primary strategy for fully utilizing the eight vector pipes of a single MSP is parallelism through outer loops and pipelined operations.

Each MSP has 2 MB of cache memory, and the cache has more than enough single-stride bandwidth to saturate the vector units of the MSP. The cache is needed because the bandwidth to the main memory is not enough to saturate the vector units without data reuse—memory bandwidth is roughly half the saturation bandwidth. This design represents a compromise between non-vector-cache systems, like the SX-6, and cache-dependent systems, like the IBM p690, with memory bandwidths an order of magnitude less than the saturation bandwidth. Because of its short cache lines and extra cache bandwidth, random stride scatter/gather memory access on the X1 is expected to be just a factor of two slower than stride-one access, not the factor of eight or more seen with typical cache-based systems like the IBM p690, Hewlett Packard (HP) Alpha, or Intel IA-64.

It is important to note that the Cray X1's cache-based design deviates from the full-bandwidth design model only slightly. Each X1 processor is designed to have more single-stride bandwidth than an SX-6 processor; it is the yet-higher peak performance that creates the imbalance. A relatively small amount of data reuse, which most modern scientific applications do exhibit, could enable a very high percentage of peak performance to be realized, and worst-case data access could still provide double-digit efficiencies.

Four MSPs and a flat, shared memory of 16 GB form a Cray X1 node. The memory banks of a node provide 200 GB/s of bandwidth, enough to saturate the paths to the local MSPs and service requests from remote MSPs. Each bank of shared memory is connected to a number of banks on remote nodes, with an aggregate bandwidth of roughly 50 GB/s between nodes. This is a remarkable number, for it represents a byte per flop of interconnect bandwidth per computation rate, compared to 0.25 bytes per flop on the Earth Simulator and less than 0.1 bytes per flop expected on an IBM p690 with the maximum number of Federation switch connections. Comparisons with other existing systems are even more unbalanced.

The collected nodes of an X1, eventually containing up to 4096 processors, will have a single system image. A single four-processor X1 node behaves like a traditional SMP, but each processor has the additional capability of directly addressing memory on any other node. Remote memory accesses go directly over the X1 interconnect to the requesting processor, bypassing the local cache. This mechanism is more scalable than traditional shared memory, but it is not appropriate for shared-memory programming models, like OpenMP, outside of a given four-processor node. This remote memory access mechanism is an excellent match for distributed-memory programming models, particularly those using one-sided put/get operations, and it is expected to provide very low latencies and unprecedented inter-processor bandwidths.

In large configurations, the Cray X1 nodes are connected in an enhanced 3D torus, or wraparound mesh. Despite the remarkable bandwidth expected per connection, this topology has relatively low bisection bandwidth compared to crossbar-style interconnects, such as those on the NEC SX-6 and IBM SP.

Whereas, bisection bandwidth scales as the number of nodes, $O(n)$, for crossbar-style interconnects, it scales as the 2/3 root of the number of nodes, $O(n^{2/3})$, for a 3D torus. Despite this theoretical limitation, mesh-based systems, such as the Intel Paragon, the Cray T3E, and ASCI Red, have scaled best to thousands of processors.

The primary benefit of full-bisection interconnects is scheduling flexibility. A single job on a mesh-based topology may achieve better performance if scheduled on contiguous nodes, while a job on a full-bisection network can be scheduled on any available nodes. Jobs on mesh-based systems may also be unable to start despite the availability of enough processors because of fragmentation of those processors. This problem was ameliorated on the Cray T3E through automated job migration; jobs were moved to create larger groups of contiguous processors. Additional benefits of implementing such job migration include system-initiated gang scheduling and checkpoint-restart. These capabilities will also be available in the Cray X1. The main reason to prefer contiguous allocation of processors on the X1 is that it enables improved hardware support for remote address translation.

Because of the tightly coupled parallelism of vector processors, high memory bandwidth, high-bandwidth/low-latency interconnect, and scalable systems software, the Cray X1 has the potential to provide more capability for scientific computation in the near term than any other system available or in late-stage development, including the SX-6. Because of the short vector pipes, multiple vector units, memory caching, and distributed memory of the X1 design, the software optimizations needed for efficiency on the X1 are expected to be similar to the optimizations needed for clusters of traditional SMPs. For example, modern high-frequency commodity processors, like the IBM Power4, have multiple floating-point units with deep pipelines and aggressive memory prefetching. They require moderately large inner loops with no dependencies, just like X1 processors. The expected raw performance and efficiency of the X1 are much higher, however.

4. EVALUATION COMPONENTS

4.1 MICROBENCHMARKS

The objective of microbenchmarking is to characterize the performance of the underlying architectural components of the Cray X1. Both standard benchmarks and customized benchmarks will be used. The standard benchmarks allow component performance to be compared with other computer architectures. The custom benchmarks will permit the unique architectural features of the Cray X1 (distributed vector pipes, and cache and global memory) to be tested with respect to the target applications.

The architectural-component evaluation will assess the following:

- Scalar and vector arithmetic performance, including varying vector length and stride and identifying what limits peak computational performance.
- Memory-hierarchy performance, including cache, local memory, shared memory, and global memory. These tests will utilize both System V shared memory and the SHMEM primitives, as well as UPC and Fortran co-arrays. Of particular interest is the performance of the shared memory and global memory, and how remote accesses interact with local accesses.
- Task and thread performance, including performance of thread creation, locks, semaphores, and barriers. Of particular interest is how explicit thread management compares with the implicit control provided by OpenMP, and how thread scheduling and memory/cache management (affinity) perform.
- Message-passing performance, including intra-node and inter-node Message Passage Interface (MPI) performance for one-way (ping-pong) messages, message exchanges, and aggregate operations (broadcast, all-to-all, reductions, barriers); message-passing hotspots and the effect of message passing on the memory subsystem are of particular interest.
- System and I/O performance, including a set of tests to measure OS overhead (context switches), virtual memory management, low-level I/O (serial and parallel), and network (TCP/IP) performance. System calls made by an X1 application node are implemented by rescheduling the thread for execution on an operating system (OS) node. The performance of this mechanism may have significant implication for the design of I/O- or network-intensive applications. Because of their importance to many application communities, the performance of parallel I/O in HDF5 and in parallel versions of NetCDF will also be studied, as well as the performance of standard MPI I/O benchmarks.
- The microbenchmarking activity will also utilize and evaluate the Cray hardware performance monitors, clocks/timers, and profiling tools.

The deliverables for this aspect of the evaluation are the following:

- results of a standard set of benchmarks and comparisons with other computer architectures,
- performance metrics of the underlying architectural components, including arithmetic operations (scalar and vector), memory hierarchy (cache, local memory, global memory), message passing, task/thread management, and I/O primitives, and
- simple analytical models of component performance.

4.2 KERNELS AND PERFORMANCE OPTIMIZATION

As described previously, the Cray X1 has a number of unique architectural features. Unlike traditional vector machines, the vector pipes are independent within the X1 MSP. As such, the compiler must be able to identify parallelism for "multistreaming" in addition to vectorization, and to exploit memory access locality in order to minimize cache misses and maximize bandwidth. Coding style is anticipated to be as important as the compiler in achieving good single-processor performance. Early indications are that coding styles that perform well on the NEC vector systems do not necessarily perform well on the X1.

Kernels extracted from or representative of the target application codes will be used to examine the performance impact of a variety of coding styles. For example, index and loop orderings and other aspects of the memory-access patterns will be carefully examined. Kernel benchmarking will include an evaluation of the Cray X1 compiler, and the evaluation team will work closely with Cray to identify and resolve problems.

As described earlier, the Cray X1 is a collection of 4-processor, shared-memory nodes with a network that supports globally addressable memory. How best to program for the hierarchical parallelism represented by clusters of SMP nodes is still an area of open research in parallel computing. The addition of both multistreaming and vectorization makes it even less clear which one of the many shared- and distributed-memory programming paradigms are most appropriate for a given application. Cray is currently planning to support both MPI and OpenMP programming paradigms, as well as the SHMEM one-sided communication library and the Co-Array Fortran and UPC parallel programming languages. System V shared-memory, Multi-Level Parallelism (MLP), and Global Arrays are also possible programming models for the Cray X1.

The Cray system has multiple options for both inter-node and intra-node communication. Inter-node communication can be through MPI-1 two-sided, point-to-point primitives, MPI-1 collective communication operators, MPI-2 one-sided messaging, or SHMEM one-sided messaging. The same paradigms can be used for communication within a node, or each can be combined with OpenMP, POSIX threads, or System V shared memory in a hierarchical, hybrid approach. Alternatively, the exact mechanism used for inter-processor communication can be hidden behind the parallel syntax of Co-Array Fortran, UPC, or Global Arrays.

An important task in evaluating the system will be assessing the functionality and measuring the costs associated with each programming approach. Each approach will be optimized separately, for example, in evaluating the many different communication protocols supported by MPI-1 primitives and comparing point-to-point implementations of the collective communication operations with the vendor-supplied implementations. For each approach, the associated peak performance for a set of low-level and kernel codes will be determined.

The goal is not just to choose the best programming paradigm, assuming that one is clearly better than the others, but also to evaluate what can be gained by rewriting codes to exploit a given paradigm. Because of their complexity, many of the important scientific and engineering application codes will not be rewritten unless significant performance improvement is predicted. The evaluation data will also allow us to interpret the fairness of using a given benchmark code implemented using, for example, pure MPI-1 or OpenMP, when comparing results between platforms.

The evaluation of programming paradigms is a complicated task. The work will be staged, optimizing and evaluating the programming paradigms used by the target application codes first, typically hybrid MPI/OpenMP. As performance issues are identified, efforts on the less common approaches may be accelerated. Some codes still have legacy support for SHMEM one-sided messaging, however, and

libraries may be able to take advantage of the more system-specific approaches without requiring modifications to the application codes. Therefore, evaluation of alternative paradigms will be given attention from the very beginning of the evaluation project.

The kernel-based evaluation is driven by the choice of kernels, which is determined in turn by the application codes. Whenever possible, standard kernels will be used, e.g., codes from the NAS or PARKBENCH benchmark suites or routines from the PETSc library, but profile data from the application codes will be the ultimate determiner. For example, it may be necessary to use code fragments extracted directly from the application code in some cases.

The deliverables for this aspect of the evaluation are as follows:

- Evaluation of a variety of programming styles appropriate for use in the target application codes, identifying both those well-suited for the Cray X1 and those that should be avoided.
- Performance benchmarks for each paradigm, and recommendations about how to optimize performance. The benchmarks and optimization analysis will be performed for inter-node and intra-node communications separately, as they are likely to differ even within a common paradigm.
- Cross-paradigm comparisons using application-specific kernels, and paradigm recommendations for application codes. Certain of the application codes, for example, the Community Atmospheric Model (CAM) used in climate research, already supports MPI/OpenMP hybrid implementations. Some of the other paradigms will require significantly more effort to evaluate fairly, and the earlier results will be used to guide this effort, concentrating on the most promising choices.

4.3 APPLICATION EVALUATION AND BENCHMARKING

The performance and efficiency of applications relevant to the DOE Office of Science in the areas of global climate, fusion, materials, chemistry, and astrophysics will be evaluated. In addition to measures of performance and scalability common to evaluations of microprocessor-based MPP systems, the extent to which Cray compilers and tools can effectively vectorize full application codes will be investigated. The extent to which localized tuning can improve vectorization and efficiency will also be investigated.

The evaluation team will work closely with the principal investigators leading the Scientific Discovery through Advanced Computing (SciDAC) application teams to identify the leaders and participants in the application-evaluation efforts. As described above, initial steps in each domain will be the detailed understanding of selected kernels and critical aspects of the system. At the same time, existing implementations of the full applications (described below) will be ported to the machine. The initial emphasis will be upon correct functioning of the code, and, subsequently, upon sequential and then parallel performance and scalability. Maturing compilers and system software will require that this work be performed in close collaboration with Cray. Once a satisfactory port of the existing software has been established, the scientific application evaluation teams will proceed with more detailed examination of the applications and, as necessary, reformulate and re-implement selected algorithms either for performance or for evaluation of different programming paradigms, as described above. Finally, once the installed computer system has reached the necessary size, and the system and the application software has attained the required levels of robustness and performance, these teams will embark upon an extensive evaluation of the computer as a scientific production resource by performing full-scale simulations.

The following applications will be initial targets for evaluation.

4.4 CLIMATE SCIENCE: CAM, CLM and POP

The Community Atmospheric Model (CAM) is the atmospheric component of the Community Climate System Model (CCSM), the primary model for global climate simulation in the U.S. and the target of the climate SciDAC project, "Collaborative Design and Development of the Community Climate System Model for Terascale Computers." The prominent dynamics algorithm of the CAM is a semi-Lagrangian transport method in combination with a semi-implicit Eulerian spectral method. This code grew out of the spectral models developed at the National Center for Atmospheric Research (NCAR).

The CAM will be used to test the Cray X1 in a variety of ways. The dynamics of atmospheric circulation will test the bandwidth and latency of the X1 interconnect. Simulation of the physical processes in each atmospheric column will test memory bandwidth and vectorization capabilities of the compilers. These studies will provide information on the limitations to performance of non-vector code as well as the style of effective vector coding for the X1. The Community Land Model (CLM) is integrated into CAM, and it will also be used to evaluate memory bandwidth and vectorization capabilities. The transfer of data between the CLM and the atmospheric dynamics will evaluate the X1 interconnect.

After initial performance and scaling measurements of CAM, the effectiveness of targeted tuning, such as modifying physics loops for better vectorization and replacing MPI-1 calls within the dynamics with lower latency MPI-2 one-sided or SHMEM calls will be investigated. An effort to explore the resolution, scaling and throughput parameter space to understand how best to utilize the X1 for climate science will be undertaken. Target resolutions will range from the current T42 and T85 climate resolutions to high-end weather resolutions.

The Parallel Ocean Program (POP) is the ocean component of CCSM and is being developed and maintained at Los Alamos National Laboratory (LANL). The code is based on a finite-difference formulation of the three-dimensional flow equations on a shifted polar grid. In its high-resolution configuration, 1/10-degree horizontal resolution, the code resolves eddies for effective heat transport and the locations of ocean currents.

POP is expected to be amenable to vectorization; the evaluation will test this expectation. The two primary processes in POP will test the Cray X1 in different ways. The "baroclinic" process is three dimensional with limited nearest-neighbor communication and should scale well. The "barotropic" process, however, involves a two-dimensional, implicit solver and will limit scalability. This two-dimensional process will test the latency of the X1 interconnect; early POP results from the NEC SX-6 show that latency is a significant bottleneck. The effectiveness of targeted communication optimizations to minimize latency, such as replacing MPI-1 calls with MPI-2 one-sided calls or SHMEM calls will be investigated.

The target resolutions for ocean calculations will be 1/10 degree horizontal grid spacing. This resolution allows for accurate eddy transport of heat by the ocean currents as well as accurate location of features such as the Gulf Stream and other important currents. This is the resolution expected on the Japanese Earth Simulator to drive the next series of scientific discoveries.

The performance of the key kernels, components, and fully configured coupled models will be documented through reports, conference talks and by providing input to the "Matrix." Feedback will be provided to Cray and input will be sought from experts in optimization at Cray and at other research centers.

4.5 FUSION SCIENCE: AORSA, GYRO, GTC, M3D, AND TLBE

The All-Orders Spectral Algorithm (AORSA) codes solve for the wave electric field and heating in a stellerator plasma heated by radio-frequency waves. They are important applications in the fusion-sciences SciDAC project, "Numerical Calculation of Wave-Plasma Interactions in Multidimensional Systems." The computation times of AORSA2D and AORSA3D are dominated by the use of ScaLAPACK to solve large, dense systems of linear equations. ScaLAPACK shows good efficiency on many computer systems, and the same is expected on the X1. AORSA results from the NEC SX-6 show excellent efficiency using ScaLAPACK, but the results reveal that the matrix generation vectorizes poorly and requires a significant amount of time. The efficiency of the X1 for matrix generation will be evaluated.

Current AORSA development is targeting algorithms that trade smaller linear systems, the solutions of which scale as the cube of the problem size, for more expensive matrix generation, that scales as the square of the problem size. The evaluation will help determine the optimal tradeoff for the X1.

The Gyro code solves time-dependent, nonlinear gyrokinetic-Maxwell equations for electrons and ions in a plasma. It is being developed under "The Plasma Microturbulence Project," a fusion SciDAC project. Gyro uses a five-dimensional grid and propagates the system forward in time using a fourth-order, explicit, Eulerian algorithm. This application has shown good scalability on large microprocessor-based MPPs, and similar scalability is expected on the X1. The extent to which this scalability is enhanced by greater per-processor efficiency will be evaluated.

The Gyrokinetic Toroidal Code (GTC) is a three-dimensional particle-in-cell simulation code for microturbulence studies in magnetically confined plasmas. It is also supported under "The Plasma Microturbulence Project." The code solves the nonlinear Gyrokinetic Vlasov-Maxwell system of equations using particle-in-cell methods for the dynamic equations and iterative (including multi-grid) methods for the elliptic field equations, with MPI and OpenMP parallelization. The code has been running on the IBM SP at National Energy Research Scientific Computing (NERSC) using from 64 to more than 2000 processors, with a parallel efficiency of up to 98% and with only 5% of the computing time spent for inter-processor communications. There is strong interest in the evaluation of the performance of the X1 processors for gather and scatter (random access) operations for GTC.

M3D is a code for simulating the magneto-hydrodynamics (MHD) of fusion plasmas in three dimensions, and it is being developed under the SciDAC "Center for Extended Magnetohydrodynamic Modeling." The parallel version uses ideal, resistive, and two-fluid MHD models, at lower granularity but much longer simulation times than gyrokinetic models. M3D uses finite differences in the radial direction and FFTs in the toroidal and poloidal directions, and PETSc provides the elliptic solver for its quasi-implicit time-integration method. Early results on the SX-6 from the developers at Princeton have raised issues about the performance of PETSc on vector systems. The M3D evaluation will provide an opportunity to resolve these issues on the Cray X1.

Lattice Boltzmann methods provide a mesoscopic description of the transport properties of physical systems using a linearized Boltzmann equation. They offer an efficient way to model turbulence and collisions in a fluid. The TLBE application performs a 2D simulation of high-temperature plasma using a hexagonal lattice and the BGK collision operator. The TLBE simulation has three computationally demanding components: computation of the mean macroscopic variables (integration); relaxation of the macroscopic variables after colliding (collision); and propagation of the macroscopic variables to neighboring grid points (stream). The first two steps are floating-point intensive, while the third consists of data movement only. The problem is well suited for vector architectures.

4.6 MATERIALS SCIENCE: LSMS, PWSCF, AND PARATEC

The Locally Self-consistent Multiple-Scattering (LSMS) code uses a real-space, multiple-scattering, Green-function-based method for calculating the electronic structure of materials and for treating the quantum mechanical interactions between large numbers of atoms. LSMS was the first computer model that accurately captured the magnetic interactions responsible for the formation and stabilization of the local magnetic moments observed in neutron scattering data for CuNi, and it has continued to be extensively used for first-principles materials research. The code scales linearly in the number of atoms, and it is remarkably efficient and scalable on microprocessor-based MPPs. Record-setting performance using LSMS resulted in the award of the 1998 Gordon Bell Prize; the code should approach the theoretical peak performance of the X1.

The single-processor efficiency of LSMS relies on its use of BLAS3 dense-matrix operations. To scale LSMS to much larger problems, the developers are moving to sparse-matrix formulations, which typically achieve significantly lower efficiency on microprocessor-based systems. The relative advantages of the Cray X1 for these sparse formulations in large number of atom configurations will be evaluated. This work on sparse formulations will be coordinated with the work on sparse solvers in PETSC described earlier.

Complementary to the LSMS method, the evaluation will include the Plane-Wave Self-Consistent Field (PWscf) code, a leading implementation of psuedopotential algorithms for electronic-structure calculations. PWscf performs a variety of calculations in density-functional theory and density-functional perturbation theory. The developers listed support for both scalar and vector systems, along with parallelism through MPI. PWscf is open source and sees prominent use both within the DOE materials-science community and abroad; it is used by the French Atomic Energy Commission (CEA) for benchmarking.

PARATEC (PARAllel Total Energy Code) performs first principles quantum mechanical total energy calculations using pseudopotentials and a plane wave basis set. The approach is based on density functional theory (DFT) that has become the standard technique in materials science to accurately calculate the structural and electronic properties of new materials with a full quantum mechanical treatment of the electrons. Codes performing DFT calculations are some of the largest consumers of computer cycles in centers around the world with the plane wave pseudopotential approach being the most commonly used. Both experimental and theory groups use these types of codes to study properties such as strength, cohesion, growth, catalysis, magnetic, optical, and transport for materials like nanostructures, complex surfaces, doped semiconductors, and others.

4.7 CHEMISTRY: NWCHEM AND GAMESS

NWChem is the DOE's massively parallel computational-chemistry code, and it is the first such code designed for efficient parallel execution. NWChem is supported by the Office of Biological and Environmental Research as part of the operation of the Environmental Molecular Sciences Laboratory at Pacific Northwest National Laboratory (PNNL). The code provides extensive functionality for determining the electronic structure of molecules and solids, including Gaussian-based Hartree-Fock and density functional theory (DFT) with analytic second derivatives, many-body methods such as MP2 (with gradients) and CCSD(T), and plane-wave DFT with a variety of pseudo-potentials and boundary conditions with Car-Parinello dynamics. Also included is an extensive classical molecular-dynamics (MD) and free-energy capability that may be the most efficient such code currently available. The quantum and classical models may be combined to perform mixed quantum-mechanics/molecular-mechanics (QM/MM) simulations. NWChem has demonstrated efficient scaling to over 1000 processors

(MD on a Cray T3E), and the new CCSD(T) capability has been designed for efficient execution up to 10,000 processors.

NWChem is a large code (nearly 1 million lines) with a wide variety of computational kernels that vary in the demands they place upon the underlying hardware. For example, the Gaussian-based DFT and the MD codes are sensitive to the latency of remote memory access, the CCSD(T) code requires efficient local matrix multiplication, and the MP2 gradients and several other modules require high-performance I/O with both sequential and random access.

In contrast to the other applications codes in the evaluation plan, NWChem uses a distributed, shared-memory programming model supported by the Global Arrays library. Point-to-point message passing is only used in third-party libraries, such as for parallel linear algebra or FFTs. The NWChem algorithms are typically based upon a non-uniform memory-access model. Shared data (in a global array) are assumed to be accessible by any process without the explicit involvement of any other process. This one-sided access is critical to realizing the efficiency and ease of programming demonstrated in NWChem. Shared memory is assumed to be more expensive to access than local memory (higher latency and lower bandwidth), and the algorithms have been designed to offset this cost by a variety of mechanisms, including loop tiling and blocking. The MD module uses spatial decomposition but still employs the Global Arrays library, rather than message passing, because the one-sided access mechanisms improve the scalability of the code through increased asynchrony of execution and easier dynamic load balancing.

A plan has been formulated for the PNNL NWChem team (led by Theresa Windus), the Global Arrays team (led by Jarek Nieplocha), and ORNL staff (led by Robert Harrison) to port and optimize NWChem for use on the X1. Because of the size and complexity of the code, this is a multi-stage project. The initial emphasis will be on porting and optimizing the supporting libraries, optimizing critical kernels, and demonstrating correct sequential and parallel functioning, including passing the full NWChem quality-assurance suite. Subsequently, attention will focus upon improving the performance of selected important modules, the overall parallel scalability, and exploring critical aspects of system performance of specific interest to NWChem (e.g., I/O and distributed, shared memory). The MD and two-electron integral modules were initially developed for Cray vector systems, and it is anticipated that they retain many favorable performance characteristics despite subsequent modifications for microprocessor-based machines. Most of NWChem has never been tuned for a vector computer; however, so much work may be necessary to reach a satisfactory base-line performance. In our favor is a wealth of experience in vector computing in computational chemistry, and, because of this experience, it is expected that nearly all of NWChem will eventually attain high performance on the X1.

As a complement to NWChem, the evaluation will also include the General Atomic and Molecular Electronic Structure System (GAMESS), an *ab initio* quantum-chemistry code that has been developed by Mark Gordon's group since 1982. It is being developed further under the chemistry SciDAC project, "Advancing Multiple-Reference Methods in Electronic Structure Theory." GAMESS is in use at well over 5,000 sites worldwide.

An important feature of GAMESS is that an increasing number of its modules can be run on scalable computers, ranging from clusters of low-cost commodity systems, to clusters of high-end workstations using advanced switching technology, to MPPs. The scalability of GAMESS is facilitated by the distributed data interface (DDI), developed by Graham Fletcher, which allows large data arrays to be distributed across all available nodes. Scalable modules in GAMESS that make use of DDI include Hartree-Fock, MCSCF, density-functional theory, closed-shell second-order perturbation theory (MP2) energies and gradients, multi-reference MP2 energies, and RHF Hessians.

4.8 ASTROPHYSICS: CACTUS AND MADCAP

One of the most challenging problems in astrophysics is the numerical simulation of Einstein's equations in order to explore the Theory of General Relativity (GR). These equations are among the most complex in physics; a set of coupled nonlinear hyperbolic and elliptic equations containing thousands of terms when fully expanded. The Albert Einstein Institute in Potsdam, Germany, developed the Cactus code to evolve these equations stably in 3D on supercomputers to simulate astrophysical phenomena with extremely high gravitational fluxes, such as the collision of two black holes and the gravitational waves that radiate from that event.

For the purpose of performance evaluation, we focus on two different implementations of the core Cactus ADM solver. The first benchmark application uses the older Fortran77-based ADM kernel (BenchADM) that dates back to a time when vector machines were more common; consequently, we expect it vectorize well. The second Cactus ADM kernel (BenchBSSN) is a newer solver and written using F90 syntax. BenchBSSN was developed specifically for microprocessor-based architectures; hence, it has many features such as more intense use of conditional statements in inner loops that may degrade performance on vector architectures.

MADCAP is a package that analyzes Cosmic Microwave Background Radiation datasets by recasting the treatment as a problem in dense linear algebra. The goal is to take any time-ordered CMB dataset, construct the maximum likelihood pixelized sky temperature map, and from this extract the maximum likelihood angular power spectrum. The algorithms employed scale with the number of pixels cubed in flop count and squared in memory requirement. To date we have managed to analyze datasets of up to 100,000 pixels on NERSC's 3000 processor IBM SP, but existing and forthcoming datasets will increase this by up to 2 orders of magnitude.

4.9 DELIVERABLES

The deliverables for application evaluation and benchmarking are the following:

- performance results for each application, with comparisons to other systems,
- performance analyses, describing strengths and weaknesses of the X1 relative to other systems, and
- results of tuning experiments, identifying effective tuning techniques and expectations for performance improvement. The performance results will be documented through reports, conference talks, and input to the "Matrix." These public forums will allow continuous monitoring of the status of the evaluation project and status of the X1.

4.10 I/O AND SYSTEM SOFTWARE EVALUATION AND BENCHMARKING

The correctness and usability of the system software are essential components in providing a fully functioning computer system. Large systems suitable for both development and production use make significant demands on this software. The system software of the Cray X1 provides a variety of unique features, including a single system image, checkpoint and restart, and job migration.

The job-migration feature is a necessity on the X1 because of the requirement that a job must run on a contiguous set of processors to use the X1 interconnect most effectively. Typical job mixes will create distributed gaps, such that enough processors may be available for a given job, but the job cannot start because the processors are not contiguous. With job migration, jobs in progress may be moved to

different processors to create larger holes. The integration of this feature with the job-management software and its effectiveness at maximizing utilization of the full system will be evaluated.

The single system image will make many aspects of system administration much easier and more efficient than on clusters of SMPs. It has important implications on the scalability and fault tolerance of the system as a whole, however. With large, parallel systems, the ultimate performance of the system can be limited by the reliability and serviceability of the hardware and software. If the system fails frequently or takes a long time to repair, then applications are stalled.

The mean time between failure (MTBF) and mean time to repair (MTTR) for the Cray X1 will be determined for the ORNL environment. The system tools for installing and updating system software, error detection and reporting capabilities of the system, and system monitoring tools will be evaluated. An evaluation of the resilience of the full system to failures of individual components will be performed with the goal helping Cray to improve the tolerance of the full system to the most common component and system-software failures.

In addition to unique system-software features, the Cray X1 has a unique I/O architecture. All I/O travels over Fiber Channel (FC) links, with "local" disks attached through FC controllers, and external networks attached through servers that translate from Gigabit Ethernet (GigE) to IP over FC. Configuration variables include the numbers of FC links, FC controllers, RAID arrays, disks per RAID, networking servers, and GigE links. Additional options come from tuning the parameters of XFS, the high-performance file system delivered with the X1, parameters such as number of stripes, stripe width, and block size. In addition to the I/O load from application codes themselves, checkpoint and restart can impose extreme requirements on file systems. The file systems used in the evaluation will be configured and tuned in different ways to look for anomalies and performance problems. This work complements the more application-specific evaluations described above.

Finally, both performance of individual codes and total throughput of the system depend heavily on how well the machine integrates into the total HPC environment. How the Cray system links with external file systems, such as HPSS and NFS, data-analysis servers, and high-bandwidth wide-area networks will be investigated. Again, the Cray X1 has unique configuration variables, including the number of FC links and FC-to-GigE servers.

The deliverables for the I/O and systems software aspect of the evaluation are:

- functionality and performance of the job-management systems,
- MTBF and MTTR statistics for the X1, as well as a description of all problems and Cray's responses,
- functionality and performance of the scalability and fault tolerance of the system,
- data showing the performance of XFS under a variety of configuration and usage patterns, including checkpoint and restart, and
- recommended tuning parameters and hardware configurations for connecting the Cray X1 to HPSS, NFS, and high-bandwidth wide-area networks.

4.11 SCALABILITY EVALUATION

Determining the scalability of a system from a given instantiation of an architecture is not a simple exercise. A simpler approach is to attempt to identify system components whose performance will adversely affect scalability.

As part of this latter approach, extensive studies of the message-passing layers of the X1 will be conducted, investigating performance behavior under a range of load conditions. Tests will be used that mimic common behavior, such as all-to-all communication, and that create hotspots, with all processors contending for a common resource. Similar studies will be performed for the I/O system.

For example, the out-of-core extension of ScaLAPACK to stress the parallel I/O system will be used. This software is written to use either multiple files local to each processor (when run on a cluster) or a large shared file on a parallel file system. It solves systems of linear equations and eigenvalue problems that are several times larger than the total available real memory. Large blocks (hundreds of megabytes) of data are transferred in a regular pattern between disks and memory. Since nearly all-available memory is dedicated for computation, it is unlikely that disk caching will be effective.

The emphasis here is different from benchmarking, although the evaluations are similar. The behavior will be measured when scaling problem size, number of processors, or other hardware or software parameters. Then trends are sought that, if extrapolated, would cause performance difficulties for larger systems. If such trends are discovered, the cause will be determined, and whether it is an impediment to scalability will be ascertained. Corrective action for problems identified will be worked with Cray.

Once obvious showstoppers have been eliminated, an attempt will be made to extrapolate performance for well-understood kernels and application codes. This requires codes for which performance models have been constructed. Fortunately, there are a number of codes for which such models exist or for which performance can be bounded from below by characterizations of message-passing behavior or memory pattern accesses. By using a range of problem sizes, these models can be parameterized, validated against existing empirical data, and then used to estimate performance on a larger system. While the quantitative accuracy of the predictions will be difficult to assess, scaling problems indicated by this analysis should be accurate. The advantage of using models in this study is that the cause of poor scalability will be obvious from the model.

The deliverables for this aspect of the evaluation are hotspot analyses for:

- inter-node and intra-node communication,
- shared-memory access, and
- parallel I/O,

as well as:

- trend analyses for both standard and optimized versions of selected communication and I/O patterns,
- trend analyses for standard and application-specific parallel kernels, and
- predicted scalability (or lack thereof) of selected kernels and application codes from performance models and lower bounds.

5. RELATION TO DOE SCIENCE PROJECTS

DOE's science projects—in accelerator physics, biological science, chemical science, environmental science, fusion science, materials science, nuclear physics, and particle physics—will be direct beneficiaries of the knowledge gained from this evaluation project. Initially, selected participants from the major SciDAC scientific-application projects will be involved in the evaluation of the Cray X1 either by actively working with the X1 or by acting as a liaison between the evaluation team and the broader DOE science community. In addition, since the evaluation plan addresses questions of performance and programming paradigms using selected benchmarks from each application, scientists will discover firsthand the impact of the X1 on their applications and will share this information with colleagues in other disciplines. The critical performance issues explored as part of this evaluation plan, including the discovery of specialized coding and optimization techniques, will result in a clearly defined path of code modifications and design decisions needed to adapt the SciDAC scientific applications to the X1 architecture.

Profound implications for the tools and methods projects of the SciDAC Integrated Software Infrastructure Centers will also result from the evaluation process, since vector architectures have been mostly neglected for the past decade in the U.S. Without representative hardware, most computer-science research and tool development has targeted commodity scalar processors and clusters. Methods research has pushed the scalability issues toward massive, distributed-memory configurations. While much of the research will be relevant and applicable to a Cray X1 with thousands of processors, the richness of the Cray software environment and high-bandwidth interconnect will shift the balance. This evaluation will give important information to tool designers and algorithm researchers about the relevance of current designs and implementations.

Finally, as this evaluation includes aspects of the production environment supplied by Cray, it will be of great interest to computational scientists throughout the DOE laboratory complex and to the broader scientific community. Reliability, robustness, and ease of use are relevant questions for those intending to use the X1 as an "extraordinary tool for extraordinary science." This evaluation will determine how to effectively manage large scientific calculations and extract information and eventually knowledge. This is the foundation of scientific discovery through advanced computing.

5.1 ACKNOWLEDGEMENTS

The authors wish to thank Mark Gordon (Ames Laboratory), William Tang (Princeton Plasma Physics Laboratory), and Warren Washington (National Center for Atmospheric Research) for reading and commenting on the evaluation plan, and to the attendees of the application workshops for their help in identifying the important application codes. The authors also wish to thank C. Hamby for preparation of the manuscript and D. McCoy for editorial assistance.

6. REFERENCES

1. T Zacharia, A. Geist, S. Scott, et al., "Reasserting U.S. Leadership in Scientific Computation," July 4, 2002.

APPENDIX A. WORKSHOPS

This evaluation plan was the outcome of a number of meetings with both HPC system vendors and application experts over the past 9 months. Highlights from several of these meetings are given below.

CLIMATE MODELING SYSTEM WORKSHOP

The Climate Modeling System Workshop was held at IBM Research in Yorktown Heights, June 10-12, 2002, and included representatives from IBM, ORNL, and NCAR. The topics of discussion were the computational requirements of the next generation of climate model, along with what hardware and system configurations and performance would be sufficient to meet those needs. Attendees are listed below by organization.

IBM	NCAR	ORNL
Jim Abeles	Tony Craig	David Bernholdt
Ravi Arimili	Al Kellie	Buddy Bland
Carol Crothers	Rich Loft	Ed D'Azevedo
Mark Dean		John Drake
Kevin Gildea		Tom Dunigan
Anshul Gupta		Trey White
Taffy Kingscott		Thomas Zacharia
Dave McQueeney		
Jamshed Mirza		
Pretap Pattnaik		
Bill Pulleyblank		
Craig Stunkel		
Lloyd Treinish		
Dick Treumann		
Bob Walkup		

ORNL CRAY X1 TUTORIAL

The ORNL Cray X1 Tutorial was held on November 7, 2002 at ORNL (**http://www.ccs.ornl.gov/CrayX1/Tutorial/index.html**). ORNL staff described the acquisition and evaluation plans, and Cray staff made presentations on the X1 hardware, tools, performance, and optimization strategies. Attendees are listed below by project or organization.

Name	Affiliation

Application Projects

Genome to Life: Mass-Spec Data Analysis (**Biology**)
Fridman, Tema JICS/UT/ORNL

Large Scale Genome Annotation (**Biology**)
LoCascio, Phil ORNL

Advanced Methods for Electronic Structure (**Chemistry**)
Harrison, Robert ORNL

A Geodesic Climate Model with Quasi-Lagrangian Vertical Coordinates (**Climate**)

Dazlich, Donald Anthony, Jr.	Colorado State
Fowler, Laura	Colorado State

Collaborative Design and Development of the Community Climate System Model (**Climate**)

Branstetter, Marcia	ORNL
Drake, John	ORNL
Hernandez, Jose	ORNL
Hoffman, Forrest	ORNL
Jones, Phil	LANL
Taft, Jim	LANL

CEMM: Center for Extended Magnetohydrodynamics Modeling (**Fusion**)

Chen, Jin	PPPL
Jardin, Stephen C.	PPPL

Collisional and Turbulent Transport in Fusion Plasmas (**Fusion**)

Spong, Don	ORNL

Numerical Computation of Wave-Plasma Interactions in Multi-dimensional Systems (**Fusion**)

Hirshman, Steve	ORNL
Jaeger, Fred	ORNL
Luszczek, Piotr	UT

Plasma Microturbulence Project (**Fusion**)

Ethier, Stephane	PPPL

Advanced Computing for 21st Century Accelerator Science and Technology (**HENP**)

Guetz, Adam	SLAC

TSI: Terascale Supernova Initiative (**HENP**)

Cardall, Christian Y.	ORNL/UT
D'Azevedo, Ed	ORNL
Eijkhout, Victor	UT/ICL
Hayes, John	UCSD
Hix, William	UT/ORNL
Messer, Bronson	UT/ORNL
Mezzacappa, Tony	ORNL
Myra, Eric	SUNY at Stony Brook
Rao, Nagi	ORNL
Saied, Faisal	NCSA/UIUC
Stoitcheva, Gergana	ORNL
Strayer, Michael	ORNL

Magnetic Materials Bridging Basic and Applied Science (**Materials**)

Eisenbach, Markus	ORISE/ORNL
Janotti, Anderson	ORNL
Stocks, Malcolm	ORNL
Ujfalussy, Balazs	UT/ORNL

Computer Science, Middleware, and Mathematics Projects

CCA: Common Component Architecture

Bernholdt, David	ORNL
Kohl, Jim	ORNL
Leininger, Matt	Sandia

CCTTSS: The Center for Component Technology for Terascale Simulation Software

Elwasif, Wael	ORNL

Center for Programming Models for Scalable Computing

Apra, Edo	EMSL - PNNL
Tipparaju, Vinod	PNNL

DOE Science Grid and Earth System Grid

Chanchio, Kasidit	ORNL

DOE Science Grid

Skouson, Gary	PNNL

Evaluation of Modern Parallel Vector Systems

Oliker, Lenny	LBNL

Fundamental Algorithms for Advanced Application in Science and Technology

Fann, George	ORNL

PERC: Performance Evaluation Research Center

Dunigan, Tom	ORNL
Gao, Xiaofeng	UCSD/SDSC
Hollingsworth, Jeff	Univ. Maryland
Hovland, Paul D.	Argonne
McCracken, Michael O.	UCSD/SDSC
Mendes, Celso	UIUC
Moore, Shirley	UT
Norris, Boyana	Argonne
Strohmaier, Erich	NERSC/LBNL
Shelton, Bill	ORNL
Terpstra, Dan	UT
Vetter, Jeff	LLNL
Worley, Pat	ORNL
You, Haihang	UT
Zhang, Ying	UIUC
Zhou, Min	UT

Scientific Data Management

Burris, Randy	ORNL
Thakur, Rajeev	Argonne

TOPS/PERC/TSI

Dongarra, Jack	UT

TOPS: Terascale Optimal PDE Solvers

Balay, Satish	Argonne
Gropp, Bill	Argonne
Grigori, Laura	LBNL
Keyes, David E.	ODU
Tarafdar, Arijit	ODU

TSTT: Terascale Simulation Tools and Technology Center

Oh, Wonho	Brookhaven
Trease, Harold	PNNL

Scalable Systems Software

Geist, Al	ORNL
Lusk, Rusty	Argonne
Mugler, John	ORNL
Naughton, Thomas	ORNL

Other

ASCI Applications for Sandia

Cole, Ben	Sandia

Center for Computational Sciences

Alexander, Richard	ORNL
Bland, Buddy	ORNL
Fahey, Mark	JICS/UT/ORNL
Fahey, Rebecca	ORNL
Halloy, Christian	JICS/UT/ORNL
White, James B., III (Trey)	ORNL
Wong, Kwai	JICS/UT/ORNL

Computational Facility for Reacting Flow Science

Ray, Jaideep	Sandia

Computer Science and Mathematics Division

Nichols, Jeff	ORNL

Cray X1 Applications Tuning

Hoelzeman, Larry	Cray Inc.
Kiefer, Dave	Cray Inc.
Levesque, John	Cray Inc.
Maschhoff, Kristyn	Cray Inc.
Perry, Steve	Cray Inc.
Schwarzmeier, Jim	Cray Inc.
Wichmann, Nathan	Cray Inc.

Earth Simulator Evaluation

Foley, Dennis	NCI-Frederick

Joint Institute for Computational Sciences
Dunning, Thom ORNL

Materials Science
Meunier, Vincent ORNL
Smirnov, Andrei ORNL

Multiple SciDAC projects
Johnson, Fred DOE
Johnson, Gary DOE
Office of Science
Turner, Dave Ames Laboratory

Other
de Jong, Bert PNNL

CRAY-FUSION ULTRASCALE COMPUTING WORKSHOP

The Cray-Fusion UltraScale Computing Workshop was held at the Garden Plaza Hotel in Oak Ridge, Tennessee, February 3-5, 2003 (**http://www.csm.ornl.gov/meetings/fusion2.html**). Topics of discussion were the characterization of models/codes in terms of physics objectives and models; characterization of performance extensions needed for making significant advances; identification of code/code kernel/algorithm for analysis on a range of platforms; and development plans for carrying out this task.

DOE SCIDAC WORKSHOP: PORTING CCSM TO THE CRAY X1

The DOE SciDAC Workshop: Porting CCSM to the Cray X1, was held at the National Center for Atmospheric Research in Boulder, Colorado, February 6, 2003 (**http://www.csm.ornl.gov/meetings/climate2.html**). The goal of the workshop was to identify and coordinate efforts in the porting of the Community Coupled System Model (CCSM) to the Cray X1, including discussions of vectorization and software engineering issues.

COMPUTATIONAL MATERIALS SCIENCE: EARLY EVALUATION OF THE CRAY X1

The Computational Materials Science: Early Evaluation of the Cray X1 was held at the Hyatt Regency in Austin, Texas, March 2, 2003 (**http://www.csm.ornl.gov/meetings/materials2.html**). The goals of the workshop were to follow up on ultra-simulation initiative white papers; provide a prioritized list of application codes that will be ported to the Cray X1; and provide a list of names and research projects affiliated with these codes.

APPENDIX B. USER FACILITIES

Oak Ridge National Laboratory is the home of 18 highly sophisticated experimental user facilities including the Center for Computational Sciences. These research laboratories are designed to serve not only our staff scientists and engineers, but also researchers from universities, industry, foreign institutions, and other government laboratories. They simultaneously advance national research and development and fulfill the U.S. Department of Energy (DOE) missions by minimizing unnecessary duplication of effort, promoting beneficial scientific interactions, and making the most effective use of costly and, in many cases, unique equipment.

The diverse and sophisticated research conducted by our staff scientists, coupled with the availability of unique resource equipment, is attracting a growing number of guest researchers. In FY 2001, there were more than 4,000 researchers from over 380 organizations that utilized the user facilities. The organizations with user agreements in place are listed below. The most recent agreement is with Cray.

A. Finkl & Sons, UC-96-074
AAB Building Systems, Inc., UR-98-111
ABB C-E Services, Inc., UC-96-083
Activated Metals & Chemicals, UC-02-221
Adtech Nepth Inc., UC-96-145
Advanced Ceramics Corporation, UC-98-011
Advanced Ceramics Research, UC-93-057
Advanced Energy, UR-01-105
Advanced Engineered Materials, LLC, UC-95-091
Advanced Refractory Tech, Inc., UA-93-011
Advanced Technology Materials, Inc., UC-96-023
AECL Research, UA-91-053
Agriboard Industries, UR-97-101
AHT, Inc., UC-97-093
AKZO Nobel Chemicals, Inc., UC-99-148
Alabama A&M University, UA-93-056
Alcoa Technical Center, UA-91-045
Alcoa, Inc., UC-02-2024
AlliedSignal Inc., Ceramic Components, UA-88-013
Allison Engine Company, UA-90-001
Alloy Engineering & Casting Company, UC-98-005
ALLVAC, UC-02-192
Alpha Optical Systems, UA-92-046
Aluminum Company of America, UA-90-013
Alvord-Polk, Inc., UC-02-155
Alzeta Corporation, UA-92-001
Amercom Inc., UA-91-076
AMERCORD, INC., UC-98-121
American Boarts Crushing Company, Inc., UC-00-190
American Iron and Steel Institute, UR-99-116
American Magnetics, Inc., UR-98-021
American Matrix, Inc., UA-87-007
American Superconductor Corporation, UA-90-026
Ametek Specialty Metals, UC-02-199
Anaconda Foundry Fabrication Co., Inc., UC-97-024

Aoyama Gakuin University, UA-98-096
Appalachian State University, UA-95-127
Appliance Industry/ Government CFC Replacement Consortium Incorporated, UR-00-035
Applied Nano Metrics, Inc., UC-98-101
Applied Thin Films, Inc., UC-02-239
ARCO Aluminum, Inc., UC-98-039
ARDE, Inc., UC-99-004
Arizona State University, UA-94-032
Army Research Lab, UR-97-053
Atlantic Research Corporation, UC-00-009
Auburn University, UA-88-003
Australian Geological Survey Organization, UC-96-115
Australian National University, UA-97-055
Australian Nuclear Science and Technology Organization, UC-97-122
B&W Nuclear Technologies, UA-94-033
Babcock & Wilcox / McDermott Technology, Inc., UA-91-092
Belcan Corporation, UC-01-103
Berea College, UA-92-032
Bethlehem Steel Corporation, UC-96-150
Bicron-NE, UR-99-136
BMR Geology and Geophysics (Australia), UA-91-091B
Bosch Braking Systems, UC-98-002
Boston University - LMES, YNP-94-018
Brigham & Women's Hospital, UC-99-128
Brown University, UA-89-025
Browne Technology, Inc., UC-97-099
BTR Sealing Systems, UC-99-165
Busek Co. Inc., UC-97-085
Calgon Corporation, UC-96-162
California Institute of Technology/Jet Prop. Laboratory, UA-95-125
California State University at LA, UA-91-014
Camel Manufacturing Company, UR-03-245
Canadian Office of Public Works and Government Services, UR-98-090
CarboMedics, Inc., UA-89-031
Carborundum Company, UA-88-024
Carlisle Syntec, Inc., UR-98-082
Carnegie Mellon University, UA-92-088
Carpenter Technology Corporation, UC-02-216
Carroll Kenneth Johnson, UC-97-143
Case Western University, UA-92-093
Caterpillar, Inc. (Technology Center), UA-92-008
Cavin Consulting Services, UA-94-006
CC Metals and Alloys, Inc., UC-02-188
CDH Energy Corporation, UR-99-184
CEA, UA-95-069
Central Research Institute of Electric Power Industry, UC-99-180
Ceradyne, Inc., UC-94-011
Ceramaspeed, Inc., UC-02-240
Ceramic Magnetics, Inc., UA-92-083
Ceramics Process Systems Corporation , UA-88-018
CERCOM, Incorporated, UA-94-014
Certainteed Corporation, UA-89-024

Certech, Inc., UC-98-063
Chand Kare Technical Ceramics, UA-93-014
Chand Metallurgical, UC-02-165
Charles University , UA-97-062
Chiroscience R&D Inc./Darwin Molecular, UR-97-140
Chrysler Corporation, UA-90-037
Church and Dwight Company Inc., UA-91-059
Cincinnati Machine Company, UC-99-129
Citation Corporation, UC-98-119
Clark University, UA-97-031
Clarkson University, UA-93-051
Clemson University, UA-88-009
CMP Industries, Inc., UC-96-027
Colorado School of Mines, UA-96-098
Columbia Falls Aluminum Company, UC-97-004
Commonwealth Aluminum, UC-02-212
Component Composite Coatings International LLC, UC-02-200
Computational Mechanics Corporation, UR-95-070
Concurrent Technologies, Inc., UA-92-002
Congoleum, UC-99-009
Constellation Technology Corporation, UR-01-131
Coors Ceramics Company, UA-89-009.1
Cornell University, UA-90-031
Corning Inc., UA-92-037
Cray, Inc., UF-03-277
Cree Research, Incorporated, UC-94-079
CTI, Inc., UC-99-151
Ctr Nat'l dela Recherche Scientifique, UA-91-098
Cummins Engine Company, Inc., UA-90-004
Dallas Optical Systems, Inc., UC-01-107
Dana Corporation, UC-96-107
Daresbury Laboratory, UC-96-158
Dartmouth College, UA-87-003
Deere & Company, UC-02-156
Deformation Control Technology, Inc., UC-98-122
Deloro Stellite Group, Ltd., UC-02-184
Delphi Corporation, UC-02-218
Denison University, UA-95-081
Detroit Diesel Corporation, UA-90-019
DG Trim Products, UA-89-012
Doehler-Jarvis Technologies, Inc., UA-96-121
DOW Chemical Company, UC-97-067
Dow Corning Corporation (Midland), UA-88-014
Drexel University, UA-97-063
DRI Institute, UR-99-122
Duke University, UC-95-095
Duraloy Technologies, Inc., UC-02-157
Duro-Last Roofing, Inc., UR-98-077
Dynamet Technology, UC-00-020
E. I. du Pont de Nemours, UA-91-025
E. I. du Pont de Nemours (Fluorochemicals), UA-91-087
E. R. Johnson Associates, Inc., UC-96-060

Eagle Racing, UC-00-078
East Tennessee State Univ. - Johnson City, UA-94-040
Eastern Kentucky State, UA-93-067
Eastman Chemical Company, UA-91-002
Eastman Kodak Company, UC-95-077
Eaton Corporation, UA-91-006
Eberline Instruments, UR-97-055
Ecole Polytechnique de Montreal, UC-99-008
Edison Welding Institute, UA-96-024
EG&G ORTEC, UC-99-143
EmeraChem, UC-02-191
Energy Conversion Devices, Inc., UA-89-014
Energy Recovery, Inc., UC-96-133
ENERMODAL, UR-97-138
Engelhard Corporation, UA-89-032
Engineering Mechanics Corporation of Columbus, UC-99-190
Ensci, Inc., UC-98-036
Environment Canada, UA-99-155
Environmental Technology & Education Center (ETEC), UR-97-108
Equistar Technology Center, UC-99-160
ETH Zentrum, UC-99-158
Exxon Research & Engineering Company, UC-94-082
F. W. Gartner Thermal Spray Company, UC-02-237
Federal Highway Administration, UC-97-008
Federal Mogul Power Train Systems, UC-98-061
Firestone Building Products Co., UR-98-075
Fisk University, UA-95-080
Florida A&M University, UA-96-095
Florida Atlantic University, UA-91-029
Florida International University, UA-91-018
Florida Solar Energy Center, UA-93-010
Florida State University, UA-91-078
Florida Tile Industries, UA-93-052
FLUENT Incorporated, UR-96-071
FMC Corporation, UC-96-137
FMC Naval Systems Division, UA-91-065
Ford Motor Company, UA-90-008
Forged Performance Products , UC-96-068
Foseco Morval, UC-02-222
Foster Miller, Inc., UA-89-011
Furman University, UA-01-122
Galvalume Sheet Producers of North America, UR-99-120
GANIL, UC-99-170
General Electric Aircraft Engines, UA-92-010
General Electric Company, UC-95-032
George Washington University, The, UA-99-133
Georgetown University Medical Center, UA-99-134
Georgia Institute of Technology, UA-88-026
Global Building Systems, Inc., UF-02-228
GM Corporation/Delco Remy Division, UA-94-007
GM Powertrain Group, UC-98-033
GM R&D Center, UC-94-031

Goal Line Company, UC-00-0-15
Golden Technologies Company, UC-95-009
Gonzaga University, UA-98-095
Goodyear Tire & Rubber Company, UA-90-034
Gray-Syracuse-Inc., UC-02-207
Great Lakes Research Div. of HRD, UA-87-008
GSI - Darmstadt, UC-99-172
GTE Laboratories, Inc., UA-89-015
Guidance & Control Systems-Litton Systems, Inc., UA-89-017
H. B. Fuller Company, UC-02-196
H. R. DeSelm, UA-99-171
Harvard University, UA-92-027
Haynes International, UC-99-147
Hilcoit Corporation, UC-99-131
Hiroshima University, UA-95-114
Hoskins Manufacturing Company, UC-97-079
Howard University, UA-95-108
Howmet, UC-95-097
Huntington Alloys/Special Metals Div., UC-99-186
Hydrogen Microplasmatron Technologies, LLC., UR-00-004
HY-Tech Research Corporation, UC-97-115
IAP Research, Inc., UC-02-167
IBACOS, UR-97-002
IBM Almaden Research Center, UA-91-020
ICYNENE Inc., UR-96-118
II-VI Incorporated, UC-01-139
Illinois Institute of Technology, UA-90-010
Illinois Tool Works (ITW), UC-02-162
IMTech Company, UA-89-034
INCO Alloys, Inc., UC-97-130
Ind. Tech Res. Inst/Taiwan, UA-94-034
Indiana University, UA-98-040
Indiana University, UR-98-059
Inductoheat, UC-02-158
INRAD, Inc., UA-91-042
Institut Laue Langevin - France, UA-91-082
Institute for Defense Analyses, UA-89-010
Institute for Nuclear Physics & Engineering, UC-97-114
Institute of Macromolecular Chemistry, UA-97-061
Institute of Paper Science and Technology, UC-96-154
Instituto Nacional de Investigaciones Nucleares, UA-97-015
Insulating Concrete Form Association, UR-99-168
Intellutran, UC-02-203
International Paper, UC-96-084
Intonet Technologies, UC-99-132
IntraSpec, Inc., UC-96-093
Ionic Atlanta, Inc., UA-89-008
Iowa State University, UA-91-095
Istituto ISM del CNR (Italy), UA-92-014
Istituto Nazionale Di Fisica Nucleare - LNS, UA-96-119
J & L Specialty Steel, UC-97-129
J. A. Martin, UA-93-021

Jawaharlal - India, UA-91-075
Jeffrey Chain Corporation, UC-95-137
John Carroll University, UA-92-082
John Mansville, UR-98-078
Johns Hopkins University, UA-89-001
JPS Elastomerics Corporation, UR-98-085
Kansas State University, UA-96-061
Kennametal, Inc., UC-96-111
Kent State University, UA-90-021
Krispin Technologies, Inc., UC-97-144
Kryton International, Inc., UR-03-246
Kyocera Industrial Ceramics Corp., UC-95-064
Lehigh University, UA-89-019
Leroy A. Landers, UC-95-083
Lexmark, YNP-94-001
Lincoln Electric, UC-96-164
Lincoln University, UA-96-134
Lockheed Martin Michould Space Systems, UC-98-047
Lockheed Martin Skunk Works, UR-96-035
Logan Aluminum, UC-00-081
LOTEC, Inc., UA-92-045
Loughborough Univ. of Tech. (UK), UA-91-080
Louisiana State University, UA-90-029
LTV Steel Company, UC-97-124
Lucent Technologies, UA-91-096
Ludlum Measurement Inc., UR-97-044
Lund University, UA-98-112
M. J. Schiff & Associates, Inc., UC-99-189
Machining Research Inc., UC-96-053
Magnequench Technical Center, UC-01-094
Manchester University (UK), UA-91-093
Mantic Corporation, UC-96-091
Marquette University, UA-89-018
Maryville College, UA-94-028
MascoTech Forming Technologies, UC-99-139
Massachusetts Inst. of Technology, UA-90-023
Material Physics Research, UC-00-007
Material Technologies, Inc., UC-96-085
Materials Focus, Inc., UC-97-098
Materials Modifcation, Inc., UC-99-005
Materials Resources International, UC-00-006
Materials Technologies of Virginia, UC-95-104
Matrix Solar Technologies, Inc, UC-02-163
MATTEC, LLC, UC-02-235
McDonnell Douglas Aerospace, UA-90-005
McGill University - Canada, UA-92-016
McMaster University - Canada, UA-91-012
Mead Research, UC-00-089
Membrane Technology Research, UA-92-064
Metal Building Manufacturers Association, UR-99-117
Metal Construction Association, UR-01-124
Metal Construction Association, UR-99-118

Metalspray North America LLC, UC-99-124
Metal-Tech of Tennessee,LLC, UC-02-208
Metamic LLC, UC-03-247
Michigan State University, UA-90-030
Michigan Technological University, UA-89-023
Microbial Insight, Inc., UC-96-135
Microcoating Technologies, UC-98-100
MicroMet Technology, Inc., UC-99-137
Milacron, Inc., UC-99-130
MINCO Acquisition Company, UC-99-176
Miniature Precision Bearings, UA-89-027
Minnesota Mining and Manufacturing, UC-97-117
Mississippi College, UA-93-037
Mississippi State University, UA-92-056
MMPact, Inc., UC-01-147
Modern Alloying Technologies, LLC, UC-99-174
Mogas Industries, UC-02-238
Monarch Tile, Inc., UA-90-009
Monash University, UA-91-073
Monofrax, Inc., UC-99-111
Motorola, UC-96-002
Motorola Energy Systems Group, UC-99-166
Mount Holyoke College, UA-91-074
MRC, Division of Praxair Surface Technologies, UC-01-145
MTI-Emory Ford, UC-02-178
Multi Phase Services, Inc., UC-01-151
N A Technologies, UC-01-104
NAHB Research Center, UR-99-006
Nano Instruments, Inc., UC-95-005
Nanopowder Enterprises, Inc., UC-99-113
NASA Glenn Research Center/Formerly NASA Lewis, UC-95-087
NASA Langley Research Center, UG-NP-006
NASA/Marshall Space Flight Center, UC-01-128
NASA/Marshall Space Flight Center, UC-99-194
NASA-Jet Propulsion Laboratory, UC-02-183
Nat'l Inst of Standards & Technology, UR-95-007
National Coil Coaters Association, UR-99-119
National Electronics Manufacturing Initiative, Inc. (NEMI), UC-02-217
National Forge Company, UC-02-154
National High Magnetic Field Laboratory, UA-98-049
National Highway Traffic Safety Administration, UR-96-114
National Sun Yat-Sen University, UA-97-043
Naval Post Graduate School, UG-NP-005
NEC Research Institute, Inc., UC-97-034
New Jersey Institute of Technology, UA-94-008
New Mexico State University, UA-97-131
New Mexico Tech, UA-99-001
Noranda Magnesium, Inc., UC-01-093
Norfolk State University, UC-99-144
North Carolina Agri. & Tech. State Univ., UA-91-009
North Carolina State University, UA-88-002
Northeastern University, UA-97-035

Northrop Corporation, UA-93-055
Northwest Numerics & Modeling, Inc., UC-02-166
Northwestern University, UA-90-033
Norton/TRW Ceramics, UA-90-011
Nuclear & Aerospace Materials Corporation, UA-88-023
Nuclear Energy Agency, The, UC-95-039
NY State College of Ceramics at Alfred University, UA-88-016
Ohio State University, UA-91-003
Ohio University, UA-95-112
Oklahoma State University, UA-87-006
Oregon Graduate Institute of Science&Tech, UA-94-041
Oregon State University, UA-93-034
Osram Sylvania/University of Massachusetts at Lowell, UC-98-017
Owens Corning Technical Center, UC-97-017
Oxyrase, Inc., UA-91-044
Parker Abex NWL, UC-99-123
Park-Ohio Transportation Group, UC-97-134
PCC Airfoils, Inc., UC-99-183
Pennsylvania State University, UA-88-025.1
Philip Morris Incorporated, UC-94-057
Philips Analytical Inc., UC-00--084
Phone-Poulenc, Inc., UC-97-096
Phygen, Inc., UC-99-140
Plasma Coatings, Inc., UC-98-064
Poco Graphite Inc., UC-01-106
Pohang University of Science and Technology, UA-99-152
Polytechnic University, UA-94-081
Portland State University, UA-97-008
Porvair Fuel Cell Technology, UC-00-087
PPG Industries, Inc., UC-94-078
Pratt and Whitney, UA-92-022
Praxair, Inc., UC-02-195
Precious Metals Corporation, UC-02-220
Princeton University, UA-92-054
Procter & Gamble Company, UA-88-001
Protecto Wrap Company, UR-02-159
Purdue University, UR-96-039
Purdue University Calumet, UA-96-037
Pusan National University, UA-95-119
Pyrotek, Inc., UC-01-133
Quadrax Corporation, UA-91-027
Quantum Peripherals, UC-96-009
Queen's University of Belfast, UA-97-118
RAND Corporation, UR-01-132
RASTRA of the Americas, UR-98-094
RCF Seals, UC-95-084
RCMA, UR-97-060
Refractory Composites, Inc., UC-97-050
Refractory Testing Associates, UA-90-020
ReMaxCo Technologies, Inc., UA-90-007
Rensselaer Polytechnic Institute, UA-91-010
Republic Technologies International, UC-00-001

Reynolds Metal Company, UR-95-004
Rhenium Alloys, Inc., UC-97-123
Rice University, UA-88-020
Richard Knof McMullan, UC-98-054
RIKEN (The Institute of Physical and Chemical Research), UA-96-066
Robert Hageman, UR-97-078
Rochester Institute of Technology, UA-91-001
Rodel, Inc., UC-94-075
Rohm and Haas (formerly Morton), UC-99-162
Rohr, Inc., UC-94-083
Rolls Royce, Inc., UC-94-074
Ronald K. McConathy, UA-99-193
Rutgers University, UA-89-020
Rutherford Appleton Laboratory (UK), UA-91-081
Saint-Gobain Norton, UA-88-017
Samsung Advanced Institute of Tech., Dept. of Supercomputer Applications, UR-96-032
Sandusky International, UC-96-051
Sarnafil, Inc., UR-98-084
Savcor Process Oy, UC-00-090
SB&TD Business Systems, UA-90-003
Schuller International, Inc., UC-96-066
SCI-TEC, Inc., UC-02-206
Scot Forge, UC-02-185
Seagate Technology, UC-99-007
Secat, Inc., UC-00-037
Selee Corporation, UA-88-015
SENES Oak Ridge, Inc., UR-95-110
Siemens Westinghouse Power Corporation, UC-96-079
Single Ply Roofing Institute (SPRI), UR-98-073
SIPA, UR-02-172
Smelter Service Corporation, UC-98-052
Smith & Nephew, UC-97-071
Smith International, Inc., UC-99-185
Soil and Land Use Technologies (SALUT), UC-97-046
Solar Turbines, Inc., UA-91-024
South Dakota State University, UA-91-007
Southern Illinois University, UA-88-012.1
Southern Research Institute, UC-96-088
Southern University, UA-89-033
Southern University Baton Rouge, UA-94-016
Southwest Research Institute, UC-96-094
Space and Naval Warfare Systems Center, UR-98-109
Special Metals Corp. (formerly Inco Alloys Int. Inc.), UC-97-130
SPI/SPFD, UR-96-159
Spirax Sarco, Inc., UC-02-175
SRI International, UC-95-086
Staffordshire University, UA-97-113
Stanford University, UA-91-023
StarFire Systems, Inc., UC-02-205
State University of New York at Stony Brook(SUNY), UA-97-088
Steven Winter Associates, Inc., UR-95-094
Stevens Institute of Technology, UA-91-091A

Stone & Webster Engineering Corp, UC-96-140
Stoody Company, UC-97-058
Structural Insulated Panel Association, UR-97-018
Stryker Howmedica Osteonics, UC-01-102
Sullivan Mining Corporation, UA-89-016
Sulzer Metco, Inc., UC-97-051
Sunstrand Corporation, UA-90-018
Surface Metallurgical, Inc., UC-02-201
Surmet Corporation, UC-95-101
Syngenta Agribusiness Biotechnology Research, UR-01-121
Synterials, Inc., UC-01-130
T. J. Watson Research Center, UC-98-104
TDA Research, Inc., UC-96-069
Technical University Graz, UA-98-107
Technical University of Denmark, UA-89-005
Technische Universitat Wien (Austria), UA-93-029
Technology for Energy Corporation(TEC), UC-96-121
Teledyne Allvac, UA-90-022
Tenneco Automotive, UC-99-135
Tennessee Center for Research and Development, UA-88-030
Tennessee State University, UA-95-075
Tennessee Technological University, UA-91-085
Texas A&M University, UA-93-026
Textron Specialty Materials-Div. Avco Corp., UA-89-013
Thermacore, Inc., UA-89-003
Thermal Ceramics, UC-98-057
Third Millennium Technologies, Inc., UA-89-028
Thixomat, UC-98-023
Timken Company, The, UC-99-141
Torrington Company, UA-92-023
Torrington Company, The, UR-99-127
Tosoh SMD, Inc. , UA-92-015
Tower Automotive, UC-98-067
TPL Incorporated, UC-99-121
TTE Diecasting, UC-98-106
Tufts University, UA-95-023
Tuskegee University, UA-90-012
Tycom Corporation, UC-98-103
U. K. Software Services, Inc. (UKSS, Inc.), UC-99-115
U. S. Naval Academy, UG-NP-002
U.S. Bureau of Mines Albany Research Center, UG-NP-001
U.S. Food and Drug Administration, UR-97-103
UES, Inc., UA-89-006
Ultramet, UA-93-060
Uniform Metal Technologies, LLC, UC-97-126
United Defense LP, UC-96-120
United States Army Research Laboratory, UC-97-053
United Technologies Pratt & Whitney, UA-90-006
Universal Energy Systems, Inc., UA-89-006
Universidad de Sonora, UA-02-165
Universidade Federal De Pernambuco, UA-95-138
Universita' degli Studi di Messina, UA-97-106

Universitat de Barcelona, UA-98-022
Universitat des Saarlandes, UA-98-042
University "La Sapienza" of Rome, UA-97-109
University of Aarhus Dept. of Chemistry, UA-99-167
University of Akron, UA-90-024
University of Alabama, UA-92-092
University of Alabama - Huntsville, UA-99-179
University of Alberta, UA-01-123
University of Arizona, UA-89-021.1
University of Arkansas, UA-02-194
University of Birmingham (UK), UA-92-020
University of Brighton, UA-97-080
University of Bristol (UK), UA-91-083
University of British Columbia, UA-99-112
University of California - Davis, UA-98-099
University of California - Irvine, UA-94-015
University of California @ Santa Barbara, UA-95-124
University of California at Los Angeles , UA-89-029
University of California at San Diego, UA-91-013
University of California at Santa Cruz, UA-90-016
University of California, Berkeley, UA-95-024
University of Cambridge, UA-94-018
University of Central Florida, UA-97-059
University of Cincinnati, UA-89-022
University of Colorado, UA-99-154
University of Connecticut, UA-91-041
University of Dayton, UA-91-026
University of Delaware, UA-90-014.1
University of Delhi, UC-94-087
University of Denver, UA-89-030
University of Edinburgh, UA-97-001
University of Florida, UA-89-004
University of Frankfurt, UA-98-041
University of Hawaii (Honolulu), UA-92-085
University of Houston, UA-91-077
University of Illinois, UA-88-010
University of Kentucky, UA-90-015
University of Leeds, UA-99-161
University of Leuven, UC-02-234
University of Liverpool, UA-95-090
University of Louisville, UA-98-004
University of Maine, UA-97-006
University of Manitoba, UC-02-232
University of Maryland, UA-91-052.1
University of Massachusetts, UA-91-089
University of Melbourne, UC-94-043
University of Memphis, UA-95-093
University of Memphis, UR-95-141
University of Michigan, UA-88-019.1
University of Minnesota, UA-88-029
University of Mississippi, The, UA-01-127
University of Missouri-Columbia, UA-90-036

University of Missouri-Rolla, UA-88-008
University of Montana, UA-97-121
University of Nebraska-Lincoln, UA-95-085
University of Nevada, Reno, UA-95-078
University of New Mexico, UA-88-011
University of New Orleans, UA-98-013
University of North Carolina at Charlotte, UA-99-149
University of North Carolina at Charlotte, UA-99-149
University of North Carolina, Chapel Hill, UA-90-035
University of North Dakota, UA-96-036
University of North Texas, UA-94-076
University of Notre Dame, UA-92-029
University of Oklahoma, UA-92-081
University of Oxford, UA-92-052
University of Palermo, UA-91-097
University of Pennsylvania, UA-92-026
University of Pittsburgh, UA-91-035
University of Puerto Rico, UA-91-005
University of Rochester, UA-92-055
University of Salford, UA-93-022
University of Sao Paulo, UA-92-087
University of South Carolina, UA-90-017
University of Southern California, UA-88-005
University of Stockholm/Manne Siegbahn Laboratory, UA-94-042
University of Surrey, UA-97-100
University of Sydney, UA-96-142
University of Tennessee, UA-87-001
University of Tennessee, UR-95-001
University of Texas (Arlington), UA-91-040
University of Texas at Austin, The, UA-97-068
University of Texas at El Paso, UA-97-133
University of Texas-Pan American, UA-01-144
University of Tokyo, UC-02-231
University of Toledo, UA-99-145
University of Toronto, UA-91-016.1
University of Tsukuba, UA-96-122
University of Utah, UA-89-002
University of Virginia, UA-95-118
University of Warwick (UK), UA-92-019
University of Washington, UA-92-044
University of Waterloo, UR-01-100
University of Western Australia, UA-92-053
University of Wisconsin, UA-91-011
Valenite, Inc., UA-94-009
VAMISTOR Corporation, UC-96-064
Vanderbilt University, UA-88-006.1
VG IsoTech, UA-93-043
Virginia Polytechnic & State University, UA-87-002
Visteon Corporation, UC-01-143
W. H. Porter, Inc., UR-02-215
W. R. Grace & Co-Conn, UC-95-022
Wagner Castings Company, UC-97-070

Walford Technologies(formerly IAC Oak RIdge), UC-02-160
Warsaw University, UA-99-169
Washington State University, UA-91-008
Washington University, UA-92-057
Waukesha Electric Systems, UC-96-070
Wayne State University, UA-95-073
Weirton Steel Corporation, UC-98-026
West Homestead Engineering and Machine Company, UC-02-173
West Virginia University, UA-92-017
Western Environmental Corporation, UR-96-087
Western Kentucky University, UA-98-117
Western Michigan University, UA-96-082
Westinghouse Electric Corporation (Science & Technology Center), UA-91-017
Weyerhaeuser Company, UC-96-073
Wichita State University, UC-98-088
William Thomas Pope, UC-99-003
Wright State University, UA-92-028
X-Ray Instrumentation Associates, UC-99-150
X-Ray Optical Systems, Inc., UC-01-152
Yale University, UA-94-031

INTERNAL DISTRIBUTION

1. F. H. Akers, Jr.
2. A. S. Bland
3. J. Dongarra
4. J. B. Drake
5. T. Dunigan, Jr.
6. T. H. Dunning, Jr.
7. A. Geist
8. R. G. Gilliland
9. W. F. Harris
10. R. J. Harrison
11. J. A. Nichols
12. W. J. Madia
13. T. E. Mason
14. D. McCoy
15. L. Riedinger
16. J. B. Roberto
17. J. B. White, III
18. P. H. Worley
19. T. Zacharia
20. Central Research Library
21. ORNL Laboratory Records – RC
22. ORNL Laboratory Records - OSTI

EXTERNAL DISTRIBUTION

23. R. Orbach, SC-1 Forrestal Building, U.S. Department of Energy, 1000 Independence Avenue, S.W., Washington, D.C. 20585
24. J. F. Decker, SC-1 Forrestal Building, U.S. Department of Energy, 1000 Independence Avenue, S.W., Washington, D.C. 20585
25. C. E. Oliver, , SC-30/GTN, U.S. Department of Energy, 1000 Independence Avenue, S.W., Washington, D.C. 20585
26. W. M. Polansky, SC-32/GTN, U.S. Department of Energy, 1000 Independence Avenue, S.W., Washington, D.C. 20585
27. G. Johnson, SC-30/GTN, U.S. Department of Energy, 1000 Independence Avenue, S.W., Washington, D.C. 20585
28. A. Laub, SC-1/GTN, U.S. Department of Energy, 1000 Independence Avenue, S.W., Washington, D.C. 20585
29. M. H. Wright, Courant Institute of Mathematical Sciences, New York University, 251 Mercer Street, New York, NY 10012
30. J. W. D. Connolly, Center for Computational Sciences, University of Kentucky, 325 McVey Hall 0045, Lexington, KY 40506
31. H. E. Kulsrud, Center for Communications Research, 805 Bunn Drive, Princeton, NJ 08540-3699
32. J. P. Dahlburg, General Atomics, P.O. Box 85608, San Diego, CA 92186
33. R. C. Giles, Center for Computational Science, Boston University, 3 Cummington Street, Boston, MA 02215
34. E. B. Stechel, Ford Motor Company, Powertrain Operations & Engine Engineering, 21500 Oakwood Blvd., MD#3, Room CB064, Dearborn, MI 8121-2053
35. K. R. Sollins, Massachusetts Institute of Technology, Cambridge, MA 02139
36. J. C. Meza, High Performance Computing Research, Lawrence Berkeley National Laboratory, One Cyclotron Road, MS: 50B-2239, Berkeley, CA 94720
37. G. J. McRae, Department of Chemical Engineering, Massachusetts Institute of Technology, 66-372, Cambridge, MA 02139
38. W. A. Lester, Jr., Department of Chemistry, University of California – Berkeley, 221B Hildebrand, Berkeley, CA 94720-1460

39. W. M. Washington, National Center for Atmospheric Research, P.O. Box 3000, Boulder, CO 80307
40. S. Wolff, CISCO Systems, 5116 Rockwood Parkway, NW, Washington, DC 20016-1916
41. J. Grosh, Office of the Deputy Undersecretary, U. S. Department of Defense, 3010 Defense Pentagon, Washington, DC 10301-3010
42. D. Nelson, National Coordination Office for Information Technology Research and Development, 4201 Wilson Blvd., Suite II405, Arlington, VA 22230
43. F. C. Johnson, SC-31/GTN, U.S. Department of Energy, 1000 Independence Avenue, S.W., Washington, D.C. 20585
44. C. H. Romine, SC-31/GTN, U.S. Department of Energy, 1000 Independence Avenue, S.W., Washington, D.C. 20585
45. M .A. Scott, SC-31/GTN, U.S. Department of Energy, 1000 Independence Avenue, S.W., Washington, D.C. 20585
46. W. H. Miner, Jr., SC-30/GTN, U.S. Department of Energy, 1000 Independence Avenue, S.W., Washington, D.C. 20585
47. T. D. Ndousse-Fetter, SC-31/GTN, U.S. Department of Energy, 1000 Independence Avenue, S.W., Washington, D.C. 20585
48. G. R. Seweryniak, SC-31/GTN, U.S. Department of Energy, 1000 Independence Avenue, S.W., Washington, D.C. 20585
49. J. R. vanRosendale, SC-31/GTN, U.S. Department of Energy, 1000 Independence Avenue, S.W Washington, D.C. 20585
50. D. Hitchcock SC-30/GTN, U.S. Department of Energy, 1000 Independence Avenue, S.W., Washington, D.C. 20585
51. A. A. Patrinos, SC-70/GTN, U.S. Department of Energy, 1000 Independence Avenue, S.W., Washington, D.C. 20585
52. J. W. Elwood, SC-74/GTN, U.S. Department of Energy, 1000 Independence Avenue, S.W., Washington, D.C. 20585
53. D. Bader, SC-74/GTN, U.S. Department of Energy, 1000 Independence Avenue, S.W., Washington, D.C. 20585
54. P. M. Dehmer, SC-10/GTN, U.S. Department of Energy, 1000 Independence Avenue, S.W., Washington, D.C. 20585
55. W. J. Stevens, SC-14/GTN, U.S. Department of Energy, 1000 Independence Avenue, S.W., Washington, D.C. 20585
56. W. H. Kirchhoff, SC-14/GTN, U.S. Department of Energy, 1000 Independence Avenue, S.W., Washington, D.C. 20585
57. S. A. Eckstrand, SC-55/GTN, U.S. Department of Energy, 1000 Independence Avenue, S.W., Washington, D.C. 20585
58. S. P. Rosen, SC-20/GTN, U.S. Department of Energy, 1000 Independence Avenue, S.W., Washington, D.C. 20585
59. V. White, SC-32, U.S. Department of Energy, 1000 Independence Avenue, S.W., Washington, D.C. 20585
60. S. Szykman, Office of Science and Technology Policy, 1650 Pennsylvania Avenue, N.W. Washington, DC 20502-0001
61. T. J. Mackin, Office of Science and Technology Policy, 1650 Pennsylvania Avenue, N.W., Washington, DC 20502
62. J. Parriott, The Office of Management and Budget, 725 17th Street, N.W., Washington, DC 10503
63. H. Simon, Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720
64. E. Blythe, Virginia Polytechnic Institute and State University, Blacksburg, VA 24061-0531
65. L. Peters, Pacific Northwest National Laboratory, P. O. Box 999, Richland, WA 99352
66. C. Liotta, School of Chemistry and Biochemistry, Georgia Institute of Technology, Atlanta, GA 30332

67. R. Kendall, Ames Laboratory, Iowa State University, Ames, IA 50011
68. T. M. Kaarsberg, U.S. House of Representatives, 390 Ford House Office Building, Washington, DC 20515
69. Timothy E. Valentine, U. S. Senate, 703 Hart Senate Office Building, Washington, D.C. 20510
70. J. F. Bagley, Battelle Washington Operations, 901 D Street, 370 L'Enfant Promenade, S.W., Suite 900, Washington, DC 20024-2115
71. C. Desmond, Battelle Washington Operations, 901 D Street, 370 L'Enfant Promenade, S.W., Suite 900, Washington, DC 20024-2115
72. R. Townsend, Oak Ridge Associated Universities, P. O. Box 117, Oak Ridge TN 37831-0117
73. J. Rottsolk, Cray, Inc., 411 First Avenue South, Suite 600, Seattle, WA 98104
74. D. Kiefer, Cray, Inc., 1050 Lowater Road, P. O. Box 6000, Chippewa Falls, WI 54729-0080
75. S. Scott, Cray, Inc., Engineering Building, 1050 Lowater Road, Chippewa Falls, WI 54729-0078
76. L. Hoelzeman, Cray, Inc., 98 Jay Drive, Madison, AL 325758
77. W. Gropp, Argonne National Laboratory, 9700 S. Cass Ave., Argonne, IL 60439
78. D. Keyes, Old Dominion University, 232 Education Bldg., Norfolk, VA 23529-0162
79. R. Stevens, Argonne National Laboratory, 9700 S. Cass Ave., Argonne, IL 60439
80. M. Gordon, Ames Laboratory, 201 Spedding, Ames, IA 50011-3020
81. W. Tang, Princeton Plasma Physics Laboratory, P.O. Box 451, Princeton, NJ 08543
82. B. Gorda, Lawrence Berkeley National Laboratory, 1 Cyclotron Road Mail Stop 50A1148, Berkeley, CA 94720
83. L. Oliker, Lawrence Berkeley National Laboratory, 1 Cyclotron Road Mail Stop 50B2239, Berkeley, CA 94720
84. P. Muzio, Army High Performance Computing Research Center, 1200 Washington Avenue South, Minneapolis, MN 55415
85. R. Rood, NASA/Goddard, Code 930, Bldg. 38, RM W230H, Greenbelt, MD 20771
86. V. Kumar, Department of Computer Science and Engineering, 4-192, EE/CSci Building, University of Minnesota, Minneapolis, MN 55455