OAK RIDGE NATIONAL LABORATORY

MANAGED BY UT-BATTELLE FOR THE DEPARTMENT OF ENERGY

P.O. Box 2008
Oak Ridge, TN 37831-6163
Phone: (865) 574-4897
Fax: (865) 574-4839
Email: zachariat@ornl.gov

March 2, 2006

Dr. Anil Deane
U.S. Department of Energy
SC-21.1, Germantown Building
1000 Independence Ave., S.W.
Washington, DC 20585

Dear Dr. Deane:

Program Announcement LAB 06-04: Scientific Discovery through Advanced Computing (SciDAC): Science Application Partnership

Enclosed please find the proposal entitled, "Performance Engineering for the Next Generation Community Climate System Model," for your consideration.

This work addresses the need for the Community Climate System Model to use existing computing systems more efficiently and to be optimized for next generation massively parallel systems. This will require improving scalability to thousands of processors and carrying out comprehensive performance analyses. The work will include (1) improving scalability of the atmosphere and ocean component models with respect to processor count; (2) investigating and optimizing performance at high spatial resolution; (3) improving scalability of the full coupled model through improved load balance; and (4) porting to and optimizing on next generation high performance computing systems.

The total request is $500,000 per year for three years.

Questions regarding this proposal should be directed to the Principal Investigator, Patrick H. Worley, 865-574-3128 or e-mail: worleyph@ornl.gov.

Sincerely,

Thomas Zacharia

Thomas Zacharia
Associate Laboratory Director
Computing and Computational Sciences

Enclosures

c: W. C. Lin J. A. Nichols T. Zacharia

DOE F 4650.2
Rev. (10-03)
(All Other Editions Are Obsolete)

Office of Science (SC)

OMB Control No.
1910-1400
(OMB Burden Disclosure Statement on Back)

Face Page

TITLE OF PROPOSED RESEARCH:
Performance Engineering for the Next Generation Community Climate System Model

1. CATALOG OF FEDERAL DOMESTIC ASSISTANCE #

2. CONGRESSIONAL DISTRICT:
Applicant Organization's District: 2nd & 3rd Districts
Project Site's District: 2nd & 3rd Districts

3A. I.R.S. ENTITY IDENTIFICATION OR SSN:
621788235

3B. DUNS Number:

4. AREA OF RESEARCH OR ANNOUNCEMENT TITLE/#:
LAB 06-04: Scientific Discovery through Advanced Computing

5. HAS THIS RESEARCH PROPOSAL BEEN SUBMITTED TO ANY OTHER FEDERAL AGENCY?
☐ YES ☑ NO

PLEASE LIST

6. DOE/OER PROGRAM STAFF CONTACT (if known):
Dr. Anil Deane

7. TYPE OF APPLICATION:
☑ New ☐ Renewal
☐ Continuation ☐ Revision
☐ Supplement

8. ORGANIZATION TYPE:
☐ Local Govt. ☐ State Govt.
☐ Non-Profit ☐ Hospital
☐ Indian Tribal Govt. ☐ Individual
☑ Other ☐ Inst. of Higher Educ.
☐ For-Profit
☐ Small Business ☐ Disadvan. Business
☐ Women-Owned ☐ 8(a)

9. CURRENT DOE AWARD # (IF APPLICABLE):
None

10. WILL THIS RESEARCH INVOLVE:
10A. Human Subjects ☑ No ☐ If yes
Exemption No. ______ or
IRB Approval Date ______
Assurance of Compliance No: ______
10B. Vertebrate Animals ☑ No ☐ If yes
IACUC Approval Date ______ or
Animal Welfare Assurance No: ______

11. AMOUNT REQUESTED FROM DOE FOR ENTIRE PROJECT PERIOD $ 1500000.00

12. DURATION OF ENTIRE PROJECT PERIOD:
07/01/06 to 06/30/09
MM/DD/YY MM/DD/YY

13. REQUESTED AWARD START DATE
7/1/06
MM/DD/YY

14. IS APPLICANT DELINQUENT ON ANY FEDERAL DEBT?
☐ Yes (attach an explanation) ☑ No

15. PRINCIPAL INVESTIGATOR/PROGRAM DIRECTOR
NAME Patrick H. Worley
TITLE Senior Research Staff
ADDRESS P.O. Box 2008
Oak Ridge, TN 37831-6016

PHONE NUMBER 865-574-3128

SIGNATURE OF PRINCIPAL INVESTIGATOR/ PROGRAM DIRECTOR
(please type in full name if electronically submitted)
Date March 2, 2006

PI/PD ASSURANCE: I agree to accept responsibility for the scientific conduct of the project and to provide the required progress reports if an award is made as a result of this submission. Willful provision of false information is a criminal offense. (U.S. Code, Title 18, Section 1001).

16. ORGANIZATION'S NAME Oak Ridge National Laboratory
ADDRESS P.O. Box 2008
Oak Ridge, TN 37831-6163

CERTIFYING REPRESENTATIVE'S
NAME Thomas Zacharia
TITLE Associate Laboratory Director
PHONE NUMBER (865) 574-4897

SIGNATURE OF ORGANIZATION'S CERTIFYING REPRESENTATIVE
(please type in full name if electronically submitted)
Date March 2, 2006

CERTIFICATION and ACCEPTANCE: I certify that the statements herein are true and complete to the best of my knowledge, and accept the obligation to comply with DOE terms and conditions if an award is made as the result of this submission. A willfully false certification is a criminal offense. (U.S. Code, Title 18, Section 1001).

DOE Form 5120.2 (Page 1)
(1-96)

U.S. DEPARTMENT OF ENERGY
OAK RIDGE OPERATIONS
FIELD WORK PROPOSAL

PROGRAM: KJ Computational and Technology Research

1. WORK PROPOSAL NO. ERKJD13	2. REVISION NO. 0	3. DATE PREPARED 04-15-2006	12
4. WORK PROPOSAL TITLE: Performance Engineering for the Next Generation Community Climate		5. BUDGET AND REPORTING CODE KJ 01 02 00 0	
6. WORK PROPOSAL TERM BEGIN: 07-01-2006 END: 06-30-2009	PATENT STATUS This proposal is being transmitted in advance of patent review for evaluation purposes only. No further dissemination or publication shall be made without prior approval of the Assistant General Counsel for Patents, DOE.		7. Is This Work Proposal Included in the Institutional Plan? ☒Yes ☐No
NAME: (Last, First, MI) (Phone Number) 8. HEADQUARTERS/OPERATIONS OFFICE PROGRAM MANAGER: Deane, Anil (301)903-7774	11. HEADQUARTERS ORGANIZATIONS: **Office of Science**		14. DOE ORGANIZATION CODE: **SC**
9. OPERATIONS OFFICE WORK PROPOSAL REVIEWER: Lin, Wayne C. (865)576-0639	12. FIELD OFFICE: **Oak Ridge Operations**		15. DOE ORGANIZATION CODE: **ON**
10. CONTRACTOR WORK PROPOSAL PRINCIPAL INVESTIGATOR(S)/MANAGER: Worley, Patrick H. (865)574-3128	13. CONTRACTOR NAME: **Oak Ridge National Laboratory** **Managed by UT-Battelle, LLC** **For the U.S. Department of Energy** **Post Office Box 2008** **Oak Ridge, TN 37831**		16. DOE CONTRACTOR CODE: **41**

17. WORK PROPOSAL DESCRIPTION (Approach, anticipated benefits in 200 words or less)

This proposal describes a Science Application Partnership aligned with the Scientific Discovery through Advanced Computing projects "A Scalable and Extensible Earth System Model for Climate Change Science" and "Performance Engineering Research Center." Under the former the Community Climate System Model is being extended to become a first-generation Earth System Model that treats the coupling between the physical, chemical, and biogeochemical processes in the climate system. To continue to achieve the computational throughput of today's simulations will require additional optimization and scalability to thousands of processors. We plan to improve scalability of the atmosphere and ocean models with respect to processor count; investigate and optimize performance at high spatial resolution; improve scalability of the full coupled model through improved load balance; and port to and optimize on next generation high performance computing systems. This work is a collaborative effort between Argonne, Lawrence Livermore, and Oak Ridge National Laboratories.

18. CONTRACTOR WORK PROPOSAL MANAGER: (Name and Phone No.) Thomas Zacharia (865)574-4897 Signature: Date: 04-15-2006	19. OPERATIONS OFFICE REVIEW OFFICIAL (Signature) (Date)

20. DETAIL ATTACHMENTS: (See instructions for page 3)

☐ a. Facility Requirements	☐ e. Approach	☐ i. NEPA Requirements	☒ m. ES&H Considerations
☐ b. Publications	☐ f. Technical Progress	☐ j. Milestones	☒ n. Human/Animal Subjects
☐ c. Purpose	☐ g. Future Accomplishments	☐ k. Deliverables	☒ o. Other (Specify)
☐ d. Background	☐ h. Relationships To Other Projects	☐ l. Perform Measures/Expectations	

DOE Form 5120.2 (Page 2)
(1-96)

WORK PROPOSAL REQUIREMENTS FOR OPERATING/EQUIPMENT OBLIGATIONS AND COSTS

PROGRAM: KJ Computational and Technology Research

CONTRACTOR NAME:	WORK PROPOSAL TITLE:		
UT-BATTELLE, LLC	Performance Engineering for the Next Generation Community Climate System Model		
	WORK PROPOSAL NO.	REVISION NO.	DATE PREPARED
	ERKJD13	0	04-15-2006

21. STAFFING (in staff years)	FY 2006	FY 2007	FY 2008		FY 2009	FY 2010	TOTAL TO COMPLETE
			REQUEST	AUTHOR.			
a. SCIENTIFIC / OTHER DIRECT - ORNL	0.1	0.3	0.3		0.2		
b. OTHER DIRECT - OTHER SITES							
c. TOTAL DIRECT	0.1	0.3	0.3		0.2		
22. OPERATING EXPENSE (in Thousands)							
a. TOTAL OBLIGATIONS	27	108	108		81		
COSTS:							
1) WAGE POOL AND ORG. BURDEN	18	74	72		53		
2) MATERIALS AND SERVICES	1	3	4		4		
3) SUBCONTRACTS AND CONSULTANTS							
4) INDIRECT COSTS	8	31	32		24		
b. TOTAL COSTS	27	108	108		81		
23. EQUIPMENT (in Thousands)							
a. EQUIPMENT OBLIGATIONS							
b. EQUIPMENT COSTS							

24. MILESTONE SCHEDULE (TASKS:)	DOLLARS (in Thousands)		SCHEDULE (DATE)	
	PROPOSED	AUTHORIZED	PROPOSED	AUTHORIZED
Analyze Community Atmosphere Model (CAM) performance on representative architectures. Quantify nature of scalability bottlenecks in terms of problem size, processor count, and computer system.	27		09/06	

25. REPORTING REQUIREMENTS (DESCRIPTION:)

Results will be reported in periodic highlights to the U.S. Department of Energy and in peer-reviewed journals and conference proceedings.

DOE Form 5120.2 (Page 3)
(1-96)

WORK PROPOSAL REQUIREMENTS FOR OPERATING/EQUIPMENT OBLIGATIONS AND COSTS

PROGRAM: KJ Computational and Technology Research

CONTRACTOR NAME:	WORK PROPOSAL TITLE:		
UT-BATTELLE, LLC	Performance Engineering for the Next Generation Community Climate System Model		
	WORK PROPOSAL NO. ERKJD13	REVISION NO. 0	DATE PREPARED 04-15-2006

20. DETAIL ATTACHMENT CONTINUED:

m. ES&H Considerations

Paper studies only.

n. Human/Animal Subjects

No human or animal subjects involved.

DOE Form 5120.2 (Page 3)
(1-96)

WORK PROPOSAL REQUIREMENTS FOR OPERATING/EQUIPMENT OBLIGATIONS AND COSTS

PROGRAM: KJ Computational and Technology Research

CONTRACTOR NAME:	WORK PROPOSAL TITLE:		
UT-BATTELLE, LLC	Performance Engineering for the Next Generation Community Climate System Model		
	WORK PROPOSAL NO. ERKJD13	REVISION NO. 0	DATE PREPARED 04-15-2006

20. DETAIL ATTACHMENT CONTINUED:

o. Other

(1) OBLIGATIONS FOR OPERATING EXPENSES-Budget Authority (B/A)

	Obligation Estimates FY 2006	FY 2007	FY 2008
Cost (B/O) Estimates	27	108	108
Less: Uncosted Balance (--) at 10/01			
Plus: Commitments for Continued Operations			
Outstanding Commitment Balance 10/08			
TOTAL OBLIGATIONS--CHANGE	27	108	108

(2) CAPITAL EQUIPMENT OBLIGATIONS AND COSTS - None

Performance Engineering for the Next Generation Community Climate System Model

A Proposal Submitted to the DOE Office of Science

Program Announcement:	LAB 06-04
Program Area:	Science Application Partnership: Computer Science
Program Office:	Office of Advanced Scientific Computing Research
Technical Contact:	Dr. Anil Deane

Applicant

Institution	*Principal Investigator*
Oak Ridge National Laboratory	Patrick H. Worley
PO Box 2008, MS 6016	(865)574-3128
Oak Ridge, TN 37831-6016	worleyph@ornl.gov

Field Work Proposal ERKJD13

Participating Institutions / Senior Personnel

***Lead PI:* Patrick H. Worley, Oak Ridge National Laboratory**

Institution	*Senior Personnel*
Lawrence Livermore National Laboratory (LLNL)	Arthur Mirin
Argonne National Laboratory (ANL)	Robert Loy

Science Application Partner / Science Liaisons

A Scalable and Extensible Earth System Model for Climate Change Science

***Lead Science Application PI:* John B. Drake, Oak Ridge National Laboratory**

Institution	*Senior Personnel*
Oak Ridge National Laboratory	John Drake
Los Alamos National Laboratory	Philip Jones
Lawrence Livermore National Laboratory	Phillip Cameron-Smith
Argonne National Laboratory	Robert Jacob

Projected Funding Request: $500K per year for three years.

TABLE OF CONTENTS

1 Abstract

We propose a Science Application Partnership aligned with the science application project *A Scalable and Extensible Earth System Model for Climate Change Science* (SEESM) and with the center for enabling technology *Performance Engineering Research Center* (PERC). SEESM will extend the Community Climate System Model (CCSM) to become a first-generation Earth System Model that treats the coupling between the physical, chemical, and biogeochemical processes in the climate system. It will explore new science and will run at higher resolution than today's models. Both the improved physical processes and the increase in resolution come at the cost of significant increase in computational complexity.

In order to continue to achieve the computational throughput of present-day simulations, CCSM will need to be modified to use existing computing systems more efficiently and optimized for next-generation massively parallel systems. This will require improving scalability to thousands of processors and carrying out more comprehensive performance analysis in order to attain that goal. We plan to (1) improve scalability of the atmosphere and ocean component models with respect to processor count; (2) investigate and optimize performance at high spatial resolution; (3) improve scalability of the full coupled model through improved load balance; and (4) port to and optimize on next generation high performance computing systems.

2 Narrative

2.1 Background and Significance

The goal of the DOE Climate Change Prediction Program (CCPP), including the SciDAC2 Science Application project *A Scalable and Extensible Earth System Model for Climate Change Science* (SEESM), is:

To determine the range of possible climate changes over the 21st century and beyond through simulations using a more accurate climate system model that includes the full range of human and natural climate feedbacks with increased realism and spatial resolution.

Over the next 5 years, SEESM intends to create a first generation Earth system model based on the Community Climate System Model (CCSM) that treats the coupling between the physical, chemical, and biogeochemical processes in the climate system. The model will include comprehensive treatments of the processes governing well-mixed greenhouse gases, natural and anthropogenic aerosols, the aerosol indirect effect and tropospheric ozone for climate change studies. It will also include improved representation of carbon and chemical processes, particularly for treatment of greenhouse gas emissions and aerosols. These improvements all come at the cost of significant increases in computation, minimally 3-5 times as costly per horizontal grid point. SEESM will also be running experiments at much higher grid resolutions, in order to validate the accuracy of the new processes and to investigate specific science questions, which involves an increase in computational cost that is at least linear in the increased size of the computational grid.

During the integration of new methods and new chemical and biogeochemical processes, the model will need to continue to perform well on the available high performance computing systems in order to maximize the length of simulations and number of ensembles that can be used for assessment products. The major assessments by the International Panel on Climate Change (IPCC) occur roughly every six years, and simulations from CCSM have been used extensively in the third and fourth IPCC reports. For the fifth assessment (AR5), the target is to use the Earth-system version of CCSM that predicts the co-evolution of physical and chemical climate. In order to complete the full suite of desired assessment runs, it will be important to achieve the same computational rate that was achieved in previous assessments. This is a very aggressive goal given the planned modifications to CCSM.

The computer resources available for computational science within the DOE Office of Science have increased dramatically with the installation of the 1024 processor Cray X1E and the 5296 processor Cray XT3 at Oak Ridge National Laboratory (ORNL), the 1024 node IBM BlueGene/L (BG/L) at Argonne National Laboratory (ANL), and the 888 processor IBM POWER5 cluster and the 640 processor Advanced Micro Devices (AMD) Opteron cluster at Lawrence Berkeley National Laboratory. This trend is expected to continue for a number of years, with the goal of installing a system with petaflop performance in 5 to 10 years. This potentially will provide the necessary resources to meet the substantial needs of SEESM, if the CCSM can take advantage of them. A petaflop system will necessarily employ tens of thousands of processors. Two major issues face the CCSM code structure as we move from systems with 1000 to 10,000 and 50,000 processors. First is parallel scalability and the second is load balancing. A number of practical issues also remain, as evidenced by the fact that that the memory requirements are currently too large to be able to use systems such as the IBM BlueGene/L system.

2.2 Preliminary Studies

2.2.1 Problem Description

The CCSM is currently made up of atmosphere, ocean, land, and sea ice *component* models. The atmosphere component is the Community Atmosphere Model (CAM3). The ocean component is the Parallel Ocean Program (POP). The land component is the Community Land Model (CLM3). The sea ice model is the Community Sea Ice Model (CSIM4). In a coupled run these components execute concurrently (on separate processors) and communicate via a coupler component (CPL6). The atmosphere model is the most computationally expensive component currently, and is likely to become even more dominant as additional physical processes are introduced. The ocean is the second most expensive component. When allocating processors to the CCSM components, typically as many as possible are given to the atmosphere. Processors are then allocated to the other components so as to minimize the time components are idle waiting for data from the coupler. Note that the cost of a component is a function of the grid resolution used in the component. For some experiments it may be appropriate to use a much larger grid for the ocean than for the atmosphere, in which case the roles of the ocean and the atmosphere in this configuration process would switch.

The land and sea ice models and the coupler are essentially two-dimensional surface models, and, while they need to run efficiently on parallel platforms, there are intrinsic limits to their scalability. In contrast, the atmosphere and ocean models are both three-dimensional models and potentially can achieve high degrees of parallelism. To achieve the goals of SEESM, high degrees of parallelism and computational efficiency will be required of both ocean and atmosphere components.

The ocean model supports a two-dimensional domain decomposition, which allows thousands of processors to be used even for the relatively small 1-degree grid resolution (approximately 120,000 horizontal grid points and 40 vertical levels) used in current CCSM simulations. However, performance can still suffer from scaling problems due to load imbalances and the overhead of a linear system solution required at the ocean surface, that is, a two-dimensional subsystem that needs to be solved every timestep of the simulation. A new version of POP has been developed that addresses some of these issues, including a subblock algorithm to eliminate some of the load imbalances and an option to run the two-dimensional linear system solution on a subset of the processors [Jones2005]. This will be incorporated into CCSM within the current year. The efficacy of these modifications for the one-degree grid is not clear, and a careful performance analysis is required to understand the scalability of this new version of POP. As new physical processes are introduced into POP, and as the grid resolution increases, the performance characteristics of the model will change. SEESM is also developing an entirely new ocean model, one incorporating a different vertical coordinate, to replace POP.

The atmosphere model is itself made up of two major subcomponents: the physical parameterizations (*physics*) and the dynamical core (*dycore*). CAM comes with three different dycores, one of which is selected at compile-time. The current released version of CCSM uses a spectral Eulerian dycore. This dycore supports only a one-dimensional domain decomposition, severely limiting scalability. For example, at the current production resolution using ~32000 horizontal grid points, no more than 128 processors can be used in the dynamics. In contrast, a two-dimensional decomposition is supported in the physics, and many more processors can be used there. On a shared memory parallel computer (SMP), or a cluster of SMPs, OpenMP parallelism can be used to provide more processors to the physics than can be used by the dynamics. It has been shown that a two-dimensional decomposition is possible for the spectral

Eulerian dycore, but the efficiency of such an implementation is in doubt due to interprocessor communication overhead [Drake1995; Foster1997]. Experimenting with such a new implementation has been deprecated as the climate scientists are currently moving to one of the other dycores for production, a semi-Lagrangian finite-volume (FV) dynamical core [Lin2004]. The FV dycore is currently utilized with only ~13000 horizontal grid points, but more recently has been run at 1.0-degree resolution (~52000 horizontal grid points) and 0.5-degree resolution (~220000 horizontal grid points). The FV dycore can be domain-decomposed in two dimensions [Mirin2005], but the number of subdomains is limited to at most 480 at 1.0-degree resolution with 26 vertical levels, and limited to 960 for the 0.5-degree resolution with 26 vertical levels. As with the spectral Eulerian dycore, OpenMP can be used to exploit more processors in the physics than in the dynamics. This is illustrated in Fig. 1, where systems not using OpenMP can use at most 960 processors, while those supporting OpenMP can use more or, like the IBM p575 system, can use fewer MPI processes for a given number of processors and achieve better scalability.

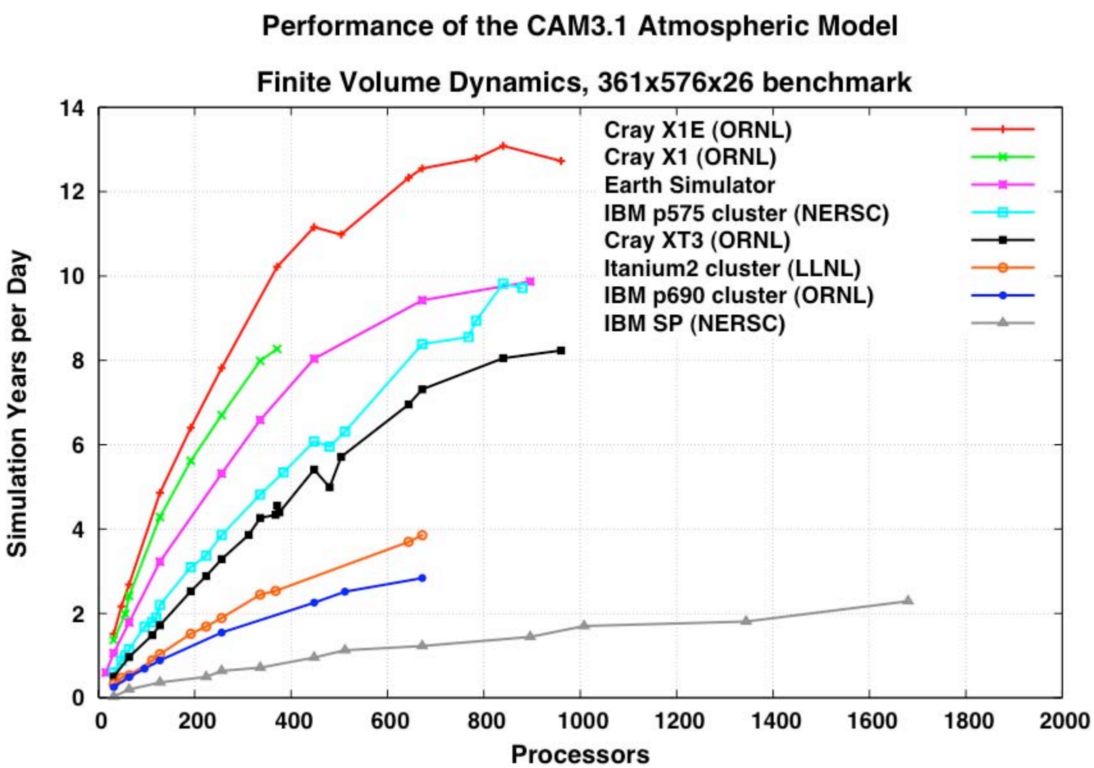

Fig. 1. Performance Scalability of CAM.

While the number of processors used in the full CCSM is the sum of those used in the individual component models, it is not productive to use many more processors than twice the number used by the atmosphere model. Ultimately, the scalability of the climate models has been limited by the relatively modest mesh sizes and accompanying domain decomposition methodologies. With the expected increase in the cost per horizontal grid point due to the inclusion of new physical processes, these limitations are no longer acceptable if CCSM is to make effective use of next generation petaflop systems.

The IBM BG/L system has many architectural features in common with a number of the proposed petaflop systems. The BG/L system is optimized for bandwidth, scalability and the ability to handle large amounts of data while consuming a fraction of the power and floor space required by today's fastest systems. Every processor on BG/L contains a two-element wide FPU, which results in 4 flops/cycle/processor. There are very high bandwidth paths between processor and memory (on a per flop basis), particularly for sequential access. The toroidal network is well suited to periodic boundary conditions, such as is found in global climate models, and a great deal of work went in to making the network partitionable. A simple kernel on the compute nodes and offloading of system services to I/O and service nodes results in less interference with running applications. The MPI library for BG/L was written to exploit hardware features and deliver high communication performance. The drawbacks of BG/L, compared to other contemporary platforms, are its relatively low clock speed (700MHz) and small per-node memory (512MB). The small memory will present particular challenges in porting CCSM to BG/L because CCSM components typically use global arrays on a single node for I/O and many of the physical processes planned for addition will also increase the per-node memory requirements.

As will be shown in the following sections, we are proposing a number of solutions to these problems. Much is still unknown however, and will remain unknown unless significant effort is made to measure and track the performance of the CCSM as the model evolves. The primary responsibility of measuring and tracking performance resides within the SEESM project, and we will work closely with the software engineers within SEESM to analyze this data and prioritize the performance aspects of model development. This Science Application Partnership (SAP), focusing as it does on scalability and on enabling the exploitation of the next generation of massively parallel systems, is a crucial component to the success of the SEESM research goals.

2.2.2 Collaboration and Past Accomplishments

This SAP is led by 3 researchers who have been involved in CCSM development for many years, have led analyses that have quantified the performance limitations of the current model, and have a vision of how to correct the problems.

- Patrick Worley has worked with the atmosphere model for over 15 years. He designed, implemented, and evaluated the parallel algorithms used in the spectral dynamics in a Computer Hardware, Advanced Mathematics and Model Physics (CHAMMP) funded massively parallel version of the Community Climate Model (CCM), the predecessor to CAM. For CAM, Worley designed and implemented the data structures and load balancing scheme used in CAM physics, supporting both efficient vectorization and increased scalability when used with OpenMP parallelism. Worley was also an active participant in the porting, performance analysis, and performance optimization of both CAM and POP on the Cray X1, X1E, XD1, and XT3, the IBM POWER4 and POWER5 clusters, and the SGI Altix. Worley is a co-chair of the CCSM Software Engineering Working Group. Worley was the liaison between the SciDAC science application projects in climate and the SciDAC Integrated Software Infrastructure Center *Performance Evaluation Research Center* and will be the coordinator of application engagement activities in the SciDAC2 Center for Enabling Technologies *Performance Engineering Research Center* (PERC).
- Arthur Mirin has worked in the area of high-performance computing for global climate modeling since the beginnings of the CHAMMP program over 15 years ago. He developed the two-dimensional domain decomposition methodology for the finite-volume dynamical core, significantly increasing the scalability of CAM with FV. He also led the vectorization of the FV dycore. He has extensive experience with communication paradigms (e.g., one-sided MPI-2, Shmem, and hybrid programming models). Mirin coordinated the simulation and was lead author of the work that won the Gordon Bell

Award for Best Performance in 1999. He has worked in large-scale computing for nearly 37 years and is presently leader of the Scientific Computing Group in the Center for Applied Scientific Computing.

- Raymond Loy has over 10 years of experience with parallel message-passing performance and load-balancing of structured, block-structured, and unstructured computations in a variety of codes including CCSM and the FLASH code from the Advanced Simulation and Computing (ASC)/Alliances Center for Astrophysical Thermonuclear Flashes. He is the author of Autopack, a tool for MPI-based programs. Autopack is a library providing automatic message-packing support for an application, reducing the burden on the application scientist, and easing platform-specific tuning. An early user of BG/L, he has ported several versions of POP to BG/L and has assisted in the BG/L port of other applications at several BlueGene Consortium Application Workshops. He has performed message-passing performance analysis of POP and the full CCSM on various platforms and has contributed to development of the Model Coupling Toolkit [Larson2005], which is used in CPL6 [Craig2005]. Loy has participated in a Terascale Simulation Tools and Technologies working group as well as the Common Component Architecture Forum.

All three members of the team have close working relationships with the computational scientists in SEESM, and have prior experience working with the ocean, sea ice, land, atmosphere, and coupler developers. For this SAP, the science application partners are John Drake (ORNL) for the atmosphere, Phil Jones (Los Alamos National Laboratory) for the ocean, Philip Cameron-Smith (LLNL) for the atmospheric chemistry, and Robert Jacob (ANL) for the coupler and coupler-level interactions within the full CCSM.

2.2.3 Purpose

The main purpose of this work is to analyze and optimize CCSM so that it can simulate tomorrow's science at the same throughput as CCSM simulates today's science. This will be accomplished through increased scalability, performance analysis and improvement, extension to new architectures, and application and exploitation of performance prediction techniques.

2.3 Research Design and Methods

As indicated in Section 2.2.1, we are concerned with CCSM performance and scalability when introducing new physics, new problem sizes, and new architectures. Here we describe the approaches we expect to take to address these concerns. We begin with a description of approaches for the individual component models, in particular, the atmosphere. However, a number of the issues we will be facing may not become clear until the CCSM model development is more advanced, and a careful performance analysis and tracking, in coordination with internal SESSM efforts, is an important aspect of this proposal. This will require the exploitation of sophisticated tools and techniques for performance instrumentation and analysis. To this end, we plan to work with the SciDAC2 Center for Enabling PERC.

2.3.1 Model Performance and Scalability

2.3.1.1 CAM Scalability

Performance Instrumentation and Analysis

We will begin with the performance analysis and performance modeling of CAM, to determine a performance baseline before the new physical processes are introduced. As CAM evolves, we will periodically update this analysis, using it to predict or identify scalability and other performance problems early in the development process. It is important that state-of-the-art performance tools, such as Integrated Performance Monitoring (IPM) [IPM], Message Passing Interface (MPI) Parallel Environment (MPE)/Jumpshot [Chan2002], or Tuning and Analysis Utilities (TAU) [Mohr1994, Malony2004], be used in this activity. For example, we recently used IPM to compare communication patterns in 1-D and 2-D decompositions of CAM (see Fig. 2), where the diagram shows latitudinal border communications (main diagonal), vertical geopotential communications (parallel to and offset from main diagonal) and transpose communications (diamond pattern).

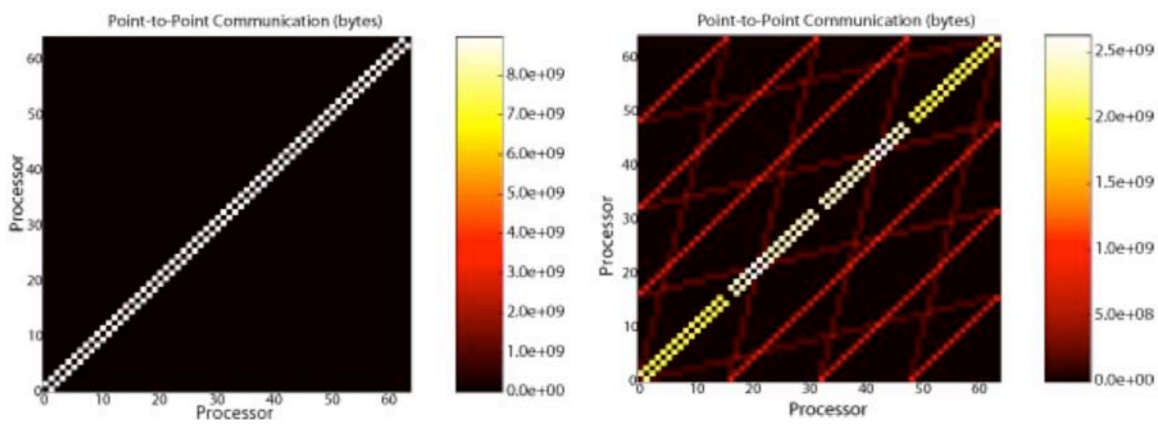

Fig. 2. Communication patterns in CAM as measured by the IPM tool.

The results from this analysis, showing over double the traffic with the 1-D decomposition (top) versus 2-D (bottom), were in contradiction to our conventional wisdom that the transposes associated with the 2-D decomposition technique would dominate. This reinforces the importance of such tools.

Much of the instrumentation and analysis will take place within the SEESM project. We are primarily consumers of these data, but will also advise on the process. Performance models and other prediction techniques are our primary contribution to this aspect of the work. The performance analysis process is iterative: posing questions, collecting data, analyzing data, reformulating questions. As such, performance models can be extremely useful in encapsulating current knowledge about performance and directing future experiments. In particular, the cost of the performance experiments can be high, and performance models also allow extrapolation to scenarios that cannot be easily tested, such as new problem sizes or new architectures. While performance models have proven their worth in architectural evaluations and algorithm evaluations, current modeling techniques make them less useful for development and optimization of codes as complex as CAM. As such, CAM performance modeling is a higher risk (and high payoff) activity than the other work described in this proposal. While modeling will not be the highest priority within the project, we do plan on pursuing the activity in collaboration with our PERC colleagues at the San Diego Supercomputing Center, at Lawrence Livermore National Laboratory (LLNL), and at ORNL.

Optimization of CAM in Chemistry-Dominated Regime

While many of the proposed modifications to the physical processes in CAM will have significant impacts on performance and scalability, here we focus on the atmospheric chemistry. The performance impact of atmospheric chemistry is better understood than that for some of the other proposed changes, and the issues and solutions are generally applicable.

Up until now, climate simulations using CAM, while modeling many important physical effects, have only marginally taken into account chemical processes. This has been due largely to the expense involved with modeling chemistry along with climate. Next-generation climate simulations require inclusion of these important chemical mechanisms. From a performance standpoint this will require analysis of regimes containing hundreds of advected constituents (instead of only a few). Also, the balance of computation involved in columnar processes (e.g., model physics, chemistry) versus the dynamical core will put us into a whole new regime, thereby requiring us to take a fresh look at parallelization methodology and scaling. For example, Fig. 3 describes the scaling of the dynamics versus the physics for the current model on the Cray X1E and on the IBM p575 cluster. Here, dynamics dominates for even small processor counts, and does not scale as well as the physics. With the introduction of atmospheric chemistry, the physics will become much more expensive.

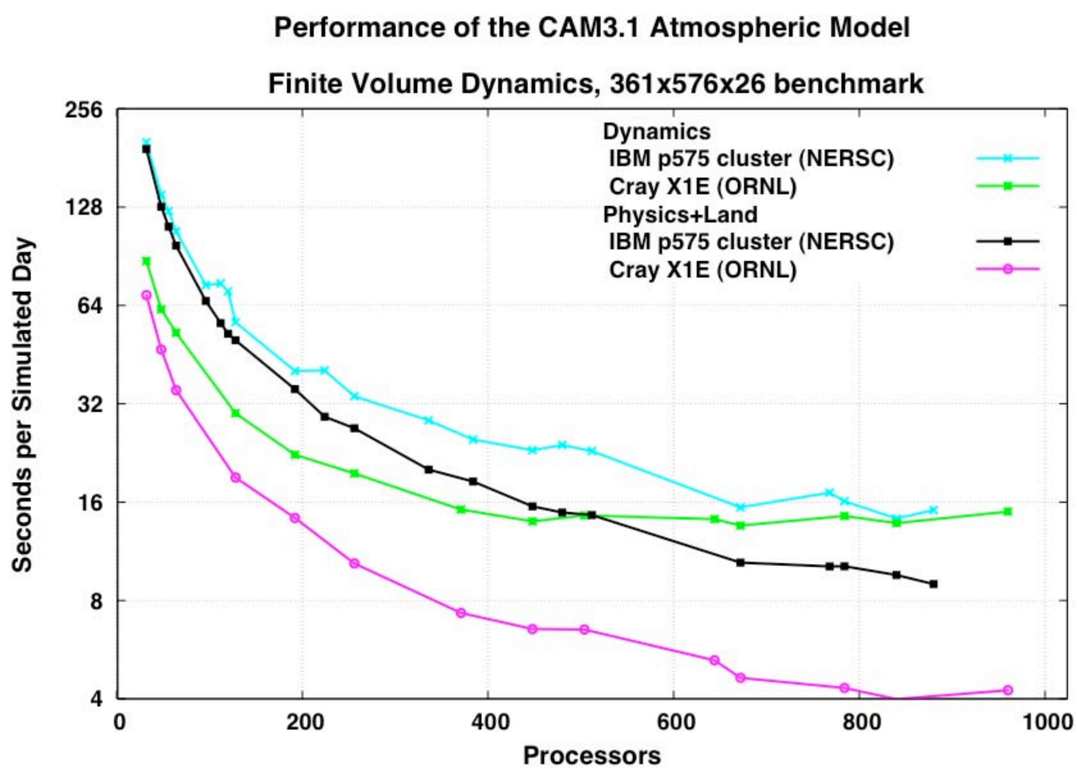

Fig. 3. Scalability of CAM dynamics versus physics.

Preliminary scaling tests show that, with the current CAM model, advecting 100 tracers more than quadruples the execution time. There is a memory cost as well to store the additional 3-D tracer arrays, and that can become a limiting factor at moderate resolutions on certain architectures. It is therefore all the more important to be able to amortize the memory costs by scaling to a very large number of processors. We propose to in effect carry out a three-dimensional decomposition of the tracer advection phase of the dycore, with the third dimension being the tracer index. Based on our experience with the multi-two-dimensional domain decomposition in the FV dycore, we expect the increased communication costs of the tracer decomposition to be manageable. The tracer advection phase also involves some border communication. We will aim to hide this border communication by communicating small groups of tracers using nonblocking primitives. We will investigate optimal sizes of those blocks to maximize the overlap with computation.

In the chemistry-dominated regime, the highest priority will be given to optimizing performance of the chemical and physical processes, which operate along vertical columns independent of one-another. Much effort has already gone into parallelizing the columnar processes using a general chunking methodology [Worley2005]. That technology generally scales well to large numbers of processes and in terms of process count should match up well with the tracer decomposition. If the current two-dimensional decomposition approach is not sufficient, much of the atmospheric chemistry is independent between grid points, not just between columns, and a full three-dimensional decomposition just for the atmospheric chemistry may be required. One concern though is with vector machines, where at fine decomposition, the work per computational process might result in inefficient vector lengths. This will need to be considered in the overall scaling strategy.

The chemistry-dominated regime also requires a different approach to load-balance. There is presently a choice of how to decompose the horizontal mesh among processes. The default choice has been to use a physics decomposition identical to that of the dynamics. However, such a decomposition is generally not load-balanced due to variations in computational work associated with solar radiation. We have also provided several alternatives that take load-balancing into account in the decomposition. These alternatives require an additional transpose between the dynamics and physics decompositions, While the communication costs associated with load balancing are typically smaller than the savings brought about by improved load-balance, the performance enhancement due to load balancing is sensitive to computer system, processor count, and problem size. For example, on the Cray X1E using the spectral Eurlerian dycore and the production problem size, load balancing improves overall CAM performance by 15% for all processor counts [Worley2006]. In contrast on an IBM p690 cluster, load balancing is only useful for smaller processor counts. In the chemistry-dominated regime, however, we expect the equation to change and for load-balanced decompositions to be much superior in all situations.

Up to now the alternative decompositions have been from among several fixed choices. With the added costs of chemistry, refining the load-balance will be all the more important. We plan on generalizing the decomposition choices as well as implementing a dynamic load-balancing scheme. The dynamic decomposition choice will depend on statistics over previous time steps. We will also need to consider decomposition startup costs in determining the frequency with which to alter decompositions.

The scenario being considered involves four or five domain decompositions – two for the main dynamics (already implemented in the FV dycore), one for the tracer advection, and one or two for the chemistry and physics. We see no way to increase the process count associated with the primary dycore evolution using the currently available dycores, beyond that which will come naturally from increased computational grid resolution. Therefore, we will need to allow different phases of the evolution to use different numbers of processes. This is supported in the current atmosphere model through the use of OpenMP parallelism. We will need to adapt and generalize this approach to also work in an MPI-only environment.

Extension to Cubed Sphere Grid

Under the direction of S-J Lin, one of the inventors of the FV dycore, work is under way to extend the dynamical core to the cubed sphere grid [Rancic1996]. This is motivated by the well-known drawbacks of the polar singularity of the latitude/longitude grid. The cubed sphere version of the dycore is expected to be available in early CY 2007. An alternative dycore being developed using a discontinuous Galerkin (DG) discretization on the cubed sphere will also be available in CAM in the same time frame. Elimination of the polar singularity will allow decomposition in all three dimensions – latitude, longitude and vertical. We plan to collaborate with SEESM on a performance analysis and modeling of the cubed sphere version of FV and DG as soon as they become available. If, as expected, the cubed sphere implementation allows scaling to a much larger process count, we will immediately begin experiments to determine how best to exploit this scalability in the context of CAM and a chemical atmosphere.

2.3.1.2 Scalability of Other Component Models

As mentioned in Section 2.2.1, a new version of the ocean model, POP 2, has recently been developed that addresses many of the known scalability and performance problems of the current production version. Until the performance of this new ocean model is carefully analyzed, it is difficult to identify what more needs to be done. However, preliminary performance instrumentation indicates that there may be an opportunity for further optimizations. For example, MPE has been used to analyze the intra-model message passing performance of POP 2.0 (see Fig. 4). Results suggest possible efficiency gains through use of Autopack [Loy2000], a message-passing library developed at ANL that transparently packs small messages into fewer larger ones for more efficient transport by MPI. Performance models of POP have been developed by two different research groups [Kerbyson2005, Snavely2003]. We will examine using these to identify performance optimization opportunities. As these models do not incorporate the new POP2 features, we will also update the models or generate a new performance model that allows us to determine performance and scalability bottlenecks.

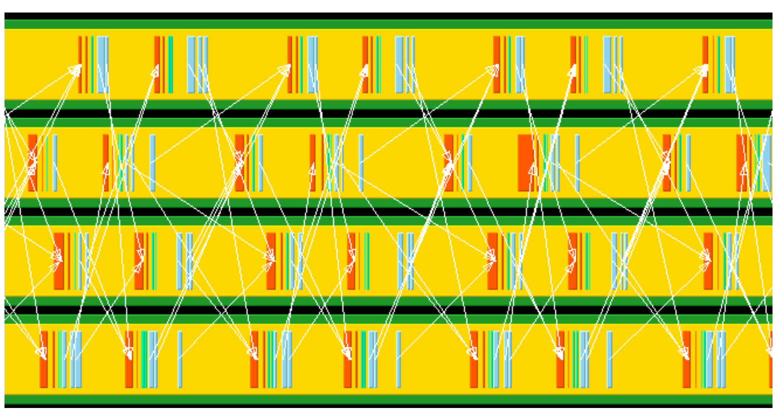

Fig. 4. Jumpshot visualization of the POP baroclinic solver running a test case on 4 processors. The yellow background is the boundary condition calculation within the baroclinic solver (green). Colored vertical bars represent time spent in various MPI calls: MPI_Isend() (light blue), MPI_Irecv() (turquoise), and MPI_Waitall() (red). Direct arrows connect MPI send and receive operations on source and target processors, respectively.

The surface (land, ice) models will also be undergoing rapid development, and their performance will need to be tracked as well (as part of SEESM). We do not expect to focus on the surface models during the duration of this proposal but will adjust our efforts as necessary if the performance analysis indicates that they have become significant limiters to scalability, and if our expertise in modeling and parallel algorithm development are deemed by the Science Application leads to be of use.

2.3.1.3 CCSM Modeling and Configuration Optimization

Optimizing the allocation of processors to CCSM is very challenging. It is a combination of five separate parallel codes having an asymmetric communication pattern, with the associated communication barriers preventing optimal allocation of processors based solely on component throughput. Deciding on a CCSM configuration has been accomplished largely by trial and error.

More sophisticated approaches utilizing performance tools have been tried recently. For example, the MPE tool has been used in an attempt to gain a better understanding of the interaction between CCSM components. While MPE can automatically log all MPI activity, CCSM has been manually instrumented to provide specific information about message passing activity between the different model components, facilitating determination of the optimum load balance (see Fig. 5). This type of work will continue within SEESM, representing an improvement over past practice.

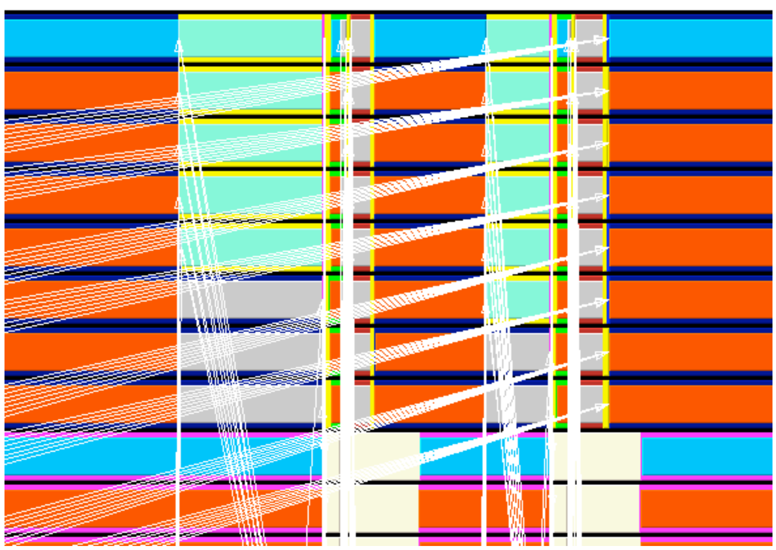

Fig. 5. Each horizontal bar depicts the timeline of a single processor. The color indicates what state the processor is in. As shown in the example above, states may be nested. Arrows denote messages being sent between processors.

Although MPE provides some insight into the workings of CCSM, a much more comprehensive performance modeling effort appears to be needed. Modeling the full CCSM with the goal of optimizing the configuration process has different accuracy requirements than identifying scalability issues in the component models. The full CCSM is also much too complicated a code to model adequately within the time frame of this project using standard techniques. Instead we will use a technique undergoing development in PERC, namely that of artificial neural networks (ANNs) trained on a sampling of performance data over a parameter space [Ipek2005]. The criteria will consist of machine and configuration parameters. The idea will be to choose judicious sets of parameters at which the code will be evaluated, hoping to extrapolate results to parameter ranges of interest. B. de Supinski of LLNL, one of the pioneers of the ANN approach, has agreed to work with us at no cost. The neural network approach avoids the challenge of representing the inner workings of the code but provides little insight into those inner workings. Hence we will augment the ANN approach with data from more traditionally based performance models of the component models, to provide an integrated performance modeling capability for CCSM.

2.3.2 CCSM Portability and Exploitation of New Architectures

A second focus of this proposal is preparing CCSM for porting to the next generation of massively parallel computer systems. The initial targets of the Cray X1E, Cray XT3, IBM p575 SMP cluster, and IBM BG/L cover most of the currently available architectures. Porting to the BG/L is the most challenging of these four, and a successful port to the BG/L will be the most useful preparation for future systems, requiring as it does both increased scalability and improved memory and I/O management. In the rest of this section we focus exclusively on the challenges and approaches to porting and optimizing the CCSM on the BG/L system.

2.3.2.1 Porting CCSM to BG/L

A main topic to address for the CCSM port to BlueGene/L will be global arrays associated with parallel I/O. All CCSM component models perform I/O to create large history and restart files. Binary mode file writes may be performed in parallel for POP2, though the current implementation of this feature has not been ported successfully to all target platforms. The preferred format for history output is NetCDF [NetCDF] and support in POP2 for NetCDF is currently only serial. A version of CAM that uses ZioLib [Yang2003] and Parallel NetCDF [Li2003] is available and we will transfer these capabilities to CCSM when porting to BG/L. Parallel-NetCDF will then be introduced in other components as necessary to improve scalability and performance.

Under the first SciDAC program, the Los Alamos standalone versions of POP 1.4.3 and POP 2.0 were ported to BG/L and good scaling was observed (see Fig. 6). POP2 will become the default ocean model in the development version of CCSM within this calendar year. Additional work may be required to parallelize diagnostic calculations that are added to POP by the CCSM ocean model working group.

A port of CAM/CLM has begun at the National Center for Atmospheric Research (NCAR) by programmers in the Scientific Computing Division. They found parallel I/O was necessary for CAM's binary restart files. We will join this effort and provide additional resources to ensure any required code changes are integrated with the development branches of CAM and CLM. The performance profiling and alternate parallelization schemes outlined above for CAM will ported to BG/L as soon as they have been validated. (To ensure that the parallel algorithm work is incorporated into the production versions of the component models as soon as possible, we will focus most of the initial development work of the non-BG/L specific modifications on the current production platforms: Cray X1E, Cray XT3, and IBM p575 cluster.)

Once the standalone versions of POP and CAM are ported, we shall port the simplest coupled system in CCSM, the coupler interacting with 4 "dead" models, to BG/L. The dead models are small parallel programs that send simple two-dimensional sine waves to the coupler for data fields. Since all modes of CCSM, full, data and dead models, use the same coupler "hub", this system will allow us to study the performance of the hub on BG/L. The coupler's underlying library, the Model Coupling Toolkit [Larson2005], has already been ported to BlueGene/L.

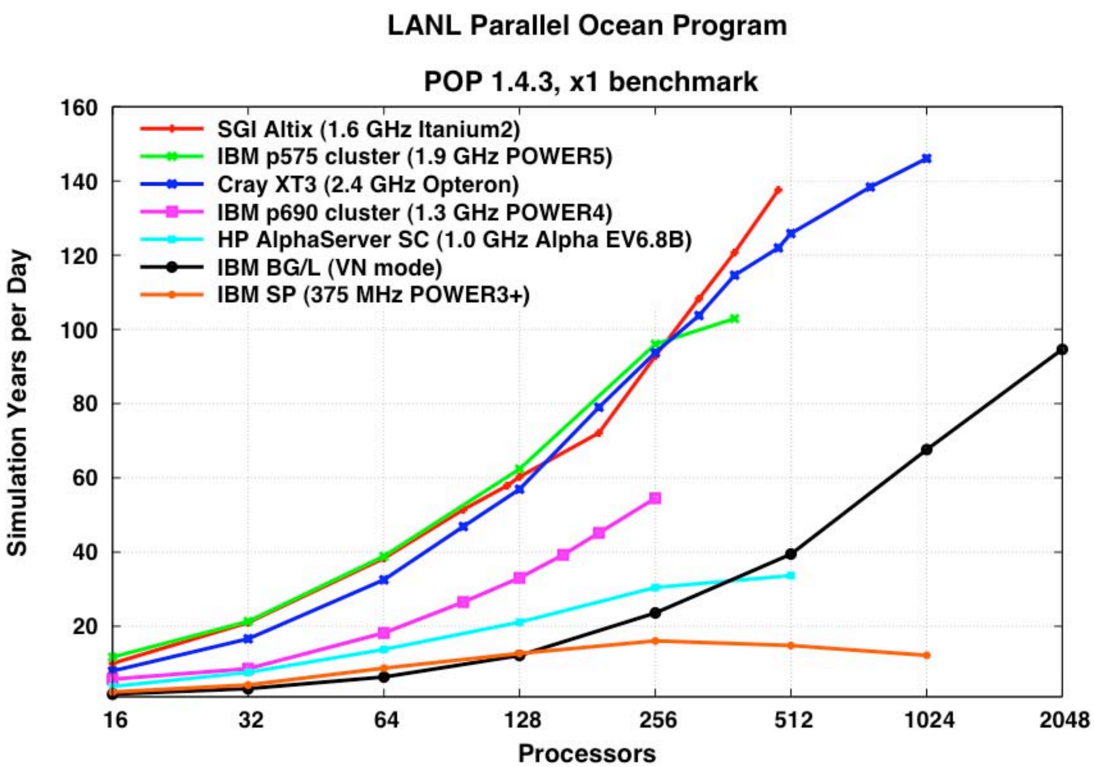

Fig. 6. Performance of POP 1.4.3 "x1" Benchmark on the ANL BG/L. Scaling for virtual node (VN) mode is shown in comparison to other platforms.

BG/L system software has the limitation that it can only load and run a single executable, unlike other MPP systems that can load a different MPI executable on each node of a run. While CCSM components are traditionally run in multi-executable configuration, work is currently underway by CCSM participants to produce a single-executable version of CCSM. Completion of this work will be supported by SEESM. If a single executable is not available when we are ready to try the dead-model case, we will assist this effort.

Once a single-executable, dead-model case is running on BlueGene, we will proceed to the all-data model case. Data models send real fields read from files to the coupler but perform no internal calculations other than possibly time interpolation. The data-models are currently not parallel but they are necessary when running with at least one full model (dead and full models can not be run together). The next step after the all-data model case is working will be to add one or two full models at a time starting with the "G-case" which uses an active ocean and ice model but data land and atmosphere models. The final goal will be a fully coupled case

Once CCSM is working and scaling reasonably well on the BG/L system at ANL, we will perform enhanced scalability tests of the large LLNL BlueGene system. The purpose will be not only to better understand CCSM scaling to very large processor count, but also to understand the applicability of the BG/L architecture to a broader set of algorithms.

The next generation BlueGene system, BG/P, is expected to have increased performance over BG/L through a processor clock speedup of 20 percent and increased inter-node communication bandwidth. Each node will be 4-way SMP, rather than BG/L's 2-way nodes which have only limited ability for intra-node communication due to lack of cache coherency. While BG/P is expected to have approximately twice the memory per processor, compared to other massively parallel systems it remains low in its ratio of memory per processor owing to the basic design goal of balancing heat generation, package density, and clock rate. BG/P will offer the possibility of shared memory parallelism as well. We will attempt to run with hybrid parallelism provided the system support is there. We expect BG/P systems to be available near the end of this project and will port our BG/L-capable version of CCSM to this platform.

2.4 Consortium Arrangements

2.4.1 Overall Management

Patrick Worley is the principal investigator for the project, and will determine overall project direction and milestones. Arthur Mirin has primary responsibility for the investigation of CAM dycore scalability and CCSM configuration optimization. Raymond Loy has primary responsibility for the port of CCSM to the BG/L system. Worley has primary responsibility for CAM physics scalability and POP scalability analysis. All three researchers have prior experience working with each other, and we expect the collaboration to be close on all of these areas.

The SAP is motivated and guided by the needs of the SEESM project. Worley will discuss progress and goals at least monthly with the SEESM principal investigator John Drake (ORNL), and members of the SAP will confer with other members of SEESM as needed. In particular, Phil Jones (LANL) is the liaison between the SAP and SEESM for the ocean model, Phillip Cameron-Smith (LLNL) is the liaison between the SAP and SEESM for atmospheric chemistry, and Robert Jacob (ANL) is the liaison between the SAP and SEESM for the coupler and coupler-level interactions within the full CCSM.

Mirin and Worley are also members of SEESM, with responsibilities for the short term performance engineering needs of SEESM, as described in the SEESM proposal, and are well-positioned to determine scalability issues during future CCSM development within SEESM. Worley is the liaison between the SAP, SEESM, and the PERC center for enabling technology, Worley will provide guidance on the application of PERC technology to the CCSM, and provide knowledge of CCSM performance issues to PERC researchers. Finally, Worley is also a co-chair of the CCSM Software Engineering Work Group, which will enable the SAP to stay coordinated with the activities of other CCSM software engineering activities that impact CCSM performance and scalability.

2.4.2 Software Management and Software Engineering

The SAP contributions to the CCSM are via the SEESM project, and so utilize and depend on the same software management as SEESM. The follow description is from the SEESM proposal: *Mariana Vertenstein at NCAR manages access to the development team repository. Each component model is quality assured by a gatekeeper or change review board. Permission to commit changes to the development trunk is, however, a critical item for the rapid advancement of the model by researchers on this proposal. This permission was granted and exercised in SciDAC 1. The distribution outside this consortium of code under development is strictly prohibited, though researchers engaged with the CCSM Working groups have the ability to manage branches and perform simulations using the development code. Released code is publicly available and freely distributed.*

2.4.3 Project Plans and Milestones

The following represent project milestones by fiscal year. The lead institution for the given milestone is listed in parentheses. However in most cases, more than one institution will be contributing to the work.

FY06 (last 3 months)

1. Work with SEESM to analyze current CAM performance on representative architectures. Quantify nature of scalability bottlenecks in terms of problem size, processor count, and computer system. (ORNL)
2. Characterize performance impact of additional tracers on CAM performance. (LLNL)
3. Port CAM/CLM to BG/L at ANL using NCAR port. Analyze performance and identify current bottlenecks. (ANL)

FY07

1. Construct performance model of current CAM physics. (ORNL)
2. Work with SEESM to analyze load imbalances introduced by atmospheric chemistry. (ORNL)
3. Optimize physics load-balance in chemistry-dominated regime using static load balancing. (ORNL)
4. Work with SEESM to analyze performance of POP2 on Cray X1E, Cray XT3, and IBM p575 cluster. (ORNL)
5. Optimize CAM at large tracer count, including decomposition over tracers. (LLNL)
6. Optimize chemical mechanisms, including decomposition over vertical levels. (LLNL)
7. Eliminate global arrays and implement parallel I/O throughout CCSM components. (ANL)
8. Optimize CAM/CLM on BG/L. (ANL)
9. Port CICE to BG/L. (ANL)

FY08

1. Update/generate POP2 performance model, compare with empirical performance characterization, and use to predict scalability. (ORNL)
2. Construct performance model of CAM physics with atmospheric chemistry. (ORNL)
3. Implement dynamic load-balance capability in chemistry-dominated regime. (ORNL)
4. Extend capability to vary process count over various phases of computation to include atmospheric chemistry and tracer advection. (LLNL)
5. Develop and implement ANN performance model for CAM and POP. (LLNL)
6. Port full CCSM to BG/L. (ANL)
7. Optimize CCSM on BG/L. (ANL)
8. Validate CCSM simulated climate on BG/L. (ANL)

FY09 (first 9 months)

1. Generate performance model for CAM, including support for varying processor count between different computational phases. Use model to predict scalability and optimal processor allocations. (ORNL)
2. Analyze and optimize DG dycore for cubed sphere grid, for both current
3. and high resolutions. (ORNL)
4. Analyze and optimize FV dycore for cubed sphere grid, for both current and high resolutions. (LLNL)
5. Develop and implement ANN performance model for full CCSM; Demonstrate CCSM configuration optimization using the model. (LLNL)
6. Update CCSM port to include new physics and new parallel algorithms. (ANL)
7. Port CCSM to BG/P. (ANL)
8. Determine optimal load balance and optimize communication performance of CCSM on BG/P. (ANL)
9. Perform extended scalability tests on LLNL BlueGene system. (LLNL)

2.4.4 Other SciDAC Interactions

As described throughout section 2.3 and again in section 2.4.1, this SAP expects to work closely with the SciDAC Center for Enabling Technology *Performance Engineering Research Center*. From PERC we expect guidance on performance tools and on performance analysis, optimization, and prediction techniques. In particular, PERC involvement will be important in the CCSM configuration optimization research. In turn, the SAP and SEESM will provide insight into the performance needs of CCSM, accelerating the development of PERC modeling and optimization technology. The SAP and performance engineering funded directly by SEESM will provide the necessary mechanism for PERC to be able to work with us.

3 Literature Cited

[Chan2002] Chan, A.,W. Gropp, and E. Lusk. "Scalable Log Files for Parallel Program Trace Data." Argonne National Laboratory Technical Report, ANL/MCS-TM-256, 2002.

[Craig2005] Craig, A. , R. Jacob, B. Kaufmann, T. Bettge, J. Larson, E.Ong, C. Ding and Y. He. "CPL6: The New Extensible, High-Performance Parallel Coupler for the Community Climate System Model." *Int'l. Jour. High Performance Computing Applications* **19** (3), 2005, pp. 309-327.

[Drake1995] Drake, J. B., I. T. Foster, J. G. Michalakes, B. Toonen and P. H. Worley. "Design and Performance of a Scalable Parallel Community Climate Model." *Parallel Computing*, **21** (10), October 1995, pp. 1571-1592.

[Ipek2005] Ipek, E., B.R. de Supinski, M. Schulz and S.A. McKee. "An Approach to Performance Prediction for Parallel Applications." *Proc. of Euro-Par 2005*, Lisbon, Portugal, Aug. 2005.

[Jones2005] Jones, P. W., P. H. Worley, Y. Yoshida, J. B. White III and J. Levesque. "Practical performance portability in the Parallel Ocean Program (POP)." *Concurrency - Practice and Experience* **17** (10): 1317-1327 (2005).

[Kerbyson2005] Kerbyson, D. J. and P.W. Jones. "A Performance Model of the Parallel Ocean Program." *Int'l. Jour. High Performance Computing Applications* **19** (3), 2005, pp. 261-276.

[Foster1997] Foster, I. T. and P. H. Worley. "Parallel algorithms for the spectral transform method." *SIAM J. Sci. Stat. Comput.* **18** (3), 1997, pp. 806-837.

[IPM] "Integrated Performance Monitor." http://ipm-hpc.sourceforge.net/.

[Larson1995] Larson, J., R. Jacob and E. Ong. "The Model Coupling Toolkit: A new Fortran90 Toolkit for Building Multi-physics Parallel Coupled Models." *Int'l. Jour. High Performance Computing Applications* **19 (**3), 2005, pp. 277-292.

[Li2003] Li, J., W-K. Liao, A. Choudhary, R. Ross, R. Thakur, W. Gropp, R. Latham, A. Siegel, B. Gallagher and M. Zingale. ``Parallel netCDF: A Scientific High-Performance I/O Interface." *Proc. of Supercomputing Conference (SC03)*, November, 2003.

[Lin2004] Lin, S-J., "A vertically Lagrangian Finite-Volume Dynamical Core for Global Models." *Mon. Wea. Rev.*, **132** (3), 2004, 2293.

[Loy2000] Loy, R. M., "Autopack Version 1.3 User Manual." Argonne National Laboratory, 2000.

[Malony2004] A. Malony and S. Shende. "Overhead Compensation in Performance Profiling." *Proc. of Euro-Par Conference*, LNCS 3149, Springer, 2004, pg 119-132.

[Mohr1994] Mohr, B., D. Brown and A. Malony. “TAU: A Portable Parallel Program Analysis Environment for pC++.” *Proceedings of CONPAR 94 - VAPP VI*, University of Linz, Austria, LNCS 854, Sept. 1994, pg. 29-40.

[Mirin2005] Mirin, A. A. and W. B. Sawyer. “A Scalable Implementation of a Finite-Volume Dynamical Core in the Community Atmospheric Model.” *Int'l. Jour. High Performance Computing Applications* **19** (3), 2005, pp. 203-212.

[NetCDF] “NetCDF.” http://www.unidata.ucar.edu/software/netcdf/ .

[Oliker2005] Oliker, L., J. Carter, M. Wehner, A. Canning, S. Ethier, G. Bala, A. Mirin, D. Parks, P. Worley, S. Kitawaki and Y. Tsuda. “Leading Computational Methods on Scalar and Vector HEC Platforms.” *Proc. Supercomputing 2005 Conference (SC05)*, Seattle (2005).

[Rancic1996] Rancic, M., R.J. Purser and F. Mesinger, 1996. “A global shallow-water model using an expanded spherical cube: gnomic versus conformal coordinates.” *Q. J. R. Met. Soc.* **122**, pp. 959-982.

[Snavely2003] Snavely, A., X. Gao, C. Lee, L. Carrington, N. Wolter, J. Labarta, J. Gimenez and P. Jones. “Performance Modeling of HPC applications.” *Proc. Parallel Computing 2003*, Dresden, Germany, 2003.

[Worley2005] Worley, P. H. and J. B. Drake. “Performance Portability in the Physical Parameterizations of the Community Atmosphere Model.” *Int'l. Jour. High Performance Computing Applications* **19** (3), 2005, pp. 187-201.

[Worley2006] Worley, P. H. “Benchmarking using the Community Atmospheric Model.” *Proc. 2006 SPEC Benchmark Workshop*, Austin, TX (2006).

[Yang2003] Yang, W.-S. and C. Ding (2003). “ZioLib: A parallel I/O library.” Lawrence Berkeley National Laboratory Technical Report, LBNL-53521, 2003.

4 Budget and Budget Explanation

4.1 Budgets

Proposed Budget by Collaborator				
Institution	FY006	FY007	FY008	FY009
ORNL	$27,000	$108,000	$108,000	$81,000
Argonne National Laboratory	$56,000	$224,000	$224,000	$168,000
Livermore National Laboratory	$42,000	$168,000	$168,000	$126,000
Total Proposed Budget	$125,000	$500,000	$500,000	$375,000

Cover Page

Title of Proposed Project:

Performance Engineering for the Next Generation
Community Climate System Model
FWP # 57648

Office of Science Program Announcement Title/#:

Lab 06-04
Scientific Discovery through Advanced Computing

Name of Laboratory:

Argonne National Laboratory

Principal Investigator(s):

Raymond M. Loy, Senior Software Developer
Argonne National Laboratory
Mathematics and Computer Science Division
9700 So. Cass Avenue - Bld.221
Argonne, IL 60439
(T) 630-252-7205 (Fax) 630-252-6104
rloy@mcs.anl.gov

Official signing for Laboratory:

Rick L. Stevens, Associate Laboratory Director-PBCS
Argonne National Laboratory
Mathematics and Computer Science Division
9700 So. Cass Avenue - Bld.221
Argonne, IL 60439
(T) 630-252-3378, (Fax) 630-252-6333
stevens@mcs.anl.gov

Requested funding for Argonne National Laboratory for Proposed Project:

Year 1	$	224,000
Year 2	$	224,000
Year 3	$	224,000
Total:	$	672,000

Duration of Entire Project Period:

07/01/2006 to 06/30/2009

Use of human subjects in proposed project: No
Use of vertebrate animals in proposed project: No

Signature of PI, Date of Signature:

PI's electronic signature on file — March 6, 2006

Signature of Official, Date of Signature:

March 6, 2006

DOE F 4620.1	**U.S. Department of Energy**	OMB Control No.
(04-93)	**Budget Page**	1910-1400
All Other Editions Are Obsolete	(See reverse for Instructions)	OMB Burden Disclosure Statement on Reverse

ORGANIZATION The University of Chicago, Operator of Argonne National Laboratory				**Budget Page No:** pg 1 of 4 Yr 1 of 3	
PRINCIPAL INVESTIGATOR/PROJECT DIRECTOR Raymond M. Loy				Requested Duration: 12 (Months) FWP # 57648	
A. SENIOR PERSONNEL: PI/PD, Co-PI's, Faculty and Other Senior Associates (List each separately with title; A.6. show number in brackets)	DOE Funded Person-mos.			Funds Requested	Funds Granted
	CAL	ACAD	SUMR	by Applicant	by DOE
1. Raymond M. Loy, PI	9.00			$159,119	
2.					
3.					
4.					
5.					
6. ( ) OTHERS (LIST INDIVIDUALLY ON BUDGET EXPLANATION PAGE)					
7. (1) TOTAL SENIOR PERSONNEL (1-6)	9.00			$159,119	
B. OTHER PERSONNEL (SHOW NUMBERS IN BRACKETS)					
1. ( ) POST DOCTORAL ASSOCIATES					
2. ( ) OTHER PROFESSIONAL (TECHNICIAN, PROGRAMMER, ETC.)					
3. ( ) GRADUATE STUDENTS					
4. ( ) UNDERGRADUATE STUDENTS					
5. ( ) SECRETARIAL - CLERICAL					
6. ( ) OTHER					
TOTAL SALARIES AND WAGES (A+B)				$159,119	
C. FRINGE BENEFITS (IF CHARGED AS DIRECT COSTS)					
TOTAL SALARIES, WAGES AND FRINGE BENEFITS (A+B+C)				$159,119	
D. PERMANENT EQUIPMENT (LIST ITEM AND DOLLAR AMOUNT FOR EACH ITEM.)					
TOTAL PERMANENT EQUIPMENT					
E. TRAVEL 1. DOMESTIC (INCL. CANADA AND U.S. POSSESSIONS)				$2,800	
2. FOREIGN					
TOTAL TRAVEL				$2,800	
F. TRAINEE/PARTICIPANT COSTS					
1. STIPENDS (Itemize levels, types + totals on budget justification page)					
2. TUITION & FEES					
3. TRAINEE TRAVEL					
4. OTHER (fully explain on justification page)					
TOTAL PARTICIPANTS ( ) TOTAL COST					
G. OTHER DIRECT COSTS					
1. MATERIALS AND SUPPLIES				$2,981	
2. PUBLICATION COSTS/DOCUMENTATION/DISSEMINATION					
3. CONSULTANT SERVICES					
4. COMPUTER (ADPE) SERVICES					
5. SUBCONTRACTS					
6. OTHER					
TOTAL OTHER DIRECT COSTS				$2,981	
H. TOTAL DIRECT COSTS (A THROUGH G)				$164,900	
I. INDIRECT COSTS (SPECIFY RATE AND BASE) Section H. Direct cost X Aggregate rate of: 35.840%					
TOTAL INDIRECT COSTS				$59,100	
J. TOTAL DIRECT AND INDIRECT COSTS (H+I)				$224,000	
K. AMOUNT OF ANY REQUIRED COST SHARING FROM NON-FEDERAL SOURCES					
L. TOTAL COST OF PROJECT (J+K)				$224,000	

DOE F 4620.1
(04-93)
All Other Editions Are Obsolete

U.S. Department of Energy

Budget Page

(See reverse for Instructions)

OMB Control No.
1910-1400
OMB Burden Disclosure
Statement on Reverse

ORGANIZATION The University of Chicago, Operator of Argonne National Laboratory				**Budget Page No:** 2 of 4 Yr. 2 of 3	
PRINCIPAL INVESTIGATOR/PROJECT DIRECTOR Raymond M. Loy				Requested Duration: 12 (Months) FWP # 57648	
A. SENIOR PERSONNEL: PI/PD, Co-PI's, Faculty and Other Senior Associates (List each separately with title; A.6. show number in brackets)	DOE Funded Person-mos.			Funds Requested	Funds Granted
	CAL	ACAD	SUMR	by Applicant	by DOE
1. Raymond M. Loy, PI	8.80			$164,322	
2.					
3.					
4.					
5.					
6. ( ) OTHERS (LIST INDIVIDUALLY ON BUDGET EXPLANATION PAGE)					
7. (1) TOTAL SENIOR PERSONNEL (1-6)	8.80			$164,322	
B. OTHER PERSONNEL (SHOW NUMBERS IN BRACKETS)					
1. ( ) POST DOCTORAL ASSOCIATES					
2. ( ) OTHER PROFESSIONAL (TECHNICIAN, PROGRAMMER, ETC.)					
3. ( ) GRADUATE STUDENTS					
4. ( ) UNDERGRADUATE STUDENTS					
5. ( ) SECRETARIAL - CLERICAL					
6. ( ) OTHER					
TOTAL SALARIES AND WAGES (A+B)				$164,322	
C. FRINGE BENEFITS (IF CHARGED AS DIRECT COSTS)					
TOTAL SALARIES, WAGES AND FRINGE BENEFITS (A+B+C)				$164,322	
D. PERMANENT EQUIPMENT (LIST ITEM AND DOLLAR AMOUNT FOR EACH ITEM.)					
TOTAL PERMANENT EQUIPMENT					
E. TRAVEL 1. DOMESTIC (INCL. CANADA AND U.S. POSSESSIONS)					
2. FOREIGN					
TOTAL TRAVEL					
F. TRAINEE/PARTICIPANT COSTS					
1. STIPENDS (Itemize levels, types + totals on budget justification page)					
2. TUITION & FEES					
3. TRAINEE TRAVEL					
4. OTHER (fully explain on justification page)					
TOTAL PARTICIPANTS ( ) TOTAL COST					
G. OTHER DIRECT COSTS					
1. MATERIALS AND SUPPLIES				$578	
2. PUBLICATION COSTS/DOCUMENTATION/DISSEMINATION					
3. CONSULTANT SERVICES					
4. COMPUTER (ADPE) SERVICES					
5. SUBCONTRACTS					
6. OTHER					
TOTAL OTHER DIRECT COSTS				$578	
H. TOTAL DIRECT COSTS (A THROUGH G)				$164,900	
I. INDIRECT COSTS (SPECIFY RATE AND BASE) Section H. Direct cost X Aggregate rate of: 35.840%					
TOTAL INDIRECT COSTS				$59,100	
J. TOTAL DIRECT AND INDIRECT COSTS (H+I)				$224,000	
K. AMOUNT OF ANY REQUIRED COST SHARING FROM NON-FEDERAL SOURCES					
L. TOTAL COST OF PROJECT (J+K)				$224,000	

DOE F 4620.1
(04-93)
All Other Editions Are Obsolete

U.S. Department of Energy

Budget Page

(See reverse for Instructions)

OMB Control No.
1910-1400
OMB Burden Disclosure
Statement on Reverse

ORGANIZATION The University of Chicago, Operator of Argonne National Laboratory				**Budget Page No:** 3 of 4 Yr 3 of 3	
PRINCIPAL INVESTIGATOR/PROJECT DIRECTOR Raymond M. Loy				Requested Duration: 12 (Months) FWP # 57648	
A. SENIOR PERSONNEL: PI/PD, Co-PI's, Faculty and Other Senior Associates (List each separately with title; A.6. show number in brackets)	DOE Funded Person-mos.			Funds Requested	Funds Granted
	CAL	ACAD	SUMR	by Applicant	by DOE
1. Raymond M. Loy, PI	8.50			$164,530	
2.					
3.					
4.					
5.					
6. ( ) OTHERS (LIST INDIVIDUALLY ON BUDGET EXPLANATION PAGE)					
7. (1) TOTAL SENIOR PERSONNEL (1-6)	8.50			$164,530	
B. OTHER PERSONNEL (SHOW NUMBERS IN BRACKETS)					
1. ( ) POST DOCTORAL ASSOCIATES					
2. ( ) OTHER PROFESSIONAL (TECHNICIAN, PROGRAMMER, ETC.)					
3. ( ) GRADUATE STUDENTS					
4. ( ) UNDERGRADUATE STUDENTS					
5. ( ) SECRETARIAL - CLERICAL					
6. ( ) OTHER					
TOTAL SALARIES AND WAGES (A+B)				$164,530	
C. FRINGE BENEFITS (IF CHARGED AS DIRECT COSTS)					
TOTAL SALARIES, WAGES AND FRINGE BENEFITS (A+B+C)				$164,530	
D. PERMANENT EQUIPMENT (LIST ITEM AND DOLLAR AMOUNT FOR EACH ITEM.)					
TOTAL PERMANENT EQUIPMENT					
E. TRAVEL 1. DOMESTIC (INCL. CANADA AND U.S. POSSESSIONS)					
2. FOREIGN					
TOTAL TRAVEL					
F. TRAINEE/PARTICIPANT COSTS					
1. STIPENDS (Itemize levels, types + totals on budget justification page)					
2. TUITION & FEES					
3. TRAINEE TRAVEL					
4. OTHER (fully explain on justification page)					
TOTAL PARTICIPANTS ( ) TOTAL COST					
G. OTHER DIRECT COSTS					
1. MATERIALS AND SUPPLIES				$370	
2. PUBLICATION COSTS/DOCUMENTATION/DISSEMINATION					
3. CONSULTANT SERVICES					
4. COMPUTER (ADPE) SERVICES					
5. SUBCONTRACTS					
6. OTHER					
TOTAL OTHER DIRECT COSTS				$370	
H. TOTAL DIRECT COSTS (A THROUGH G)				$164,900	
I. INDIRECT COSTS (SPECIFY RATE AND BASE) Section H. Direct cost X Aggregate rate of: 35.840%					
TOTAL INDIRECT COSTS				$59,100	
J. TOTAL DIRECT AND INDIRECT COSTS (H+I)				$224,000	
K. AMOUNT OF ANY REQUIRED COST SHARING FROM NON-FEDERAL SOURCES					
L. TOTAL COST OF PROJECT (J+K)				$224,000	

DOE F 4620.1	**U.S. Department of Energy**	OMB Control No.
(04-93)	**Budget Page**	1910-1400
All Other Editions Are Obsolete	(See reverse for Instructions)	OMB Burden Disclosure Statement on Reverse

ORGANIZATION The University of Chicago, Operator of Argonne National Laboratory				**Budget Page No:** 4 of 4 3-Yr. ANL Total Project	
PRINCIPAL INVESTIGATOR/PROJECT DIRECTOR Raymond M. Loy				Requested Duration: 36 (Months) FWP # 57648	
A. SENIOR PERSONNEL: PI/PD, Co-PI's, Faculty and Other Senior Associates (List each separately with title; A.6. show number in brackets)	DOE Funded Person-mos.			Funds Requested	Funds Granted
	CAL	ACAD	SUMR	by Applicant	by DOE
1. Raymond M. Loy, PI	26.30			$487,971	
2.					
3.					
4.					
5.					
6. ( ) OTHERS (LIST INDIVIDUALLY ON BUDGET EXPLANATION PAGE)					
7. (1) TOTAL SENIOR PERSONNEL (1-6)	26.30			$487,971	
B. OTHER PERSONNEL (SHOW NUMBERS IN BRACKETS)					
1. ( ) POST DOCTORAL ASSOCIATES					
2. ( ) OTHER PROFESSIONAL (TECHNICIAN, PROGRAMMER, ETC.)					
3. ( ) GRADUATE STUDENTS					
4. ( ) UNDERGRADUATE STUDENTS					
5. ( ) SECRETARIAL - CLERICAL					
6. ( ) OTHER					
TOTAL SALARIES AND WAGES (A+B)				$487,971	
C. FRINGE BENEFITS (IF CHARGED AS DIRECT COSTS)					
TOTAL SALARIES, WAGES AND FRINGE BENEFITS (A+B+C)				$487,971	
D. PERMANENT EQUIPMENT (LIST ITEM AND DOLLAR AMOUNT FOR EACH ITEM.)					
TOTAL PERMANENT EQUIPMENT					
E. TRAVEL 1. DOMESTIC (INCL. CANADA AND U.S. POSSESSIONS)				$2,800	
2. FOREIGN					
TOTAL TRAVEL				$2,800	
F. TRAINEE/PARTICIPANT COSTS					
1. STIPENDS (Itemize levels, types + totals on budget justification page)					
2. TUITION & FEES					
3. TRAINEE TRAVEL					
4. OTHER (fully explain on justification page)					
TOTAL PARTICIPANTS ( ) TOTAL COST					
G. OTHER DIRECT COSTS					
1. MATERIALS AND SUPPLIES				$3,928	
2. PUBLICATION COSTS/DOCUMENTATION/DISSEMINATION					
3. CONSULTANT SERVICES					
4. COMPUTER (ADPE) SERVICES					
5. SUBCONTRACTS					
6. OTHER					
TOTAL OTHER DIRECT COSTS				$3,928	
H. TOTAL DIRECT COSTS (A THROUGH G)				$494,700	
I. INDIRECT COSTS (SPECIFY RATE AND BASE)					
TOTAL INDIRECT COSTS				$177,300	
J. TOTAL DIRECT AND INDIRECT COSTS (H+I)				$672,000	
K. AMOUNT OF ANY REQUIRED COST SHARING FROM NON-FEDERAL SOURCES					
L. TOTAL COST OF PROJECT (J+K)				$672,000	

Lab 06-04: Scientific Discovery through Advanced Computing
Performance Engineering for the Next Generation
Community Climate System Model
Raymond M. Loy, PI
FWP # 57648
Budget Explanation

A-C Salaries and Fringe Benefits

Argonne National Laboratory is a government-owned facility operated by the University of Chicago. As a contractor for the Department of Energy, Argonne National Laboratory must comply with DOE general policies and procedures on budgeting and accounting. The Laboratory's costing procedures are based on the assumption that all costs incurred will be recovered. The costing procedures use standard rates, which are used throughout the Laboratory on a consistent basis and uniformly applied to all work supported by the Department of Energy and other federal agencies.

Standard rates are established at the beginning of the fiscal year for each research division, and are monitored and revised as necessary. All labor costs are distributed using standard rates which are developed by the laboratory's budget office for each major payroll classification within the lab. The division-wide rates are based on pay bands (salary ranges) and fringe benefits (35.2% for a regular staff and clerical, and 11% for post/pre doctoral appointees), plus a factor for divisional overhead and for paid absences. Graduate and undergraduate students costs include housing allowance and fringe benefits(7.65%). Effort is escalated each year by a rate provided by the Argonne Budget Department.

The prinicipal investigator for this proposal is:	Raymond M. Loy	
The PI's effort charged per year to this proposal is :		9.00 man-months

E Travel

Domestic: $1.4 K per trip/escalate 4.5% per yr.
Foreign: $2.6 K per trip/escalate 4.5% per yr.

G Other Direct Costs

1. Materials and Supplies:
Hardware/software maintenance, software, low-end computers (<$5k), computer and misc.supplies.
2. Publication Costs:
Books/literature, subscriptions, publishing costs related to research.

I Indirect Costs

Standard rates are also developed for Laboratory General and Administrative (G&A) expense. The procedures for distributing Laboratory G&A and program expense is applied on the basis of the total cost of the work performed. The following indirect rates are provisional and have been estimated for each fiscal year budget period:
PBCS Program Expenses @ 3.7%
Laboratory G&A:
Common Support @ 27.3%
Service Centers @ 21.3%
Equipment/Subcontracts@ 8.1%
G&A Burden @ 2.9%

Argonne' s indirect rates are continuously reviewed and audited by
Cognizant Federal Agency:
Martin Straka 630-252-7724 Department of Energy-Chicago Operations Office

DOE F 4650.2
Rev. (10-03)
(All Other Editions Are Obsolete)

Office of Science (SC)

Face Page

OMB Control No.
1910-1400
(OMB Burden Disclosure Statement on Back)

TITLE OF PROPOSED RESEARCH:
Performance Engineering for the Next Generation Community Climate System Model

1. CATALOG OF FEDERAL DOMESTIC ASSISTANCE #
81.049

2. CONGRESSIONAL DISTRICT:
Applicant Organization's District: 10th
Project Site's District: 10th

3A. I.R.S. ENTITY IDENTIFICATION OR SSN:
95-6031193

3B. DUNS Number:

4. AREA OF RESEARCH OR ANNOUNCEMENT TITLE/#:
SciDAC 2 Lab 06 04

5. HAS THIS RESEARCH PROPOSAL BEEN SUBMITTED TO ANY OTHER FEDERAL AGENCY?
☐ YES ☒ NO

PLEASE LIST

6. DOE/OER PROGRAM STAFF CONTACT (if known):
Anil Deane (301) 903-1465

7. TYPE OF APPLICATION:
☒ New ☐ Renewal
☐ Continuation ☐ Revision
☐ Supplement

8. ORGANIZATION TYPE:
☐ Local Govt. ☐ State Govt.
☒ Non-Profit ☐ Hospital
☐ Indian Tribal Govt. ☐ Individual
☐ Other ☐ Inst. of Higher Educ.
☐ For-Profit
☐ Small Business ☐ Disadvan. Business
☐ Women-Owned ☐ 8(a)

9. CURRENT DOE AWARD # (IF APPLICABLE):

10. WILL THIS RESEARCH INVOLVE:
10A. Human Subjects ☒ No ☐ If yes
Exemption No. or
IRB Approval Date
Assurance of Compliance No:
10B. Vertebrate Animals ☒ No ☐ If yes
IACUC Approval Date or
Animal Welfare Assurance No:

11. AMOUNT REQUESTED FROM DOE FOR ENTIRE PROJECT PERIOD $ 504,000.00

12. DURATION OF ENTIRE PROJECT PERIOD:
07/01/06 to 06/30/09
MM/DD/YY MM/DD/YY

13. REQUESTED AWARD START DATE
07/01/06
MM/DD/YY

14. IS APPLICANT DELINQUENT ON ANY FEDERAL DEBT?
☐ Yes (attach an explanation) ☒ No

15. PRINCIPAL INVESTIGATOR/PROGRAM DIRECTOR
NAME Arthur Mirin
TITLE Principal Investigator
ADDRESS 7000 East Avenue, L-561
Livermore, CA 94550
PHONE NUMBER (925) 422-4020

16. ORGANIZATION'S NAME Regents of the University of California
ADDRESS Lawrence Livermore National Laboratory
7000 East Avenue
Livermore, CA 94550

CERTIFYING REPRESENTATIVE'S
NAME Dona L. Crawford
TITLE Computation, Associate Director
PHONE NUMBER (925) 422-1985

SIGNATURE OF PRINCIPAL INVESTIGATOR/ PROGRAM DIRECTOR
(please type in full name if electronically submitted)
Date

PI/PD ASSURANCE: I agree to accept responsibility for the scientific conduct of the project and to provide the required progress reports if an award is made as a result of this submission. Willful provision of false information is a criminal offense. (U.S. Code, Title 18, Section 1001).

SIGNATURE OF ORGANIZATION'S CERTIFYING REPRESENTATIVE
(please type in full name if electronically submitted)
Date

CERTIFICATION and ACCEPTANCE: I certify that the statements herein are true and complete to the best of my knowledge, and accept the obligation to comply with DOE terms and conditions if an award is made as the result of this submission. A willfully false certification is a criminal offense. (U.S. Code, Title 18, Section 1001).

NOTICE FOR HANDLING PROPOSALS
This submission is to be used only for DOE evaluation purposes and this notice shall be affixed to any reproduction or abstract thereof. All Government and non-Government personnel handling this submission shall exercise extreme care to ensure that the information contained herein is not duplicated, used, or disclosed in whole or in part for any purpose other than evaluation without written permission except that if an award is made based on this submission, the terms of the award shall control disclosure and use. This notice does not limit the Government's right to use information contained in the submission if it is obtainable from another source without restriction. This is a Government notice, and shall not itself be construed to impose any liability upon the Government or Government personnel for any disclosure or use of data contained in this submission.

PRIVACY ACT STATEMENT
If applicable, you are requested, in accordance with 5 U.S.C., Sec. 562A, to voluntarily provide your Social Security Number (SSN). However, you will not be denied any right, benefit, or privilege provided by law because of a refusal to disclose your SSN. We request your SSN to aid in accurate identification, re erral and review of applications for research/training support for efficient management of Office of Science grant/contract programs.

DOE F 4620.1 (04-93) All Other Editions Are Obsolete	**U.S. Department of Energy**	OMB Control No. 1910-1400 OMB Burden Disclosure Statement on Reverse

Budget Page

(See reverse for Instructions)

Year 1

ORGANIZATION Lawrence Livermore National Laboratory					**Budget Page No:** 1	
PRINCIPAL INVESTIGATOR/PROJECT DIRECTOR Art Mirin					Requested Duration: 12 (Months) Year 1	
A. SENIOR PERSONNEL: PI/PD, Co-PI's, Faculty and Other Senior Associates (List each separately with title; A.6. show number in brackets)	DOE Funded Person-mos.				Funds Requested	Funds Granted
	CAL	ACAD	SUMR		by Applpicant	by DOE
1. Art Mirin	4.44				$53,210	
2.						
3.						
4.						
5.						
6 ( ) OTHERS (LIST INDIVIDUALLY ON BUDGET EXPLANATION PAGE) TSM 110-119						
7. ( ) TOTAL SENIOR PERSONNEL (1-6)	4.44				$53,210	
B. ( ) OTHER PERSONNEL (SHOW NUMBERS IN BRACKETS)						
1. ( ) POST DOCTORAL ASSOCIATES						
2. ( ) OTHER PROFESSIONAL (TECHNICIAN, PROGRAMMER, ETC.)						
3. ( ) GRADUATE STUDENTS						
4. ( ) UNDERGRADUATE STUDENTS						
5. ( ) SECRETARIAL - CLERICAL						
6. ( ) OTHER						
TOTAL SALARIES AND WAGES (A+B)					$53,210	
C. FRINGE BENEFITS (IF CHARGED AS DIRECT COSTS)					$22,614	
TOTAL SALARIES, WAGES AND FRINGE BENEFITS (A+B+C)					$75,824	
D. PERMANENT EQUIPMENT (LIST ITEM AND DOLLAR AMOUNT FOR EACH ITEM.)						
TOTAL PERMANENT EQUIPMENT						
E. TRAVEL		1. DOMESTIC (INCL. CANADA AND U.S. POSSESSIONS)			$1,500	
		2. FOREIGN				
TOTAL TRAVEL					$1,500	
F. TRAINEE/PARTICIPANT COSTS						
1. STIPENDS (Itemize levels, types + totals on budget justification page)						
2. TUITION & FEES						
3. TRAINEE TRAVEL						
4. OTHER (fully explain on justification page)						
TOTAL PARTICIPANTS ( )		TOTAL COST				
G. OTHER DIRECT COSTS						
1. MATERIALS AND SUPPLIES						
2. PUBLICATION COSTS/DOCUMENTATION/DISSEMINATION					$500	
3. CONSULTANT SERVICES						
4. COMPUTER (ADPE) SERVICES					$900	
5. SUBCONTRACTS						
6. OTHER					$7,573	
TOTAL OTHER DIRECT COSTS					$8,973	
H. TOTAL DIRECT COSTS (A THROUGH G)					$86,297	
I. INDIRECT COSTS (SPECIFY RATE AND BASE)						
TOTAL INDIRECT COSTS					$81,684	
J. TOTAL DIRECT AND INDIRECT COSTS (H+I)					$167,981	
K. AMOUNT OF ANY REQUIRED COST SHARING FROM NON-FEDERAL SOURCES						
L. TOTAL COST OF PROJECT (J+K)					$167,981	

DOE F 4620.1
(04-93)
All Other Editions Are Obsolete

U.S. Department of Energy
Budget Page
(See reverse for Instructions)

OMB Control No.
1910-1400
OMB Burden Disclosure
Statement on Reverse

Year 2

ORGANIZATION Lawrence Livermore National Laboratory				Budget Page No: 2	
PRINCIPAL INVESTIGATOR/PROJECT DIRECTOR Art Mirin				Requested Duration: 12 (Months) Year 2	
A. SENIOR PERSONNEL: PI/PD, Co-PI's, Faculty and Other Senior Associates (List each separately with title; A.6. show number in brackets)	DOE Funded Person-mos.			Funds Requested	Funds Granted
	CAL	ACAD	SUMR	by Applpicant	by DOE
1. Art Mirin	4.32			$53,324	
2.					
3.					
4.					
5.					
6 ( ) OTHERS (LIST INDIVIDUALLY ON BUDGET EXPLANATION PAGE) TSM 110-119					
7. ( ) TOTAL SENIOR PERSONNEL (1-6)	4.32			$53,324	
B. ( ) OTHER PERSONNEL (SHOW NUMBERS IN BRACKETS)					
1. ( ) POST DOCTORAL ASSOCIATES					
2. ( ) OTHER PROFESSIONAL (TECHNICIAN, PROGRAMMER, ETC.)					
3. ( ) GRADUATE STUDENTS					
4. ( ) UNDERGRADUATE STUDENTS					
5. ( ) SECRETARIAL - CLERICAL					
6. ( ) OTHER					
TOTAL SALARIES AND WAGES (A+B)				$53,324	
C. FRINGE BENEFITS (IF CHARGED AS DIRECT COSTS)				$22,663	
TOTAL SALARIES, WAGES AND FRINGE BENEFITS (A+B+C)				$75,987	
D. PERMANENT EQUIPMENT (LIST ITEM AND DOLLAR AMOUNT FOR EACH ITEM.)					
TOTAL PERMANENT EQUIPMENT					
E. TRAVEL 1. DOMESTIC (INCL. CANADA AND U.S. POSSESSIONS)				$1,500	
2. FOREIGN					
TOTAL TRAVEL				$1,500	
F. TRAINEE/PARTICIPANT COSTS					
1. STIPENDS (Itemize levels, types + totals on budget justification page)					
2. TUITION & FEES					
3. TRAINEE TRAVEL					
4. OTHER (fully explain on justification page)					
TOTAL PARTICIPANTS ( ) TOTAL COST					
G. OTHER DIRECT COSTS					
1. MATERIALS AND SUPPLIES					
2. PUBLICATION COSTS/DOCUMENTATION/DISSEMINATION				$500	
3. CONSULTANT SERVICES					
4. COMPUTER (ADPE) SERVICES				$900	
5. SUBCONTRACTS					
6. OTHER				$7,342	
TOTAL OTHER DIRECT COSTS				$8,742	
H. TOTAL DIRECT COSTS (A THROUGH G)				$86,229	
I. INDIRECT COSTS (SPECIFY RATE AND BASE)					
TOTAL INDIRECT COSTS				$81,787	
J. TOTAL DIRECT AND INDIRECT COSTS (H+I)				$168,016	
K. AMOUNT OF ANY REQUIRED COST SHARING FROM NON-FEDERAL SOURCES					
L. TOTAL COST OF PROJECT (J+K)				$168,016	

DOE F 4620.1
(04-93)
All Other Editions Are Obsolete

U.S. Department of Energy

Budget Page

(See reverse for Instructions)

OMB Control No.
1910-1400
OMB Burden Disclosure
Statement on Reverse

Year 3

ORGANIZATION Lawrence Livermore National Laboratory				Budget Page No: 3	
PRINCIPAL INVESTIGATOR/PROJECT DIRECTOR Art Mirin				Requested Duration: 12 (Months) Year 3	
A. SENIOR PERSONNEL: PI/PD, Co-PI's, Faculty and Other Senior Associates (List each separately with title; A.6. show number in brackets)	DOE Funded Person-mos. CAL	ACAD	SUMR	Funds Requested by Applpicant	Funds Granted by DOE
1. Art Mirin	4.20			$53,416	
2.					
3.					
4.					
5.					
6 ( ) OTHERS (LIST INDIVIDUALLY ON BUDGET EXPLANATION PAGE) TSM 110-119					
7. ( ) TOTAL SENIOR PERSONNEL (1-6)	4.20			$53,416	
B. ( ) OTHER PERSONNEL (SHOW NUMBERS IN BRACKETS)					
1. ( ) POST DOCTORAL ASSOCIATES					
2. ( ) OTHER PROFESSIONAL (TECHNICIAN, PROGRAMMER, ETC.)					
3. ( ) GRADUATE STUDENTS					
4. ( ) UNDERGRADUATE STUDENTS					
5. ( ) SECRETARIAL - CLERICAL					
6. ( ) OTHER					
TOTAL SALARIES AND WAGES (A+B)				$53,416	
C. FRINGE BENEFITS (IF CHARGED AS DIRECT COSTS)				$22,702	
TOTAL SALARIES, WAGES AND FRINGE BENEFITS (A+B+C)				$76,118	
D. PERMANENT EQUIPMENT (LIST ITEM AND DOLLAR AMOUNT FOR EACH ITEM.)					
TOTAL PERMANENT EQUIPMENT					
E. TRAVEL 1. DOMESTIC (INCL. CANADA AND U.S. POSSESSIONS)				$1,500	
2. FOREIGN					
TOTAL TRAVEL				$1,500	
F. TRAINEE/PARTICIPANT COSTS					
1. STIPENDS (Itemize levels, types + totals on budget justification page)					
2. TUITION & FEES					
3. TRAINEE TRAVEL					
4. OTHER (fully explain on justification page)					
TOTAL PARTICIPANTS ( ) TOTAL COST					
G. OTHER DIRECT COSTS					
1. MATERIALS AND SUPPLIES					
2. PUBLICATION COSTS/DOCUMENTATION/DISSEMINATION				$500	
3. CONSULTANT SERVICES					
4. COMPUTER (ADPE) SERVICES				$900	
5. SUBCONTRACTS					
6. OTHER				$7,141	
TOTAL OTHER DIRECT COSTS				$8,541	
H. TOTAL DIRECT COSTS (A THROUGH G)				$86,159	
I. INDIRECT COSTS (SPECIFY RATE AND BASE)					
TOTAL INDIRECT COSTS				$81,850	
J. TOTAL DIRECT AND INDIRECT COSTS (H+I)				$168,009	
K. AMOUNT OF ANY REQUIRED COST SHARING FROM NON-FEDERAL SOURCES					
L. TOTAL COST OF PROJECT (J+K)				$168,009	

DOE F 4620.1
(04-93)
All Other Editions Are Obsolete

U.S. Department of Energy

Budget Page

(See reverse for Instructions)

OMB Control No.
1910-1400
OMB Burden Disclosure
Statement on Reverse

Years 1 - 3

ORGANIZATION Lawrence Livermore National Laboratory				Budget Page No: 4	
PRINCIPAL INVESTIGATOR/PROJECT DIRECTOR Art Mirin				Requested Duration: 36 (Months) Summary - All Years	
A. SENIOR PERSONNEL: PI/PD, Co-PI's, Faculty and Other Senior Associates (List each separately with title; A.6. show number in brackets)	DOE Funded Person-mos. CAL	ACAD	SUMR	Funds Requested by Applpicant	Funds Granted by DOE
1. Art Mirin	12.96			$159,950	
2.					
3.					
4.					
5.					
6 ( ) OTHERS (LIST INDIVIDUALLY ON BUDGET EXPLANATION PAGE) TSM 110-119					
7. ( ) TOTAL SENIOR PERSONNEL (1-6)	12.96			$159,950	
B. ( ) OTHER PERSONNEL (SHOW NUMBERS IN BRACKETS)					
1. ( ) POST DOCTORAL ASSOCIATES					
2. ( ) OTHER PROFESSIONAL (TECHNICIAN, PROGRAMMER, ETC.)					
3. ( ) GRADUATE STUDENTS					
4. ( ) UNDERGRADUATE STUDENTS					
5. ( ) SECRETARIAL - CLERICAL					
6. ( ) OTHER					
TOTAL SALARIES AND WAGES (A+B)				$159,950	
C. FRINGE BENEFITS (IF CHARGED AS DIRECT COSTS)				$67,979	
TOTAL SALARIES, WAGES AND FRINGE BENEFITS (A+B+C)				$227,929	
D. PERMANENT EQUIPMENT (LIST ITEM AND DOLLAR AMOUNT FOR EACH ITEM.)					
TOTAL PERMANENT EQUIPMENT					
E. TRAVEL 1. DOMESTIC (INCL. CANADA AND U.S. POSSESSIONS)				$4,500	
2. FOREIGN					
TOTAL TRAVEL				$4,500	
F. TRAINEE/PARTICIPANT COSTS					
1. STIPENDS (Itemize levels, types + totals on budget justification page)					
2. TUITION & FEES					
3. TRAINEE TRAVEL					
4. OTHER (fully explain on justification page)					
TOTAL PARTICIPANTS ( ) TOTAL COST					
G. OTHER DIRECT COSTS					
1. MATERIALS AND SUPPLIES					
2. PUBLICATION COSTS/DOCUMENTATION/DISSEMINATION				$1,500	
3. CONSULTANT SERVICES					
4. COMPUTER (ADPE) SERVICES				$2,700	
5. SUBCONTRACTS					
6. OTHER				$22,056	
TOTAL OTHER DIRECT COSTS				$26,256	
H. TOTAL DIRECT COSTS (A THROUGH G)				$258,685	
I. INDIRECT COSTS (SPECIFY RATE AND BASE)					
TOTAL INDIRECT COSTS				$245,321	
J. TOTAL DIRECT AND INDIRECT COSTS (H+I)				$504,006	
K. AMOUNT OF ANY REQUIRED COST SHARING FROM NON-FEDERAL SOURCES					
L. TOTAL COST OF PROJECT (J+K)				$504,006	

Performance Engineering for the Next Generation Community Climate System Model

Budget Justification

A. Senior Personnel

The PI, Arthur A. Mirin, will analyze and optimize performance of the next generation Community Climate System Model. The work will include performance modeling using both analytic and artificial neural network approaches, scaling to thousands of processors, and extending to the chemistry-dominated regime, alternative meshes, and high resolution.

B. Other Personnel

Bronis deSupinski will act as a PERC liaison at no cost, particularly vis-a-vis the artificial neural network approach.

C. Fringe Benefits

The Laboratory's Payroll Burden Rate is 42.5% and is applied to the non-leave standard salary of all Laboratory employees, including overtime. PostDocs are charged a 35% Payroll Burden, and students are charged 9.5%.

D. Equipment

None planned.

E. Travel

The anticipated travel is 1-2 trips per year for one person to interact with other researchers doing related work. Travel cost is estimated at $4,500.

F. Trainee/Participant Costs

N/A.

G. Other Costs

2.) Publication costs for technical review and release of publishing project results is anticipated at $1,500.
4) Computer services, as needed on the project, are estimated at $2,700.
6.) Office space is estimated at $22,056.

H. Total Direct Costs

$258,685

I. Indirect Costs

Total Indirect Costs are estimated at $245,321. LLNL rate amounts and their definitions are explained in Attachment A. Note that rates are applied in a specified order and not all taxes apply to each direct cost.

Attachment A

Lawrence Livermore National Laboratory
Current FY2006 Rates

Indirect Cost Pool	Rate (%)	Allocation Base/Rate Determination
Organization Personnel Charge (OPC): Computation - Associate Director's Office	19.09%	Distribution of specific Organization's personnel management costs to users of the Organization Personnel Charge accounts. The rates vary by the Organization providing the service.
Program Management Charge (PMC): Computation - Associate Director's Office Computation - Program	 4.50% 8.10%	A distribution of costs associated with managing and administering direct funded Programs within a Directorate. PMC is allocated on a value-added cost input base to Direct Operating, LDRD, Capital Equipment,
General & Administrative (G&A):	31.50%	G&A is allocated on a value-added base, which is total operating costs less direct materials, subcontracts, and the Electricity Recharge. Supplemental Labor is included included in the base. G&A is applied to Direct Operating, Capital Equipment and Construction accounts.
Strategic Mission Support (SMS):	7.00%	SMS costs include institutional strategic planning, institutional capabilities, outreach, and special employees. Applied to direct operating, WFO, and capital equipment accounts using a value-added base.
Institutional General Purpose Equipment (IGPE):	0.80%	The IGPE allocation is for capital equipment of a general use or institutional nature that benefits multiple cost objectives and is required for general-purpose site-wide needs. It is allocated on a total-cost base, and is not applied to DOE major items of equipment, general plant projects, line item construction or Work for Others - DOE.
Institutional General Plant Projects	0.65%	The IGPP allocation is for new construction projects that cost less than $5M and are of a general institutional nature benefiting multiple cost objectives and required for general-purpose site-wide needs.
Laboratory Directed Research & Development (LDRD): Operating Calculation - An assessment applied to Laboratory operating costs to support exploratory research and development.	6.38%	Rate is applied against total capital equipment (excluding line items) and operating costs, excluding Federal Administrative Charge.

DOE F 4620.1
(04-93)
All Other Editions Are Obsolete

U. S. Department of Energy
Budget Page
(See reverse for Instructions)
(Amounts in Thousands)

OMB Control No.
1910-1400
OMB Burden Disclosure
Statement on Reverse

ORGANIZATION
OAK RIDGE NATIONAL LABORATORY

Budget Page No: FY2006

PRINCIPAL INVESTIGATOR/PROJECT DIRECTOR
Patrick Worley

Requested Duration: 3 (Months)

A. SENIOR PERSONNEL: PI/PD, Co-PI's, Faculty and Other Senior Associates (List each separately with title; A.6. show number in brackets)	DOE Funded Person-mos. CAL	ACAD	SUMR	Amounts in Whole Dollars Funds Requested by Applicant	Funds Granted by DOE
1. Patrick Worley	1.0			10,165	
2.					
3.					
4.					
5.					
6. ( ) OTHERS (LIST INDIVIDUALLY ON BUDGET EXPLANATION PAGE)					
7. (1) TOTAL SENIOR PERSONNEL (1-6)	1.0			10,165	
B. OTHER PERSONNEL (SHOW NUMBERS IN BRACKETS)					
1. ( ) POST DOCTORAL ASSOCIATES					
2. ( ) OTHER PROFESSIONAL (TECHNICIAN, PROGRAMMER, ETC.)					
3. ( ) GRADUATE STUDENTS					
4. ( ) UNDERGRADUATE STUDENTS					
5. ( ) SECRETARIAL - CLERICAL					
6. ( ) OTHER (CRAFTS)					
TOTAL SALARIES AND WAGES (A+B)				10,165	
C. FRINGE BENEFITS (IF CHARGED AS DIRECT COSTS)				3,558	
TOTAL SALARIES, WAGES AND FRINGE BENEFITS (A+B+C)				13,722	
D. PERMANENT EQUIPMENT (LIST ITEM AND DOLLAR AMOUNT FOR EACH ITEM.)					
TOTAL PERMANENT EQUIPMENT					
E. TRAVEL 1. DOMESTIC (INCL. CANADA AND U.S. POSSESSIONS)					
2. FOREIGN					
TOTAL TRAVEL				1,000	
F. TRAINEE/PARTICIPANT COSTS					
1. STIPENDS (Itemize levels, types + totals on budget justification page)					
2. TUITION & FEES					
3. TRAINEE TRAVEL					
4. OTHER (fully explain on justification page)					
TOTAL PARTICIPANTS ( ) TOTAL COST					
G. OTHER DIRECT COSTS					
1. MATERIALS AND SUPPLIES					
2. PUBLICATION COSTS/DOCUMENTATION/DISSEMINATION					
3. CONSULTANT SERVICES					
4. COMPUTER (ADPE) SERVICES					
5. SUBCONTRACTS					
6. OTHER Division Organization Burden and Labor Burden				4,433	
TOTAL OTHER DIRECT COSTS				4,433	
H. TOTAL DIRECT COSTS (A THROUGH G)				19,156	
I. INDIRECT COSTS (SPECIFY RATE AND BASE) G&A 35.0%, Legacy Tax 4.8% Management Fee 2.90%					
TOTAL INDIRECT COSTS				7,852	
J. TOTAL DIRECT AND INDIRECT COSTS (H+I)				27,007	
K. AMOUNT OF ANY REQUIRED COST SHARING FROM NON-FEDERAL SOURCES					
L. TOTAL COST OF PROJECT (J+K)				27,007	

DOE F 4620.1 (04-93) *All Other Editions Are Obsolete*	**U. S. Department of Energy** **Budget Page** (See reverse for Instructions) (Amounts in Thousands)	OMB Control No. 1910-1400 OMB Burden Disclosure Statement on Reverse

ORGANIZATION OAK RIDGE NATIONAL LABORATORY				**Budget Page No:** FY2007	
PRINCIPAL INVESTIGATOR/PROJECT DIRECTOR Patrick Worley				Requested Duration: 12 (Months)	
A. SENIOR PERSONNEL: PI/PD, Co-PI's, Faculty and Other Senior Associates (List each separately with title; A.6. show number in brackets)	DOE Funded Person-mos.			Amounts in Whole Dollars Funds Requested	Funds Granted
	CAL	ACAD	SUMR	by Applicant	by DOE
1. Patrick Worley	3.8			41,124	
2.					
3.					
4.					
5.					
6. ( ) OTHERS (LIST INDIVIDUALLY ON BUDGET EXPLANATION PAGE)					
7. (1) TOTAL SENIOR PERSONNEL (1-6)	3.8			41,124	
B. OTHER PERSONNEL (SHOW NUMBERS IN BRACKETS)					
1. ( ) POST DOCTORAL ASSOCIATES					
2. ( ) OTHER PROFESSIONAL (TECHNICIAN, PROGRAMMER, ETC.)					
3. ( ) GRADUATE STUDENTS					
4. ( ) UNDERGRADUATE STUDENTS					
5. ( ) SECRETARIAL - CLERICAL					
6. ( ) OTHER (CRAFTS)					
TOTAL SALARIES AND WAGES (A+B)				41,124	
C. FRINGE BENEFITS (IF CHARGED AS DIRECT COSTS)				14,599	
TOTAL SALARIES, WAGES AND FRINGE BENEFITS (A+B+C)				55,723	
D. PERMANENT EQUIPMENT (LIST ITEM AND DOLLAR AMOUNT FOR EACH ITEM.)					
TOTAL PERMANENT EQUIPMENT					
E. TRAVEL 1. DOMESTIC (INCL. CANADA AND U.S. POSSESSIONS)					
2. FOREIGN					
TOTAL TRAVEL				3,000	
F. TRAINEE/PARTICIPANT COSTS					
1. STIPENDS (Itemize levels, types + totals on budget justification page)					
2. TUITION & FEES					
3. TRAINEE TRAVEL					
4. OTHER (fully explain on justification page)					
TOTAL PARTICIPANTS ( ) TOTAL COST					
G. OTHER DIRECT COSTS					
1. MATERIALS AND SUPPLIES					
2. PUBLICATION COSTS/DOCUMENTATION/DISSEMINATION					
3. CONSULTANT SERVICES					
4. COMPUTER (ADPE) SERVICES					
5. SUBCONTRACTS					
6. OTHER Division Organization Burden and Labor Burden				18,013	
TOTAL OTHER DIRECT COSTS				18,013	
H. TOTAL DIRECT COSTS (A THROUGH G)				76,737	
I. INDIRECT COSTS (SPECIFY RATE AND BASE) G&A 34.0%, Legacy Tax 4.8% Management Fee 2.90%					
TOTAL INDIRECT COSTS				31,275	
J. TOTAL DIRECT AND INDIRECT COSTS (H+I)				108,012	
K. AMOUNT OF ANY REQUIRED COST SHARING FROM NON-FEDERAL SOURCES					
L. TOTAL COST OF PROJECT (J+K)				108,012	

DOE F 4620.1 *(04-93)* *All Other Editions Are Obsolete*	**U. S. Department of Energy** **Budget Page** (See reverse for Instructions) (Amounts in Thousands)	OMB Control No. 1910-1400 OMB Burden Disclosure Statement on Reverse

ORGANIZATION OAK RIDGE NATIONAL LABORATORY				**Budget Page No:** FY2008	
PRINCIPAL INVESTIGATOR/PROJECT DIRECTOR Patrick Worley				Requested Duration: 12 (Months)	
A. SENIOR PERSONNEL: PI/PD, Co-PI's, Faculty and Other Senior Associates (List each separately with title; A.6. show number in brackets)	DOE Funded Person-mos. CAL	ACAD	SUMR	Amounts in Whole Dollars Funds Requested by Applicant	Funds Granted by DOE
1. Patrick Worley	3.6			40,771	
2.					
3.					
4.					
5.					
6. ( ) OTHERS (LIST INDIVIDUALLY ON BUDGET EXPLANATION PAGE)					
7. (1) TOTAL SENIOR PERSONNEL (1-6)	3.6			40,771	
B. OTHER PERSONNEL (SHOW NUMBERS IN BRACKETS)					
1. ( ) POST DOCTORAL ASSOCIATES					
2. ( ) OTHER PROFESSIONAL (TECHNICIAN, PROGRAMMER, ETC.)					
3. ( ) GRADUATE STUDENTS					
4. ( ) UNDERGRADUATE STUDENTS					
5. ( ) SECRETARIAL - CLERICAL					
6. ( ) OTHER (CRAFTS)					
TOTAL SALARIES AND WAGES (A+B)				40,771	
C. FRINGE BENEFITS (IF CHARGED AS DIRECT COSTS)				14,678	
TOTAL SALARIES, WAGES AND FRINGE BENEFITS (A+B+C)				55,449	
D. PERMANENT EQUIPMENT (LIST ITEM AND DOLLAR AMOUNT FOR EACH ITEM.)					
TOTAL PERMANENT EQUIPMENT					
E. TRAVEL 1. DOMESTIC (INCL. CANADA AND U.S. POSSESSIONS)					
2. FOREIGN					
TOTAL TRAVEL				4,000	
F. TRAINEE/PARTICIPANT COSTS					
1. STIPENDS (Itemize levels, types + totals on budget justification page)					
2. TUITION & FEES					
3. TRAINEE TRAVEL					
4. OTHER (fully explain on justification page)					
TOTAL PARTICIPANTS ( ) TOTAL COST					
G. OTHER DIRECT COSTS					
1. MATERIALS AND SUPPLIES					
2. PUBLICATION COSTS/DOCUMENTATION/DISSEMINATION					
3. CONSULTANT SERVICES					
4. COMPUTER (ADPE) SERVICES					
5. SUBCONTRACTS					
6. OTHER Division Organization Burden and Labor Burden				17,893	
TOTAL OTHER DIRECT COSTS				17,893	
H. TOTAL DIRECT COSTS (A THROUGH G)				77,342	
I. INDIRECT COSTS (SPECIFY RATE AND BASE) G&A 35.0%, Legacy Tax 2.9% Management Fee 2.50%					
TOTAL INDIRECT COSTS				30,710	
J. TOTAL DIRECT AND INDIRECT COSTS (H+I)				108,052	
K. AMOUNT OF ANY REQUIRED COST SHARING FROM NON-FEDERAL SOURCES					
L. TOTAL COST OF PROJECT (J+K)				108,052	

DOE F 4620.1 (04-93) *All Other Editions Are Obsolete*	**U. S. Department of Energy** **Budget Page** (See reverse for Instructions) (Amounts in Thousands)	OMB Control No. 1910-1400 OMB Burden Disclosure Statement on Reverse

ORGANIZATION OAK RIDGE NATIONAL LABORATORY	Budget Page No: FY2009
PRINCIPAL INVESTIGATOR/PROJECT DIRECTOR Patrick Worley	Requested Duration: 9 (Months)

A. SENIOR PERSONNEL: PI/PD, Co-PI's, Faculty and Other Senior Associates (List each separately with title; A.6. show number in brackets)	DOE Funded Person-mos. CAL	ACAD	SUMR	Amounts in Whole Dollars Funds Requested by Applicant	Funds Granted by DOE
1. Patrick Worley	3.5			30,684	
2.					
3.					
4.					
5.					
6. ( ) OTHERS (LIST INDIVIDUALLY ON BUDGET EXPLANATION PAGE)					
7. (3) TOTAL SENIOR PERSONNEL (1-6)	3.5			30,684	
B. OTHER PERSONNEL (SHOW NUMBERS IN BRACKETS)					
1. ( ) POST DOCTORAL ASSOCIATES					
2. ( ) OTHER PROFESSIONAL (TECHNICIAN, PROGRAMMER, ETC.)					
3. ( ) GRADUATE STUDENTS					
4. ( ) UNDERGRADUATE STUDENTS					
5. ( ) SECRETARIAL - CLERICAL					
6. ( ) OTHER (CRAFTS)					
TOTAL SALARIES AND WAGES (A+B)				30,684	
C. FRINGE BENEFITS (IF CHARGED AS DIRECT COSTS)				11,046	
TOTAL SALARIES, WAGES AND FRINGE BENEFITS (A+B+C)				41,731	
D. PERMANENT EQUIPMENT (LIST ITEM AND DOLLAR AMOUNT FOR EACH ITEM.)					
TOTAL PERMANENT EQUIPMENT					
E. TRAVEL 1. DOMESTIC (INCL. CANADA AND U.S. POSSESSIONS)					
2. FOREIGN					
TOTAL TRAVEL				3,000	
F. TRAINEE/PARTICIPANT COSTS					
1. STIPENDS (Itemize levels, types + totals on budget justification page)					
2. TUITION & FEES					
3. TRAINEE TRAVEL					
4. OTHER (fully explain on justification page)					
TOTAL PARTICIPANTS ( ) TOTAL COST					
G. OTHER DIRECT COSTS					
1. MATERIALS AND SUPPLIES					
2. PUBLICATION COSTS/DOCUMENTATION/DISSEMINATION					
3. CONSULTANT SERVICES					
4. COMPUTER (ADPE) SERVICES					
5. SUBCONTRACTS					
6. OTHER Division Organization Burden and Labor Burden				13,001	
TOTAL OTHER DIRECT COSTS				13,001	
H. TOTAL DIRECT COSTS (A THROUGH G)				57,731	
I. INDIRECT COSTS (SPECIFY RATE AND BASE) G&A 35.0%, Legacy Tax 2.9% Management Fee 2.50%					
TOTAL INDIRECT COSTS				23,280	
J. TOTAL DIRECT AND INDIRECT COSTS (H+I)				81,011	
K. AMOUNT OF ANY REQUIRED COST SHARING FROM NON-FEDERAL SOURCES					
L. TOTAL COST OF PROJECT (J+K)				81,011	

DOE F 4620.1
(04-93)
All Other Editions Are Obsolete

U. S. Department of Energy
Budget Page
(See reverse for Instructions)
(Amounts in Thousands)

OMB Control No.
1910-1400
OMB Burden Disclosure
Statement on Reverse

ORGANIZATION: OAK RIDGE NATIONAL LABORATORY

Budget Page No: YRS 1 - 3

PRINCIPAL INVESTIGATOR/PROJECT DIRECTOR: Patrick Worley

Requested Duration: 36 (Months)

A. SENIOR PERSONNEL: PI/PD, Co-PI's, Faculty and Other Senior Associates (List each separately with title; A.6. show number in brackets)	DOE Funded Person-mos. CAL	ACAD	SUMR	Amounts in Whole Dollars Funds Requested by Applicant	Funds Granted by DOE
1. Patrick Worley	11.9			122,744	
2.					
3.					
4.					
5.					
6. ( ) OTHERS (LIST INDIVIDUALLY ON BUDGET EXPLANATION PAGE)					
7. (1) TOTAL SENIOR PERSONNEL (1-6)	11.9			122,744	
B. OTHER PERSONNEL (SHOW NUMBERS IN BRACKETS)					
1. ( ) POST DOCTORAL ASSOCIATES					
2. ( ) OTHER PROFESSIONAL (TECHNICIAN, PROGRAMMER, ETC.)					
3. ( ) GRADUATE STUDENTS					
4. ( ) UNDERGRADUATE STUDENTS					
5. ( ) SECRETARIAL - CLERICAL					
6. ( ) OTHER (CRAFTS)					
TOTAL SALARIES AND WAGES (A+B)				122,744	
C. FRINGE BENEFITS (IF CHARGED AS DIRECT COSTS)				43,881	
TOTAL SALARIES, WAGES AND FRINGE BENEFITS (A+B+C)				166,625	
D. PERMANENT EQUIPMENT (LIST ITEM AND DOLLAR AMOUNT FOR EACH ITEM.)					
TOTAL PERMANENT EQUIPMENT					
E. TRAVEL 1. DOMESTIC (INCL. CANADA AND U.S. POSSESSIONS)					
2. FOREIGN					
TOTAL TRAVEL				11,000	
F. TRAINEE/PARTICIPANT COSTS					
1. STIPENDS (Itemize levels, types + totals on budget justification page)					
2. TUITION & FEES					
3. TRAINEE TRAVEL					
4. OTHER (fully explain on justification page)					
TOTAL PARTICIPANTS ( ) TOTAL COST					
G. OTHER DIRECT COSTS					
1. MATERIALS AND SUPPLIES					
2. PUBLICATION COSTS/DOCUMENTATION/DISSEMINATION					
3. CONSULTANT SERVICES					
4. COMPUTER (ADPE) SERVICES					
5. SUBCONTRACTS					
6. OTHER Division Organization Burden and Labor Burden				53,340	
TOTAL OTHER DIRECT COSTS				53,340	
H. TOTAL DIRECT COSTS (A THROUGH G)				230,965	
I. INDIRECT COSTS (SPECIFY RATE AND BASE) G&A 35.0%, Legacy Tax 2.9% Management Fee 2.50%					
TOTAL INDIRECT COSTS				93,117	
J. TOTAL DIRECT AND INDIRECT COSTS (H+I)				324,082	
K. AMOUNT OF ANY REQUIRED COST SHARING FROM NON-FEDERAL SOURCES					
L. TOTAL COST OF PROJECT (J+K)				324,082	

ORNL Budget Explanation

Budget Pages

Cost estimates presented in the "budget pages" of this proposal have been reclassified in order to be comparable to proposals submitted by other research institutions. At the Oak Ridge National Laboratory (ORNL), costs are collected and reported in accordance with approved Department of Energy (DOE) accounting guidelines. Although costs have been reclassified in this proposal, integrity has been maintained in total and between direct versus indirect costs.

A. (1-7) Senior Personnel

The ORNL's cost accounting system utilizes wage pools based upon salary ranges. For purposes of this budget, the wage pool cost estimate is divided by the fringe benefits rate. The labor component is being reported in Item A and the fringe component is being reported in Item C.

The list of senior personnel participating in this project is as follows:

Patrick Worley is the Principal Investigator (PI) for the project. He will oversee the activities at ORNL, at Lawrence Livermore Naitonal Laboratory, and at Argonne National Laboratory, and will consult with the Science Application PI John Drake on project direction and results. Worley has primary responsibility for the work on CAM physics scalability and on POP scalability.

C. Fringe Benefits

Fringe Benefits for ORNL employees are estimated to be 35.1% of labor costs for FY 2006, 35.5% for FY2007 and 36% for FY2008 and out years.

D. Permanent Equipment

None

E. (1-2) Travel

Worley will travel to meet with other members of the collaboration once in FY06, three times in FY07, four times in FY08, and three times in FY09. Estimated cost per domestic travel is $1000 and includes plane fare, housing, meals, registration, and other allowable costs under government per diem rules.

G.6 Other - Organization Burden Administration

Use of cost collection centers in ORNL R&D divisions is the approved method for collection and distribution of organization burden costs. These accounts are established to collect costs associated with an R&D division. The types of costs which can be charged to organization burden cost collection centers are unfunded paid hours; division administration; and general materials/service costs, including, but not limited to telecommunications, space, utilities, word processing, and copying which are not directly attributable or chargeable to R&D projects. Division Administration costs include: *(i)* managerial, technical, and administrative oversight; and *(ii)* support personnel such as facilities and operations, environmental, safety, and health, finance and budget, quality, and health physics provided for the general benefit of a division.

For ORNL staff, the labor and fringe components have been estimated and reported in items A - C. For Post-BS subcontractors, the subcontract costs have been reported in Item B.1. For ORNL staff and Post-BS subcontractors, the organization burden component has been estimated and is being reported in Item G.6.

Inclusion of these costs is necessary to provide a full accounting of estimated cost for the project period. All cost will be collected and reported in ORNL's cost accounting system.

I. Indirect Costs

Full General & Administrative (G&A), Legacy Charge, and Management Fee are assessed on ORNL labor costs (Items A, C, and G.6), Materials and Supplies, and Equipment less than $35,000 unit value. Full G&A is estimated to be 35.0% for FY2006, 34.0% for FY2007 and 36.50% for FY2008, with an estimated 3% increase each year after that for additional fiscal years. Legacy Charge is estimated to be 4.8% for each year. Management Fee is estimated to be 2.9% for FY2006, 3.2% for FY2007 and 3.5% for FY20081% each year.

Non-DOE-contractor subcontract costs are assessed Subcontract G&A and Management Fee. Subcontract G&A is estimated to be 1.1% each year. Management Fee is estimated to be 2.9% for FY2006, 3.2% for FY2007 and 3.5% for FY20081% each year.

Travel costs are assessed Travel G&A and Management Fee. Travel G&A is estimated to be 7.0% each year. Management Fee is estimated to be 2.9% for FY2006, 3.2% for FY2007 and 3.5% for FY20081% each year.

4.2 Tasks and Milestones

Oak Ridge National Laboratory

Oak Ridge National Laboratory is the nominal lead of this project, however, this is truly a close collaboration between all three institutions. Oak Ridge National Laboratory will participate in virtually all areas of this project. ORNL will have primary responsibility for the performance analysis on the current production platforms: Cray XT3, Cray X1E, and IBM p575 cluster. ORNL will also lead the CAM physics scaling studies and optimization, and POP scalability analyses.

In the first year of this project, ORNL will take primary responsibility for quantifying the current status of CAM and POP2 scalability for a variety of grid resolutions. In collaboration with PERC, ORNL will also lead the generation of a performance model of the CAM physics. ORNL will then determine the performance impact of the atmospheric chemistry on CAM physics, and modify the current static load balancing algorithms accordingly. ORNL will work with LLNL in optimizing the performance of chemical mechanisms, especially with respect to platform-specific analysis and optimization.

In the second year of the project, ORNL will collaborate with PERC and use the POP2 performance characterization to either update one of the existing performance models or generate a new performance model to include the new features found in POP2. This model will then be used to examine POP2 scalability for a range of problem resolutions and architectural assumptions. ORNL will work with LLNL in introducing additional remapping phases into CAM, supporting different processor counts and different work distributions in CAM dynamics, CAM tracer advection, CAM physical parameterizations, and in CAM atmospheric chemistry. ORNL will also add atmospheric chemistry to the performance model for CAM physics. ORNL will lead the introduction of dynamic load balancing to CAM physics, examining both statistical and model-based load-balancing mechanisms. ORNL will work with LLNL to evaluate the ANN performance models for CAM and POP, comparing them with alternative performance models.

In the final year of the project, ORNL will work with PERC to generate a complete model for CAM, for both physics and dynamics, and including support for varying processor count between different computational phases. ORNL will then use the model to investigate scalability and optimal processor allocations. ORNL will collaborate with LLNL in analyzing the performance and scalability of cubed sphere-based dynamical cores, with LLNL leading the analysis of the FV dycore and ORNL leading the analysis of a Discontinuous Galerkin (DG) dycore. ORNL will also work with LLNL in the CCSM configuration optimization work, and with ANL in the incorporation of the new physics and parallel algorithms into the BG/L port of CCSM.

Lawrence Livermore National Laboratory

Lawrence Livermore National Laboratory will participate in virtually all areas of this project. LLNL, through Art Mirin, will have primary responsibility for liaison with PERC artificial neural network activities (Bronis deSupinski) and with core application atmospheric chemistry work (Philip Cameron-Smith). While ANL has lead responsibility for BlueGene port and optimization, LLNL, through its 1024-node and 65536-node BG/L systems, will provide valuable insight into the BlueGene-related tasks.

In the first year of this project, LLNL will take primary responsibility for optimizing the atmospheric component CAM at large tracer count. This will include implementing an additional decomposition direction over tracers as well as minimizing communication overhead. LLNL will also take the lead in optimization of chemical mechanisms, including decomposition over vertical levels. LLNL will contribute to the ORNL-led activity of optimizing load-balance in the chemistry-dominated regime. While ANL will be responsible for the main BG/L port, LLNL will contribute assistance as needed. LLNL will also contribute to the interpretation of CAM performance, particularly vis-à-vis communication and load-balance aspects involving the 2-D decomposition of the finite-volume dycore.

In year 2, LLNL will take lead responsibility for the ability to vary process count over various phases of the computation. We expect to use possibly 5 different domain decompositions at the same time. LLNL will contribute to the ORNL-led activity of dynamic load-balance in the chemistry-dominated regime. LLNL will have primary responsibility for interacting with PERC to utilize the artificial neural network approach for modeling performance in CAM. LLNL will continue to contribute to BG/L-related activities, as needed.

In year 3, LLNL will take lead responsibility in optimizing the FV dycore on the cubed sphere grid (the porting work will have been covered under the core application proposal). This will involve interaction with S-J Lin of GFDL and Bill Putman of NASA/GSFC. This activity will go hand-in-hand with the ORNL-led optimization of the DG dycore on the cubed sphere grid. LLNL will continue to interact with PERC on the ANN-approach to performance modeling, with a focus on the coupled model (CCSM). LLNL will also perform a limited number of extended scalability tests of CCSM on its 65536-node BG/L.

Argonne National Laboratory

The full peer-reviewable proposal is being submitted by Oak Ridge National Laboratory as the lead institution.

At Argonne, this project addresses several of the challenges involved in improving the performance of the Community Climate System Model (CCSM). We will lead the porting of CCSM to the Blue Gene/L (BG/L) platform and its successors. We will then optimize the performance of CCSM, including its components and utilities, on BG/L and similar massively parallel platforms that may be available in the next three years such as BG/P, which is expected to have a higher speed, more memory and more processors per node.

In the first year of the project, Argonne will take the port of CAM initiated at NCAR and analyze its performance for the FV dynamical core and identify current bottlenecks. This will be done in collaboration with LLNL. Argonne will also transfer the version of CAM that supports parallel NetCDF for history writes and parallel binary output for restart writes to BG/L and work to get these changes in to the development repository. Argonne will begin the implementation of parallel I/O in other components as needed and also port the CCSM sea ice model to BG/L.

In the second year, Argonne will complete the porting of all CCSM component models to BG/L and assemble the first coupled system. The first coupled system will consist of CCSM "dead" models and the coupler. We expect to take advantage of a single-executable CCSM by this time and will aid that effort if it is not yet completed. A fully coupled system, with all active components, will be completed and optimized on BG/L. We will also conduct a validation run for the simulated climate. The addition of parallel I/O to all necessary components will be completed. We shall perform an initial analysis of the benefits of message aggregation in the parallel performance of CCSM.

In the final year, we will port the pre-release version of CCSM4 to BG/L including new parallel algorithms developed as part of this SAP by LLNL and ORNL. We expect a BlueGene/P system to be available at this time and will port the BlueGene/L version of CCSM to this platform and optimize its performance and load balance. Our analysis of communications performance improvements through message aggregation will be completed.

5 Other Support of Investigators

Patrick H. Worley / Oak Ridge National Laboratory

Current and Pending Support (Effort will be adjusted if pending proposals are funded.)

Current Support

Title	*Source*	*Period*	*Award/Yr*	*Effort*
PERC-2 High-End Computer System Performance: Scalable Science and Engineering	DOE	7/04 – 6/06	$323K	50%
Collaborative Design and Development of the Community Climate System Model for Tera-Scale Computers	DOE	7/01 – 6/06	$850K	40%
Center for Plasma Edge Simulation	DOE	10/05 – 9/10	$260K	10%

Pending Support

Title	*Source*	*Period*	*Award/Yr*	*Effort*
Performance Engineering Research Center for Enabling Technology	DOE	7/06 – 6/11	$500K	30%
A Scalable and Extensible Earth System Model for Climate Change Science	DOE	7/06 – 6/11	$900K	50%
Performance Engineering for the Next Generation Community Climate System Model SAP	DOE	7/06 – 6/09	$108K	30%
PIC Methods on Leadership Class Computers SAP	DOE	7/06 – 6/09	$230K	20%
Framework Application for Core Edge Transport Simulations	DOE	7/06 – 6/11	$120K	10%

Arthur A. Mirin / Lawrence Livermore National Laboratory

Current and Pending Support (Effort will be adjusted if pending proposals are funded.)

Current Support

Title	*Source*	*Period*	*Award/Yr*	*Effort*
Dynamic Data-Driven Event Reconstruction for Atmospheric Releases	LLNL	10/03 – 9/06	$646K	35%
Detection and Attribution of Regional Climate Change	LLNL	10/04 – 9/06	$140K	25%
Collaborative Design and Development of the Community Climate System Model for Tera-Scale Computers	DOE	7/01 – 6/06	$700K	20%

Pending Support

Title	*Source*	*Period*	*Award/Yr*	*Effort*
A Scalable and Extensible Earth System Model for Climate Change Science	DOE	7/06 – 6/11	$563K	50%
Performance Engineering for the Next Generation Community Climate System Model SAP	DOE	7/06 – 6/09	$168K	40%

Raymond Loy / Argonne National Laboratory

Current and Pending Support (Effort will be adjusted if pending proposals are funded.)

Current Support

Title	*Source*	*Period*	*Award/Yr*	*Effort*
Collaborative Design and Development of the Community Climate System Model for Tera-Scale Computers	DOE	7/01 – 6/06	$700K	100%

Pending Support

Title	*Source*	*Period*	*Award/Yr*	*Effort*
Performance Engineering for the Next Generation Community Climate System Model SAP	DOE	7/06 – 6/09	$224K	100%

6 Resumes

Dr. Patrick H. Worley
Computer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, TN 37831-6016. *Telephone*: (865) 574-3128; *Fax*: (865) 576-5491
http://www.csm.ornl.gov/~worley; *E-mail*: worleyph@ornl.gov

Education
Indiana University, Bloomington, IN, B.A. Computer Science and Mathematics, 1980
Stanford University, M.S.. Computer Science, 1983
Stanford University, Ph.D. Computer Science, 1988

Professional Experience
Senior Research Staff, 2003 – present, Oak Ridge National Laboratory
Research Staff, 1987 – 2002, Oak Ridge National Laboratory

Selected Publications Relevant to Proposal

1. P. Worley. *Benchmarking using the Community Atmospheric Model*. In Proceedings of the 2006 SPEC Benchmark Workshop, Austin, TX, January 23, 2006.
2. L. Oliker, J. Carter, M. Wehner, A. Canning, S. Ethier, A. Mirin, G. Bala, D. Parks, P. Worley, S. Kitawaki and Y. Tsuda. *Leading Computational Methods on Scalar and Vector HEC Platforms*. In Proceedings of the ACM/IEEE Conference on High Performance Networking and Computing (SC05), Seattle, WA, November 12-18, 2005.
3. P. Worley and J. Drake. *Performance Portability in the Physical Parameterizations of the Community Atmospheric Model*. International Journal for High Performance Computer Applications, 19 (3), August 2005, pp. 187-202.
4. P. Worley, J. Candy, L. Carrington, K. Huck, T. Kaiser, G. Mahinthakumar, A. Maloney, S. Moore, D. Reed, P. Roth, H. Shan, S. Shende, A. Snavely, S. Sreepathi, F. Wolf, and Y. Zhang. *Performance Analysis of GYRO: A Tool Evaluation*. In Proceedings of the 2005 SciDAC Conference, San Francisco, CA, June 26-30, 2005.
5. P. Worley, S. Alam, T. Dunigan, Jr., M. Fahey and J. Vetter. *Comparative Analysis of Interprocess Communication on the X1, XD1, and XT3*. In Proceedings of the 47th Cray User Group Conference, Knoxville, TN, May 16-19, 2005.

Synergistic Activities
Worley's research interests include parallel algorithms for scientific computing and performance evaluation of parallel applications and computer systems. Worley currently leads Oak Ridge National Laboratory's participation in the Department of Energy (DOE) SciDAC project in performance evaluation (Performance Evaluation Research Center). He is also a lead performance engineer for the National Science Foundation and DOE Community Climate System Model (CCSM), the principal investigator for the Performance Evaluation and Analysis Consortium End Station at the National Center for Computational Sciences Leadership Computing Facility, and a co-chair of the Software Engineering Working Group for the CCSM.

Recent professional activities include: organizing committee for the SIAM (Society for Industrial and Applied Mathematics) Conference on Parallel Processing for Scientific Computing: 2004, 2005; Tutorials Committee, SC03 and SC04 (2003, 2004); ACM SIGMETRICS Electronic Bulletin Board moderator (1996-2004); Secretary of the SIAM Activity Group on Supercomputing: 2001-2002.

Dr. Arthur A. Mirin
Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, Livermore, CA 94551. *Telephone*: (925) 422-4020; *Fax*: (925) 423-2993.
http://www.llnl.gov/casc/people/mirin; *E-mail*: mirin@llnl.gov.

Education
Ph.D, Mathematics, University of California, Berkeley, 1974.
A.B., Mathematics, University of California, Berkeley, 1969.

Professional Experience

11/03–present:	Leader, Scientific Computing Group, Center for Applied Scientific Computing, Lawrence Livermore National Laboratory (LLNL).
5/96–10/03	Computational Physicist, Center for Applied Scientific Computing, LLNL.
1985–4/96	Leader, Computational Physics Group, NERSC, LLNL.
1969–1984	Computational Physicist, LLNL.

Research Interests
Scientific computing, high-performance computing, climate and atmospheric modeling.

Honors:
Gordon Bell Award for Best Performance, 1999 (coordinated award-winning simulation).

Selected Publications and Presentations

Mirin, A.A., P.H. Worley, W.B. Sawyer, L. Oliker, D. Parks and M.F. Wehner, "Performance Intercomparison of Community Atmosphere Model on High-End Computing Platforms," Twelfth SIAM Conference on Parallel Processing for Scientific Computing, San Francisco (2006). UCRL-ABS-215859.

Mirin, A.A. and W.B. Sawyer, "A Scalable Implementation of a Finite-Volume Dynamical Core in the Community Atmospheric Model," *Int'l. Jour. High Performance Computing Applications*, 19, No. 3 (2005), 203. UCRL-JRNL-206816.

Oliker, L., J. Carter, M. Wehner, A. Canning, S. Ethier, G. Bala, A. Mirin, D. Parks, P. Worley, S. Kitawaki and Y. Tsuda, "Leading Computational Methods on Scalar and Vector HEC Platforms," Proc. Supercomputing 2005 Conference, Seattle (2005). UCRL-CONF-212184.

Mirin, A.A., R.H. Cohen, B.C. Curtis, W.P. Dannevik, A.M. Dimits, M.A. Duchaineau, D.E. Eliason, D.R. Schikore, S.E. Anderson, D.H. Porter, P.R. Woodward, L.J. Shieh and S.W. White, "Very High Resolution Simulation of Compressible Turbulence on the IBM-SP System, (**Gordon Bell Award for Performance, 1999**), *Supercomputing 99 Conference*, Portland, OR, Nov. 13-19, 1999. UCRL-JC-134237.

Raymond M. Loy
Mathematics and Computer Science Division
Argonne National Laboratory
Argonne, IL 60439

Phone: 630-252-7205
Fax: 630-252-6424
E-mail: rloy@mcs.anl.gov

Education

Ph.D. Computer Science, Rensselaer Polytechnic Institute, 1998
M.S. Computer Science, Rensselaer Polytechnic Institute, 1989
B.S. Columbia University, School of Engineering, 1987

Professional Experience

2004-present	Senior Software Developer, Argonne National Laboratory, Mathematics and Computer Science Division
2000-2004.1	Research Associate, University of Chicago, ASCI Center for Astrophysical Thermonuclear Flashes
1998-2000	Postdoctoral Appointee, Argonne National Laboratory, Mathematics and Computer Science Division
1998	Postdoctoral Research Associate, Rensselaer Polytechnic Institute, Scientific Computation Research Center

Selected Publications

- "Theoretical Cost Comparison of Remote Visualization Strategies," *Int. J. of Information, Special Issue on Parallel and Distributed Computing*, 6:3, Jul., 2003, (with L. Freitag).
- "Comparison of Remote Visualization Strategies for Interactive Exploration of Large Data Sets," *Proc. 15th International Parallel and Distributed Processing Symposium (IPDPS 2001)*, Apr. 23-27, 2001 (with L. Freitag).
- "Software for the Parallel Adaptive Solution of Conservation Laws by Discontinuous Galerkin Methods," in B. Cockburn, G. Karniadakis, and S.-W. Shu, eds., *Discontinuous Galerkin Methods Theory, Computation and Applications*, 11:113–124, Springer, Berlin, 2000 (with J. Flaherty, M. Shephard, and J. Teresco).
- "Adaptive, Multiresolution Visualization of Large Data Sets using a Distributed Memory Octree," *Proc. SC99*, Nov. 13-19, 1999 (with L. Freitag).
- "Distributed Octree Data Structures and Local Refinement Method for the Parallel Solution of Three-Dimensional Conservation Laws," in *Grid Generation and Adaptive Algorithms, The IMA Volumes in Mathematics and its Applications*, 113:113-134, 1999 (with J. Flaherty, M. Shephard, M. Simone, B. Szymanski, J. Teresco, and L. Ziantz).

Selected Presentations

- "POP Performance on ANL BG/L," *SciDAC CCSM Consortium Meeting*, Mar. 17, 2005.
- "Autopack: A Library for Efficient Message-Passing," poster session, *ADAPT '03: Adaptive Methods for Partial Differential Equations and Large-Scale Computation*, Oct. 11-12, 2003.

7 Description of Facilities and Resources

This project will rely on access to several high-performance computing systems in order to demonstrate results and techniques on multiple platforms. These include the Cray X1E and Cray XT3 at Oak Ridge National Laboratory, the IA64 cluster and BG/L machines at Lawrence Livermore National Laboratory, the IBM BG/L at Argonne National Laboratory, and the IBM Power5 cluster and AMD Opteron cluster at NERSC.

Argonne National Laboratory

Personnel associated with this proposal will have access to facilities at Argonne National Laboratory, and in particular to facilities associated with the Mathematics and Computer Science Division at Argonne.

Argonne National Laboratory has computing and networking facilities located in the Mathematics and Computer Science Division. These resources include major parallel computing clusters, visualization systems, advanced display environments, collaborative environments, and high-capacity network links.

As one of the nine participants in the NSF's Distributed Terascale Facility, Argonne operates the TeraGrid's visualization facility. The entire TeraGrid is a 13.6 TF grid of distributed clusters using Intel McKinley processors with over 6 TB of memory and greater than 600 TB of disk space. The full machine is distributed between NCSA, SDSC, Caltech, the Pittsburgh Computer Center, Purdue, Indiana University, the Texas Advanced Computing Center, and U Chicago/Argonne. The individual clusters are connected by a dedicated 40 Gb/s link that acts as the backbone for the machine. The Argonne component of the machine consists of 16 dual IA-64 nodes for computation, 96 dual Pentium IV nodes with G Force Ti 4600 graphics accelerators for visualization, and 20 TB of storage.

A second supercomputer at Argonne, which is available to researchers for production computing, is "Jazz". This Linux system, which has achieved a sustained teraflop, ranks among the 50 fastest computers in the world. Jazz has 350 compute nodes, each with a 2.4 GHz Pentium Xeon with 1.5GB of RAM. The cluster uses Myrinet 2000 and Ethernet for interconnect and has 20 TB of on-line storage in PVFS and GFS file systems.

In addition, Argonne has a cluster dedicated for computer science and open source development called "Chiba City". Chiba City has 512 Pentium-III 550MHz CPUs for computation, 32 Pentium-III 550 CPUs for visualization and 8 TB of disk. Chiba City is unique testbed that is principally used for system software development and testing.

Argonne's most recent addition to its supercomputing facilities is a one-rack IBM Blue Gene/Light. The system includes a 2048-processor compute node with a peak performance of 5.7 teraflops.

Another facility available to researchers is the recently constructed wireless sensor network research and deployment laboratory. The lab includes Mica2 motes with a wide range of sensors, including weather boards and GPS. The motes, StarGate gateway nodes, servers, and a digital image capture device allow researchers to develop and test deployments.

Argonne also is a participant in the I-WIRE project, which links to the TeraGrid and StarLight, as well as linking facilities at Argonne to various research institutions in Illinois.

Argonne has substantial visualization devices as well, each of which can be driven by the TeraGrid visualization cluster, by Chiba City, or by a number of smaller dedicated clusters. These devices include the ActiveMural (an 11 million pixel large-format tiled display) and several smaller tiled displays such as the portable MicroMural, which has approximately 3 million pixels.

Furthermore, Argonne currently supports numerous Access Grid nodes, ranging from AG nodes in continual daily use to AG2 development nodes.

Lawrence Livermore National Laboratory

Lawrence Livermore National Laboratory (LLNL) is home to BlueGene/L and Purple, ranked numbers 1 and 3 on the TOP500 list of the world's fastest computers. Within LLNL's Computation Directorate is housed the Center for Applied Scientific Computing (CASC), the organizational home of applied mathematics and computer science research at the Laboratory. CASC has about one hundred scientific staff members and all have ready access to the supercomputing resources administered by Livermore Computing (LC), which is the computer center for LLNL. In addition to maintaining desktop workstations for staff and visiting researchers, LC maintains various large-scale computing platforms, including the 2304 and 4096 processor MCR and Thunder Linux clusters and the 131,072 processor BlueGene/L platform. These production computers provide users with a rich tool environment that includes high-performance compilers, debuggers, analyzers, editors, and locally developed custom libraries and application packages for software development. Access to such resources is provided by an LLNL Multiprogrammatic and Institutional Computing Initiative and therefore does not directly affect the requested budget for this proposal.

Lawrence Livermore National Laboratory is also known for its world class research in global climate modeling. LLNL houses the Program for Climate Model Diagnosis and Intercomparison (PCMDI), whose mission is to develop improved methods and tools for the diagnosis, validation, and intercomparison of global climate models, and to conduct research on a variety of problems in climate modeling and analysis. LLNL carries out state-of-the-art research in atmospheric modeling, atmospheric chemistry and aerosols, coupled climate / carbon cycle modeling, carbon sequestration, anthropogenic effects on climate, and the societal impacts of climate change. The Laboratory also carries out risk and emergency response management through its National Atmospheric Release Advisory Center (NARAC) and the Interagency Modeling and Atmospheric Assessment Center (IMAAC).

Oak Ridge National Laboratory

Oak Ridge National Laboratory (ORNL) has appropriate computers, printers, peripherals, networking equipment, and security for carrying out this task. The Center for Computational Sciences (CCS) houses the ORNL computing platforms. The CCS was established in 1992 and is a designated User Facility. Primary CCS systems include the following:

- Jaguar: a 5,296 processor Cray XT3 system providing a peak performance of over 25 teraflops and over 10 TB of memory. Planned upgrades of Jaguar are to 100 TF in 2006 and to 400 TF in 2007.
- Phoenix: a Cray X1E, with 1,024 multistreaming vector processors (MSPs) and 2 TB of globally addressable memory. Each MSP has 2 MB of cache, and four MSPs form a node with 8 GB of shared memory. Memory bandwidth is very high, up to half the cache bandwidth. The interconnect functions as an extension of the memory system, offering each node direct access to memory on other nodes at high bandwidth and low latency. The peak performance of Phoenix is 18.5 teraflops.
- OIC: ORNL Institutional Cluster is a collection of eight SGI Xeon clusters providing 640 dual processor nodes and almost 10 TF of peak performance.

- Cheetah: a 27-node IBM Power-4 system. Each Power-4 node of Cheetah has thirty-two 1.3 GHz Power4 processors. Twenty of the nodes have 32 GB of memory, five nodes have 64 GB of memory and two nodes have 128 GB of memory. The peak performance of Cheetah is 4.5 teraflops.
- Ram: a 256-processor SGI Altix with 2 TB of shared memory. Each processor is the Intel Itanium2 1.5 GHz processor. The full system runs a single Linux image and the large shared memory facilitates analysis of very large data sets. The peak performance of Ram is 1.5 teraflops.

Access to the Cray X1E, XT3, and SGI Altix will be available via the National Center for Computational Sciences (NCCS) end station in climate and/or the end station for performance evaluation and analysis.

ORNL has a professional, experienced operational and engineering staff comprised of groups in HPC Operations, Technology Integration, User Services, and Scientific Computing. The ORNL computer facility is staffed 24 hours a day, 365 days a year to provide for continuous operation of the center and for immediate problem resolution. On evenings and weekends, the operators provide first-line problem resolution for users with additional user support and system administrators on-call for more difficult problems.

ORNL is connected to every major research network at rates of 10 gigabits per second or greater. Connectivity to these networks is provided via optical networking equipment owned and operated by ORNL that runs over leased fiber optic cable. This equipment has the capability of simultaneously carrying either 192 10-gigabit per second circuits or 96 40-gigabit per second circuits and connects the CCS computing facility to major networking hubs in Atlanta and Chicago. Currently, only 16 of the 10-gigabit circuits are committed to various purposes, allowing for virtually unlimited expansion of the networking capability. ORNL will be expanding the current TeraGrid connection from 10 to 30 gigabits per second in the near future. Currently, the connections into ORNL include: TeraGrid, Internet2, ESnet, and Cheetah at 10 gigabits per second as well as UltraScienceNet and National Lambda Rail at 20 gigabits per second. ORNL operates the Cheetah research network for NSF and the UltraScience Net research network for DOE.

The CCS local-area network is a common physical infrastructure that supports separate logical networks, each with varying levels of security and performance. Each of these networks is protected from the outside world and from each other with access control lists and network intrusion detection. Line rate connectivity is provided between the networks and to the outside world via redundant paths and switching fabrics. A tiered security structure is designed into the network to mitigate many attacks and to contain others. The new Cray system will be connected in the TeraGrid enclave to the TeraGrid Force10 E600 router via a 10 Gbps link.

ORNL has a comprehensive physical security strategy including fenced perimeters, patrolled facilities, and authorization checks for physical access. An integrated cyber security plan encompasses all aspects of computing. Cyber security plans are risk-based and separate systems of differing security requirements into enclaves of similar requirements allowing the appropriate level of protection for each system, while not hindering the science needs of the projects.