Preprint
UCRL-ID-147450

Unclassified Computing Capability

User Responses to a Multiprogrammatic and Institutional Computing Questionnaire

M. McCoy, L. Kissel

U.S. Department of Energy

Lawrence
Livermore
National
Laboratory

January 29, 2002

This work was performed under the auspices of the United States Department of Energy by the University of California, Lawrence Livermore National Laboratory under contract No. W-7405-Eng-48.

UCRL-ID-147450

Unclassified Computing Capability

User Responses to a Multiprogrammatic and Institutional Computing Questionnaire

Integrated Computing and Communications Department (ICCD)

Computation Directorate
Lawrence Livermore National Laboratory

January 29, 2002

Contents

The Multiprogrammatic and Institutional Computing (M&IC) Questionnaire

Introduction to the Questionnaire

We are experimenting with a new computing model to be applied to a new computer dedicated to that model. Several LLNL science teams now have computational requirements, evidenced by the mature scientific applications that have been developed over the past five plus years, that far exceed the capability of the institution's computing resources. Thus, there is increased demand for dedicated, powerful parallel computational systems. Computation can, in the coming year, potentially field a capability system that is low cost because it will be based on a model that employs open source software and because it will use PC (IA32-P4) hardware. This incurs significant computer science risk regarding stability and system features but also presents great opportunity. We believe the risks can be managed, but the existence of risk cannot be ignored.

In order to justify the budget for this system, we need to make the case that it serves science and, through serving science, serves the institution. That is the point of the meeting and the White Paper that we are proposing to prepare.

Many of the questions here are directed at scientists with mature and demanding parallel codes: codes that by themselves tax the full capability of the system. However, some of the questions address the void that is created between very low parallelism capacity infrastructure (that was procured with the FY01 budget) and the highest end (the 2–8 teraFLOP/s level). We could call this intermediate range the "intermediate capability" range. This middle range (8–32 processors, or 25–100 gigaFLOP/s) is called to attention because the high end has moved up, and the lowest end remains at the one to four processor level. This intermediate range probably has other characteristics, namely the need by science teams to run many such problems as part of a single study. Thus, we are talking about "capacity at the (intermediate) capability" level. It is then necessary to seek to calibrate the demand for both service regimes, and this might allow M&IC, within its limited budget, to achieve the right balance for the generation of scientific results.

Most responses to these questions will be presented in a Requirements Document (a White Paper) that will go to the Deputy Director for Science & Technology, Jeff Wadsworth. In part, the demand as reflected in this document will be used to size the system(s) and to rethink the existing M&IC service posture and rules. Size equates to cost, obviously.

Our current model for the use of this system at the high capability end by institutional projects is as follows. Efforts will be enabled for between 6–18 months and then can apply for renewal to the Council for Strategic Science and Technology (CSST). During this period, we expect that there will be some sub-periods (that sum up to perhaps 1–4 months) in which the code will be running on at least one-third of the computer much or most of the time. The rest of the time, while we assume use not to be insubstantial, we do not expect the code to be running huge, multiday problems, but to be running shorter, smaller production runs and scaling calculations. This allows for the enfranchisement of several projects simultaneously, each in a different phase of its cycle.

For the intermediate capability regime, there are a number of solutions possible, including dedicating half or more of TC2K to this kind of work. Good solutions should emerge after the responses have been read and understood.

Directions for Responding to the Questionnaire

Please respond to the following questions. Use of figures or pictures is encouraged if this is convenient. We would like to prepare a very professional document for management. We would like the answers to these questions to reflect serious thought—overestimation of requirements is not necessary, just objective and careful estimates. Please limit yourself to approximately three or four pages.

Two classes of service are covered here, highest-end and intermediate capability. It is not necessary to answer the questions specifically directed at one class of service if you are not interested in that class of service.

The Questions

To facilitate the review of the responses and to eliminate the need to repeat the questions in their entirety with each response, a brief identifying phrase is given in bold after each question number below and preceding the full text of the question. These phrases will appear with the responses in place of the full text of the question.

1. **Scientific intent and scope**
   Describe the project's scientific intention and its scope. Describe its importance to your program, directorate, or to the institution. What is the potential or current overall scientific impact of the proposed project?
2. **Project funding source, size in FTEs, PI, and major contributors**
   Describe the project's current funding source and rough number of FTEs on the team. Please name the PI or the main author and the major contributors.
3. **Code maturity level and history**
   Describe the current maturity level of the application code(s) used for this project, including level of parallelism, production status, where the code was developed, its history of use, what systems it has used in the past, and (if known) its level of efficiency at various node and processor counts. [The code should be checkpoint restartable—if the code does not currently have this capability, how long would it take to create this capability?] What capability does your project have to visualize the computational results?
4. **Anticipated scientific advance at 2–8 teraFLOP/s**
   What would be the scientific advance or breakthrough that could be achieved if this project received dedicated access to an M&IC Open Computing Facility capability platform of between 2–8 teraFLOP/s for 1–4 months? (Note: eventual size of the system is not yet determined, but would be at least 2 TF.)
5. **Anticipated scientific advance at 8–32 processor range**
   What would be the scientific advance or breakthrough that could be achieved if this project received copious access to banks of processors cooperating at the 8–32 processor range (25–100 GF)?

6. **Connection with experiment**
   How does this project support experimental science or how does this project depend on experimental data from previous (or planned) experiments?

7. **Requirements for scientific challenge at 2–8 teraFLOP/s**
   Assuming M&IC could provide your project a 2–8 teraFLOP/s resource for a specific scientific challenge, please describe your computational requirements for success, including (but not limited to):

   - Ratio of bytes of memory/flop (or how much memory per processor or how much total memory and how many total FLOP/s)
   - Ratio of bytes of disk/flop (or how much local and/or global disk would the run require)
   - What visualization requirements do you have? What visualization support would you require?
   - Tertiary storage requirements in TB
   - When you would be ready to start a 6–18 month project as described in the introduction?

8. **Requirements for scientific challenge at intermediate capacity level**
   Assuming M&IC could provide your project with "capacity at the intermediate capability level" (many 16–32 processor runs) for a specific scientific challenge, please describe your computational requirements for success, including (but not limited to):

   - Ratio of bytes of memory/flop (or how much memory per processor or how much total memory and how many total FLOP/s)
   - Ratio of bytes of disk/flop (or how much local and/or global disk would the run require)
   - What visualization requirements do you have? What visualization support would you require?
   - Tertiary storage requirements in TB
   - When you would be ready to start a project?

9. **Machine use pattern at 2–8 teraFLOP/s**
   Assuming the machine features 256–512 nodes with 2–4 processors with 3.6–4.0 gigaFLOP/s peak processors and 2–8 GB of memory and 80–160 GB of disk per node with a Quadrics interconnect (>300 MB/s), please describe how you would use this machine in the period in question in which you enjoy a preferential allocation (between 6 months and 18 months). Think in terms of the whole project effort as well as the capability science runs. What scaling or convergence studies would your project need to perform? Is there any experimental data that needs to be analyzed as part of the effort? What validation of the algorithms and/or application will need to be performed? This question requires some real (and realistic) thought, as we are using these requirements to size a system and make an honest case to the Deputy Director about the spectrum of mature unclassified efforts around the Laboratory.

10. **Machine use pattern at intermediate capacity level**
    Assuming the machine features many banks of nodes dedicated to 8–32 processor runs (think of the system architecture similar to that described in question 8 but dedicated to smaller problems, or think alternatively of TC2K dedicated to 8–32 processor runs), please describe how you would use such a system to achieve your goals. Think in terms of the whole project effort as well as the science runs. Is there any experimental data that needs to be analyzed as part of the effort? What validation of the algorithms and/or application will need to be performed? What is the approximate number of runs that need to be done and the approximate wall-clock time of each run?

11. **Memory vs compute trade-off**
    Would this machine be of use with less memory, say 1 GB per processor (.75 GB available to the user)? In other words, would you trade away memory to get more nodes or gigaFLOP/s?

12. **Comments on value of unclassified computing**
    Please take some time to describe the value of unclassified computing (including M&IC) to your efforts in the past and today. In particular, we would like to hear about programmatic impact. M&IC has not been an inexpensive investment, and we need to make the case to our Deputy Directors that this investment has enriched and enhanced this institution. Questions we get concern "return on investment." [For instance, has your involvement with M&IC allowed you to hire additional personnel? Has it enhanced your ability to meet programmatic milestones? Has it enriched the institution? (that is, have to been asked to give talks to review committees or in other forums in part based on your work with M&IC systems?] Please dwell on unmet needs (especially in the arena of access and computer time) as well as needs that have been met adequately.

The Responses

The following responses to the questionnaire are of two formats. If the responder answered specific questions, the identifying question phrase precedes the response to the question. (Note that not all responders answered every question.) If the responder did not respond specifically to questions, the response appears as submitted.

Adaptive Laser Plasma Simulator (ALPS) Project

Milo Dorr

1. **Scientific intent and scope**
   The goal of the ALPS project is to investigate the use of parallel adaptive mesh refinement (AMR) in the simulation of laser–plasma interaction (LPI). The ability to predict LPI is important for the design and analysis of laser driven fusion experiments, such as those to be performed at the National Ignition Facility (NIF). Through the use of AMR, problems with a wide range of scales can be more efficiently solved, which in turn enables larger LPI simulations to be performed in less time. In addition to developing advanced algorithmic and code technologies, we are also involved in testing these new capabilities on LPI problems of importance to the Laboratory's laser-driven high-energy-density science missions.

2. **Project funding source, size in FTEs, PI, and major contributors**
   In FY02, the ALPS project will be funded by LDRD and ASCI ITS. The total number of FTEs is 3.15. The PI is Milo Dorr. The other investigators are Xabier Garaizar, Jeff Hittinger (Computation/CASC), and James Schek (Computation/CAO). We have substantial collaborations with Richard Berger, Bruce Langdon, Bert Still, Ed Williams (DNT/X Division), Bruce Cohen, and Laurent Divol (PAT/M Division).

3. **Code maturity level and history**
   ALPS is currently a research code. Nevertheless, it is in a sufficiently mature state to perform meaningful science runs in the hands of expert users (i.e., its developers). ALPS is built using the SAMRAI (Structured Adaptive Mesh Refinement Application Infrastructure) system, which is a C++ class library under development in CASC. Parallelism is achieved through the use of SAMRAI functionality, which in turn is based on the Message Passing Interface (MPI) standard. Therefore, ALPS can (only) run on platforms to which SAMRAI has been ported, which to date have included: Sparc clusters (on which the code is developed), the Compaq Compass/Tera/TC2K clusters, and the ASCI Blue-Pacific machine. A Linux port is under way. The largest number of processors we have used with ALPS is 128 processors (32 nodes) of TC2K. We have not performed a systematic test of parallel efficiency, but it is clear that this will be highly problem-dependent. ALPS does have checkpoint restart capabilities. For visualization, we use a home-grown VTK-based client called Vizamrai to visualize our structured AMR data.

4. **Anticipated scientific advance at 2-8 teraFLOP/s**
   We are interested in such a capability to run large, 3D calculations of crossed laser beams in a plasma flow. Under certain conditions, a resonant interaction between the light beam wave and an ion acoustic wave can result in the transfer of energy from one beam to the other. Such problems are of potential importance for NIF, where beams in the inner and outer cones

will cross in regions of plasma flow sufficient to excite this energy transfer. If this resonance is not detuned or otherwise accounted for, the deposition of laser energy on the hohlraum wall could be adversely affected.

We have been running some 2D calculations to verify an analytic model and make predictions for cases where the model does not apply. Even with the help of AMR, these problems require a handful of TC2K processors. We need a much larger machine to allow us to even consider making 3D runs at the size of the experimental configuration currently being used on the Omega facility at the University of Rochester in support of the NIF Program.

5. **Anticipated scientific advance at 8–32 processor range**
   This type of capability would allow us to run a large number of scaled-down 2D crossed beam problems covering a wide range of plasma conditions and beam configurations. This study would allow us to establish a qualitative understanding of the scaling of beam energy transfer with respect to a large number of parameters. We are currently running these types of problems on TC2K, but at the average rate of one problem per day our progress has been slow.

6. **Connection with experiment**
   A goal of our use of ALPS to model crossed laser beams in a plasma flow is to provide guidance for the experimental program currently being conducted at the Omega facility at the University of Rochester in support of the NIF Program. Given the limited number of laser shots allocated for the investigation of crossed beam issues, it is clear that experiments alone will not be able to map out the entire parameter space of plasma conditions and beam configurations of potential interest. Using simulation, we hope to aid the experimental campaign by predicting which experiments will yield the most useful information. In order to do this, of course, we will utilize experimental data to validate the code.

7. **Requirements for scientific challenge at 2–8 teraFLOP/s**
   Our requirements are strongly driven by the fact that we are much more memory-limited than CPU-limited. ALPS scales reasonably well enough to compensate for a lack of memory per node by using more nodes, but our strong preference is for as much memory per node as possible. For the baseline capability run described in the response to question 9 below, assuming 512 nodes, 8 GB per node would be tight but probably sufficient. Although we are not CPU-limited, we do perform enough floating point operations that any increase in single node performance translates to a proportional decrease in wall-clock time. In other words, the faster the processors the better. The current Alpha EV67s on the TC2K have worked well; another factor of 3 increase would obviously be welcome.

   Our local disk requirements are driven by the need to save visualization and restart files every so often (the frequency is selected at run time). We also use the data in our visualization files for other types of post-processing. For the capability runs described below, we estimate that 16 GB of local (per node) disk space will be sufficient to hold the visualization files for about 256 time steps, which is a reasonable sampling of the time interval of interest. A restart file requires about five times more disk space than a visualization file, but we can dump those much less frequently so that they will not substantially increase this estimate. Based on this local disk estimate, the tertiary storage requirements for the data from one large run would be about 8 TB. Hopefully, we will be able to somehow compress the data to reduce this.

Visualization for such a large problem is uncharted territory for us. To date, we have been using a home-grown VTK-based client called Vizamrai to visualize our data. Vizamrai works on a single data file concatenated from parallel runs. Clearly, this mode of operation will have to be changed to a more distributed one. We believe that there are plans to parallelize Vizamrai, but the exact schedule has not been determined (although the schedule may be strongly influenced by the need to respond to the opportunity presented by the availability of a new, larger M&IC resource).

We estimate that we would be ready to begin a large simulation in April 2002.

8. **Requirements for scientific challenge at intermediate capacity level**
   Our "capacity at intermediate capability level" requirements would be met with a machine at the level of TC2K or slightly better. We are currently using this machine for 2D runs, for which the current memory and /nfs/tmp disk resources have been sufficient. We just need to be able to get them more frequently. Currently, we are able to make about one 7–10 hour run on a couple of nodes per day. If more capacity were available, we could start using it today.

9. **Machine use pattern at 2–8 teraFLOP/s**
   We believe that such a machine would get us very close to being able to simulate a full-scale Omega facility crossed-beam experiment. By "full-scale," we mean the use of actual-size 250-μm-diameter beams in a plasma-filled region sufficiently large to contain the interaction region of interest. The figure of merit in such a calculation would be the amount of power transferred from one beam to the other. We do not expect that a single such calculation would require more than a couple of days of dedicated time on the entire machine. However, we will undoubtedly want to make several such runs with varying plasma and laser beam parameters. Experimental data exists to which we will compare our computed solutions. As explained above, our ultimate goal is to use computations to help guide and interpret future experiments.

   There is considerable uncertainty with respect to such a large calculation. To date, the largest number of processors we have used with ALPS is 128 (32 TC2K nodes). We will therefore need to spend some time carefully ramping up from this level, which could expose new scalability issues. We depend heavily upon a number of libraries, in particular the SAMRAI system and the Hypre linear solver libraries. The resolution of any performance issues associated with the use of these systems for such a large problem will therefore require the assistance and commitment of their developers (which has been very good in the past).

10. **Machine use pattern at intermediate capacity level**
    This type of resource would allow us to run a large number of scaled-down 2D crossed-beam problems covering a wide range of plasma conditions and beam configurations. The nominal wall-clock time for a single run would be about 12 hours, but could be longer. The number of runs required is difficult to estimate, since it depends on the results of previous runs. Nevertheless, given the need to vary plasma density, velocity scale length, laser intensity, beam width and shape, the number of runs required could reach into the hundreds. This study would allow us to establish a reasonably qualitative understanding of the scaling of beam energy transfer with respect to several parameters. Quantitative predictions with actual beam configurations (i.e., the setup used for the Omega experiments) will require the 3D capability runs described above.

11. **Memory vs compute trade-off**
    No, we would not prefer such a trade. We are much more memory-limited than CPU-limited.

12. **Comments on value of unclassified computing**
    The resources provided by the M&IC have been essential for the success of our LDRD-funded project. In particular, the TC2K cluster has allowed us to test the ALPS code in ways that would not have been possible otherwise. The ASCI machines were the only other alternative, but they (in particular, Blue-Pacific) had become so heavily utilized that it was difficult to get sufficient turnaround on batch jobs to make progress. The computations we performed on the TC2K provided a critical part of our successful presentation to the UC committee that reviewed Computation Directorate research in April of this year.

    One of the goals of CASC is to collaborate with Laboratory programs on computational science projects. The availability of institutional computing resources is very important in lowering the artificial organizational and programmatic barriers to the multidisciplinary teamwork essential in pushing the scientific computing envelope. By pooling our computational assets in this way, the Laboratory can accomplish far more than the sum of its otherwise strictly programmatic efforts.

The Djehuty Project—Stellar Modeling

Dave Dearborn, Don Dossa

"Star White, Far Light
First Flop I do tonight"

Stars and 3D Modeling

The best determinations of the size, age, and composition of the universe are founded on measurements of stars, combined with an understanding that has relied on 1D computations. The importance of energy transport by convection (a 3D process) limits the value of these 1D calculations even for spherical stars. Observations assure us that our best 1D approximation of steady-state convection is flawed, and there is no real ability to model the time-dependent convection critical in helium flashes or deep mixing events that follow thermal pulses. Even worse, nearly half of the visible mass of the universe is incorporated into intrinsically 3D binary stars. Understanding integral properties such as the evolution of the elements, or the detailed pre-explosion structure of metric phenomena like type I supernovae awaits development of a true 3D stellar evolution code.

Developing a 3D code capable of modeling stars is exceedingly challenging. In spite of the inherent difficulty, the importance of obtaining a better understanding of stars has led a small number of groups to begin such work. Most current calculations are practically limited to of order 1 million zones, and cannot realistically handle whole stars (in 3D). They have begun simulations of envelope convection by modeling modest segments of a star. Naturally, they do not incorporate physics pertinent to the core (nuclear energy production), and the gravitational potential is an imposed condition. Very recently, Porter and Woodward of the University of Minnesota Laboratory for Computational Science and Engineering have completed a calculation of red giant atmospheres with 134 million zones run on 120 processors. The primary thrust of this calculation was to minimize latency time. It did not include nuclear energy generation or a realistic equation of state. Energy transport was approximated as dependent on gas pressure.

While these are important starts towards understanding convection in stars, it is clear that some problems (e.g., solar rotation) yield results that are dependant on the size of the segment that can be simulated, necessitating whole star modeling.

The LLNL code Djehuty is now the world's only operating code capable of modeling complete stars in three dimensions. It has an equation of state (EOS) that is very accurate for a large range of masses (>0.5 solar masses), Opal opacities with Alexander opacities for the lower temperatures where molecules are significant, and nuclear-reaction network for hydrogen, helium, and carbon burning. Nuclear energy production can be computed either by a tabular set of reactions, or by a very quick analytic network. Energy transport is modeled by a standard pair of coupled diffusion equations. The gravity implementation is currently complete only for spherical stars, but is adequate to begin the first major 3D study of the convective cores in massive stars. We in the course of the next month, we expect to begin testing of a Poisson solver that will enable us to study rapidly rotating and binary stars. Then, we will enter completely new territory.

We have tested stellar decomposition and parallel operation on 8, 16, 32, 64, 128, and 440 processors, and constructed large (7, 14, and 56 million zone) meshes needed for realistic stellar modeling of whole single stars. Some extremely interesting phenomena will require even larger meshes and far more CPUs than are currently available.

While the magnitude of current academic modeling efforts are very credible, our code is unique, and there are a number of astrophysical puzzles that await 3D modeling on a scale that requires operation in a massively parallel environment.

Modeled as Single Stars—Big Runs

(1) The ignition of helium in degenerate core (helium flash) occurs in most stars, and the results are strongly dependent on the efficiency of the development of convection. Stars are seen to survive this process, while models sometimes have difficulty. The truth may explain the early composition irregularities seen in some stars.

(2) Radial and nonradial pulsations on the upper Main Sequence (53 Persei stars); several other classes of pulsating star. Calculation of radially pulsating stars offers an opportunity to validate our code against a vast body of observational data as well as compare to 1D hydrodynamic calculation. From this basis, we can confidently engage the unique abilities of Djehuty to study nonradially pulsating stars. With the MACHO data set, Djehuty will break new ground modeling stars with short lifetimes and long pulsation periods (Long Period Variables).

(3) Rotationally driven circulation and mixing is the suspected cause of a number of isotope anomalies. It has been proposed to alleviate the ^{3}He overproduction of solar type stars that leads to an inconsistency with Big Bang nucleosynthesis. Reconciliation of the robust production of ^{3}He with the Big Bang will require a 3D modeling ability.

(4) Possible thermal H-flashes on first giant branch.

Modeled as Single Stars with Higher Resolution—Very Big

(1) Modeling of the solar convection zone and distribution of angular momentum. The sun is the best-known star, and, with the results of helioseismology, provides a snapshot of its current convection zone and rotational properties. Modeling this observed structure will provide considerable support for the code's calculational ability.

(2) Thermal He-shell flashes on the second (asymptotic) giant branch, the final stage of evolution of intermediate mass stars. These pulses are responsible for many of the s-process elements, and probably involve the time development of deep convective intrusions.

(3) Explorations of conditions at the boundary between an accretion disc and the star at its center. In many important stellar systems (novae, type Ia supernovae), material is transferred from one star to another through an accretion disc. The energy and angular momentum of material falling onto the star will probably be Rayleigh–Taylor unstable, and over a broader region lead to mixing patterns that establish nonspherical temperature profiles on the star.

Modeled As Single Component of a Binary System—Big to Monster Runs
(1) Tidal dissipation in convective or radiative envelopes. The ability of a tidally disturbed star to dissipate energy is fundamental for understanding stellar capture, and orbital evolution.

(2) Dynamical effects of very close encounters as a mechanism for producing the chemical inhomogeneities seen in globular cluster abundances.

Modeled as Binaries—Monster Runs
(1) The Darwin instability, by which the orbits of binary stars with a high mass ratio become unstable, and the stars merge. The importance of this phenomenon can be explored and understood once the tidal dissipation problem (above) has been explored.

(2) Roche-lobe overflow can result in a disk/hot-spot or in direct stream impact and penetration on the companion, depending on the system. Modeling these processes requires a 3D calculation that includes two stars in one calculation.

(3) When stars are in contact, they effectively have a small common envelope. Observations show that there must be a vigorous yet stable energy transport mechanism operating. How can this occur, and are all such systems stable? Probably not, as the observed structure of some stellar systems requires one star to have been embedded in the envelope of the other at some time in the past while in other cases it is clear that two stars have merged into a single object.

Code Metrics

The code exhibits quite good scaling across a broad range of MPI processes and problem size. Because the code solves time dependent problems, the code writes out the state of the problem as directed by the user at the end of certain number of cycles through the calculation. On TC2K, we routinely see the system calculate at a rate of 7–10 cycles per hour per CPU on a million zone problem. We have also seen that Frost is about four times faster than TC2K.

The other useful metric is the I/O load. When the code is writing silo files, we see disk files being created that amount to 355 MB/million zones per saved cycle.

Using the above metrics and scaling to the suite of problems we wish to run lead to some truly alarming requirements.

For the binary star system calculations, we will want to run with 100 million zones for tens of thousands of cycles. Of course, one calculation will not be sufficient, and we would want to run the problem set many times, varying the mass ratio of the stars, the orbital separation, and mass composition. Other researchers, both within LLNL and our external collaborators, have other problems to solve, such as the effect of dust grains in the atmospheres of red giants. Using the above metrics, we would expect that a single calculation using 500 CPUs on 100 million zones on TC2K for 100,000 cycles would require 2,800 hours of wall-clock time (17 weeks) and would

generate at least 3.5 TB of output. At this scale of problem, Djehuty alone could consume the entire TC2K cluster for more than its planned lifetime. This still ignores the compute requirements for many of the outstanding problems in astrophysics that we are both capable of and desirous of solving. These are world-class problems that require a world-class system.

Adaptive Mesh Refinement (AMR) Code Development Project

Jeff Greenough

1. **Scientific intent and scope**
   The mission of the AMRh code development project is to develop state-of-the-art AMR methods and apply them to problems of interest to A Program. The code includes both hyperbolic methods using high-order monotone upwind numerical methods and parabolic methods using robust multilevel iterative methods. The code is used to obtain unprecedented spatial and temporal evolution of shock-driven fluid flows and radiatively driven plasma flows. We are also associated with the Turbulence and Mixing Program within A Program. The interest here is to use the code to calculate compressible, high Reynolds number flows where all the flow length scales are resolved to the highest degree possible.

   This last activity is the one that has pushed us beyond the capabilities of the machines we have access to. For example, to obtain 3-μm zoning for a 2D diffuse cylinder interacting with a weak shock wave problem, we need approximately 40 million cells on the finest grid. This is the resolution required to resolve the Kolmogorov length scale. We need to compute the flow for 8 ms, which requires 450 coarse grid time steps. On 128 CPUs of TC2K, this will require nearly 4 months of dedicated weekend time (nearly sixteen 62.5-hour weekend segments). With 512 CPUs, we can complete the calculation in 1 month of dedicated weekends. Our current allocation, astrolas bank with 4.75% of TC2K, can only support the 128-CPU dedicated weekend runs. This problem is the simplest prototypical problem our project faces. Other ones planned will be much more expensive, particularly the ones involving radiation physics and any of the 3D ones.

   As way of background, when we use the code for a particular scientific investigation, we typically perform a series of runs where the resolution is successively increased. It is the highest resolution runs that consume the greatest amount of resources.

2. **Project funding source, size in FTEs, PI, and major contributors**
   Our project is currently funded within A Program as ASCI Research and Development. In total there are 4.5 FTEs associated with the project. Allen Kuhl is the Project Leader, and Jeff Greenough is the Technical Project Leader.

3. **Code maturity level and history**
   The current code is based on the hybrid C++/Fortran code framework developed by the Center for Computational Sciences and Engineering (CCSE) at Lawrence Berkeley National Laboratory (LBNL). The development in A Division has been under way for nearly 4 years and has focused on adding physics capabilities that include multifluid hydrodynamics, single-group and multigroup diffusion radiation. The code, called RAPTOR, is primarily a research code with a long history of large-scale high-end scientific computing. It has been used on vector supercomputers in the past and now on various MPPs, including T3E, SP2, SP3, and machines like TC2K. The hydrodynamics problems we typically run are not memory bound

but rather CPU bound. Therefore, it runs best on the machine with the fastest clock, which currently is the Alpha-based system. RAPTOR has always been checkpoint restartable.

The linear solver used in the inner iteration of the nonlinear solve in the diffusion radiation package is Hypre, developed by CASC. There are two options for multigrid solvers, SMG and PFMG. We have found that we typically need the robustness of SMG (guaranteed convergence for symmetric positive definite systems) for our applications.

The parallelization model is currently a coarse-grained parallelism where the grids on a level are distributed across the available processors. Note that the largest number of grids (and cells) is on the finest (highest spatial and temporal resolution) AMR level. A fine-grained parallelization model using Pthreads is currently under development. This will provide an additional degree of freedom for tuning AMR applications to a particular platform.

For pure hydrodynamics problems using AMR, RAPTOR has been run on up to 256 processors (SP2 and SP3) and 128 processors (Compaq Alpha) and the scaling efficiency is around 93% for each problem size and CPU doubling. The scaling for the diffusive physics is very problem dependent. The latest version of Hypre seems to scale fairly well on simple test problems at better than 80%. If the nonlinear problem is hard to solve, then the cost of the linear solution goes up. Typically though, a problem will only have a short transient where the nonlinear diffusion problem is difficult (expensive).

Commercial visualization products have historically not been of use with our AMR data sets. Hence, we have developed our own in-house visualization tools. We have begun discussions with the VIEWS group and hope to leverage some of their applications for our problems.

4. **Anticipated scientific advance at 2–8 teraFLOP/s**
   Access to a 2 TF size machine for 1–4 months would allow us to complete the first fully resolved (analogous to direct numerical simulations (DNS) for incompressible low Reynolds number flow) compressible shock-driven flows in two dimensions. This has never been demonstrated before. In addition, less resolved 3D simulations, compared to the previously described 2D ones, could be performed that surpass the resolutions currently achievable on TC2K class machines (128 CPUs for six to seven 12-hour segments). These 3D simulations would allow us to more fully resolve the turbulent breakup of vortex rings produced by shock/sphere interactions that are currently under experimental investigation (RIKSPHERE Omega Laser experiments).

5. **Anticipated scientific advance at 8–32 processor range**
   Access to banks of processors in the 8–32 processor range could serve to provide cycles for running the lower resolution runs of the refinement study and for debugging/code development.

6. **Connection with experiment**
   This project is actively supporting the Omega Laser AGEX experimental program in A Division. We (and our small user community) use the code for experimental design work.

7. **Requirements for scientific challenge at 2–8 teraFLOP/s**
   The memory available on machines like TC2K is sufficient for our current and planned computational campaigns. The local rotating disk space required would be on the order of hundreds of gigabytes per run. Tertiary storage requirements would be on the order of several terabytes. We could be ready to start a 6–18 month project as soon as the machine is

available. Machine stability is not a primary consideration if the machine can deliver the computing power advertised.

8. **Requirements for scientific challenge at intermediate capacity level**
   The numbers given for question 7 can be scaled down by about 10–100.

9. **Machine use pattern at 2–8 teraFLOP/s**
   We would perform a limited set of scaling studies utilizing up to the entire machine for hydrodynamics test problems and coupled radiation–hydrodynamics test problems. Then the attempts to obtain fully resolved simulations in two dimensions (shock/diffuse cylinder problem and the diffuse single-mode Richtmyer–Meshkov instability problem) would begin. Three-dimensional simulations of the shock/sphere Omega experiments would begin. Note that the above two computational campaigns require only hydrodynamics. High-resolution simulations of the AWE supersonic jet experiment could also begin (collaborative effort with M. John Edwards of X Division). These last simulations require coupled radiation–hydrodynamics. To go beyond 1 μm spatial resolution will be beyond the capabilities of TC2K.

10. **Machine use pattern at intermediate capacity level**
    Using the machine in this mode would be to do runs at lower resolutions (fill out the scaling studies), investigate equation-of-state (EOS) issues, and various sensitivity issues.

11. **Memory vs compute trade-off**
    One GB per processor may be a bit small for the 3D simulations. But we have some flexibility in how we generate grids and distribute them across the machine, so we could probably make it fit.

12. **Comments on value of unclassified computing**
    Unclassified computing is essential to the vitality of the Laboratory. It is our connection to the outside world. It is also our vehicle for interaction and earning the respect of the outside world. It keeps our scientists current in the latest advances in our respective fields. It definitely provides benefits to the programs by attracting top talent that can impact programmatic work. It encourages significant scientific advances that (at least for hydrodynamics and radiation-hydrodynamics numerical algorithms) advances programmatic computational capabilities.

    There currently is no resource of the class of TC2K for classified computing, and ASCI White is unavailable to all but a select few research teams. This is a significant computational shortfall that is strongly affecting our code project in its programmatic role. To fully realize my previous comments on this question, the programs need to provide comparable resources to those provided by M&IC.

Monte Carlo Methods for Solving the Schrödinger Equation for Many-Fermion Systems

Malvin Kalos

1. **Scientific intent and scope**
   This research seeks to establish and make practical exact Monte Carlo methods for solving the Schrödinger equation for many-fermion systems. We have chosen to use atomic and

molecular systems to study electronic structure, and systems of ^{3}He atoms as examples of a fermion fluid.

2. **Project funding source, size in FTEs, PI, and major contributors**
   The project is currently funded by LDRD (LW). There are about three FTEs working on it. The PI is Malvin Kalos.

3. **Code maturity level and history**
   The applications codes are fairly mature; versions have been running here for 2–3 years. They have been used mainly on SPs here, but are capable of running well on any distributed memory system, including Linux clusters. They can be scaled to very large numbers of processors. They have checkpoint restart capability.

   We have used no advanced visualization capabilities.

4. **Anticipated scientific advance at 2–8 teraFLOP/s**
   Although the codes are relatively mature, there remains much methodological development to be done. As indicated above, we are studying small molecules and fermion fluids. Quite recently, we have achieved speedups by orders of magnitude in the fluid calculations. We need to apply the same or similar methods to the molecular calculations. The pace of methods development is governed very much by access to powerful resources for testing and optimization.

   Our intermediate aim is the water dimer done by an exact all-electron method. That may be within reach with 2–8 teraFLOP/s for a few months.

5. **Anticipated scientific advance at 8–32 processor range**
   Ability to do clusters of water molecules, plus hundreds of atoms of ^{3}He. The physical chemical properties of water are at the heart of much chemistry and biology.

   Exact equilibrium properties of systems of about a hundred hydrogen nuclei plus electrons.

6. **Connection with experiment**
   This project is truly ab initio, requiring only fundamental constants. It serves as a check on experiments, and will lead to replacement of some experiment by computational methods.

7. **Requirements for scientific challenge at 2–8 teraFLOP/s**
   We need about 100 MB per processor, a few times 10^{19} floating point operations, about 1 GB local disk per node, and a few terabytes total storage. We have no advanced visualization requirements as yet.

8. **Requirements for scientific challenge at intermediate capacity level**
   The numbers given for question 7 can be scaled down by about 10–100. We can be ready to start in a month.

9. **Machine use pattern at 2–8 teraFLOP/s**
   See answer to question 7 above.

   We would need to do some tuning: parameters of the method and codes can be varied to achieve best scalar and parallel performance. A few days of validation of the parallel performance might be necessary. The methods and programs are well validated now, but for the larger systems we expect to treat, some comparison with experiment on similar systems will be invaluable.

10. **Machine use pattern at intermediate capacity level**
    We can make progress with small or intermediate systems: 10–20 particles, depending on the Monte Carlo efficiency. We would need of the order of 10 runs per system and about 50 hours of wall-clock time per run. This includes some validation and tuning.

11. **Memory vs compute trade-off**
    As indicated above, our memory requirements per node are modest. We would definitely trade away memory to get more gigaFLOP/s.

12. **Comments on value of unclassified computing**
    High-performance computing is at the heart of my project. The project is unclassified and includes collaborators who do not have access to our Secure Computing Facility and whose efforts are essential to the science. Thus, the use of M&IC facilities is essential for scientific progress. As indicated above, the project is important as a purely scientific enterprise and is of substantial importance to the mission of the Laboratory. It is, I believe, world-class work, likely to have high visibility after some additional progress.

Modifying the HYDRA and Cretin Codes for Simulating Direct-Drive ICF Implosions

Steven Langer

1. **Scientific intent and scope**
   The goal of this project is to develop a capability to simulate direct-drive Inertial Confinement Fusion (ICF) implosions by modifying LLNL codes that are currently used for simulating indirect-drive implosions. The codes that will be used are HYDRA, a 3D radiation-hydrodynamics code, and Cretin, a 3D code that is used to model atomic line emission from ICF capsules. Developing this simulation capability will provide us a tool for better understanding the relative merits of indirect and direct drive. The exercise of modeling direct-drive experiments will also help to validate the codes.

2. **Project funding source, size in FTEs, PI, and major contributors**
   The current LLNL members of the project team are Steve Langer (PI), Howard Scott, and Marty Marinak. All three are employees of DNT and spend most of their time working on other projects. Howard is the author of Cretin, and Marty is the author of HYDRA. We would like to find a postdoc, uncleared designer, or a graduate student to work on this project. This would be an excellent way to learn how to use a sophisticated multiphysics simulation code. We will collaborate with designers and experimentalists from the University of Rochester's Omega laser.

3. **Code maturity level and history**
   HYDRA is a multiphysics 2D/3D radiation-hydrodynamics code that contains the necessary physics for modeling ICF ignition targets. It has been used for 7 years to model a broad range of laser-driven targets. HYDRA was developed at LLNL and currently runs on all major parallel platforms at LLNL. It has run on over 1000 processors on SKY with good efficiency. The principal technical challenge will be learning how to use HYDRA to model direct-drive implosions. We anticipate the need for some modest changes to HYDRA's laser package, but much of the effort will be spent on finding the zoning and code parameters that work best for direct drive.

Short spatial wavelength variations in the laser intensity imprint themselves onto the ICF capsule during the first 100 ps of the laser pulse. At later times heat conduction smooths out the shortest wavelength variations. The laser package in HYDRA is not yet ready to model the first 100 ps of the laser pulse. To reduce the technical risks in this project, we intend to rely on our collaborators at the University of Rochester to provide us with an initial mass perturbation equivalent to the early time laser perturbations. We will apply this perturbation to the capsule surface and drive the capsule with laser beams from which short spatial wavelength variations have been removed.

Cretin uses temperatures and densities from HYDRA as input. It runs on all major shared- and distributed-memory computers at LLNL. For the problems considered in this project, it runs best on shared-memory computers with large memories. The principal technical challenge for Cretin will be improving its efficiency on distributed-memory computers.

The results of these codes can be visualized by MeshTV, the Terascale Browser, and a 3D visualization package developed by the PI. It should be possible to visualize results with Ensight and Visit, although we have not yet tried that.

4. **Anticipated scientific advance at 2–8 teraFLOP/s**
   See the response to question 5.

5. **Anticipated scientific advance at 8–32 processor range**
   To model the full capsule surface while resolving all the relevant perturbation wavelengths would require several hundred million zones and could easily use the entire 5 TF computer. We feel, however, that such runs are not justified until some basic questions about direct-drive implosions have been answered. For example, does the capsule shell break up during the implosion? Which wavelengths contribute most strongly to the rms perturbation of the interface between the fuel and the surrounding shell? How should the ALE package in HYDRA be tuned to best resolve the imploding capsule? These questions are best addressed by running simulations covering only portions of the capsule surface. These simulations can be performed on roughly 32 processors, but we will need to run many of them.

   Scientists at the University of Rochester have been able to explain much of their data using a model in which material from the fuel and shell is mixed together. The amount of material that is mixed is a free parameter that is set by fitting the data. Similar models were used in attempting to explain Nova data 10 years ago. These models were proposed before LLNL had the ability to model capsule implosions in two and three dimensions. In the subsequent 10 years, LLNL has developed the ability to simulate capsules in which capsule surface imperfections and spatial variations in the drive x-rays are directly modeled. These simulations show that a highly distorted interface between the fuel and the capsule shell can develop, but they do not look like the earlier models in which material was “stirred together.” These models, in which we have done our best to directly model the implosion, are in good agreement with data obtained on Nova.

   The question we are most interested in resolving is whether the experimental data obtained at Rochester can be explained by simulations in which the interface between the fuel and the surrounding shell becomes highly distorted but is still recognizably an interface. The answer to this question will greatly increase our understanding of what is happening during direct-drive implosions. This knowledge will, in turn, make it possible to design better direct-drive implosions.

6. **Connection with experiment**
   We intend to make a detailed comparison of our simulations to experimental data from Omega. The data includes thermonuclear burn rates and x-ray emission (continuum and line) from a variety of different capsules. The reduction in yield compared to a spherical implosion on Omega is thought to be due to the growth of hydrodynamic perturbations. Our simulations, if successful, should provide insight into how these perturbations grow and how to mitigate their effects. We intend to validate our code mostly by comparison to experimental data, although we may also compare to other simulations.

7. **Requirements for scientific challenge at 2–8 teraFLOP/s**
   We do not plan to request this level of access in the next 1–2 years.

8. **Requirements for scientific challenge at intermediate capacity level**
   Current HYDRA simulations have run on Blue-Pacific, where there is roughly 2 GB of memory for four CPUs. We believe that this amount of memory per CPU will be adequate for the planned capsule simulations. Cretin currently runs best on shared-memory computers. We plan to request significant amounts of time on the ES45s with 32 GB of memory that was recently purchased by M&IC. We will try to improve Cretin's performance on distributed-memory computers so that we can handle larger 3D simulations.

   A HYDRA capsule simulation on 32 CPUs produces roughly 70 GB of disk files. A Cretin simulation (that fits on a 32 GB ES45) produces roughly 60 GB of disk files. We anticipate running more HYDRA simulations than Cretin simulations because of the need to learn how to run direct-drive problems. We anticipate 60 runs per year between the two codes. This will produce roughly 4 TB of data to be stored in the tape archive.

   A 32-processor HYDRA capsule simulation run on Blue-Pacific takes 45 hours to complete (1440 CPU hours). HYDRA would probably require half as many CPU hours on TC2K. A Cretin run takes roughly 4–800 CPU hours on Frost and would run faster on an ES45. Cretin does not write restart files until it has completed a time step, and there are some simulations where a time step takes more than 24 wall-clock hours. We thus may need to request time slices on the ES45s that are longer than normal batch limits.

   We could start the project immediately. The initial phase of the project will be spent investigating the best way to run direct-drive problems. It may prove necessary to make some modifications to the laser package in HYDRA to achieve a laser drive that is smooth enough that the intentionally imposed laser perturbation is large compared to the numerical fluctuations due to tracing discrete laser rays. Cretin should be able to model direct-drive implosions as soon as HYDRA can provide it with a temperature and density structure.

9. **Machine use pattern at 2–8 teraFLOP/s**
   Not applicable.

10. **Machine use pattern at intermediate capacity level**
    As noted in the response to question 8, the first phase of the project will involve finding a way to run HYDRA such that deliberately imposed laser perturbations are correctly applied to the capsule. We will need to verify that the laser properties match those of the Omega experiments. This includes the time history, the pointing of each beam, and the intensity pattern within each beam. This will involve interaction with experimentalists at Rochester. We will also investigate zoning schemes to try to resolve the growth of perturbations on the

inner and outer surfaces of the capsule shell at minimal computational cost. HYDRA uses an ALE hydro package. One of the goals will be to find an algorithm that causes regions of fine radial zoning to follow the shell as it implodes. We will also investigate which spatial wavelengths need to be modeled and consider the advantages of running a full octant vs a patch on a soccer ball. This period of the investigation may not require heavy computer use (e.g., 1D models may be sufficient to set some of the parameters for the ALE package). After this is complete, we will begin running simulations of specific Omega experiments. We expect to run several simulations for each experiment to deal with uncertainties in the exact laser drive and beam pointing.

11. **Memory vs compute trade-off**
    The memory on Blue-Pacific or TC2K should be sufficient for running the HYDRA simulations. The memory on these machines is insufficient to run our Cretin simulations. Lack of access to a 32 GB SMP would seriously impact our ability to run Cretin.

12. **Comments on value of unclassified computing**
    Collaboration with Rochester is impractical on classified computers. We cannot send electronic results from the closed network to our colleagues, and some of them may want to run our codes. Our past use of the M&IC computers has been limited to code development and a few science runs. We would like to have enough computing resources available to us that we can carry out design studies and detailed comparisons to experiments. We hope that the acquisition of a new high-end parallel computer for "large runs" will free up time on TC2K for the type of design studies we are interested in.

Z3 Project

Barbara Lasinski, Bruce Langdon, Bert Still

1. **Scientific intent and scope**
   The goal of the Z3 project is the development of a state-of-the-art modern massively parallel 3D particle-in-cell (PIC) code. PIC codes, which self-consistently solve Maxwell's equations and the Lorentz force equation, provide a fundamental, first-principles description of laser–matter interactions. The 2D, legacy serial ZOHAR suite of codes has been in use at this Laboratory for over 25 years in support of many application areas. The control and mitigation of the instabilities that arise in the laser–plasma interactions encountered in the ICF Program have been a major application area. More recently, these codes have been used to study charged particle production and laser propagation issues relevant to the fast ignitor effort.

   Simulations on today's parallel computers and those of the near future approach time and space scales that allow for more direct comparison with experiment. Large 2D and 3D PIC fundamental studies are now possible, and a carefully constructed PIC code serves as a framework for more advanced and/or reduced descriptions that further extend the parameter regimes in which these codes operate.

   PIC codes have traditionally been at the forefront in using ever newer computers. In the 1970s, our simulation suite consisted both of 1D and 2D simulations. Now, we envision a mix of 1D, 2D, and 3D simulations. The flagship simulations will be 3D, which again will be at the leading edge in taking advantage of the dramatic increase in computer speed and

memory. The insights gained will be supported by many more numerous 2D simulations, some of which can be done on either intermediate-level or smaller computing platforms. Scientific progress with PIC codes depends on access to a balanced mix of computer resources.

2. **Project funding source, size in FTEs, PI. and major contributors**
   Z3 development is one goal of the DNT LDRD entitled "Computational Methods for Collisional Plasma Physics," now in its second year, which is funded at the 1.65 FTE level. The overall project goal is to develop the computational capability to model partially collisional plasma physics. This LDRD provides for the continued development of Z3 into a modern PIC code with a wide range of diagnostics. Then a succession of collisional models are to be incorporated into this code as it evolves into a more general tool for a wide range of DNT and LLNL applications. The PI for this LDRD is Dennis Hewett; Bruce Langdon, David Larson, Barbara Lasinski, and Bert Still are the other participants.

   In addition, there is a DOE Office of Fusion Energy Sciences project whose goal is to set the basis for a more comprehensive effort on the fast ignitor. This effort is mainly experimental, but a small amount has been set aside for PIC modeling, which so far has been a key aspect of the planning and analysis of the short-pulse, high-intensity laser–plasma interaction campaign. This project, with Mike Key as the principal LLNL contact, started in the summer of 2000 and is expected to last for 3 years.

   Several other non-DNT LDRDs set aside ~10% of their proposed FY02 budgets for PIC modeling of high-intensity laser–plasma interactions relevant to their experiments. The expectation is that Z3 would be used for large 2D and 3D simulations.

   A DNT LDRD proposal is currently being formulated to study electron transport in high-density cold plasmas. This effort, for which Max Tabak is the PI, includes pure PIC modeling with Z3 to serve as a comparison point for other code efforts in the cold, high-density plasma regime.

3. **Code maturity level and history**
   Production runs using Z3 have begun. The initial application has been the propagation of short-pulse, high-intensity laser light in underdense plasma, and definite 3D effects appear in the beam breakup. For these 3D runs, we have used as many as 256 processors (64 nodes) of the S machine for several days.

   Z3 shares many computer-science aspects with the pF3D code. Both codes, written in C with MPI used for message passing, are packages under the Yorick code development umbrella, which allows for efficient code steering. In principle, these codes run on any system to which Yorick has been ported. The Linux port of Yorick has been done. Z3 itself is currently being ported to Frost, and the Linux implementation will be next.

   Z3 currently makes dump files that are used for restart and post-processing purposes. A checkpoint restart capability has not been implemented, because the dump files work well for restart purposes. Unless required, we have no plans to implement checkpointing.

   We make heavy use of the Yorick graphics for 2D plots during a run and for post-processing from both the general dump files and those generated specifically for subsequent analysis. Three-dimensional graphics, both those being implemented within Yorick and independent

programs, are still at a very early stage in Z3 visualization. This area will need much more attention as Z3 is applied.

4. **Anticipated scientific advance at 2–8 teraFLOP/s**
   At 2–8 teraFLOP/s for 1–4 months, we envision modeling a complete *f*/8 speckle, the relevant plasma volume of a NIF beam, for an interesting period of time, ~100 ps, in terms of laser–plasma interactions. Scaling from present 1D runs for the plasma conditions in high-temperature hohlraums, each simulation would take 17–20 days on a 2 TF platform. This estimate is for a minimum plasma volume; ideally we would then scale up to larger volumes and perhaps more than one *f*/8 speckle. At the one speckle size and greater, we could then for the first time model beam propagation with a first-principles code for relevant space and time scales of a realistic plasma.

   On the short-pulse, high-intensity front, we scale our 2D simulations in the overdense region to 3D and provide a first-principles prediction on the directed angular distribution and energy spectra of the charged particles produced in these laser–matter interactions. Production and transport of these particle beams is a key component of the fast-ignitor scenario and other applications.

5. **Anticipated scientific advance at 8–32 processor range**
   Here, we envision many large-scale 2D simulations. The occasional 3D simulation would have compromises on length and time scales but would enable us to begin to answer the question of how the 3D view differs from the 2D results.

   Simulations in 2D of a short-pulse, high-intensity laser beam propagating in a plasma that spans electron plasma densities from underdense to overdense would allow us to begin to assess how beam focusing and breakup in the underdense plasma affect the subsequent production of directed particle beams. To date, very little has been done on this key question.

6. **Connection with experiment**
   In the history of PIC modeling, experiment planning and analysis have always been essential components. However, the smaller space and shorter time scales of the simulations have historically meant that there needed to be one or more intermediate steps in relating the PIC results to experiment. Now, we will be able to reduce or remove these intermediate steps. For the NIF parameter regime, experiments now can isolate a single speckle and unprecedented direct comparisons on beam propagation and instability saturation could then be made with the PIC first-principles description.

   For short-pulse, high-intensity parameters, PIC with Z3 provides a 3D description of the production of charged particles that again allows a closer link to experiment than has been possible so far. To the extent that planned enhancements allow extensions into the collisional regime, the Z3 project will also have a view of the transport of these directed beams, and that again leads to a more direct relationship with ongoing experiments.

7. **Requirements for scientific challenge at 2–8 teraFLOP/s**
   Our current particle movers have 380 FLOPs per particle in 3D and 241 FLOPs in 2D (two space coordinates and three velocity coordinates). This sequence of operations is repeated for billions of particles for a 3D simulation over many ten thousands of time steps. In these simulations, the laser frequency and wavelength are resolved. And for the explicit electromagnetic field advance, the Courant condition imposes the requirement that light

travels a distance less than a cell size in one time step. The particle processing time is overwhelmingly the determining factor in how long a run takes and is what we have used to make the estimates in response to question 3.

We mainly use single precision. The 3D code stores six words per particle; the 2D code has five words per particle. The field algorithms currently set aside 20 words per node point. However, as more diagnostics are added, the number of words per node point is expected to increase and the field storage requirements will become comparable to those for the particles. For the *f*/8 speckle simulation discussed in response to question 3, one full state of the problem is a minimum of $\sim 2.5 \times 10^{10}$ single-precision words.

Because we save the state of a problem in dump files, we need disk capacity that allows for more than two full sets of the simulation. Typically, after we successfully write one set of dump files, we store the prior one. Over a long run, we may store as many as ten sets of dump files; thus, reliable fast access to storage is essential. We also keep various history, other post-processing, and graphics output for each simulation.

These estimates assume the code as it is in production today and should be regarded as lower limits. Planned enhancements, such as particle collision models, will increase the number of operations per particle.

We would be ready to start in a few months.

8. **Requirements for scientific challenge at intermediate capacity level**
   Here, we envision many 2D simulations as described in response to question 4. Based on our history of computer use, ideally we would set system sizes and run times such that each simulation lasted a few days to a week.

9. **Machine use pattern at 2–8 teraFLOP/s**
   Our use would consist of a mix of flagship 3D simulations and many 2D runs as described above. The basic algorithms have been presented by A. B. Langdon and B. F. Lasinski in "Methods in Computational Physics," Academic Press, 1976 and by C. K. Birdsall and A. B. Langdon, "Plasma Physics via Computer Simulation," McGraw Hill, 1985, Institute of Physics, 1991.

10. **Machine use pattern at intermediate capacity level**
    Here, we would investigate algorithm enhancements such as collision and ionization models as well as more elaborate boundary conditions. These are necessary studies as we seek to expand the scope of traditional PIC modeling.

11. **Memory vs compute trade-off**
    Right now we are memory bound. Based on our present experience, we opt for more memory per processor.

12. **Comments on value of unclassified computing**
    We were heavy users of the M&IC computers in the earliest days of the Compass cluster. At that time, we used our legacy serial PIC code, ZOHAR, to model short-pulse high-intensity laser–plasma interactions. This work was part of the overall experimental and theoretical campaign on the PetaWatt laser. The PIC modeling itself led to a American Physical Society invited talk and a subsequent publication, "Particle-in-Cell Simulations of Ultra Intense Laser Pulses Propagating through Overdense Plasma for Fast-Ignitor and Radiography

Applications," B. F. Lasinski, A. B. Langdon, S. P. Hatchett, M. H. Key, and M. Tabak, *Phys. Plasmas* **6**, 2041 (1999). Our work was also cited in other invited talks and publications. We did many 2D simulations that each ran for several hundred hours. The interesting and intriguing ZOHAR results form the basis of interest in the next generation PIC codes. This effort would have been impossible without the steady, reliable access to the M&IC computers and associated storage.

To make further progress in this area, we needed a massively parallel code for large 2D and 3D runs, and we turned our efforts to Z3, a completely new code. We use Blue-Pacific and the SKY machines and plan to use Frost in the very near future, but as our effort gears up we again need the steady, consistent, reliable computer access afforded by the next generation of M&IC computers and storage.

On the Compass cluster, we used a variant of the legacy ZOHAR code, BZOHAR, that focuses on ion dynamics, to study the saturation of the Brillouin instability, an important issue in the plasma physics that is expected to occur in NIF generated plasmas for inertial confinement fusion. That work resulted in several publications including "Resonantly Excited Nonlinear Ion Waves," B. I. Cohen, B. F. Lasinski, A. B. Langdon, and E. A. Williams, *Phys. Plasmas* **4**, 956 (1997) and "Suppression of Stimulated Brillouin Scattering by Seeded Ion Wave Mode Coupling," B. I. Cohen, B. F. Lasinski, A. B. Langdon, E. A. Williams, H. A. Baldis, and C. Labaune, *Phys. Plasmas* **5**, 3402 (1998).

We currently use the Tera cluster for simulations with ZOHAR and BZOHAR both for short-pulse, high-intensity studies and modeling relevant to NIF plasmas. However, many of these projects quickly reach the limits imposed by the legacy serial code and will switch to the parallel code as the planned simulations become bigger and longer.

The renewed interest in modern PIC codes is also a recruiting tool. This past summer, two graduate students (one from Princeton, one from the University of Michigan) spent the summer here developing enhancements to traditional PIC codes for their research.

Shock Propagation in Materials

Maria J. Caturla

1. **Scientific intent and scope**
   The scope of this project is to understand the effects of shocks on materials, in particular, dynamic void nucleation, phase transformations, and shock front evolution. The results of these simulations could be used as input for continuum models, in particular in relation to anisotropies generated during shock propagation, void nucleation mechanisms, and interactions of shocks with defects (dislocations, loops, grain boundaries, etc.). The goal is to gain basic understanding of the interaction of shocks with the microstructure.

2. **Project funding source, size in FTEs, PI, and major contributors**
   The funding sources for this project are ASCI, LDRD, and NIF. The PIs are B. D. Wirth and J. S. Stolken. Major contributors are M. J. Caturla (CMS), A. Kubota (CMS), and E. Bringa (CMS).

3. **Code maturity level and history**
   The main tool used is a parallel molecular dynamics code (MDCASK). This code is written in Fortran and uses MPI for message passing as well as OpenMP. It uses spatial domain decomposition for parallelization.

   MDCASK was first developed at LLNL from a serial version. It used PVM in its first version by T. Pierce and P. Marcelin on a Cray T3D (1994). It was later modified to use MPI in 1996 and in 1999 to include OpenMP.

   MDCASK scales linearly with the number of nodes as shown in Fig. 1. Figure 1 shows the time (in seconds per time step and atom) as a function of the number of processors (calculated with ASCI Blue-Pacific) for different system sizes. The black line shows perfect scaling.

   This code has been used in many different platforms: Cray T3D and T3E, Compaq Alpha, Linux cluster, IBM SP (ASCI Blue-Pacific and Frost). It was developed under LDRD and ASCI funding. The visualization tools used are those available under M&IC.

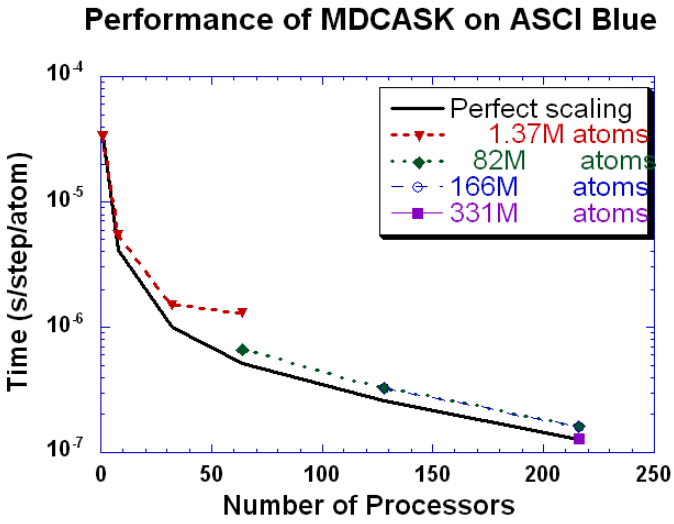

Figure 1. Time per time step and atom as a function of the number of processors computed on ASCI Blue-Pacific using MDCASK for different system sizes.

4. **Anticipated scientific advance at 2–8 teraFLOP/s**
   - Shock-induced phase transformations and structural modifications of materials
   - Shock propagation in spherical symmetries: void nucleation
   - Shock wave interaction with defects (vacancy clusters, loops, and bubbles)

5. **Anticipated scientific advance at 8–32 processor range**
   Parametric studies of void size as a function of shock pressure and pulse length in single crystals.

6. **Connection with experiment**
   These simulations are closely coupled with experiments. Laser-driven shock experiments will be performed, and interpretation of the results of those experiments will rely on the simulations explained above. Some quantities can be measured, such as the wave velocities and diffraction patterns. These will be used to validate the simulations. In many experiments, sample recovery after the shock is not possible. Simulation is the key to understanding phenomena occurring during the shock.

7. **Requirements for scientific challenge at 2–8 teraFLOP/s**

   - Memory per processor?
     0.1–0.5 GB.
   - Local disk required?
     >30 GB.
   - Visualization?
     Rasmol/Xmol/Atom TV/Renderman and similar tools.
   - Tertiary storage?
     2–3 TB.
   - Ready to start?
     Now.

8. **Requirements for scientific challenge at intermediate capacity level**
   Same as in response to question 7.

9. **Machine use pattern at 2–8 teraFLOP/s**
   (1) Spherical shocks in the presence of grain boundaries (nanocrystalline materials): we will study the effect of a grain boundary in the propagation of a shock by including small grains (~20 nm) in the material. We will look at the effect of grain boundary orientation (low and high misorientation between grains), shock pressures (at least four different values), and materials (at least two materials). A total of 144 hours will be needed for these simulations if using 256 processors.

   (2) Structural transformations induced by laser-driven shocks in different materials and the effect of surfaces: large system sizes will be necessary for these simulations to compare directly with experimental studies. Two materials will be tested and four shock pressures and four pulse shapes, a total of 384 hours with 256 processors.

   (3) Shock interaction with defects produced during self-irradiation: in this case different shock pressures and defect types will be tested. We estimate at least 48 cases will be necessary (four pressures, three defect types, four sizes), which will require 240 hours with 256 processors.

   If we assume a machine with 2048 PEs, we will require one-eighth of the machine during ~1.5 months.

11. **Memory vs compute trade-off**
    Yes.

12. **Comments on value of unclassified computing**

Using the unclassified computers we have been able to study the nucleation of voids in copper due to propagation of spherical shock waves. Figure 2 shows a slab through the center of a copper sphere. In Fig. 2(a) we show the shock wave (kinetic energy) before it reaches the center of the sphere. At this initial shock pressure (200 kbar) a void generated in the center of the sphere, as shown in Fig. 2(b).

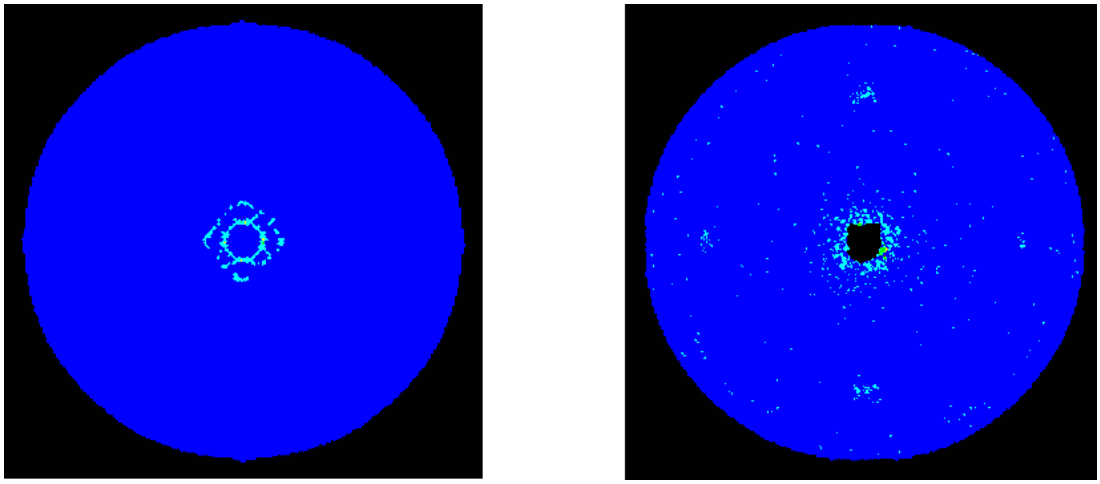

(a) (b)

Figure 2. Slab through the center of a copper sphere: (a) before shock converges in the center; (b) after the formation of a void at the center.

We have also studied the modifications produced in fused silica due to shocks. This is of interest for optics used in high-power lasers such as NIF. Figure 3 shows the results of the simulation of shocks propagating in fused silica and the changes in the structure observed, which can lead to obscuration of the optics.

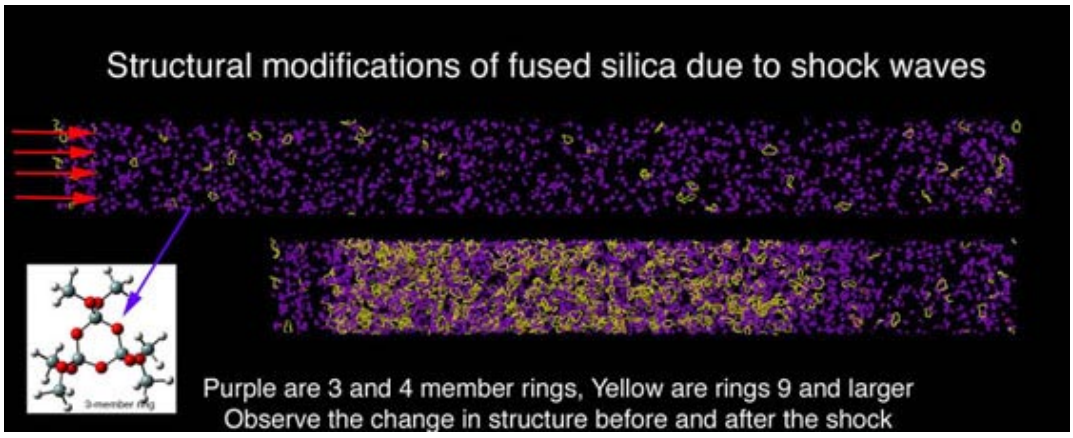

Figure 3. Results of simulations of shock waves traveling through fused silica. We show the rings of sizes 3 and 4 (in purple) and 9 and larger (in yellow) before (top) and after (bottom) the shock. Changes in density and point defects are observed that can lead to obscuration of the optics.

To advance in our understanding of shock waves propagating in complex materials it is necessary to go to larger simulation cells to be able to look at (1) size effects, (2) the effect of

grain boundaries, and (3) the effect of surfaces. The computer resources available to us now will not allow us to study these systems.

Radiation Damage of Materials

Maria J. Caturla

1. **Scientific intent and scope**
   The scope of this project is to understand and predict changes in microstructure and properties (mechanical, thermal, dimensional stability) of materials exposed to irradiation.

   This is an important problem for many areas. It is crucial for predicting safe operating lifetime of structural materials in nuclear energy technologies (both fission and fusion) as well as for predicting the long-term disposition of materials that undergo radioactive α decay (Pu in stockpile, waste). As shown in Fig. 1, the particular conditions of damage rate and defect production can differ by orders of magnitude depending on the specific problem. However, the underlying processes of damage production are similar.

   The ultimate objective is to design new, radiation-resistant structural materials for fusion energy and to predict lifetimes of radioactive materials. Physically based models are necessary to extrapolate test results to in-service conditions. For example, for fusion environments there is no neutron source available at the moment that operates under the conditions expected during operation.

2. **Project funding source, size in FTEs, PI, and major contributors**
   The funding sources for this project are the Defense Program, the Office of Fusion Energy, and ASCI. The PIs are T. Díaz de la Rubia (CMS), B. D. Wirth (CMS), and W. G. Wolfer. Major contributors are M. J. Caturla (CMS), A. Kubota (CMS), J. Marian (CMS), and J. Young (CMS).

3. **Code maturity level and history**
   The main tool used is a parallel molecular dynamics code (MDCASK). This code is written in Fortran and uses MPI for message passing as well as OpenMP. It utilizes spatial domain decomposition for parallelization.

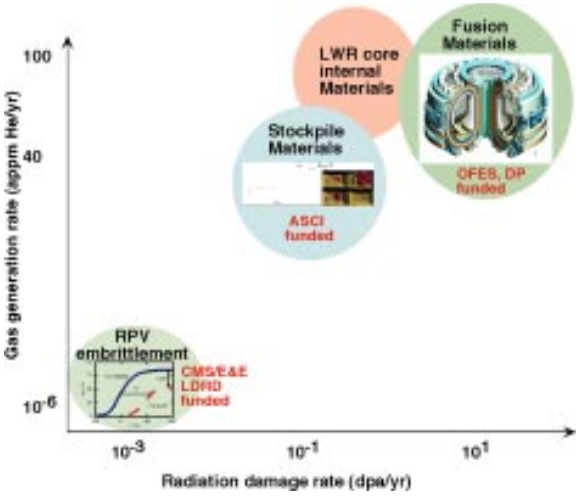

Figure 1. Damage rate as a function of gas generation rate for fusion/fission and stockpile materials.

MDCASK was first developed at LLNL from a serial version. It used PVM in its first version by T. Pierce and P. Marcelin on a Cray T3D (1994). It was later modified to use MPI in 1996 and in 1999 to include OpenMP.

MDCASK scales linearly with the number of nodes. (See Fig. 1 in the preceding response, "Shock Propagation in Materials.")

This code has been used in many different platforms: Cray T3D and T3E, Alpha, Linux cluster, IBM SP (ASCI Blue-Pacific, Frost, Red). It was developed under LDRD and ASCI funding.The visualization tools used by our group are those available under M&IC.

4. **Anticipated scientific advance at 2–8 teraFLOP/s**
   - Identify primary damage in Pu due to 85-keV recoils (self-decay): an α decay in Pu produces a recoil of 85 keV energy. We will be able to study the damage produced by such recoil with interatomic potentials developed specifically for Pu.
   - Determine damage overlap and high-dose defect production in glasses and metals.
   - Determine defect–grain boundary interaction mechanisms (important to understanding denuded defect zones observed in metals and the potential radiation resistance of nanocrystals).
   - Quantify the impact of defect clusters on mechanical strength over a wide range of defect cluster types and materials systems.
5. **Anticipated scientific advance at 8–32 processor range**
   The particular problems described here require the computer power of question 4.
6. **Connection with experiment**
   Closely coupled experiments and simulations provide opportunities for both experimental validation of simulation results and physical insight into interpreting complex experiments.

   Examples: (1) We have used in-situ straining in TEM to validate atomistic details of dislocation–defect interaction in Cu, predicted by large-scale MD simulations performed on M&IC computers. (2) Interpretation of defect migration kinetics from complex isochronal annealing experiments of Pu was made possible by coupled MD/kinetic Monte Carlo simulations.
7. **Requirements for scientific challenge at 2–8 teraFLOP/s**
   - Memory per processor?
     0.1–0.5 GB.
   - Local disk required?
     >30 GB.
   - Visualization?
     Rasmol/Xmol/Atom TV/Renderman and similar tools.
   - Tertiary storage?
     2–3 TB.
   - Ready to start?
     Now.
8. **Requirements for scientific challenge at intermediate capacity level**
   Same as response to question 5.
9. **Machine use pattern at 2–8 teraFLOP/s**
   (1) Radioactive decay: 85-keV cascades with expensive interatomic potentials (such as MEAM). Requires ~200 hours per job with 256 PEs. Several cases will be needed for statistics ~3 months. Low-energy simulations are already being done using the M&IC resources.

(2) Radiation damage of glasses: parameters are energy and overlap distance. A total of ~20 hours per job will be needed with a total of at least 60 cases, or ~2 months. Database to be used in defect evolution in glasses due to irradiation. Low-energy simulations are already being done using the M&IC resources.

(3) Radiation damage of metals: interaction of damage with grain boundaries. Response of nanocrystalline materials to radiation damage. Wide parameter range: grain boundary orientation, grain size, and energy of radiation. Requires some small-scale test cases and analysis of data to be stored; ~350 hours with 1024 PEs per job ~5 months.

If we assume a machine with 2048 processors, the simulations described above will require approximately (1) one-eighth of the machine for 3 months, (2) one-eighth for 2 months, and (3) one-half for 5 months.

10. **Machine use pattern at intermediate capacity level**
    Same as response to question 5.

11. **Memory vs compute trade-off**
    Yes.

12. **Comments on value of unclassified computing**
    Using the unclassified computers, we have been able to study very successfully the effect of single high-energy impacts in different materials, related to both the ASCI and Fusion Materials Programs. We have studied the defects produced by such events, as shown in Fig. 2(a), as well as the interaction of dislocations with those defects produced during irradiation, as shown in Fig. 2(b).

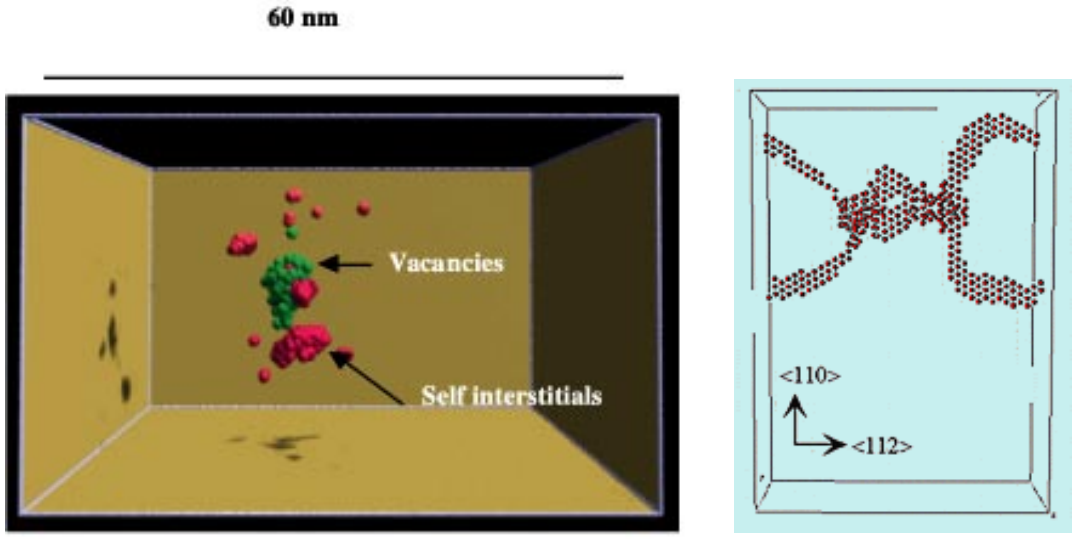

Figure 2. (a) Defects produced by a 20-keV Au atom in Au; (b) interaction of a dislocation with a cluster of vacancies produced during irradiation in Cu.

We have studied the effect of shock propagation in glasses (related to damage of optics used in NIF), and radiation damage in optics, for fusion reactors.

We have developed and implemented new interatomic potentials in our parallel molecular dynamics code. In particular, we have implemented an empirical interatomic potential for Pu developed at Los Alamos by M. Baskes. Using this potential we have looked at damage evolution in Pu: defect mobility and defect production. These are milestones for our ASCI effort on Pu aging.

Our work has been presented at the Science Day, UC National Security Review and over 10 International Conferences. It has produced over 20 papers, including Nature and Phys. Rev. B during 2000 and 2001. It has also allowed us to hire a new postdoctoral fellow, Eduardo Bringa, as well as promote people in our group from postdoc to Term positions (Alison Kubota).

However, the computer resources available today will not let us reach the next level of complexity in these simulations (damage overlap, higher energies, complex potentials, etc.) as explained in question 4. Figure 3 shows the estimated time necessary to do different system sizes, using ASCI Blue-Pacific, and assuming that we can run one 8-hour job per day with 64 processors. If we want to do a simulation with 10 million atoms (which is not enough for much of the work presented here) and time scales of hundreds of picoseconds, we will have to wait 1 month before getting results of a single job with the resources available to us now.

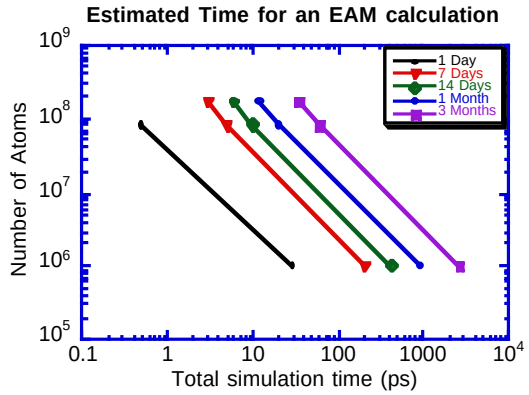

Figure 3. Estimated time for an EAM calculation assuming running one job per day (8 hours, 64 PEs) or two jobs per week (8 hours, 256 PEs) (based on experience).

Modeling of Biological Cells for Bioterrorism and Health

Carl Melius

The signaling, regulatory, and metabolic pathways controlling the biological cell are the key to understanding how life works. The understanding of cell functioning is important not only to our human health, but also to our national security, both human and agricultural. Understanding the functioning of microbial pathogens at the cellular level is critical to our biosecurity mission. By manipulating the metabolic, signaling, and regulatory pathways of the microbial cell, we can discover novel ways to shut off bacterial toxin production, defeat genetically engineered organisms, develop new antimicrobial drugs, and develop more robust diagnostic sensors for detecting bioterrorist attacks. Understanding the signaling and metabolic pathways of eukaryotic cells and their intercellular communication in tissues will help us to develop better targeting

agents for fighting cancer, develop better drug delivery systems, and counter pathogenic attacks on the body.

The post-genomics revolution has transitioned from sequencing the gene into "systems biology," the understanding of how the cells carry out their functions. DOE's Genomes-to-Life proposal (see http://www.DOEGenomesToLife.org) provides an excellent example of this new grand challenge in systems biology and identifies the important roles that computing will play. The computational challenge is not just how to handle the vast amounts of data being generated to describe the cell, but how to make sense of it all. The goal of systems biology is to determine the relationship and functionality of the proteins and metabolites involved in the cell signaling pathways, the metabolic pathways, and the other regulatory pathways. The complex network of reaction pathways involved in these pathways is referred to as the wiring or circuit diagram of the cell.

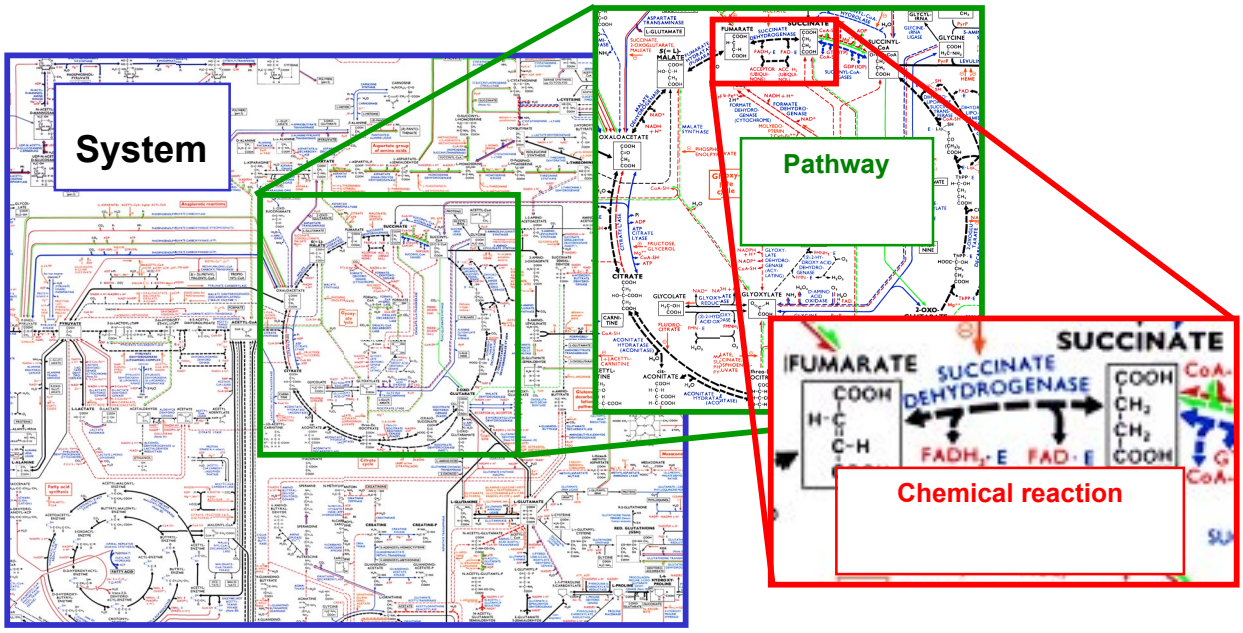

Figure 1. An example of the metabolic wiring diagram of a cell, involving interacting metabolic pathways made up of chemical reactions (see http://ca.expasy.org/cgi-bin/search-biochem-index).

The complexity of the circuit diagrams for these pathways is too great to be predicted without a detailed computational model of the cell. The computational cell model involves the solution of coupled differential equations representing the chemical reactions and transport for the metabolic, signaling, and regulatory pathways. To treat these reaction–diffusion systems of equations, we have initiated two LDRDs at LLNL to model the cell. One involves the cellular processes in a bacterial cell, representative of a biological pathogen. The second involves the cellular processes in an epithelial cell, the body's first line of defense to the outside world. The modeling of the cell will make use of LLNL's ALE3D code. ALE3D is a 3D, unstructured, arbitrary Lagrangian–Eulerian multiphysics code. ALE3D can treat reaction chemistry and diffusion as well as material transport. ALE3D is the major workhorse for modeling and simulation of physical processes within LLNL's ASCI Program. LLNL has a huge investment in implementing the code on massively parallel machines. In particular, ALE3D has been implemented on the ASCI Blue-Pacific and ASCI White architectures. Because ALE3D is

parallelized via message passing (MPI) and can run threaded, we expect that porting it to the new computer architecture will be straightforward. Furthermore, this effort fits in with our longer-term effort to have this cell-modeling code implemented on the IBM BlueGene/L machine when it comes on board. For the cell modeling, we will use a subset of the ALE3D code package, without the hydrodynamic features necessary for weapons simulation. This cell modeling code will be the working code platform for the cell modeling team.

The approach to modeling the cell within ALE3D is to compartmentalize the different functionalities involved with the metabolic, signaling, and regulatory pathways. Nature has evolved this compartmentalization approach to handle the scalability issue. This compartmentalization approach is also well suited for addressing computational scalability, dealing with hundreds of thousands of processors. Cells are compartmentalized into a multitude of organelles (e.g., mitochondria, Golgi, endoplasmic reticulum, nucleus, lysosomes and other assorted vesicles), each with their own internal compartmentalization. The ALE3D code is well suited to handle this spatial compartmentalization. The biological system to be modeled (a single cell or a multitude of cells) can be spatially discretized on an ALE3D structured or unstructured grid. Each compartment is given its own material characterization, with its own concentrations of chemical species. To resolve the material characterization of each of the subcompartments, the cell is subdivided into hundreds of grid points on a side, corresponding to tens to hundreds of millions of grid elements per cell. The total number of grid elements is the number of grid elements per cell multiplied by the number of cells to be modeled. The memory requirements per processor are roughly proportional to the number of grid points times the number of chemical species per grid point plus some working arrays (the grid elements are processed sequentially on a given processor).

Besides the ability to use ALE3D to model the intercellular and intracellular physical and chemical processes of the cell, the LLNL code development effort also provides an integrated suite of additional packages that are needed to model the cell. These packages, representing a heavy investment by ASCI as well as by ASCR, provide the pre- and post-processing codes necessary to define the cell model and to visually process the data results. These pre- and post-processing codes are integrated with ALE3D so that they will interface well with the new machine architecture. This is particularly true of the post-processing, which must access the results of the cell simulation on the various processors as the integration proceeds and send the results to the visualizers. The visualization tools being developed for ALE3D are well suited to handle the material-based compartmentalization of the 3D grid and follow the spatial–temporal evolution of the signaling, metabolic, and regulatory pathways within the various compartments of the cell.

While ALE3D is an excellent starting point to model the 3D reaction–diffusion processes of the cell, it is limited in its ability to treat the complex chemical processes carried out by the coordinated interactions of proteins. These protein machines (multiple proteins working together as a single unit) carry out much of the regulatory and signaling functionality of the cell. Examples include transporters and signal transducers that move chemicals and information across the cell membrane and protein complexes that process gene expression, replication, and repair. These molecular-scale calculations will be treated by other researchers in separate but coordinated proposals. At the cellular level, the results of these protein-machine calculations will be treated in the ALE3D cell model as chemical species within a grid element.

To handle the multitime scales, multiconcentration scales, and multicomplexity scales of these protein-machine chemical species, we are developing an object-oriented, modular, embeddable framework in which the chemical reactions and chemical species are computational objects of varying complexity. We will be designing and implementing object-oriented classes for the various chemical species, in which complex biological chemical species are derived from simpler chemical species. Simple chemical species objects will represent molecules, such as an amino acid or a hormone. More complex chemical species objects will be developed to handle proteins, which can have state (e.g., active/inactive, phosphorylated, methylated, as well as having spatial state). Increasingly more complex chemical species objects will be developed to handle a collection of proteins (a protein machine) working as a unit, such as a membrane transporter or a signal regulator. The more complex chemical species objects will be embeddable, containing their own sets of chemical species. Examples of more complex chemical species objects would be a transport vesicle involved in endocytosis/exocytosis, an organelle (e.g., the endoplasmic reticulum, a mitochondrion, etc.) or even a complete cell. The complex chemical species will have the ability to have their own chemical integrators and will be able to restrict the chemical species that are exposed to the external world. A hybrid master chemical integrator, combining both continuous–deterministic and discrete–stochastic Monte Carlo integrators, is being developed to solve the differential equations for these complex chemical species. The development of this multitiered, modular framework will proceed in parallel with the ALE3D cell-modeling effort.

Our emphasis is on the high-performance computing, making use of the new architectures, ultimately leading to IBM's BlueGene/L. The cell model, based on the ALE3D code, will incorporate the scalable system software being developed by the LLNL computer science team to run on the various high-end platforms. The cell model will also make use of the various data storage, access systems, and visualization capabilities of the overall system being developed by LLNL's computer science team.

Cell Modeling Applications

To demonstrate the different capabilities of our cell model, we will address several model systems, involving both microbial cells and human cells. To test the 3D reaction–diffusion capability of the ALE3D code, we will model human epithelial cell tissue.

Epithelial cells are a 2D layer of interconnected cells. Each of the epithelial cells is polarized, that is, it has top (apical) and bottom (basolateral) sides, composed of different membrane proteins. The purpose of the epithelial cells is to form a protective layer between the inside of the body and the outside, such as in the lungs, intestines, breast, skin, and kidneys. The epithelial cells regulate the transport of chemicals, ions, and water between the inside and the outside of the body. The transport regulation is controlled by coupled sets of intra-cellular signaling pathways, arising in both the apical and basolateral membranes. In addition, the 2D array of epithelial cells are coupled to each other by series of gap junctions, which allow intercellular communication. Most of the inter-cellular signaling is by diffusion of messenger molecules and proteins. We will investigate the transport of simple molecules across the cell and its membranes, along with the signaling proteins within the apical and basolateral membranes regulating these transport processes.

The modeling of epithelial cells should represent a good test of the capabilities of the ALE3D code. However, the ALE3D code is limited in its ability to handle complex chemical species and reaction mechanisms. In addition to the ALE3D code, we will use the object-oriented, modular framework to test the complex chemical species built up from the protein machines. We will use the chemotaxis process of microbial cells as our test case for treating these complex chemical objects. Chemotaxis is the process by which cells move in response to external chemical stimuli. The mechanism involves a two-component signaling pathway, in which the signal transducer undergoes dimerization, phosphorylation, and multiple methylation steps. The signaling and regulatory pathways will be coupled with the metabolic pathways, using the hybrid deterministic–stochastic integrator. The signaling and regulatory processes also involve collective behavior due to quorum sensing of the bacterial cells. We will therefore model the chemical reactions involved with the mobility, assembly, and interaction of multiple cells leading to their collective behavior.

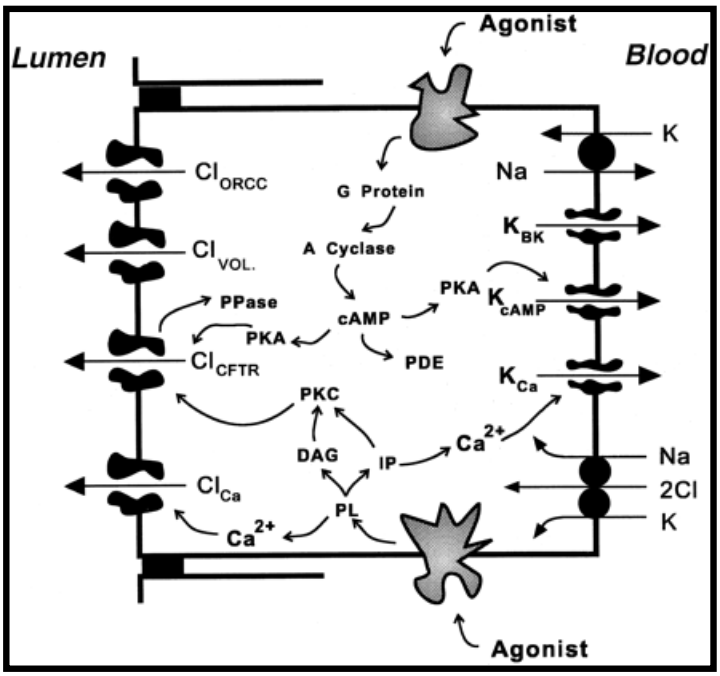

Figure 2. Diagram of an epithelial cell, indicating ion transport and regulatory pathways.

Computer Time Request

The cell modeling initiative is a new effort that fits well with the time scale for the acquisition of the new machine. The development of our cell model should make us fully operational in FY03, consistent with the acquisition time. We anticipate the need for 150 hours of total machine time per year starting in FY03 for full-scale simulations. Furthermore, we anticipate 20–30 hours of partial machine time per year, scattered throughout each year for code development and validation. Since our cell model is heavily based on the ALE3D code, we expect to leverage off the general ALE3D code porting and validation effort being carried out within the overall LLNL effort. We do not anticipate any unique requirements of our code not already present within the ALE3D code.

First-Principles Molecular Dynamics for Understanding Fundamental Biochemical Interactions

Andrew Quong, Chris Mundy

Molecular targeting, a technique to engineer antibodies to detect and ultimately kill toxic cells, is an important research area for national security and for medicine. An important aspect of engineering the antibody is to design a chelator that will bind to a specified metal ion. The demise of the toxic cell is mediated in part by its interaction with this bound metal species.

These chelating agents need not be exotic. Researchers are experimenting with common chelators such as ETDA and its interaction with metals such as indium and cobalt (see for example *J. Nucl. Med.* **28**, 83–90 (1987)). One of the challenges is to design a chelator/metal pair that will remain intact over many days, reducing the chance that the unbound metals (albeit in trace amounts) will cause damage to healthy tissue/cells. One other aspect of receptor–ligand binding is an active area of research at LLNL. For example, an LDRD-sponsored project is exploring ways to extract toxic metals such as beryllium from the body. The challenge here is to design the ligand to bind tightly to the beryllium but not to elements such as calcium and iron that are required for good health in trace amounts. There are also projects aimed at designing high-affinity ligands for both national security (single molecule detection) and health (molecular targeting).

Because molecular-level modeling in the area of ligand–receptor binding is in its infancy, an opportunity is presented to set precedents in this area using state-of-the-art computing platforms and algorithms that are present at LLNL. The need for large-scale first-principles calculations is due to the accuracy required for determining the energetics of these processes. The first-principles molecular dynamics code we run is highly optimized to run on massively parallel architectures (including all platforms at LLNL). We will be using the code for the initial thrust of this research effort, which will be dedicated to understanding ligand–receptor interactions of Tiron and beryllium to support the experimental efforts. In order to obtain an accurate picture of the energetics of receptor–ligand binding as it occurs in the body, the inclusion of solvating water is imperative. This will put us into the state of the art in terms of number of atoms simulated by first-principles molecular dynamics.

Future directions of research in this area will be to include the effects of the proteins that make up the antibody to move towards studying prototype molecular targeting agents that are studied in the laboratory. The large numbers of atoms that are needed to describe these systems render standard methods of first-principles molecular dynamics useless. In order to circumvent this problem we have proposed a methodology that will allow us to use both accurate electronic structure in the region where bonding is important (e.g., the metal–ligand interaction) and approximate quantum mechanical models to treat the large number of atoms in the protein scaffolding. The aforementioned approximate schemes are efficient (O(*N*)), accurate (polarizable), and will form a seamless interface with first-principles methods.

Another active area of research at LLNL is aimed at elucidating a molecular picture of radiation damage to DNA bases. Currently, most of the effort in this area is aimed at characterizing the biological end-points (e.g., strand breaks), and little is known about the actual molecular mechanisms that give rise to damaged DNA bases. Building on our previous results of guanine interacting with an OH-radical, we plan to use first-principles molecular dynamics to study the dimers of DNA bases to include the effects of hydrogen bonding in stabilizing radicals in DNA bases. Our major computational effort in this area will be to study a crystal of dimerized DNA bases via a first-principles approach to help elucidate experimental data on the effects of the condensed phase on radical stabilization.

Computer Time Request

The computational projects outlined above are supporting major areas of research at LLNL. The calculations are also very ambitious and require the largest computing capabilities. We anticipate the need for 250 hours of total machine time per year starting in FY03 for full-scale simulations. The proposed effort fits well with the time scale for the acquisition of the new machine.

Furthermore, we anticipate 30–40 hours of partial machine time per year, scattered throughout each year, for code development and validation.

Computational Biology Project

Felice Lightstone, Eric Schwegler, Michael Colvin

1. **Scientific intent and scope**
   The goal of the Computational Biology project is to combine LLNL's expertise in biology, advanced simulation methods, and high-performance computing to develop a laboratory core competency in computational biology. The primary focus of this project involves the use of state-of-the-art, first-principles molecular dynamics (FPMD) simulations to examine biologically relevant systems. The use of FPMD enables very accurate dynamical descriptions of biological phenomena including chemical mechanisms, enzyme catalyzed reactions, protein–protein interactions, and DNA–protein interactions.

   Currently, we are studying the cyclization reaction of phosphoramide mustard and hydroxide attack on dimethyl phosphate. These systems were chosen because of their interest in BBRP, but also because FPMD can uniquely answer the mechanistic questions where other methods cannot. Because we use explicit solvent in our FPMD simulations, our systems can contain several negatively charged species, which need to be solvated correctly. Also, we can study concerted reactions where protonation or deprotonation are concerted throughout the bond making and breaking process.

2. **Project funding source, size in FTEs, PI, and major contributors**
   The Computational Biology group (Michael Colvin, PI) has been collaborating with Physics and Computations by using JEEP as a tool to answer biological questions. This current project has been funded by LLNL LDRD SI funding since 1998. In September 2001 we received funding from the DOE. At present, our group has three FTEs, one Lawrence Fellow, and one postdoc.

3. **Code maturity level and history**
   The primary code we use for our simulations is the JEEP code, developed and maintained by François Gygi at CASC. The JEEP code is written in C++ MPI/OpenMP and has been ported and extensively tested on all of the major parallel computing facilities available at the Lab, including ASCI Blue-Pacific and White and the Compaq Compass/Tera/TC2K clusters. JEEP currently supports full restart capabilities in addition to the ability to receive DPCS signals to prematurely shut down a calculation without any loss of data (typical grace times are on the order of 10 minutes). The code has been shown to scale well on a large number of processors and has been used to perform simulations that are many times larger than what is currently feasible by any other research group in the world. At its current status we consider JEEP to be a mature code that is routinely used for production-level calculations. We have used the code for a large number of simulations that include liquid water, ion solvation, the dissociation of water in the liquid state, the conformational dynamics of dimethyl phosphate, phosphate hydrolysis, and the activation of anticancer drugs. Visualization of the simulation results is currently done with OpenDX and with a set of programs developed by François Gygi based on the VTK toolkit.

4. **Anticipated scientific advance at 2–8 teraFLOP/s**
   The majority of the problems we are interested in involve large, extended molecular systems. With our current nondedicated access to the unclassified computing facilities we are limited in the size of systems that we can directly simulate. For example, a typical model system might involve 200 atoms that are simulated for time scales on the order of 10 ps. Dedicated access to a 2–8 TF computer would enable us to perform simulations on larger model systems that more closely resemble the biological system that we are ultimately interested in.

   For example, we are currently performing simulations of phosphate hydrolysis on dimethyl phosphate. With dedicated access we could make our model system much larger. One of the main scientific focuses of the BBRP is DNA repair. In particular one enzyme, Ape1, repairs abasic sites (sites missing nucleotide bases) that can spontaneously arise in DNA and are caused by radiation damage. Ape1 recognizes the abasic site and cleaves the DNA backbone (phosphate hydrolysis) such that another enzyme can remove the abasic site and replace it with a "good" DNA base. Currently, solely because of computational limitations, we can only represent the DNA as a simple molecule, dimethyl phosphate, with no sugar or base. Also, we cannot even begin to create a model enzyme active site. Given a 2–8 teraFLOP/s platform for 1–4 months, we would be able to create a model enzyme active site and include a small piece of DNA to be cleaved. This would finally move FPMD simulations into the realm of solving a realistic biological system and would be the largest such simulation ever performed. Figure 1 shows a proposed model enzyme active site of Ape1 with 554 atoms. The DNA is shown as solid balls, and the protein is shown in sticks. The phosphate to undergo hydrolysis is shown as orange and red solid balls in the active site of the enzyme.

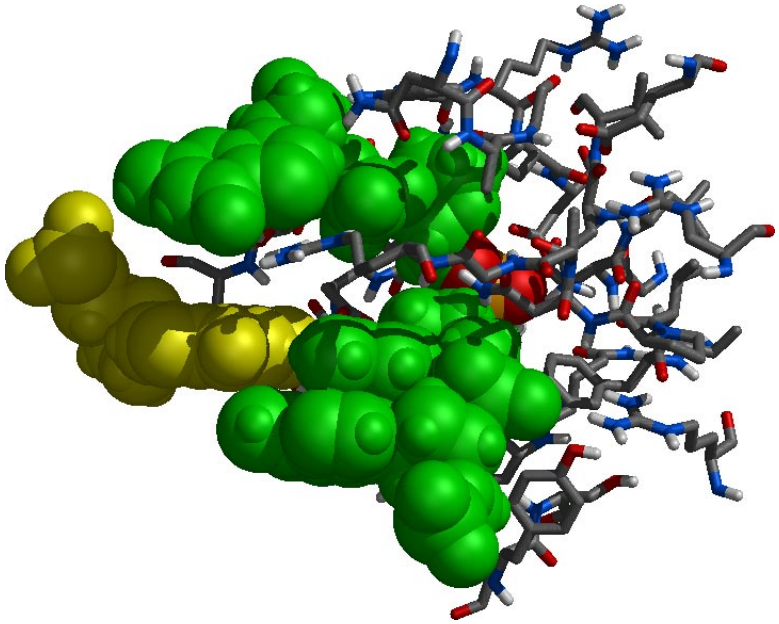

Figure 1. A proposed model enzyme active site of Ape1 with 554 atoms. The DNA is shown as solid balls, and the protein is shown in sticks. The phosphate to undergo hydrolysis is shown as orange and red solid balls in the active site of the enzyme.

5. **Anticipated scientific advance at 8–32 processor range**
   If access to a large number of 16–32 processor banks was provided, we could greatly increase the number of simulations we are currently running. This would make it possible for

us to examine the differences between a number of variations in an anticancer drug, or to sample reaction coordinates in a model system such as phosphate hydrolysis. In each case, a reaction pathway of a proposed mechanism is simulated. Along this pathway, 10–20 points are chosen to run 3-ps simulations each (10–20 independent simulations) in order to determine the free energy profile of the reaction pathway. Thus, the free energy of activation can be calculated. Each reaction mechanism can be studied and compared to other reaction mechanisms.

6. **Connection with experiment**
   The project is specifically designed to answer long-standing questions in the biological and chemical communities that traditional experimental approaches have failed to answer successfully. Because of the complexity of the problems, our simulations will complement the existing experimental data and together will answer the mechanistic questions or suggest new experiments to be performed. We anticipate that the results of the simulations can also be used as a predictive tool and would be validated by future experiments.

7. **Requirements for scientific challenge at 2–8 teraFLOP/s**
   For a large-scale calculation, our simulations typically run with up to 2 GB/node of memory. Since a small fraction of the calculation must be computed on a single node, adding processors does not reduce this requirement. Depending on the run-time length of the queues that are available, the simulations do not depend on particularly fast I/O, but do require a large global file system to save restart files and trajectories. A 1-TB global file system with 40 GB of local disk space is sufficient for most calculations. Access to a larger global file system would make it possible to save a more detailed description of the electronic structure of the system during the dynamics and would facilitate the post-processing and visualization of the simulation data. Visualization of the simulation results could be done with OpenDX along with a suite of VTK-based programs developed by François Gygi. Additional visualization support for making high quality animations of the large volume of data that will be produced may need to be investigated. For long-term storage, we will need up to 5 TB of storage space. A large-scale simulation that would run within these limits could be started now.

8. **Requirements for scientific challenge at intermediate capacity level**
   For an "intermediate capability level" calculation, our simulations could be performed with 1 GB/node of memory. The majority of the simulations could be run on 16–32 nodes. Depending on the final configuration of the machine, it would probably not be feasible to use fewer than 16 nodes. Simulations of this size typically require 2–5 GB of globally accessible disk space per simulation. The visualization requirements are identical to those of the larger scale simulations. Our group is currently performing simulations of this size, and additional simulations could be started now.

9. **Machine use pattern at 2–8 teraFLOP/s**
   The application we are using for this project, JEEP, has been extensively tested with large-scale simulations. For a suitably large system, the code is capable of efficiently using thousands of processors. Although not routine, calculations of this size are entirely feasible and could be started now with only a small number of complications. With resources of this size, the model of the Ape1 active site described above could be simulated, which would provide important information about the mechanism of phosphate hydrolysis in a realistic

environment. This simulation would be among the largest FPMD calculations performed to date.

10. **Machine use pattern at intermediate capacity level**
    With dedicated access to banks of 16–32 processors we could efficiently perform many simulations of systems of 200 to 300 atoms. Computer resources of this type would be ideal for a number of the projects we are currently working on. For example, in the calculation of a free energy profile for a reaction in solution we need to perform 10 to 15 independent simulations along a reaction coordinate. Each of these simulations need to be 3–4 ps in length and could be computed with approximately 1 month of dedicated access per simulation (all simulations can be run simultaneously). For a project such as the activation of phosphoramide mustard anticancer drugs this would enable us to develop a deeper understanding of how variations of the drug affect its activity and could be directly verified by experimental measurement.

11. **Memory vs compute trade-off**
    Since a small fraction of the calculation needs to be computed on a single node (the minimum memory size is 2 GB/node), reducing the amount of memory will put a limit on the size of simulations we can perform, no matter how many processors are available. We would prefer not to trade memory for an increase in the number or speed of the nodes. In addition to memory, we should mention that these simulations depend quite heavily on fast communication. A fast interconnect such as Quadrics (>300 MB/s) should be viewed as an essential component of the machine.

12. **Comments on value of unclassified computing**
    Unclassified computing is absolutely essential to the success of the projects performed by the Computational Biology group. Without access to large unclassified computing resources, we would be severely limited in the system sizes that we could study. For our current projects, we would like to investigate many of the enzymatic mechanism questions that can only be answered within a FPMD model. Also, we would like to increase the size of our simulations to be more biologically relevant (500 atoms and up). To achieve these goals, we are currently limited primarily by computer time. The return on investment is that we will be able to advertise that we can do the very large biological simulations because the lab has provided us the computer time we need.

    Other institutions (such as NIH) are very interested in this relatively new computational tool as applied to biology. However, as noted in the review of one of our recent grant applications, the reviewer's primary complaint focused on the biological relevance because of the short simulation times and small system sizes that we are currently limited to. These problems can only be resolved with access to more computer time. In addition, the reviewers have doubted the commitment of computer time by the LLNL. This problem can also be resolved by specifying that computational biology projects will have a commitment to dedicated computational resources at the Lab.

Green's Function Molecular Dynamics (GFMD) Project

Lin H. Yang

1. **Scientific intent and scope**
   The Green's Function Molecular Dynamics project is based on quantum-based atomistic simulations of materials properties in transition metals contributing to the LLNL program on multiscale modeling of strength and failure. The GFMD code uses parallel time-dependent Green's function techniques to treat the boundary conditions and quantum-based model generalized pseudopotentials theory (MGPT) for interatomic interactions. The ability to predict dislocation properties is essential to provide direct input into higher-length-scale multiscale simulations, including dislocation mobility for microscale dislocation-dynamics simulations of plasticity.

2. **Project funding source, size in FTEs, PI, and major contributors**
   In FY02, the GFMD project is funded by ASCI Dynamics of Metals (DofM) and LDRD (High Pressure SI). The total number of FTEs at the current level is three. The main author of the code is Lin Yang (PAT, H Division). Other contributors include John A. Moriarty (PAT, H-Division), who is the co-PI for the ASCI DofM Program and the High Pressure SI project, and Per Soderlind (PAT, H Division).

3. **Code maturity level and history**
   GFMD is currently a production code with ongoing efforts for better efficiency and more features to deal with various dislocation geometries. The 35,000-line source is written in C with MPI library calls to decompose the problem over the multiple nodes of a massively parallel system. Therefore, the GFMD code is transferable across different platforms that include Compaq Compass/Tera/TC2K clusters, Linux clusters, and the ASCI Blue-Pacific machine. The largest number of processors we have used with GFMD is 256 processors (64 nodes) of the ASCI Blue-Pacific machine and IBM SP P3 (64 SMP nodes) machine at DOD's Aeronautical Systems Center Major Shared Resource Center (ASC MSRC). We use AVS/Express as our primary visualization tool. GFMD is coupled with AVS interface network for real-time data visualization.

   GFMD has checkpoint restart capabilities and parallel I/O (home-grown C read-write routines and MPI-IO) for file access. Duplicated files are generated during production runs that allow backtracking if the system fails during I/O.

4. **Anticipated scientific advance at 2–8 teraFLOP/s**
   We are interested in such a capability to run large 3D calculations of dislocation–dislocation interactions that involve 2–3 millions atoms distributed over 500,000 structured cells.

   We have been running some scaled-down (~200,000 atoms) 3D calculations on TC2K and IMB SP P3 machines to check the scalability and accuracy of our calculations. In order to get meaningful physics out of our simulations, we need a much larger machine to run realistic 3D simulations of at least 2–3 million atoms distributed over 500,000 cells. This type of problem will account for about 20% of our total computational needs.

5. **Anticipated scientific advance at 8–32 processor range**
   This type of capability would allow us to run a long-time simulation in nanosecond time scale (a million time steps) for screw dislocation kink-pair configurations. This study would allow us to build an accurate database for screw dislocation mobility that is the key input into

much larger-length-scale simulations. We are currently running this type of calculation on the TC2K and ASCI Blue-Pacific machines, but with an average of one data point per week will not be sufficient to meet our programmatic milestones set for the ASCI DofM project. We need to have at least two data points per day in order to have any significant impact to the LLNL program on multiscale modeling of strength and failure of materials.This type of capability will account for about 50% of our total computer needs.

6. **Connection with experiment**
   A goal of the LLNL program on multiscale modeling of strength and failure of materials is to provide quantum-level prediction of strength and failure of materials under extreme (high-temperature and high-pressure) conditions that are impossible to conduct in the laboratory. In order to ensure that our simulation code is accurate and reliable for such a mission, we will need to utilize the experimental data to validate our GFMD code. Some experiments such as in-situ transmission electron microscopy (TEM) and high resolution electron microscopy (HREM) have been established to provide a direct link between our simulations and experiments.

7. **Requirements for scientific challenge at 2–8 teraFLOP/s**
   Our requirements are strongly driven by the fact that we are much more CPU-limited than memory-limited. A faster single-CPU performance means better turnaround for our code. For this type of capacity, a ratio of 1/5 (bytes of memory/flop) is sufficient to meet our needs.

   Our global disk requirements are driven by the need to save visualization and restart files for a preset time interval. For this type of capability, we estimate that 5 GB global disk space will be sufficient to hold the visualization files for each time interval, and for a typical capability science run a total of 1000 time intervals is necessary. A restart file requires about 6 GB global disk space, and there are two restart files for one such run. Based on this estimate, the global storage requirements for the data from one large run would be about 5 TB. Data compression is not implemented in the current version of GFMD code but is under development.

   We have been using AVS/Express as our run-time visualization tools. However, we have no experience of visualizing data in such size and would require much help from M&IC to explore other options.

   We estimate that we would be ready to begin such large simulations in July 2002.

8. **Requirements for scientific challenge at intermediate capacity level**
   Our requirements for such capacity would be met with a machine at the level of TC2K or IBM SP P3. We are currently running this type of simulation on a daily basis, and they will account for about 30% of our future computer needs. A 1/2 ratio for bytes of memory/flop and a 1/1 ratio for bytes of disk/flop are sufficient for this type of capacity.

   We can handle visualization using AVS/Express.

   Our code is ready for such tasks.

9. **Machine use pattern at 2–8 teraFLOP/s**
   We think that such a machine would give us sufficient computer power to simulate a full-scale 3D modeling of dislocation junctions that have been observed in TEM experiments. This type of capacity will make our simulation capability the world-class leader in the field of atomistic simulations of dislocation–dislocation interactions and multiscale modeling of

strength of materials. This type of simulation will account for about 10% of the whole project effort, and it is an important step to have direct comparison to experiments.

Although our code is designed and structured to attack this type of problems, we have very little knowledge about the scaling and convergence of our algorithm on such a machine. We should be able to validate our simulations and algorithms by direct comparison with TEM and HREM experiments.

10. **Machine use pattern at intermediate capacity level**
    This type of resource would allow us to run a large number of simulations for various dislocation configurations such as kink formations. The nominal wall-clock time for a single run would be about 8–12 hours. This study would allow us to build accurate database for a screw dislocation mobility model that is the key input into much larger-length-scale simulations. Using these resources would allow us to have two data points per day to meet the future milestones set for the LLNL program on multiscale modeling of strength and failure of materials. This type of capability will account for about 50% of our total computer needs.

11. **Memory vs compute trade-off**
    Yes, we would trade away memory for more nodes or single-CPU performance. We are much more CPU-limited than memory-limited.

12. **Comments on value of unclassified computing**
    The resources provided by the M&IC have been essential for the success of our ASCI DofM and LDRD funded High Pressure SI projects. In particular, the TC2K and Linux clusters have allowed us to test the code and make important contributions to our presentations in various ASCI and LDRD review meetings. The ASCI machines have provided very reliable resources for our computational needs, but they had become so heavily utilized that it was difficult to get sufficient turnaround on batch jobs to make significant progress. Because of this, we had to rely on other resources such as the one at ASC MSRC to meet our unclassified computing needs. We are hoping that M&IC can enhance its current capacity so that the waiting time in the queuing systems can improve significantly through this proposal.

Boundary Plasma Turbulence Code (BOUT) Project

Xu Xueqiao

1. **Scientific intent and scope**
   Magnetic fusion devices exploit a basic asymmetry of plasma in a strong magnetic field: the transport of plasma across magnetic field lines is greatly inhibited relative to that along magnetic field lines. In particular, in toroidal magnetic fusion devices, such as tokamaks, field lines in the core of the device lie at least approximately on closed toroidally shaped surfaces called flux surfaces, giving rise to good confinement of particles and heat. However, on the periphery of such devices, there inevitably exist magnetic field lines that are in direct contact with bounding surfaces. In this region, known as the scrape-off layer, the aforementioned asymmetry is manifested through a rapid loss channel (along magnetic field lines) for heat to the surrounding environment. Relative to the core plasma, this region tends to have steep radial gradients in temperature and density and to be relatively cold. The low

temperature coupled with proximity to bounding surfaces results in a relatively high concentration of neutral gas and impurities.

The boundary region tends to be characterized by large (in terms of percentage of equilibrium values) turbulent fluctuations. The presence of such fluctuations in the closed-flux-surface region leads to large radial edge transport and low plasma pressure at the edge, which is unfavorable for global confinement. However, this large turbulence in the scrape-off layer outside the last closed flux surface also leads to a less localized plasma–wall interaction, which is generally favorable for reducing peak heat loads and impurity influx. This edge turbulence has been observed for many years on nearly every toroidal confinement device, but is still not understood from first principles.

The BOUT project is building on the current generation of the boundary plasma turbulence code, the BOUT code. BOUT models the 3D electromagnetic boundary plasma turbulence that spans the separatrix using a set of fluid moment equations with the neoclassical closures for plasma vorticity, density, ion and electron temperature and parallel momentum, and with proper sheath boundary conditions in the scrape-off layer. The unique BOUT simulation capabilities have led to fundamental discoveries and worldwide recognitions. BOUT simulation results have been used for benchmarking with experimental measurements for boundary plasma turbulence research on various magnetic fusion devices, such as DIII-D at GA, C-Mod at MIT, and NSTX at PPPL. BOUT will be used to uncover the basic physical mechanisms of important edge phenomena and to predict edge plasma behavior for evaluation and optimization of future devices. The code will be applied to key issues including the role of nondiffusive and/or large-event-dominated transport, the transition to the enhanced high-confinement mode, edge-localized modes (ELMs), core density-limit phenomena, and characterization of the loss and fueling channels through the edge plasma.

An important improvement of the physical model that will be made in BOUT is the incorporation of neutral transport and the associated neutral effects on both plasma turbulence and plasma particle, momentum, and energy balances. We will initially use a fluid neutral description since this allows the neutral dynamics to be followed on the same time scale as that of the plasma, and we have extensive experience in fluid neutral modeling. This extension will allow self-consistent modeling of plasma and neutral particle transport at the edge and is expected to be a crucial ingredient for the study of the density limit.

2. **Project funding source, size in FTEs, PI, and major contributors**
   The BOUT project is funded by the MFE Theory and Computation Program. There are about two FTEs on the team at LLNL. The main author of BOUT is Xueqiao Xu of PAT/MFE, and the major contributors are Ron Cohen, Bill Nevins, Tom Rognlien, and Marv Rensink at LLNL. There are also a number of collaborators in the U.S. and around the world.

3. **Code maturity level and history**
   The BOUT code solves for plasma fluid equations in three dimensions. The complication is that the simulation domains have to include different regions due to the magnetic configuration in the boundary, such as the region somewhat inside the separatrix and extending into the scrape-off layer; the private flux region is also included. A finite-difference method is used, and the resulting difference equations are solved with a fully implicit Newton–Krylov solver, CVODE/PVODE. BOUT is a parallelized code based on domain decomposition that uses the Message Passage Interface (MPI) system.

BOUT has been tested and run on Sunbert, TC2K, and the Linux cluster at Livermore Computing. BOUT has also been developed and run on the NERSC Cray T3E and IBM SP. The same code runs on a one-processor workstation or on parallel platforms.

The parallelized version of BOUT works well with a poloidal domain decomposition, giving a factor of 13 speedup for 15 processors on the Sun Wildfire and a very encouraging factor of 69 speedup for 60 processors on the T3E-900. The supralinear speedup on the T3E is likely due to the better utilization of cache memory for the larger number of processors. Most recently, we have extended this case to 120 processors on the T3E, and find the data on the same offset linear curve.

4. **Anticipated scientific advance at 2–8 teraFLOP/s**
   For a tokamak, the most studied of toroidal confinement devices, three major types of phenomena are routinely observed: edge transport barriers, strong ELMs, and a density limit. Details of these events are given in the following description of an experimental discharge. (1) As the core heating power is increased in some regimes, spontaneous shear flow is believed to suppress the level of turbulence at the edge, leading to an edge transport barrier and transition from "L" (low) to the more favorable "H" (high)-mode energy confinement regime. (2) Simultaneously, there often occurs a new phenomenon of relatively regular bursts of energy and particle loss from the edge, known as ELMs. (3) With additional fueling beyond this level, one soon encounters first the so-called critical density limit, where enhanced and strongly intermittent edge plasma transport results in large plasma flux to the first wall and eventually disrupts the whole discharge.

   An important improvement of the physical model that will be made in BOUT is the incorporation of neutral transport and the associated neutral effects on the plasma turbulence and the plasma particle, momentum, and energy balances. This extension will allow self-consistent modeling of plasma and neutral particle transport at the edge and is expected to be a crucial ingredient for the study of the density limit. Since these simulations require multimonth runs under the current computing capabilities, we have hesitated to start the efforts. However, with dedicated access to an M&IC Open Computing Facility capability platform of 2–8 teraFLOP/s for 1–4 months, these goals seem within reach.

5. **Anticipated scientific advance at 8–32 processor range**
   We are able to investiagte the first event described above for a simpler magnetic geometry, such as circular cross section tokamaks and scaling calculations. The evolution of the plasma profiles during the time of an L–H transition can be followed within a turbulence code that evolves the background profiles.

6. **Connection with experiment**
   BOUT contains much of the relevant physics for the pedestal barrier problem for the experimentally relevant X-point divertor geometry and has been a unique tool for understanding the experimental measurements. Encouraging results have been obtained when using measured plasma profiles in current generations of major U.S. fusion devices such as DIII-D, C-Mod, and NSTX. The resistive X-point mode has been identified in an X-point divertor geometry. Comparison of the shifted-circle vs X-point geometry show the different dominant modes and turbulence fluctuation levels. The poloidal fluctuation phase velocity shows experimentally observed structure across the separatrix in many fusion devices. The fluctuation phase velocity is larger than $E \times B$ velocity. The quasi-coherent mode is believed

to be responsible for the high-energy confinement (H-mode), yet acceptably low particle (impurity) confinement in the Alcator C-Mod high-density plasma regime. The experimentally measured dispersion and mode stability is in good agreement with the resistive ballooning X-point mode predicted by the BOUT code. A strong poloidal asymmetry of particle flux in the proximity of the separatrix may explain the paradox of the JET probe measurement of the particle flux when comparisons of the limiter vs divertor experiments had been made. Our L–H transition with simple sources added shows transitions with resistive X-point modes dominating the L-mode. The levels of turbulence are similar to experimental measurements.

The BOUT code has been successfully benchmarked against experiments in DIII-D, Alcator C-Mod, and NSTX, and it should be ready to make predictions about the edge plasma behavior in future devices such as ITER-FEAT, FIRE, or ARIES. These predictions should be extremely valuable for the design of these machines and for the progress of magnetic fusion science.

7. **Requirements for scientific challenge at 2–8 teraFLOP/s**
   For a typical run now on TC2K with a resolution $64 \times 64 \times 64$, BOUT needs 550–700 MB memory per processor, 45 GB total memory, and about 2 TF.

   For a 2–8 TF resource, we anticipate for a typical run with a resolution $128 \times 128 \times 128$ and 10 times longer for a so called "density limit" run, BOUT needs 1.5 GB memory per processor, 400 GB total memory, and about 200 TF.

   BOUT only uses global disk now on TC2K. For one run now it needs 80 GB global disk. For the proposed run, it needs 800 GB global disk.

   BOUT uses our own visualization tools, xgrafix and IDL. We would like get some supports to view 3D data and generate animations.

   For 800 GB per run, 100 runs will need 80 TB.

   We would be ready within 1–2 months to start a 6–18 month project.

8. **Requirements for scientific challenge at intermediate capacity level**
   In this case, we will have to scale down our problems as we describe above.

11. **Memory vs compute trade-off**
    No.

12. **Comments on value of unclassified computing**
    Unclassified computing (including M&IC) has been a valuable resource for the BOUT project. The project has had a banner year, as measured by the number of invited papers presented in major international meetings (five), in work by the groups who are collaborating with members of the BOUT project (GA, UCSD, Lodestar, IPP/Germany), and by the clamor of experimentalists at GA, MIT and PPPL (including the Lab director at MIT) for more BOUT attention. Without the resource from unclassified computing (including M&IC), it would not be possible to develop the BOUT project. With increased interest in BOUT, we will be working with OFES/DOE to secure funding for a postdoc to work on the BOUT project.

Ab Initio Nuclear Structure Project

Erich Ormand

1. **Scientific intent and scope**
   The goal of the Ab Initio Nuclear Structure project is to attempt a first-principles description of the structure of light nuclei. By this, we mean to address the question: Do we really know how nuclei are put together? In particular, we wish to determine if our knowledge of the fundamental interactions between pairs of nucleons is sufficient to describe the rich and complex structure observed in nuclei. This is a topic of fundamental importance in nuclear physics, and our work will represent a significant improvement in our understanding of nuclei. Furthermore, these first-principles descriptions of nuclear structure will enhance our ability to accurately calculate cross sections for nuclear reactions. This is important for the nuclear data effort for SBSS within PAT.

2. **Project funding source, size in FTEs, PI, and major contributors**
   In FY02, the Ab Initio Nuclear Structure project will be funded by LDRD and PDRP. At present, the total number of FTEs is two: the PI is Erich Ormand, and the Co-PI is Petr Navratil. A proposal has been submitted to DOE Office of Science to support an additional postdoctoral researcher in FY03. We have active collaborations with academic researchers: Bruce Barrett (Univ. of Arizona), James Vary (Iowa State Univ.), Wick Haxton (Univ. of Washington), Calvin Johnson (San Diego State Univ.), and Ettiene Caurier (CNRS, Strasbourg).

3. **Code maturity level and history**
   Our project has used several "home grown" computer codes. Our approach utilizes effective interaction theory within the framework of the nuclear shell model. The goal of effective interaction theory is a renormalization of the true interaction that enables "exact" solutions in manageable model spaces. The most important feature of effective interaction theory is that the effective interactions have two-, three-, four-, etc., body components even if the true interaction is only two-body in character. Although the computational complexity increases dramatically as the number of clusters in the effective interaction increases, convergence to exact results is improved.

   We have developed computer programs to evaluate the effective interaction and to carry out the many-body shell-model calculations. Both aspects of our project are computationally intensive. The procedure for the effective interaction requires exact solutions for two- and three-body problems, and a transformation from intrinsic coordinates into a form suitable for the shell-model calculations. In particular, we have developed a Fortran 90 computer code called MANYEFF for solving A = 2, 3, 4 nucleon systems interacting by modern realistic interactions including the real three-body interaction. In addition to finding exact bound solutions using a large harmonic oscillator Jacobi coordinate basis, the code computes the effective interaction from these solutions, which is then transformed by supporting codes to a form suitable for the many-body calculations performed in the Slater determinant single-particle coordinate basis. The code is parallelized using OpenMP and allows for restarts. The many-body calculations amount to finding the eigenvalues of large matrices (dimensions of the order 10^6–10^7), and they utilize the Lanczos algorithm, which is based on the application of the Hamiltonian matrix to an arbitrary basis vector. This leads to a tri-diagonal form with eigenvalues converging to the lowest (and most interesting) values. We have used two

programs—MFD, developed at Iowa Sate and modified at LLNL, and REDSTICK, entirely developed at LLNL. The codes are written in Fortran 90, are parallelized using the Message Passing Interface (MPI) standard, and have been verified. The primary difference between the two codes is that while MFD calculates the matrix to be diagonalized, stores it on disk, and retrieves it for each iteration, REDSTICK applies the matrix to each vector "on the fly" using an efficient algorithm to essentially recalculate the matrix each time. Consequently, the primary difference between the codes is in their need for temporary disk space storage. REDSTICK is currently undergoing an improvement that will lead to a six-fold improvement in performance and will likely be better suited for future applications. Both codes have been used extensively on ASCI Blue-Pacific, while MFD has been ported to an 8-node SGI system. Code portability has been maintained as each has been tested on serial versions of Sun, Linux-Intel (Portland-Group compiler), and Linux-Alpha machines. A port to a Linux-Alpha Beowulf cluster is under way. Calculations utilizing up to 100 processors (25 nodes) on ASCI Blue-Pacific have been carried out. Limited tests of code efficiency have been carried out, with REDSTICK exhibiting approximately $N^{0.8}$ scalability over the range $2 \leq N \leq 6$. Both REDSTICK and MFD have checkpoint restart capabilities, which is a feature that has been used extensively on ASCI Blue-Pacific.

At present, we have not utilized any visualization capabilities. We are investigating ways to display our wave functions in coordinate space. In addition, we have developed codes for post-processing of our obtained many-body wave functions that allow to evaluate transition densities that serve as input for the nuclear reaction cross section calculations. These codes are written in Fortran 90, parallelized using MPI, and allow for restarts).

4. **Anticipated scientific advance at 2–8 teraFLOP/s**
This type of capability would allow us to pursue at least two possible avenues. Both are extraordinarily difficult from the computational point of view and are complementary. In particular, our choices are to either increase the size of our model space or to include higher clusters in the effective interaction. In the past year, we have demonstrated convergence to exact solutions to within 500 keV (or about 1%) for the binding energy for the nuclei ^{6}Li, ^{8}Be, and ^{10}B utilizing three-body effective interactions in a model space including excitations across four major shells. Dedicated access as described above would enable us to extend these calculations across six major shells, which would improve convergence to approximately 200–300 keV and guarantee our ability to complete calculations for all relevant nuclei up to ^{16}O.

An alternative approach would be to go up to four-body clusters in the effective interaction. Due to computational limitations, the model space would be limited to excitations across four shells. It is likely that the four-body clusters would lead to convergences at the level of approximately 100 keV and may, in fact, be essential for describing a class of excited states in ^{12}C and ^{16}O that are thought to be α-particle clusters. Including four-body interactions would represent a significant computational challenge and would only be possible with dedicated access to significant computational resources.

At present, we are fully capable of achieving the first approach given sufficient computational resources. This would represent a significant advancement in first-principles descriptions of nuclear structure. We still need to work on developing converged four-body solutions to evaluate the four-body interaction, and would likely be able to follow this approach in January 2003.

5. **Anticipated scientific advance at 8–32 processor range**
   This type of capability would enable us to carry out extensive calculations with three-body effective interactions for most nuclei targeted in the study. In addition, we are testing various prescriptions for the nucleon–nucleon interaction. A promising new development is based on quantum chromodynamics and effective field theories. Early indications are that these interactions may be easier to work with in our formalism, and considerable development and testing is needed. Hence, this would enable us to expand our research effort.

6. **Connection with experiment**
   The goal of the Ab Initio Nuclear Structure project is to describe the structure of light nuclei and compare the results with experimental data. In addition, the interpretation of some experiments, such as tests of the weak interaction using β decay, requires corrections due to nuclear structure. Our project will enable us to explore a vast array of nuclear phenomena that are verifiable with experiment. Our project does not require input from future experiments in order to proceed.

7. **Requirements for scientific challenge at 2–8 teraFLOP/s**
   Our requirements are strongly driven by both CPU and memory limitations. As we increase either the model space or the number of clusters in the effective interaction, our CPU and memory requirements increase. For our planned calculations utilizing three-body effective interactions, our production runs would require 2 GB memory/processor. The disk storage requirements for the on-the-fly code are not that great, as the matrix is essentially recalculated. These runs would require of the order 0.5 TB disk space. For the alternative code, MFD, the calculations require vast disk space. For our largest planned calculations, about 2 TB of disk space may be needed. In addition, the proposed high-capacity machine would allow us to perform calculations using the real three-body interactions as well as the induced four-body effective interactions. This would result in a realistic and close-to-convergence description of light nuclei, including clustering effects. The obtained wave functions then would allow a realistic description of nuclear cross sections for light nuclei. We are ready to start the project within a few months.

8. **Requirements for scientific challenge at intermediate capacity level**
   The "capacity at intermediate capability level" runs would allow us to perform nuclear structure calculations using the real three-body interactions and the induced three-body effective interactions for the p-shell nuclei in larger basis spaces than we are currently able to use with the machines we have access to. This would result in an improved convergence and more realistic wave functions suitable for the description of nuclear cross sections. We would require 1 GB/processor and access to 1 TB of disk storage for temporary data. We are ready to start the project now.

9. **Machine use pattern at 2–8 teraFLOP/s**
   As stated, starting January 2002, we will be ready to carry out full-scale calculations with three-body effective interactions spanning six major shells. Consequently, our primary focus would be to carry out production runs for nuclei of mass 8–18. Our codes have already been validated on smaller machines using smaller basis spaces. Scaling tests would have to be performed on the massively parallel machine. However, our current experience with our codes, based on their performance on Blue-Pacific using up to 100 processors, suggests a reasonable scaling; thus, there would be a substantial benefit from such a machine for our project. Given that we have only been able to use 100 processors on Blue-Pacific, we have

not been able to fully determine the scaling limits. The primary algorithm that needs to be investigated is the reorthogonalization of the Lanczos vectors, which must be done after each successive iteration. It is possible that improvements in this algorithm may be needed in order to achieve better scalability.

10. **Machine use pattern at intermediate capacity level**
    A machine of this type would be an important auxiliary for a project. In particular, we would utilize it to run the OpenMP codes for effective interaction computations using a single shared-memory node. About 1 GB of memory per processor and at least four processors per node would be optimal for these calculations.

11. **Memory vs compute trade-off**
    No, we would not prefer such a trade. We are much more memory-limited than CPU-limited.

12. **Comments on value of unclassified computing**
    The resources provided by the M&IC have been essential for the success of our LDRD-funded project. In particular, the Compass, Tera, and Linux clusters has allowed us to run the MANYEFF code in ways that would not have been possible otherwise. The ASCI machines (in particular, Blue-Pacific) have become crucial for our many-body calculations using the REDSTICK and MFD codes. A few issues regarding current computational capacity have limited our ability to push forward. These range from memory limitations, such as only up to 1 GB on the Linux cluster (which we need for the effective interactions), to the limited turnaround of ASCI Blue-Pacific. In general, Blue-Pacific is best suited for smaller-scale runs (20 nodes, less than 2 hours). It is often more efficient to use the restart capability two or three times than to use the 5-hour queue. In addition, Blue-Pacific is closed to dedicated runs on weekends, which further limits our turnaround time. Finally, the Blue-Pacific configuration limits our runs to approximately 350 MB/processor in MPI, which is inadequate for the full-scale project that we have outlined. Nonetheless, M&IC and ASCI resources have provided the single most important tool for us, and our project would not be possible without it. It has enabled us to achieve a significant improvement in nuclear theory that we will now be able to apply data to needs of SBSS. Indeed, these resources at LLNL provide a unique capability that has given us a competitive advantage.

The JEEP and Quantum Simulations Projects

François Gygi, Giulia Galli

This response combines the requests for institutional, unclassified computing from the JEEP project (Computation/François Gygi, PI) and the Quantum Simulations (QS) project (PAT/Giulia Galli, PI). The JEEP project focuses on code development for high-performance ab-initio simulations of matter at the microscopic level; the QS project focuses on applications of ab-initio codes to the study of a variety of systems, ranging from fluids under pressure, solid surfaces, and nanostructures to systems of biological interest.

1. **Scientific intent and scope**
   The main goals of these project are:

   - Predict physical and chemical properties of matter with great accuracy, using advanced quantum simulation techniques, e.g., state-of-the-art, first-principles molecular dynamics (FPMD) codes.

- Investigate properties of condensed systems (e.g., fluids and solids) that are not directly accessible to experiments.
- Interpret and complement experiment in close connection with experimentalists by taking advantage of high-performance computing.
- Establish LLNL as a strong player in the field of FPMD simulations worldwide.

The combination of expertise in high-performance software development, a growing group of expert code users in various LLNL directorates, and the availability of large supercomputers will potentially make the Laboratory a unique place where the most ambitious molecular simulations can be run

At present, ab-initio simulations are applied to three main areas:

- **Molecular fluids under pressure (HPF)**, e.g., hydrogen and fluid mixtures present in high-explosive mixtures. The major funding of this area has been LDRD for about 3 years. Programmatically relevant results obtained within LDRD projects have recently attracted ASCI funding which, since April 2001, constitutes half of HPF's funding. The HPF part of the project has important connections to DNT programs (hydrogen, high explosives) and so far is the part of the whole project which is most programmatically relevant. We are building a collaboration with the NIF Program regarding NIF early light (NEL) unclassified experiments; in particular we are devising simulations that may help guide some NEL experiments.
- **Computational biology (CB)**. The major funding is LDRD with some OBES contribution very recently. All of the project's computational biology activity has been done within the SI led by M. Colvin in BBRP.
- **Semiconductor nanostructures (SN)**. The major funding is LDRD with some OBES contribution. We are using this project to build collaboration with NAI programs in the area of biodetectors.

Research on new programming models and numerical algorithms is being pursued within the JEEP project in the Computation Directorate in order to extend the range of feasible FPMD simulations.

So far the scientific impact of the combined QS and JEEP projects has been remarkable. In collaboration with Colvin's SI, the project has led to 32 refereed papers published in major scientific journals during 1999–2001. Members of the QS group (QSG) in PAT and the JEEP project in Computation have had about 45 invited talks to international workshops or conferences and are actively engaged in several international collaborations with universities and other laboratories.

2. **Project funding source, size in FTEs, PI, and major contributors**
   The project started in May 1998. The PAT QSG was formed in January 2000 and the Computation JEEP project was formally started in FY99. At present, 14 scientists belong to the QSG (of whom 10 were hired after the project and the group were started): two FTEs, five Flexible Terms, two Lawrence Fellows and five postdocs. The JEEP project has two FTEs in the Computation Directorate.
3. **Code maturity level and history**
   JEEP—A C++ MPI/OpenMP FPMD code used in production in several research projects. It was initially developed in 1995 by F. Gygi at the Swiss Federal Institute of Technology in

Lausanne, Switzerland, and further developed by Gygi since 1998 at CASC. JEEP has been ported to the largest platforms available at LLNL, notably ASCI SST and ASCI White, on which the largest FPMD calculations to date were run. On these large problems, the code made efficient use of a large number of processors (up to 3840 CPUs on ASCI SST, 2640 on ASCI White). JEEP has a checkpoint restart capability. It also implements a signal caching feature based on the DPCS functions provided at LLNL, which allows for a graceful interruption (i.e., saving a restart file) of the simulation by an operator, by the batch queuing system, or by the user. Currently, JEEP has been ported to the Compaq Compass/Tera/TC2K clusters, ASCI Blue-Pacific, ASCI Frost/White and Linux i686 platforms. JEEP currently has 17 users within LLNL in the PAT, BBRP, and Computation Directorates.

MGMol—An experimental C++/MPI/OpenMP FPMD code based on a multigrid, finite-difference implementation. It is developed and maintained at CASC by J. L. Fattebert (JEEP project). This is a research code not yet used in production. It is being developed with the aim of reducing the computational cost of FPMD to linear (O(*N*)) scaling using a domain-decomposition technique. MGMol has been ported to IBM SP, TC2K, and Linux i686 platforms.

CASINO—A Quantum Monte Carlo Fortran 90 code for condensed systems and molecules. Parallelization: MPI. Development started at Cambridge Univ. (UK) and since February 2000 has been carried out at PAT by R. Hood and A. Williamson in collaboration with Cambridge.

QMCMol—A Quantum Monte Carlo Fortran 90 code for molecules. Parallelization: MPI. Development started at the Univ. of Illinois and since February 2000 has been carried out at PAT by J. Grossman in collaboration with North Carolina State Univ.

CUPID—A Path Integral Monte Carlo C/C++ code for hydrogen. Parallelization: MPI. Development started at the Univ. of Illinois and since July 2000 has been carried out at PAT by B. Militzer in collaboration with the Univ. of Illinois.

The five codes run on the IBM and Compaq and Linux clusters and are checkpoint restartable. Each of the five codes represents a state-of-the-art code in its own field and most of these codes are far superior to other main codes in the field in terms of robustness and performance. In particular, JEEP is the only C++ QMD parallel code available in the world, to our knowledge.

Visualization tools used within the project include Data explorer and AVS (installed on LC machines). Visualization software based on the VTK toolkit is also developed within the JEEP project. This allows us to produce high-resolution movies for Power Wall displays.

4. **CPU requirements for representative major calculations**
In the following we are reporting CPU requirements for some representative projects. Each of the calculations listed below led to one or more important publications; one of the calculations contributed to a DOE Award of Excellence.

   All of the calculations for the HPF part of the project were carried out on classified machines using CPU leftover. In order to consolidate these types of investigations and make them part of a steady effort, we do need access, and some dedicated access, to unclassified institutional computing.

We note that the CPU requirements cannot be trivially scaled down to obtain "some" results, instead of "many results." Simulations like shocks in deuterium, high-pressure water, nanostructures in solution, and all of the computational biology runs do need at least the CPU requirements listed below in order to give meaningful results, which have both a programmatic and a scientific impact.

High Pressure Fluids (HPF)
Shocked deuterium: wall-clock time for six shocks = 15 day dedicated run on 165 ASCI White nodes (2630 CPUs) (JEEP code; completed).

Water at ambient conditions and under pressure: wall-clock time for ~6 ps, 100 molecules, 15–30 GPa = 4.9 months on ~50 ASCI SKY nodes (200 CPUs); wall-clock time for ~14 ps, 54 molecules, 1–2 GPa = 2.8 months on ~50 ASCI SKY nodes (JEEP code, completed).

Runs similar (in terms of CPU cost and scientific impact of results) to those for water have been done for hydrogen fluoride.

Semiconductor Nanostructures (SN)
Optical gaps in Si nanocrystals (Quanto Monte Carlo codes): wall-clock time for ~10 clusters = ~1 month on 32 TC2K nodes (128 CPUs).

Si nanocrystals in water (JEEP code): wall-clock time for ~6 ps = ~5 months on ~32 ASCI Blue-Pacific nodes (128 CPUs).

Computational Biology (CB)
DMP/H_2O, Mg, Na/H_2O, H_2O dissociation: wall-clock time for ~(6 + 9 + 18) = 33 ps = ~5 months on ~50 ASCI Blue-Pacific nodes. (JEEP code; projects almost completed).

Nitrogen mustard and Mg/H_2O dissociation reaction: total estimated wall-clock time of these ongoing projects: 9 months on ~36 ASCI Blue-Pacific nodes (JEEP code; two-thirds of project completed).

5. **Major ongoing and planned projects**
In the following, major ongoing projects are listed for each area of focus. We note that all ongoing projects on unclassified machines have been slowed down (40–50%) compared with 1 year ago, greatly compromising the efficiency and overall impact of our work.

High Pressure Fluids (HPF)

Hydrogen
Shocked deuterium: on ASCI White; ASCI funded.
Metallization of hydrogen: on Blue-Pacific and SKY; ASCI (+LDRD) funded.

Low-Z Fluids
Water under pressure (≥1 Mbar): on SKY; LDRD funded
HF under pressure (≥1 Mbar): on SKY; ASCI + LDRD funded
Oxygen under pressure (20–100 Mbar): on TC2K and Blue-Pacific; LDRD funded.

Semiconductor Nanostructures (SN)
Optical gaps in Si nanocrystals: on Blue and TC2K. LDRD funded.
Surface reconstructions of Si nanocrystals: on Blue-Pacific, TC2K and NERSC (LBNL); LDRD funded.
Si nanocrystals in H_2O: on Blue-Pacific; LDRD funded.

Computational Biology (CB)
Nitrogen mustard and Mg/H_2O dissociation reaction: on Blue-Pacific and TC2K; LDRD funded.

In the following, major planned projects are listed for each area of focus. We note that with the current unclassified computer resources, no new project can be started, in spite of new hires joining the QSG and the JEEP project, without seriously compromising ongoing runs.

High Pressure Fluids (HPF)

Hydrogen:
Shocked low-Z fluid, e.g., He and H/He mixtures: on ASCI-White; ASCI funded.

Low-Z Fluids:
Liquid/solid interfaces in water under pressure: ⇒ *new LDRD*.

Semiconductor Nanostructures (SN)
Optical gaps for dots other than Si (e.g., CdSe) and arrays of dots: ⇒ *LDRD*.
Si nanocrystals with organic passivated surfaces in H2O: ⇒ *ER + new pre-SI.*

Computational Biology (CB)
Nitrogen mustard and Mg/H2O dissociation reaction: on Blue-Pacific and TC2K; LDRD funded.
Other model anticancer drugs in H_2O.

6. **Major scientific advances if access to new GF and TF facilities granted**
Major advances expected in the three areas of focus of the QS project with 2–8 TF (a) and 2 GF (b) computers are:

High Pressure Fluids (HPF)
Microscopic simulations of shock propagation in high-explosive mixtures (a). Microscopic simulations of shock propagation in heavy metals (b). These two sub-projects could have huge programmatic impact for DNT (both A and B Divisions) and for NIF.

Semiconductor Nanostructures (SN)
Optical properties of dots in solution on the fly (a). Virtual (real-time) atomic manipulation of semiconductor nanostructures in solution (b). These sub-projects could have a huge scientific visibility and eventually will lead to close collaborations with NAI.

Computational Biology (CB)
DNA base pairs in solution (a). Drug attack to DNA in solution (b). This sub-project would also have huge scientific visibility.

7. **Requirements for scientific challenge at 2–8 teraFLOP/s**

- Ratio of bytes of memory/flop (or how much memory per processor or how much total memory and how many total FLOP/s)?
  ~0.5–1.0 GB/GF, but not less than 2 GB/node.
- Ratio of bytes of disk/flop (or how much local and/or global disk would the run require)?
  Total global disk space 1.0 TB. Local space: 20 GB/node.

- What visualization requirements do you have? What visualization support would you require?
  Installation of VTK libraries on the platform.
- Tertiary storage requirements in TB?
  3–5 TB.
- When you would be ready to start a 16–18 month project as described in the introduction?
  Now.

8. **Requirements for scientific challenge at intermediate capacity level**
   - Ratio of bytes of memory/flop (or how much memory per processor or how much total memory and how many total FLOP/s)?
     ~0.5–1.0 GB/GF, but not less than 2 GB/node.
   - Ratio of bytes of disk/flop (or how much local and/or global disk would the run require)?
     Total global disk space 1.0 TB. Local space: 20 GB/node.
   - What visualization requirements do you have? What visualization support would you require?
     Installation of VTK libraries on the platform.
   - Tertiary storage requirements in TB?
     3–5 TB.
   - When you would be ready to start a project?:
     Now.
9. **Machine use pattern at 2–8 teraFLOP/s**
   After initial scaling runs, the machine would be filled with one large problem (128–256 nodes) and three to four smaller problems (32–64 nodes).
10. **Machine use pattern at intermediate capacity level**
    Several 32-node jobs would be run simultaneously. This is an efficient use of resources in some problems, notably in the calculation of the free-energy profiles of some biochemical reactions, which involve several independent molecular dynamics runs. The duration of the runs is typically limited by the batch system quota. If the machine (or part of it) is dedicated, job duration is limited by the machine's MTBF. (On Frost, several JEEP runs of >10 hours were completed). The signal-caching capabilities of JEEP allow for efficient use of multiple chained jobs with minimal user intervention.
11. **Memory vs compute trade-off**
    The JEEP code has compile-time options that can be used to take advantage of more memory if it is available. However, some operations have to be carried out on single nodes, and therefore impose a lower bound on the memory available on a node. On most platforms, JEEP currently makes use of up to 1.5 GB/node. A .75 GB limit on memory would impact the largest applications planned in this project, and would likely necessitate some software modifications.
12. **Comments on value of unclassified computing**
    All of the activity of the JEEP and quantum simulation project has been and is at present unclassified. Without the use of unclassified computing, none of our projects would have

been possible. We resorted to using leftover cycles on classified machines because the cycles on the unclassified computers were not sufficient for a successful completion of many of the QS projects. In the last 2 years, the success of our projects—all of them based on high-performance computing—allowed us to hire 11 new scientists. Some of the projects initially funded by LDRD are now funded by ASCI because of the programmatic relevance of the results produced. In the last 2 years the joint JEEP/QS project has performed two DNT award-winning investigations (1999 and 2001) and received a DOE Defense Programs Award of Excellence (2000).

PHENIX/HBT Project

Ron Soltz

1. **Scientific intent and scope**
   Our scientific goal is to measure and understand the particle emission region for relativistic heavy-ion collisions (i.e., collisions of Au nuclei at light speed). The relativistic heavy-ion program is motivated by the desire to detect and characterize the quantum chromodynamics phase transition (the melting of protons and neutrons into a plasma of their constituent quarks and gluons). Measuring the particle emission region is an integral part of the search for this phase transition and is an attempt to understand the nature of the strong nuclear force during conditions that prevailed in the first microsecond after the Big Bang.

   The project uses data from the recently commissioned Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory. To measure the particle emission region we use a technique of two-particle correlations, also known as HBT, after a similar technique used to measure stellar radii. This technique takes advantage of a statistical interference that produces an enhancement at low relative momentum between two identical particles (e.g., pions, kaons, and protons). The extent of this enhancement is inversely proportional to the spatial extent of the emission region, nominally given in units of 10^{-15} m. The correlation is typically fit to a standard gaussian in one to three dimensions.

   However, this technique is limited by its inability to measure nongaussian shapes and the difficulty of accounting for emission regions that are highly explosive in nature. Traditional techniques make multiple corrections to the data before fitting to a gaussian source. Different particle species are fit separately, and the expansion velocity is derived from a series of fits that depend heavily on a set of assumptions that are not easily tested.

   Our project is developing two techniques to overcome these limitations. The first is a source-imaging code, which seeks to measure the shape of the source directly. The second is a Coulomb wave fitting code (CRAFT) that allows one to fit a variety of source shapes and source dynamics directly to the data. Both techniques are new and computationally intensive.

   An additional need for computation is to be able to run large-scale Monte Carlo simulations in order perform the very first stages of data analysis. These simulations are required to separate detector-induced effects from the physics we are trying to understand. These simulations are required for all physics analysis, not just the two-particle measurements of the LLNL group. The simulations we run are for the PHENIX experiment, one of the two large experiments collecting data at RHIC.

2. **Project funding source, size in FTEs, PI, and major contributors**
The group consists of six physicists (three staff and three postdocs) in N Division. The five experimentalists, Ron Soltz (PI), Ed Hartouni, Stephen Johnson, Mike Heffner, and Jane Burward-Hoy, are members of the High Energy Physics group. The theorist and author of the source-imaging code is David Brown, a postdoc in the Nuclear Theory and Modeling group. The group is funded by LDRD and by a grant from the DOE Office of Science.

3. **Code maturity level and history**
There are three codes used in this endeavor, each at a different level of maturity. The PHENIX Monte Carlo code (known as PISA) is a 10,000 line code that was developed over several years prior to the start of data taking in 2000. The code was initially developed on Alpha Unix (then OSF), but later migrated to Intel Linux, which has become the cost-effective platform of choice for the high-energy physics community. In 2000, our group reported the code back to Compaq Tru64 Unix for a series of successful runs on Compass/Tera. Figure 1 below is a reproduction from a *Physical Review Letters* article on the emitted transverse energy distributions in Au–Au collisions at RHIC. The background data were generated by PISA running on the Tera cluster. These simulations were generated over the winter holidays in FY00, using approximately 100 GF for a period of 2–3 weeks. PISA uses a standard GEANT visualization package.

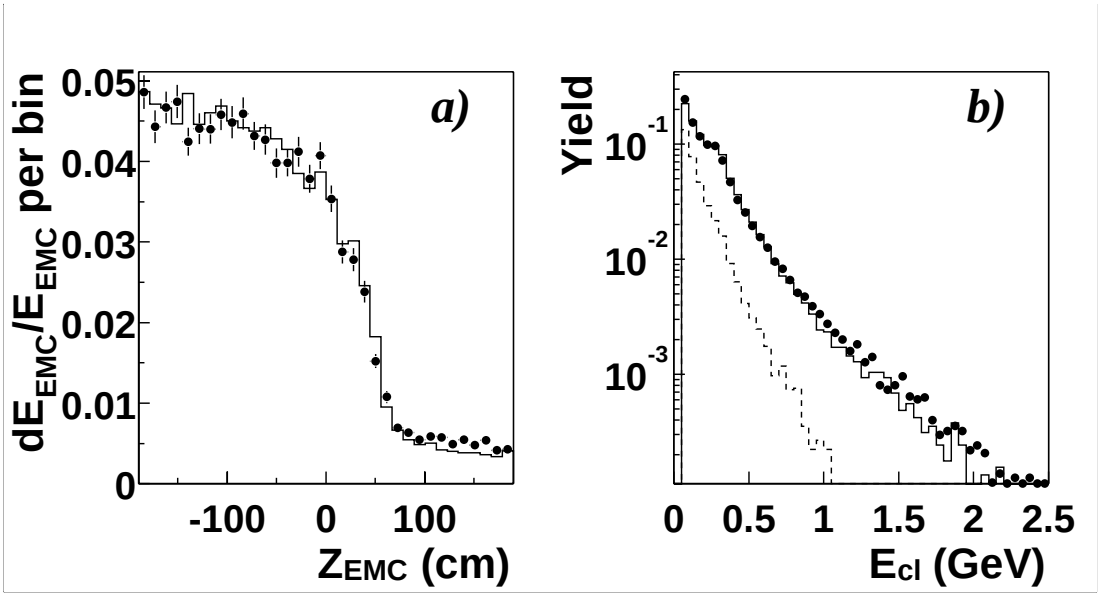

Figure 1. Simulation results (•) from Tera cluster, *Phys. Rev. Lett*. 87:052301 (2001).

Development of the imaging code began in 1997 at Michigan State University for the analysis of 1D correlation functions. Since coming to LLNL in 2000, David Brown has extended the code to accommodate analysis of correlations in three dimensions.

This new 3D capability required parallelizing the code, and it now uses OpenMP on the shared-memory Tera cluster machines. Since this parallel capability is still rather new, we are still assessing the efficiency and scalability of this code on various systems. It is known to run successfully on single-processor Linux and Tru64 PCs and Alphas. Furthermore, it has run successfully on multiprocessor Tru64 nodes. Visualization of the results is done through

standard histogram packages (PAW and ROOT). More sophisticated visualization techniques are highly desirable.

CRAFT is a collaborative effort between Mike Heffner and David Brown and has been under development for less than a year. Until we have more experience with it, estimates for computing needs are best taken from experience with the source-imaging code.

Of the three codes, only PISA is currently checkpoint restartable. Adding this capability to the other codes could be done within the times given in response to question 8.

5. **Anticipated scientific advance at 8–32 processor range**
   Access to this level of computing will enable both techniques to surpass what can be achieved from traditional analyses. For the source-imaging code, we should be able to answer questions pertaining to nongaussian source shapes and nonspherical structures in the resonance contribution to the source (the most extended region of the source), and detector effects obtained from the Monte Carlo simulations can be folded into the source inversion technique. In this way, we will remove many of the systematic errors made in the traditional data analyses of two-particle correlations and increase the accuracy of the source reconstructions.

   For the CRAFT code the primary breakthrough will be the ability to describe the pion source and expansion with a single parameterization. With additional code development, we should also be able to simultaneously fit multiple species of particles with a single parameterization. Both results will provide important constraints on current hydrodynamic models of heavy-ion collisions.

   With respect to running PISA, the scientific payoff is more difficult to predict—any physics analysis that utilizes Monte Carlo simulations stands to benefit from access to computing. The two-particle correlations are one example, but the results in Fig. 1 are from a different physics analysis, which seeks to understand the transverse energy distribution of particle produced in heavy-ion collisions. The 100 GF runs we were able to perform on Tera during the holidays would be too large to run at other times, but are well suited to this processor range.

6. **Connection with experiment**
   This project is closely coupled to experimental science. The Monte Carlo simulations are run specifically to calibrate detector effects. The source imaging and CRAFT codes are fit directly to the relativistic heavy-ion data sets.

8. **Requirements for scientific challenge at intermediate capacity level**
   A calculation with PISA typically requires 0.5 GB of memory and 1 GB of hard disk per 1 GF. Tertiary storage of approximately 2 TB is required. Calculation of simulated events could begin immediately.

   The source imaging code requires 4 GB memory and 4 GB disk per 1 GF. Calculations could begin in fall 2002. For CRAFT, we assume similar needs until shown otherwise. Large-scale calculations with CRAFT will likely begin sometime between fall 2002 and spring 2003.

10. **Machine use pattern at intermediate capacity level**
    Generating simulation data with PISA occurs in batches of 10^5 events, each requiring 100 GF (32 processors) for a period of 1 week. We estimate needing to perform approximately 5 to 10 such runs per year.

The source imaging analysis of a single data set can take several hundred CPU hours on the current LLNL Tera cluster. Initial runs would focus on a series of systematic studies designed to improve code performance to attain the highest quality source reconstruction. These studies are as follows:

- **Convergence in lmax tests**. We expand our source in spherical harmonics. We need to practice angular scale cut-off in the reconstructions.
- **Knot setting algorithm tests**. Our code has the ability to adjust the resolution in accordance with the experimental data and the kernel of the integral equation.
- **Compare resolution of types of particles**. Due to the differences in the interactions between particle pairs, we need to assess which ones give access to which features in a heavy-ion reaction.

Including initial systematic studies, and a series of ten data sets, we would need access to 25 GF (8 processors) for a total integrated time of 12 weeks. At this time, we estimate computing requirements for CRAFT to be comparable.

11. **Memory vs compute trade-off**
    This would not be a problem for PISA, nor for CRAFT. However, the source-imaging code needs the larger memory for inverting large matrices.

12. **Comments on value of unclassified computing**
    M&IC computing has been an essential ingredient in building and maintaining the Heavy Ion Physics group. The DOE Office of Science grant would not have been possible without access to LC resources for running simulations and analyzing data.

    Furthermore, the group has been directly involved in recruitment of four physicists during the last 2 years. While continuing to do basic research in the group, most group members have make significant contributions to LLNL programs in advanced radiography and nuclear reaction modeling.

    Thus far, all computing needs for our group have been met adequately, and benefit to the Lab is both significant and visible.

Microscopic Origins of Dynamic Fracture Project (MD3D)

James Belak, Robert Rudd

1. **Scientific intent and scope**
   The scope of the Microscopic Origins of Dynamic Fracture project is to model, through direct numerical simulation (molecular dynamics), the nucleation and growth of voids in ductile metals during dynamic fracture. The ultimate goal is a model of these processes suitable for continuum hydrocode simulations. Shock waves in solids produce many interesting effects. One such effect, known as spallation dynamic fracture, occurs when a shock wave reflects from a free surface and produces extreme tension inside the solid. When this tension exceeds the internal rupture strength, the solid fails by nucleating voids in ductile metals and cracks in brittle solids. The growing voids quickly link together to fracture the solid. The origin of the voids is intimately tied to the microscopic structure of the solid. Weak points in the microstructure, such as inclusions and grain boundary junctions, are the locations from which the voids are likely to nucleate. This project represents the first time in

which the dislocation mechanisms by which voids grow have been quantified. The development of constitutive models that are sensitive to material microstructure will enable the assessment of changes due to aging and remanufacture.

2. **Project funding source, size in FTEs, PI, and major contributors**
In FY02, the Microscopic Origins of Dynamic Fracture project will be funded by ASCI and LDRD (01-ERD-022). The total number of FTEs is 3.0. The PIs are James Belak and Robert Rudd. In addition, a postdoc (Eira Seppala) is dedicated to the project. The LDRD (01-ERD-022) is primarily an experimental project with several FTEs spread across PAT, CMS, and ENG. It includes a collaboration with M. Fivel (CNRS in Grenoble, France), which is facilitated through the OCF. J. Belak is the lead PI on the LDRD project.

3. **Code maturity level and history**
All of our simulations are performed using a molecular dynamics research code commonly referred to as MD3D. MD3D is a robust molecular dynamics simulation code that has undergone continuous development and usage at LLNL for the past 12 years. Is has been ported to the BBN TC2000, Cray T3D, and ASCI Blue-Pacific and is currently being ported to the TC2K and Linux clusters. MD3D has checkpoint restart capability.

   During FY01, analysis routines were added to MD3D to identify atoms participating in dislocations on the fly during a simulation. This reduced data set (a few percent) is printed often and visualized on workstations using AVS (see figure).

4. **Anticipated scientific advance at 2–8 teraFLOP/s**
We are currently running initial void sizes up to 10 nm at strain rates of 10^8–10^9 s^{-1}. The experimental strain-rates of interest are 10^6–10^7 s^{-1}. Further, though we observe many interesting dislocation processes for 10-nm voids, we expect significant size effects and need to quantify the dislocation processes for larger voids.

   The dedicated access to a 2–8 teraFLOP/s facility would enable us to simulation the behavior of a 100-nm void under tension (about 1 billion atoms) at rates of 10^8–10^9 s^{-1}. Of equal interest is the simulation of a 10-nm void at rates of 10^6–10^7 s^{-1}. This is possible by using the parallelism to simulate far fewer atoms per processor. Our previous studies suggest we will need at least 1000 atoms per processor for reasonable scalability.

   A major focus of the Dynamics of Metals program at LLNL is the bcc metal tantalum. The interatomic force models (MGPT and BOP) required to represent bonding in tantalum are 10 to 100 times more computational intensive than the simpler EAM force models we use to represent copper. The 2–8 teraFLOP/s

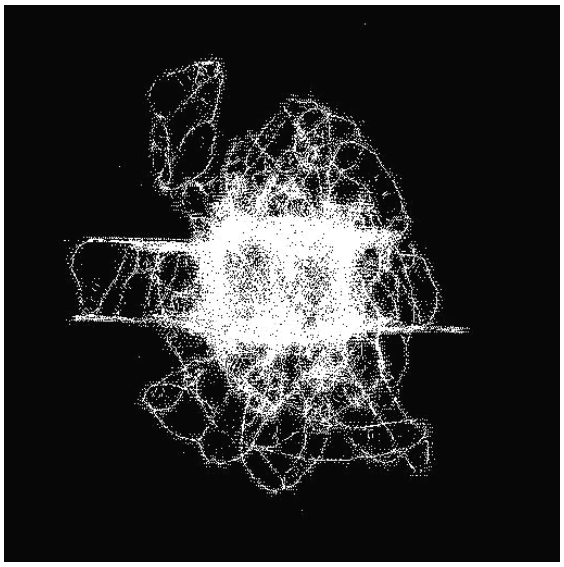

On-the-fly visualization while MD3D is running identifies atoms participating in the simulation.

capacity platform would enable us to simulate the growth of a 10-nm void in tantalum at strain rates of 10^8–10^9 s^{-1}.

5. **Anticipated scientific advance at 8–32 processor range**
   We are currently undertaking a series of simulations with this type of capability in mind. Our preliminary studies show that the dislocation mechanisms of void growth are sensitive to the stress state, i.e., relative amount of shear stress to isotropic tension; this is known as triaxiality.

   a. We plan to map out the dislocation mechanisms of void growth as a function of triaxiality.
   b. We have assumed the initial voids to be perfectly spherical. We plan to map out the dislocation nucleation as a function of void roughness. The rough regions concentrate the applied stress.

6. **Connection with experiment**
   One of us (Belak) is the lead PI on the LDRD project Microstructure Origins of Dynamic Fracture in Ductile Metals (01-ERD-022). This is primarily an experimental project and was proposed to generate experimental data to validate and focus our simulations. Our need for computational resources is not contingent upon the experiments. During FY02, void nucleation from 10-nm inclusions in copper will be experimentally studied to allow direct overlap with our molecular dynamics modeling.

   This LDRD project, in collaboration with three other LDRD-ERD projects (Becker, Campbell, and Kalantar), is currently generating an LDRD-SI proposal in the area of Dynamic Failure and Fracture in which we expect a significant need for computational resources.

7. **Requirements for scientific challenge at 2–8 teraFLOP/s**
   Our requirements are strongly driven by the fact that we are much more CPU-limited than memory-limited. Our domain decomposition algorithms scale reasonably well as long as we use at least 1000 atoms per processor. The memory requirement is about 100 bytes per atom. The 100-nm void simulation mentioned above requires 100 GB across the whole system using 1 billion atoms. The 2–8 teraFLOP/s performance would enable a strain rate of no slower than 10^8 s^{-1} in one CPU month (a total of 5×10^{18} floating point operations). This corresponds to about 100,000 molecular dynamics time steps or about 0.5 ns of simulation time.

   Our local disk requirements are driven by the need to save visualization and restart files every so often (the frequency is selected at run time). For the run described above, a snapshot of the system requires about 100 GB of disk and for an entire run we might save 10 such files for a storage requirement of about 1 TB. While the dislocation analysis files require only a few percent of this, we expect to generate a few thousand during a run. Thus our total storage needs could approach 10 TB.

   Like many other projects, visualization for such a large problem is uncharted territory for us. To date, we have been using AVS on our desktop workstations. As such, the system sizes we are proposing now are taxing the response and memory of AVS on these systems. A remote visualization application, such as a dedicated high-end graphics facility, would be of use to us.

We estimate that we would be ready to begin a large simulation in summer 2002.

8. **Requirements for scientific challenge at intermediate capacity level**
   Our physics needs for capacity runs are outlined above, and we are currently proceeding with this in mind. These systems contain 1–10 million atoms and require no more than about 1 GB of memory. Each run would require up to 1 CPU week of continuous computer time and would generate about 10 GB of data; assuming many such runs, the storage requirement is about the same as the big run, i.e., 1–10 TB. We find that AVS on the local workstations is adequate to visualize these runs.

   We estimate that we would be ready to begin a large simulation in spring 2002.

9. **Machine use pattern at 2–8 teraFLOP/s**
   Primarily, we propose to use the capability resource to extend our simulations into qualitatively new regimes of size scale, time scale, and potential model. Though the basic algorithms for parallel domain decomposition of our molecular dynamics code are well validated and a large scale simulation with one million atoms per processor on 1000 processors should scale very well, we anticipate the need to address scaling due to architecture and machine dependent problems. Thus we expect the need for brief periods of debugging (a few days) in between longer periods (a few weeks) of dedicated large-scale production simulations.

10. **Machine use pattern at intermediate capacity level**
    Our scientific exploration of dynamic fracture requires the exploration of a large parameter space. The 8–32 processor capacity resource does represent a qualitative capability over the 2–4 processor platforms of current desktop workstations. We anticipate using the resource for 32 processor runs. Considering 4 CPU months on these 32 processors, we would perform approximately fifty 1-CPU day (on 32 processors) and approximately seven 1-CPU week runs.

11. **Memory vs compute trade-off**
    Yes! We are CPU limited.

12. **Comments on value of unclassified computing**
    High-performance computing has played several important roles in achieving the success of LLNL's unclassified programs. The availability of a large capacity of computational power spread across many projects has enabled sustained scientific progress. These resources have supported the efforts of many scientists including new hires. Our ASCI project has benefited most from the availability of this capacity, and the kind of projects we have undertaken would not be possible without it. The goals of the projects have been designed to make the most of the available computational resources, and they rely on a continued growth. For example, the computational investigation of more complicated materials will require significantly more computational power. The bcc metals are only now becoming accessible by direct molecular dynamics simulations.

    On the other hand, the fact that LLNL can muster a virtually unequaled computational capability has added tremendously to the prestige of the Lab. In a few cases, this capability has permitted qualitatively new kinds of simulations. These simulations are typically not the final word on a subject, but they open the door to new investigations. In other cases they allow a validation of extrapolations that are implicit in many of the calculations.

It is our hope that LLNL's high-performance computing will continue to grow in both capacity and capability on the unclassified side.

Evolving pF3d to Develop a Predictive Laser–Plasma Interaction Modeling Capability

C. H. Still, E. A. Williams, R. L. Berger, A. B. Langdon, L. M. Divol

1. **Scientific intent and scope**
   The pF3d code is one of the principal production tools for the ICF Hohlraum Energetics Work Breakdown Statement (WBS-1), part of the theoretical component to the NIF Ignition Plan. The ability to model laser–plasma interactions (LPI), and ultimately to develop a predictive capability for LPI, is essential in modeling the energetics within a NIF ignition hohlraum, which is key to ensuring successful ignition on NIF. The code pF3d forms the basis for a LPI predictive capability.

2. **Project funding source, size in FTEs, PI, and major contributors**
   The pF3d project is funded by WBS-1 through the NIF Directorate. The main contributors are Dick Berger, Laurent Divol, Bruce Langdon, Bert Still, and Ed Williams. For the purpose of this computing proposal, the PIs are Bert Still and Ed Williams.

3. **Code maturity level and history**
   The code pF3d is based on a prior Fortran version (F3D), which played a large role leading to the successful Key Decision-1 to go ahead with NIF construction. The first parallel version of pF3d (called yf3d) was initially developed partly using ASCI money as a seed, and delivered a simulation result shown during the acceptance-period testing of the ASCI Blue-Pacific Initial Delivery machine. Since then, yf3d has been part of the simulated workload (SWL) acceptance test suite. The latest version of pF3d has all of the physics models from its Fortran predecessor as well as many new capabilities not present in the original F3D code.

   Written in ANSI-C and built upon Yorick [1], the latest version of pF3d runs on a wide variety of Unix-based systems from standalone workstations to massively parallel processors. The code produces dump files in the PDB format, and can be restarted from a save set. Production runs are typically made using 6- or 12-hour time slices on parallel machines at LLNL, and involve several restarts. Yorick provides graphics output during a run from a set of prescribed diagnostics, and additional visualization diagnostics can be made from a save set during serial or parallel post-processing.

4. **Anticipated scientific advance at 2–8 teraFLOP/s**
   The longer plasma scale lengths and higher power present on the NIF will enhance laser–plasma instabilities [2]. Definitive modeling of a single NIF inner-cone beam interacting with the relevant size plasma (from the laser entrance hole to the wall inside a NIF hohlraum) is an enormous calculation on a 64 gigazone mesh, requiring roughly 10^{19} flops. While this problem is just beginning to become feasible on the world's largest computers, a great deal of understanding can be developed via entire beam simulations of other laser facilities. We have performed a few large simulations using a Nova beam "letterbox" (an entire Nova beam in one transverse dimension and a subsection in the other transverse dimension) over the entire propagation distance in a Nova hohlraum. (Recently, we performed such a calculation on 80 nodes of the ASCI Blue SKY S-sector for just over 2

weeks.) With access to the resource described (2–8 Tflops for 1–4 months), we could perform the first simulation on the entire Nova beam volume, and perhaps even the volume of a NIF outer beam.

The ability to perform a series of runs of full Nova size beams would allow us to validate the code directly with the existing experimental database. In addition we would compare the results of these full beam-simulations to the smaller corresponding letterbox calculations and determine just how well, and when, we can extrapolate. These "capability" calculations would allow us to make more meaningful use of "capacity" resources.

5. **Anticipated scientific advance at 8–32 processor range**
In developing a predictive capability for LPI, there are some physical phenomena that must be better understood. One is the electron heat conduction. By its nature, electron heat conduction in a laser-driven hohlraum is nonlocal and nonlinear, making a computational model difficult. Another aspect is the saturation mechanisms for both stimulated Brillouin backscatter (SBS) and stimulated Raman backscatter (SRS). SBS is generated by the interaction of the laser pump with an ion-acoustic wave, while SRS is generated by the interaction with an electron plasma (Langmuir) wave. Part of understanding these saturation mechanisms involves parameter studies with pF3d, in which we vary the frequency of the light, the power in the beam, beam models (with and without smoothing techniques), and plasma conditions. Many of these studies can be carried out in 2D using the "midsize" computer system described.

   For example, we can directly compare pF3d simulation results with results from the particle-in-cell (PIC) code ZOHAR (fully kinetic, relativistic electromagnetic in $2^1/_2$D)—or its relative Bzohar (mobile ions, but Boltzmann fluid electrons), or its daughter code Z3 (3D MPP fully kinetic code)—for small problems (e.g., a single speckle) to validate the saturation models. These comparisons will require many "small" runs in 1D and 2D and some in 3D (with PIC codes and pF3d) to do parameter studies in a single speckle. Then larger pF3d simulations would be used to study the impact of those models on longer plasmas with random phase plate (RPP) beams (when many speckles are present).

6. **Connection with experiment**
pF3d simulates laser–plasma interaction experiments. Input conditions are generally taken from laser experimental data, or LASNEX models thereof, and used to produce simulated experimental results. Analyzing these results, in turn, improves physical understanding, and can be used to design new experiments. Many Nova experiments have already been modeled, and current modeling of laser experiments on Nova, Omega (at the University of Rochester), Helen (in Britain), and Lil (in France) is being used to develop understanding for NIF experimental conditions.

7. **Requirements for scientific challenge at 2–8 teraFLOP/s**
For a simulation including SBS and SRS models, we estimate 10,500 FLOPs per cell, per time step, and total memory use of approximately 220 bytes/cell. Because pF3d operates on a 3D Cartesian mesh, simulations are typically memory limited. Spatial resolution is on the order of a laser wavelength (0.351 μm typically), which leads to 3D meshes with 5×10^8 cells and up for "large" problems. Such a problem would use 110 GB of memory and require 5.3 TFLOPs per time step. Each time step is approximately 0.06 ps, and a simulation of 20–40 ps is usually performed, meaning 300–600 time steps. A restart dump consists of a

collection of problem state files (one per processor), and we typically choose to make 10 to 20 or more dump sets during the course of the simulation, depending on the expected reliability of the machine. This would anticipate a requirement of 1.1–2.2 TB of disk storage for restarts, divided among the local disks. Graphical output varies by the number of diagnostics employed during a run. A recent simulation of A Nova beam sized problem ($1024 \times 512 \times 640$ cells) created 391 MB of graphics output and 22 GB of run history data; these output files would be stored on global disk. If automatic archiving is available, the code will make use of it to reduce the local disk footprint required. If automatic archiving is not available but global disk is plentiful, the code could "archive" previous save sets to the global disk to relieve local disk usage.

Once machine access is granted, we would be ready to start almost immediately. The only known drawback at present is the need to force SIGFPE to be delivered when Yorick is compiled with the icc compiler in P7 mode.

8. **Requirements for scientific challenge at intermediate capacity level**
See the response to question 7. The mesh size for a "medium" problem (large 2D or small 3D) would range from 10^6 to 10^8 cells. This in turn would imply 10.5 GFLOPs up to 1 TFLOPs and 220 MB–22 GB total memory. Corresponding disk footprint would be 12–500 GB, or less with automatic archiving.

9. **Machine use pattern at 2–8 teraFLOP/s**
Our use would focus on 2D and 3D parameter studies. Pf3d is our primary tool to bridge between (largely theoretical) studies of microphysics and macroscale LPI experiments. As we improve our microphysical models—for instance, how do the ion-acoustic waves involved in SBS saturate?—we incorporate distillations of them in pF3d and then see what range of parameter space they can model experimentally. This requires relatively large numbers of "capacity" calculations (in 2D) to make progress. However, to connect the 2D calculations back to the experimental results requires 3D "capability" calculations (as described in the response to question 4 above).

   For the next two years (prior to NIF early light), we will be involved with LPI experiments on smaller laser facilities (Trident/Helen, Janus), where our experimentalists can get shots. A number of experimental results have already been obtained from the Trident/Helen laser. On these smaller lasers, we can simulate the entire 3D LPI problem over time scales of interest, and then directly compare to experimental results.

10. **Machine use pattern at intermediate capacity level**
Our use for this class of computing resource would focus on 2D parameter studies (see response to question 9 above). In order to connect back with experimental results, we would still need to perform code benchmarking in 3D, but 2D parameter studies would enable us to improve our microphysical models.

    As mentioned above, we anticipate modeling LPI experiments on smaller laser facilities (e.g., Trident/Helen) for the next 2 years. A Trident unsmoothed laser beam has diameter ~80 μm at best focus, and a plasma interaction length of 800 μm. Since the laser light has a wavelength of 0.53 μm, we can perform the entire 3D simulation (26.2 million cells using 5.8 GB of memory and 165.2 TF of computation for the simulation) on the described system. These simulations can be directly compared to experiment, as well as providing a basis for evaluating 2D and smaller 3D calculations.

11. **Memory vs compute trade-off**
    Reducing the memory per processor would likely result in a less useful machine for us because our pF3d simulations tend to be memory limited. Halving the memory per processor would require us to use twice as many processors, thereby increasing communication costs relative to computation, resulting in less efficient processor use.

12. **Comments on value of unclassified computing**
    For the past several years, the unclassified computing resources provided by M&IC and its predecessors have been essential in meeting our programmatic goals. PF3d modeling of Nova experiments (hohlraums and CO_2 gasbags) accomplished using M&IC resources has contributed to a number of publications ([2]–[6] to name a few). Furthermore, maintaining a cutting-edge computing environment in the unclassified arena will remain essential to us in continuing to meet our programmatic goals for several reasons. Interaction with experimentalists occurs on the unclassified networks, where their data and analysis reside, reporting physics results to our sponsors and the physics community is greatly facilitated by having our simulation results on the unclassified systems, and collaborations with our off-site scientific colleagues can only occur on the unclassified network. Having sufficient available computing cycles on the unclassified network has been, and will continue to be, of enormous value to the ICF plasma physics effort.

[1] D. Munro, "Yorick: An interpreted programming language for post-processing or scientific calculation," ftp://ftp-icf.llnl.gov/pub/Yorick/doc/index.html.

[2] C. H. Still, R. L. Berger, A. B. Langdon, D. E. Hinkel, L. J. Suter, and E. A. Williams, "Filamentation and Forward Brillouin Scatter of Entire Smoothed and Aberrated Laser Beams," *Phys. Plasmas* **7**, 2023 (2000). (This paper was presented as an invited talk to the American Physical Society/Division of Plasma Physics Meeting, November 1999.)

[3] R. L. Berger, C. H. Still, E. A. Williams, and A. B. Langdon, "On the Dominant and Subdominant Behavior of Stimulated Raman and Brillouin Scattering Driven by Nonuniform Laser Beams," *Phys. Plasmas* **5**, 4337 (1998).

[4] R. L. Berger, E. Lefebvre, A. B. Langdon, et al., "Stimulated Raman and Brillouin Scattering of Polarization Smoothed and Temporally Smoothed Laser Beams," *Phys. Plasmas* **6**, 1043 (1999).

[5] S. H. Glenzer, R. L. Berger, L. M. Divol, et al., "Reduction of Stimulated Scattering Losses from Hohlraum Plasmas with Laser Beam Smoothing," *Phys. Plasmas* **8**, 1692 (2001).

[6] S. H. Glenzer, L. M. Divol, R. L. Berger, et al., "Thompson Scattering Measurements of Saturated Ion Waves in Laser Fusion Plasmas," *Phys. Rev. Lett*. **86**, 2565 (2001).

Gas Distortion Characterization for the National Ignition Facility (NIF)

Steven Sutton

1. **Scientific intent and scope**
   One of the shot-rate limiting factors in the NIF will be the recovery of thermally driven gas distortions that result from temperature differences within the laser chain. Detailed time- and space-resolved simulations of the flow field in trapped gas volumes through which the laser

propagates are necessary to fully quantify the shot-rate capability of the facility. Fully quantifying the shot rate capacity and identifying ways to improve the shot rate will have significant impact on NIF operations.

Currently, these calculations are run on high-end workstations and take from several weeks to months for a single simulation. To meet NIF's modeling needs for the laser commissioning effort, these calculations need to be performed in much shorter clock times. To address these needs we are transitioning to parallel tools and platforms.

2. **Project funding source, size in FTEs, PI, and major contributors**
   Steve Sutton is the PI for NIF gas distortion analysis. Resources for this effort are coming from three LLNL sources. Code application is funded out of the NIF budget at the rate of 1.3 FTEs. ALE3D incompressible flow development is headed by Rose McCallen, with support from Tim Dunn and Charles Tong (ASCI). Funding for ALE3D incompressible flow development is approximately 1 FTE from ASCI and 0.5 FTE from DOE. Two of the codes being applied are commercial, so any development in these instances is being funded by those organizations.

3. **Code maturity level and history**
   Three codes are being exercised and being considered for down-selection. They are:

   - FIDAP—This code is marketed by Fluent, Inc., and is an industry standard finite-element incompressible flow code. This code is the current workhorse for our effort. The code is parallel in a shared memory environment. Discussions are under way with Fluent to determine their interest in initiating a CRADA to move to distributed-memory platforms.
   - ALE3D—The finite-element incompressible flow model being added to ALE3D is partially implemented. Current efforts focus on improving time-integration algorithms and execution efficiency. This code is fully parallelized The largest number of processors that have been used to date on the ASCI Blue-Pacific machine is 128. Tests of parallel efficiency are under way.
   - LS-DYNA—This code is marketed by Livermore Software Technology, Inc. A finite-element incompressible flow model has been added to the code and is undergoing testing at LLNL. Serial execution tests are under way, with parallel execution tests to follow shortly. This code is fully parallel in a shared memory environment. The structural mechanics portions have been successfully run on the ASCI Blue-Pacific machine.

4. **Anticipated scientific advance at 2–8 teraFLOP/s**
   The shot rate of NIF is going to be largely dictated by optical distortions from thermally driven convection currents. In addition, other distributed heat sources in the system could lead to significantly increased focal spot size. It is important for project success that a computational capability be developed and exercised to allow simulations of convection currents to be performed in a day (they currently take several weeks). This type of capability would allow us to perform calculations during NIF commissioning to address issues as they occur.

5. **Anticipated scientific advance at 8–32 processor range**
   This type of capability would allow us to run simulations for smaller geometry portions of the NIF system in the execution clock times required. For larger geometry regions, this

would still be of use but would reduce the impact that analysis would have on commissioning decisions.

6. **Connection with experiment**
   The NIF will be the major stockpile stewardship experimental facility at LLNL. This project will partially be used to address sources of and remedies for any anomalous behavior observed in the system performance. Thus, there will be a synergistic relationship between the experimental and modeling activities.

7. **Requirements for scientific challenge at 2–8 teraFLOP/s**
   In general, our requirement is driven by CPU requirements. To stabilize the flow fields generally takes ½ to 1 hour of simulation time. Thus, we need access to a multiteraFLOP machine. The solvers employed use matrix bandwith minimizing approaches to reduce storage requirements. Specific memory and CPU requirements will not be established until preliminary parallel execution studies are completed. We will begin timing studies within the next month.

8. **Requirements for scientific challenge at intermediate capacity level**
   See the response to question 7.

9. **Machine use pattern at 2–8 teraFLOP/s**
   We believe that access to such a machine would allow us to meet the NIF initial commissioning requirements. The effort can be broken into three categories:

   (1) Mesh sensitivity studies, which have been restricted to date because of the machines available. It is imperative that these be performed prior to analyzing the larger structures on NIF.

   (2) Development of a distortion database for systematic thermal conditions. The information in this database will be used in NIF propagation simulations.

   (3) Analysis of nonsystematic thermal disturbances that are identified during commissioning. This would include the effects of spot heat sources in the transport and switchyard beamtube areas.

   During the NIF commissioning effort, optical distortion measurements will be made for one or more of the laser beamlines. These will be analyzed and compared to the numerical predictions that result from this analysis effort.

10. **Machine use pattern at intermediate capacity level**
    Provided there was adequate access to several of these processor blocks, we could make progress by running several problems at the same time. Thus, this type of resource could be of benefit to our modeling efforts.

11. **Memory vs compute trade-off**
    Memory requirements are in the range of 2000 bytes/element. Thus, a 0.75 GB per processor limit should not be constraining for the problems we are running. Our major need is a higher FLOP rate.

12. **Comments on value of unclassified computing**
    We are only beginning the process of using the unclassified parallel computing resources at the Laboratory. As the next generation of computer tools becomes available, we need access to such resources in order to meet distortion characterization requirements for the NIF

commissioning process. Since the results of our analysis will be part of the NIF distortion information database, the analysis must be performed on unclassified platforms so that the information can be readily transferred.

Accurate Modeling of Electromagnetic Effects in High-Frequency Integrated Circuits

Robert M. Sharpe

1. **Scientific intent and scope**
   Accurate and timely modeling of electromagnetic phenomena is important to many DOE programs in areas as diverse as high-energy accelerators, remote sensing and nondestructive evaluation, EMI/EMC effects, and lasers and photonics. Currently, we have a strong interest in applying our expertise in computational electromagnetics towards the critical problem of electromagnetics effects in high-frequency integrated circuits (ICs). At low frequencies, ICs can be adequately modeled and designed using traditional circuit theory, which neglects distributed electromagnetic effects. At high frequencies, time-varying electromagnetic fields are distributed throughout the entire circuit, effectively coupling every component of the circuit to every other. This coupling of components via electromagnetic fields is already an important factor in RF and microwave ICs and is a potential show-stopper in the development of next-generation digital ICs with clock rates approaching 100 GHz. A computational study of electromagnetic effects within ICs would lead to physics insight and design rules that would positively impact the entire semiconductor industry.

2. **Project funding source, size in FTEs, PI, and major contributors**
   Rob Sharpe of DSED is the PI of the EIGER code suite. This team consists of three FTEs at LLNL, and six professors (and their students) from several different universities. Funding is currently provided by LDRD ER, Engineering Tech Base development, DARPA funding, and support from the DOD's HPCMO.

   Daniel White of CASC is the PI of the EMSolve code development team. This team consists of three FTEs and two students, currently funded through Lab-wide LDRD, DARPA, and UC Davis.

3. **Code maturity level and history**
   The EIGER code suite is a mature production code that was developed beginning in 1995. The original development team was LLNL, SNL, the University of Houston, and the Navy's SPAWAR System Center. The DOD's HPCMO has invested in the parallel aspects of the code since 1998. The code is a hybrid finite element/boundary element code that uses an implicit formulation to solve Maxwell's equations. There are two main compute phases in the code, filling a system matrix (either dense, sparse, or a combination) and solving the resulting system matrix. A restart feature is available for the fill algorithm. A restart is not currently available for all the linear solvers that can be used (both direct and iterative solvers can be used). The code has run at nearly 80% efficiency for 200 processors on a DOD SP2. The code runs on IBM, SGI, Compaq, and Linux platforms. Various visualization tools are available for displaying results.

   The EMSolve code is a parallel 3D unstructured grid finite element program that solves Maxwell's equations and related partial differential equations. EMSolve is a research code,

not a production code, and has no users outside of LLNL. EMSolve takes advantage of existing software components such as METIS from the University of Minnesota, PETSc from ANL, and Hypre from CASC. The efficiency and scalability of EMSolve is directly attributed to the efficiency and scalability of these underlying components. EMSolve relies on existing tools for mesh generation and visualization of results.

4. **Anticipated scientific advance at 2–8 teraFLOP/s**
   Since it is extremely difficult, if not impossible, to measure the electromagnetic fields within an IC, this project would have a tremendous impact on our understanding of electromagnetic effects inside of high frequency ICs. We will use the explicit EMSolve tools to discuss these aspects.

   Using EMSolve, a simulation of an IC with 10^{10} unknowns ($10^4 \times 10^4 \times 10^2$ mesh), and assuming 10^2 FLOPs per unknown per time step and 10^4 time steps results in 10^4 teraFLOPs. So a single run would require about 3 hours on a 1 teraFLOP/s machine. Increasing the size of the circuit by a factor of 10 would yield a 3000-hour simulation. So a 2–8 teraFLOP/s system would enable extremely detailed simulation of entire ICs.

5. **Anticipated scientific advance at 8–32 processor range**
   Based upon the above arguments, a more modest system could not be used to simulate an entire digital IC. However, much useful science could be gained from such a system. For example, analysis of electromagnetic coupling of adjacent microstrip interconnects could be studied in great detail, or electromagnetic coupling of adjacent inductors in a microwave low-noise amplifier. This type of research may not seem "sexy," but it is extremely interesting to the electronics designers in the DOE, DOD, and industry.

6. **Connection with experiment**
   The DOD is currently planning a sequence of experiments on example mixed signal circuits to which we can compare our computed results.

7. **Requirements for scientific challenge at 2–8 teraFLOP/s**
   EMSolve, being an unstructured grid finite element code, has a higher memory footprint than comparable finite difference codes. Hence 1 GB per processor, or more, is preferred.

   Not much temp disk space is required, just enough room for the input mesh, the restart files, the results. Bytes of disk = 100 × bytes of RAM should be sufficient.

   EMSolve does not perform its own visualization. Programs such as B Division's VizIt should be adequate.

   Could start 1 year from today.

8. **Requirements for scientific challenge at intermediate capacity level**
   Same as the response to question 7.

9. **Machine use pattern at 2–8 teraFLOP/s**
   A machine of this caliber is required to perform the type of simulation described in our response to question 4 above. However the problem set-up is significant, and post-simulation data reduction and analysis is also significant, so we would likely perform only a few such runs during a 6–18 month time frame. We would choose the problem parameters (mesh size, time step, etc.) such that each run would complete in 3 weeks or less. We would perform many smaller, shorter validation runs prior to each large science run.

10. **Machine use pattern at intermediate capacity level**
    Our science goal would be that as described in our response to question 5 above. Our plan would be similar to that described in our response to question 9 above, except that smaller runs are easier to set up and post-process, so we would probably get more runs in our allotted time.

11. **Memory vs compute trade-off**
    No. Go for the RAM! Modern object-oriented programs use a lot of memory, and sophisticated numerics such as adaptive mesh refinement and implicit solvers use a lot of memory per unknown.

12. **Comments on value of unclassified computing**
    I am a big user of the Compass/Tera cluster. Not only is the hardware impressive, but the set of installed software tools is equally impressive, and the system administrators and the LC-Hotline are a great support. I use these machines for all of my LDRD research, and this research has led to a $3.6 million DARPA contract. I am also supporting two Ph.D. students who use these machines for their research. I do all of my development/testing/debugging on Compass/Tera rather than on my desktop. My only criticism is the network—interactive work (such as running emacs or TotalView remotely) is painful!

High-Resolution, Nondestructive Characterization of Objects and Materials

Pat Roberson

It is important to note up front that our greatest need is a 2–8 teraFLOP/s resource for both classified and unclassified applications. Our immediate and highest priority would be a classified system. The highest capacity system that is described in the questionnaire is the type of resource that we desperately need for processing problems that we currently have.

1. **Scientific intent and scope**
   We are developing nondestructive characterization systems and techniques that require significant computational power. These needs exist due to the complexity of problems that we are dealing with and the size of data sets that are generated or acquired. One of the most significant and evolving challenges that we are facing is processing and visualizing data from large sensor arrays. We are acquiring information at very high spatial resolution over large fields of view, producing extremely large amounts of data. The Center for Nondestructive Characterization and the Nondestructive Evaluation Section/Groups are supporting many projects and programs such as the Enhanced Surveillance Campaign (ESC), NIF, and other Defense Program initiatives, and Work for Others (WFO) such as the FAA, Knolls Atomic Power Laboratory (KAPL) and many others. The significance of the work performed for these programs varies, but the common thread is our emerging ability to nondestructively characterize objects and materials at resolution and detail that has not previously been possible.

2. **Project funding source, size in FTEs, PI, and major contributors**
   We support many programs and WFO efforts. We provide imaging for ESC, NIF, FAA, KAPL, medical applications, and many other organizations. We have approximately 30 employees supporting these efforts. Please contact Pat Roberson, Harry Martz, Jim Trebes, John Kinney, or Dan Schneberk.

3. **Code maturity level and history**
   We are interested in working with four nondestructive evaluation (NDE) codes that are parallelized or running on distributed systems. These codes consist of two reconstruction codes (the Convolution Back Projection and Feldkamp), one rendering code (cell_tracer), and one radiographic/tomography simulation code (Hades). Interestingly, these codes are somewhat similar in that they are ray-tracing codes that have a high degree of nested loops and show a lot of promise for unlimited linear speedup when applied over many distributed processors. The codes we are interested in have been developed at LLNL, and we own the source. NDE has a small cluster of 22 processors (nde-ccc). We have ported the Convolution Back Projection (CBP), Feldkamp, and cell_tracer to this system, and they all run in a distributed mode and all exhibit scalability. Hades has been ported to the Compass cluster and is capable of operating in parallel on that system. Our most urgent need is for using the Feldkamp reconstruction algorithm in a distributed mode on a much larger (classified) system than the nde-ccc 22-processor cluster. We are currently developing this code to be checkpoint restartable (assuming that this means that the code can recover from being stopped or from nodes dropping out of the system). We have not developed the other codes to be checkpoint restartable but could easily do so in 2–3 person months. I am not sure about Hades, however. As far as efficiency, the CBP code is very efficient at ~93%. The rendering code, cell_tracer, is also very efficient at ~92%. However, cell_tracer produces projections around some selected orbit of an object and has a current limitation of running on only as many processors as the desired number of projections. The projections become frames of a movie. With more work, however, cell_tracer could be distributed even further but possibly with a loss of efficiency. The Feldkamp reconstruction algorithm requires transferring data during computations. Unfortunately, there is a loss in efficiency due to this requirement. It is hard to determine precisely what the efficiency is for Feldkamp as it depends on the communication hardware used in the distributed system. We have determined that this algorithm can run from 87% to 18% efficient depending on the system hardware used. We are also interested in using finite element simulations of assemblies, objects, and materials as they are made or built. This will involve the direct transfer of pixilated tomography data to node and element descriptions for simulation codes. In this effort we predict the need for very significant computational resources.

4. **Anticipated scientific advance at 2–8 teraFLOP/s**
   We have computational problems that we cannot solve without significant distributed processing. These problems are of utmost importance to the programs that we support. However, our desire is not to operate on a system for a single one to four month period, but rather to operate on a periodic basis for 1–5 days at a time for approximately 20 periods during the year. These numbers will depend on the capability of the system.

5. **Anticipated scientific advance at 8–32 processor range**
   Systems operating at 25–100 GF will not solve our problem in a reasonable way.

6. **Connection with experiment**
   This work would support both experimental work and simulation work. However, the majority of data would be experimental. These data sets would be in the range of 256 GB to 2 TB in size.

7. **Requirements for scientific challenge at 2–8 teraFLOP/s**
   Our problem requires a 2–8 teraFLOP/s resource. For simplicity, the following assumes two

distributed systems, one containing 1000 processors and the other containing 2500 processors. We currently have systems capable of producing volume images of 256 GB in size and will have 2 TB images within the next 1 to 2 years. In our reconstruction problem we have raw data that must be reconstructed from Radon space producing images in image space. In an ideal case, for optimum speed of the algorithms, we need to hold both the Radon and image space data in memory during the reconstruction. This is approximately 128 GB + 256 GB = 384 GB of data for systems we currently have and 1 TB + 2 TB = 3 TB of data for our near future systems. For a distributed system containing 1000 processors, we will require ~384 MB/processor of memory for current data sets and ~3 GB/processor of memory for future data sets. For a system containing 2500 processors we will require ~153.6 MB/processor of memory for current data sets and ~1.2 GB/processor of memory for near future data sets. We will require a 500 GB global disk for current data sets and will need a 3TB disk in the near future. We will require ~256 MB/processor (current needs) of local disk and ~2GB/processor (future needs) of local disk for a 1000 processor system and ~102.4 MB/processor (current needs) and ~800 MB/processor (future needs) of local disk. We need to analyze and visualize the image data that is reconstructed in this process. The volume images are 256 GB to 2 TB in size and will require data processing including segmentation and rendering. Real-time interactive rendering would be desired.

8. **Requirements for scientific challenge at intermediate capacity level**
   A 25–100 GF system will not be adequate for our problem.

9. **Machine use pattern at 2–8 teraFLOP/s**
   We would desire a system with 512 nodes and 4 processors with 8 GB of memory and 80 GB of disk per node. Our most urgent algorithm requires data transfers between nodes during the processing period; therefore, we would like to have the fastest communication network possible. A system of this design would meet our requirements, but we need a system that is classified and would prefer that there be both a classified and unclassified system available. We would use this system for periods of 1–5 days and approximately 20 times during the year. We are in the process of evaluating the scaling of our algorithms and have found that they do scale on distributed systems such as the one described in the questionnaire. For one of our algorithms, we do suffer some loss of efficiency due to a requirement to transfer data between nodes while processing; however, this algorithm scales and we believe we can transfer data while processing data and minimize this bottleneck. We absolutely will need to analyze and visualize our data sets after we reconstruct. These processes will require some batch type processing but will also require some interactive and iterative processing. We will be analyzing data sets that are from ¼ to 2 TB in size.

10. **Machine use pattern at intermediate capacity level**
    Our urgent needs are for a larger system than what is described in this question. However, if a smaller system were available on demand, we would use it for smaller reconstruction, simulation, analysis, and visualization efforts. It would be important that the system be available on demand because we have systems that are equivalent to this in our facility (i.e., a cluster with 22 processors).

11. **Memory vs compute trade-off**
    We desperately need the memory capacity per node. We absolutely would not trade processors for memory.

12. **Comments on value of unclassified computing**
We need both a classified and unclassified system. We would put higher priority on a classified system. We need a system with 2–8 teraFLOP/s capability. If we had this as a resource for 1–5 day periods and at 20 sessions during the year, it would support many programs and WFO efforts that we currently have. For example, we support the Enhanced Surveillance Campaign (ESC), which is a campaign that supports the Core Surveillance Program. The objective of ESC is to develop tools, techniques, and models that enable us to provide advanced capability to measure, analyze, calculate, and predict the effects of aging on weapons materials and components and to understand these effects as they impact reliability, safety, and performance of weapons that are aged beyond their originally designed lifetimes. We currently have data sets that support this campaign that we cannot process due to a lack of computational resources. We also provide imaging for NIF, the FAA, the Knolls Atomic Power Laboratory, medical applications, and many other organizations. We are running into a limit with our current computational resources and we will no longer be able to support our programs unless we get access to systems such as the one described in this questionnaire.

Seismic and Acoustic Wave Propagation at Low Source Frequency or High Geologic Velocity

Shawn Larson

1. **Scientific intent and scope**
Seismic and acoustic wave propagation using the E3D code has been applied to a number of diverse projects in several Laboratory directorates and programmatic areas (Energy and Environment, Engineering, NAI, NIF/Lasers, Physics and Applied Technologies, Computation, multiple LDRD projects, GNEM program, Defense Programs Office), as well as Laboratory-sponsored collaborations (IGPP, ISCR, CLC) with academic and other institutions. Projects and research areas utilizing E3D include earthquake hazard analysis, oil exploration, nuclear nonproliferation, response of engineered structures, ground deformation modeling using geodetic data, underground structure detection, vehicle tracking, hydroacoustics, tectonics, medical imaging, and optics.

   Seismic and acoustic wave propagation in the earth and other material is a fundamental physical phenomenon. The ability to model and characterize seismic energy is therefore critically important to a wide range of existing and future Laboratory projects involving scientific, technical, and defense-related applications.

2. **Project funding source, size in FTEs, PI, and major contributors**
Current and recent funding sources for Laboratory projects utilizing the E3D code include DOE (NN, FE), LDRD (multiple projects), and Work for Others (e.g., PG&E, NSF, Nutec Sciences Inc., Karmanos Cancer Institute, U.S. Bureau of Reclamation). It is anticipated that additional funds will be secured from other government organizations (e.g., DOD, intelligence agencies). In addition, university collaborators have joint funding from the USGS, NSF, and IGPP, while other external groups have alternative funding (USGS, DOD, industry).

   Approximately 6 to 10 FTEs work on Laboratory-sponsored projects utilizing E3D. The external effort, including collaborative projects, is estimated to be between 10 and 15 FTEs.

Shawn Larsen is the PI for the computational aspects of seismic and acoustic wave propagation using the E3D code. Numerous individuals serve as PIs on independent projects utilizing the code.

3. **Code maturity level and history**
   E3D has been under development at the Laboratory for approximately 9 years. While there has been significant funding for the scientific and technical projects that utilize the code, there has been no explicit funding for code development until this year (FY02). A licensing agreement between the Laboratory and a service company representing the oil industry (Nutec Sciences Inc.) exists for their exclusive use of the code in commercial gas and oil applications.

   The 50,000-line source is written in C, with low-level Fortran calls for efficiency. The code features several levels of parallelization, each of which can work independently or in unison. Low-level optimization exists to maximize the pipelined features of certain architectures, including vector processors. Compiler directives exist to efficiently utilize shared-memory systems. MPI library calls are used to decompose the problem over the multiple nodes of a massively parallel system, with a PVM interface to do the same over a network of workstations. A scheme to allow multiple quasi-independent simulations to run concurrently on parallel systems is under conception.

   E3D has been implemented on a number of existing and former computer systems. These include workstations (Sun, SGI, HP, PC/Linux), heterogeneous workstation networks, SMPs (Sun Enterprise, SGI Origin 2000, Compaq Alpha clusters), vector processors (Cray YMP, Cray C90, Fujitsu VPP700E), and massively parallel machines (nCUBE-2, Meiko CS2, Cray T3D, Cray T3E, IBM SP2, Compaq MPP, PC/Linux cluster). Typical single CPU speeds of 25–40% peak are observed. A performance of 1.4 GFLOP/s (58% peak) was observed on a single CPU of a Fujitsu VPP700E vector processor. A sustained speed of 80 GF/s (25% peak) was observed for a 14-hour simulation using 240 CPUs of a Compaq massively parallel processor (the single CPU speed for an equivalent problem decomposition was approximately 33% peak).

   E3D has marginal checkpoint restart capabilities. The kernel is restartable, but several I/O functions are not. It will take approximately 5 days of dedicated time to correct this deficiency. The code has run-time visualization capabilities, but sophisticated post-processing visualization utilities are not readily available.

4. **Anticipated scientific advance at 2–8 teraFLOP/s**
   Capability computing at the 2–8 teraFLOP/s level will satisfy needs in (1) underground structure detection, (2) nuclear nonproliferation, (3) oil exploration, and (4) earthquake hazard analysis.

   In general, seismic wave propagation problems are computationally limited to lower frequencies and higher geologic velocities (e.g., rock instead of soft sediment). Doubling the source frequency or reducing the geologic velocity by a factor of 2 requires an 8-fold increase in memory and a 16-fold increase in compute time. In the case of underground structure detection and nuclear nonproliferation, it is desired to increase the source frequency by a factor of 2 to 4. In the case of oil exploration and earthquake hazard analysis, it is desired to decrease the geologic velocity by a factor of 2 to 4. Such simulations will require

machine capabilities that are at least 50 to 100 times the 10–20 GF simulations that are performed today.

For example, a consortium of companies representing the oil industry recently approached the national laboratories inquiring about the feasibility of performing several (10 to 100) large-scale simulations to address current scientific problems encountered during their search for oil and gas. Several of these companies are already collaborators on DOE-sponsored LLNL research projects. The resource requirements for each simulation are approximately 0.5–1 TB of memory and 12–48 hours of compute time on a 1 TF system.

5. **Anticipated scientific advance at 8–32 processor range**
   Future seismic modeling problems, particularly those involving underground structure detection and oil exploration, will likely require hundreds if not thousands of simulations where the geologic model and other simulation parameters are slightly perturbed. This is necessary to accurately characterize the subsurface in the presence of fine-scale geologic heterogeneity. For example, the detection of underground structures is an extraordinarily difficult problem. It is possible to solve this problem with a preponderance of data or a preponderance of simulation. It is likely that very limited data will be available in real world applications, hence the need for vast intermediate capability computer resources. Each relatively small 3D simulation needed for this type of analysis currently requires several hours of compute time using a few dozen processors. It is anticipated that 500 to 1000 simulations could be performed over a 3-month period with dedicated use of 32 processors.

6. **Connection with experiment**
   Experimental data are frequently used to validate the E3D code and its application. For example, computational simulations of historic earthquakes are compared to observed seismic data to validate the geologic and fault-rupture models used to predict strong ground motions from future events. For other applications, such as underground structure detection, it is anticipated that experimental data will be used to characterize and assess stochastic models of the subsurface.

7. **Requirements for scientific challenge at 2–8 teraFLOP/s**
   It is difficult to constrain an answer to this question due to the wide range of diverse scientific problems that can be addressed if such resources are available. Problems of this size are likely to be run-time limited. Several problems could be defined that would utilize 1 TB of memory. Such problems would require approximately 12–48 hours run time on a 2 TFlop system (2 TFlop observed is equivalent to about 8 TFlops peak). The local/global disk space would be relatively small. Some problem definitions can be parameterized with exceedingly small disk requirements. Others require a full grid definition, which for a 1 TB problem in memory will require about 200 GB of disk space. Restart capabilities essentially require that the entire grid be dumped to either local or global disk locations (e.g., 1 TB space would be required, but this could be local if the same node configuration will be used on restart). It is not likely we will have the resources to pursue advanced visualization capabilities, although this is always subject to change. It is not anticipated that large tertiary storage requirements will be needed (hundreds of GB at most).

   A start date can't be determined until the likelihood of securing computer time is known (see response to question 12). There are too many variables, many that involve our interactions

with external collaborators. Depending on the project, start dates will be 2 to 6 months after it is known that there is a reasonable chance that computer time will become available.

8. **Requirements for scientific challenge at intermediate capacity level**
   Again, this is difficult to answer due to the wide range of problems being considered. In general, problems of this size are likely to be fairly well balanced between memory and run-time. A 10–15 GB problem (a relatively small 3D simulation) will require 4–8 hours compute time on a 10–20 GFlop (observed) system. The local/global disk space would be minimum. At most, approximately 5 GB of global space would be required (assuming restart capabilities wouldn't be implemented for these types of problems). Again, advanced visualization is desired, but we don't have the resources to pursue this. Tertiary storage would be small, not more than a few hundred GB.

   Projects related to oil exploration could be started almost instantaneously (in fact, we are pursuing this on a smaller scale using our existing allocated resources). The same is true for ongoing earthquake hazard research. Robust intermediate capability underground structure detection simulations could not begin until approximately 6 months.

9. **Machine use pattern at 2–8 teraFLOP/s**
   The algorithms and the applications that use the algorithms are well validated. Likewise, scaling issues are fairly well understood. More so than not, a regular grid is decomposed over an explicit finite-difference algorithm. Minimal studies need to be performed in these areas. A significant source of effort comes from the resolution of numerical instabilities. These simulations often push the bounds of simulation past what is commonly known or expected based on published literature. Numerical problems are therefore encountered that need to be understood and resolved on a case-by-case basis. Experimental and other data will be available for some problems. In general, these data will be compared to simulation. Differences between observation and calculation will primarily be attributed to inaccuracies in the model (as opposed to inaccuracies in the algorithms). If these inaccuracies are significant, the model will be refined and simulated again.

   The general approach is as follows. After a problem is formulated, a model is defined. That is, the geologic parameters that define the subsurface are defined. This model may be constructed internally, or it may be obtained from outside organizations (e.g., oil industry). The model can be synthetic or it can be a representation of a real geologic environment. In the second case, there is often experimental data associated with the model. A small number of test simulations are performed. The model and simulation parameters are modified. One or more simulations are performed. If data are available, it is often necessary to associate and compare the data with results from the simulation. In some cases, the model parameters will be modified and an iterative approach will be used to obtain a geologic model that is consistent with the data. In other cases, a stochastic representation of a suite of models is obtained, and this representation is used to define the model from a series of simulations. This methodology is applicable for underground structure detection. For some problems, data are not needed. Instead, it is desired to construct a synthetic data set that can be used to test other aspects of the problem. This is the primary need of our projects with the oil industry.

   In general, model and parameter iterations are desired; therefore, an ideal environment is one that allows the user dedicated time at several intervals during a multiyear project. For example, it is far more beneficial to have eight $^{1}/_{2}$-month-long dedicated allocations over a

period of 18 months (or 2 years), than it is to have a single 4-month-long dedicated allocation during the same period.

10. **Machine use pattern at intermediate capacity level**
    The answer is similar to the answer for question 9, although longer periods of dedicated application may be more useful (e.g., two 2-month-long dedicated applications during a 12-month interval).

11. **Memory vs compute trade-off**
    E3D problems that can be implemented on a small number of processors (e.g., less than 32) tend to be memory limited, and hence less memory per node will not be efficient. However, memory scales as the third power of problem size, while run-time scales as the fourth power of problem size. Hence, problems requiring a large number of processors (e.g., more than 128) will be run-time limited. Said in another way, capacity and intermediate capability computing requires significant memory per processor, while capability computing may be satisfied with less memory (e.g., 1 GB/processor).

12. **Comments on value of unclassified computing**
    The availability of large institutional computing resources, especially through M&IC, has been a critical and necessary component of a diverse set of wave propagation problems important to many Laboratory programs. In addition, these resources have played an invaluable role in developing new technological thrust areas at the Laboratory. The return on investment has been immeasurable.

    Scientific breakthroughs have been achieved in areas of earthquake hazard analysis, crustal deformation modeling, and oil exploration. Results from these investigations have been presented in peer-reviewed journal articles, invited talks, and in presentations at national conferences. Key deliverables have been met that satisfy ongoing programmatic requirements, such as those imposed by the nuclear nonproliferation efforts of the GNEM Program. The E3D code and institutional computing resources were used to analyze a critical design issue of NIF and were used in a collaborative effort with an external medical institute to investigate the detection of cancer in human tissue. E3D and institutional computing has and is being used in a number of LDRD projects. Scientific results from investigations made possible through M&IC resources have received local and national media coverage.

    Most important, the availability of institutional computing resources makes it possible to explore new focus areas that have the potential to develop into active and well-funded programmatic areas directed at the solution of national and international problems. For example, modeling efforts geared toward the detection and characterization of underground structures would not be possible without the capacity and capability computing resources offered through M&IC.

    Unmet needs include (1) the ability for rapid turnaround (e.g., a modest 40-CPU simulation currently remains in the batch queue for 24 hours), (2) the availability of computer time to perform those “spur of the moment” and “proof of concept” ideas that are frequently encountered in our scientific work environment, and (3) less delay between a request and approval of computer time. For example, it took over 1 year for our initial request for computer time to be granted, forcing us to seek alternative resources at Los Alamos and in Japan (this was not the fault of M&IC, which handled our requests quickly). Many projects

require multilevel coordination. It is difficult and many times impossible to proceed in any reasonable fashion when there is a significant delay in some aspect of the process.

State-of-the-Art Calculations of Flow and Dispersion of Air-Borne Hazardous Materials

Robert L. Lee, Stevens Chan

1. **Scientific intent and scope**
   We develop high-performance Computational Fluid Dynamics (CFD) models for simulating flow and dispersion of hazardous materials over urban areas. The CFD models we have developed are world class and have been used to support chemical–biological projects that are funded by both DOE and DOD. The problems of relevance to these agencies typically require calculations involving millions of grid points and the use of massively parallel computer platforms. The results of these model calculations can be used in emergency planning and response activities.
2. **Project funding source, size in FTEs, PI, and major contributors**
   The current funding sources are from DOE and DOD. The number of FTEs on the projects range from six to four scientists and computer scientists. The PIs of the projects are Robert L. Lee and Don Ermak from the Atmospheric Science Division.
3. **Code maturity level and history**
   The massively parallel code was developed in 1999 at LLNL and utilizes a high level of parallelism. It can run on a variety of platforms—ASCI White, Frost, TC2K, Compaq Alpha. We have observed that the performance of the model increases almost linearly with respect to the number of processors. We use EnSight and the MATLAB packages for visualization of the computational results.
4. **Anticipated scientific advance at 2–8 teraFLOP/s**
   We would be in a position to solicit funding via performing a number of breakthrough calculations of the 10 to 100 million grid-point range, which would be considered state-of-the-art.
5. **Anticipated scientific advance at 8–32 processor range**
   It would enable us to perform numerous parameter studies and further improve our models in both numerics and physics.
6. **Connection with experiment**
   We have successfully compared the results of the simulations with field experiments. The agreement between field and model data provides confirmation that the model is capable as a prediction tool. In addition, we could use our validated model to make calculations prior to field experiments to guide instrument setup and data collection.
7. **Requirements for scientific challenge at 2–8 teraFLOP/s**
   Not sure about the numbers. But we will certainly be able to use whatever resource is available to us to perform state-of-the-art building-scale flow and dispersion calculations. Besides sufficient memory and disk space, we would also need visualization support that will enable us to visualize our computational results efficiently for problems involving 10–100 million grid points. We could probably be ready to make such calculations in 6 months.

8. **Requirements for scientific challenge at intermediate capacity level**
   We have been doing this kind of calculations for some time and would continue to require this kind of support from M&IC.

9. **Machine use pattern at 2–8 teraFLOP/s**
   We would use the machine to make building-scale flow and dispersion calculations with very high resolution (say, ~1 m grid spacing) and an advanced turbulence model such as LES (large eddy simulation) to model atmospheric turbulence more accurately. Such results could be used to demonstrate the Lab's unique capability in this field and attract more funding from potential sponsors.

10. **Machine use pattern at intermediate capacity level**
    We have been using similar computing capability for some time and expect to require continued support from M&IC in the future.

11. **Memory vs compute trade-off**
    Yes. It would be of use for our applications.

12. **Comments on value of unclassified computing**
    The computing resource from M&IC is not only essential for us to be able to meet our programmatic goals and milestones, it also enables us to make model predictions to aid instrumentation setup, thus making field experiments more cost-effective. We have been fortunate in gaining access to the ASCI White and, more recently, Frost as beta users. This access to such world-class computing platforms enabled us to make a number of state-of-the-art calculations and has drawn much attention from our potential sponsors. This led to several projects with new funding and also enabled us to hire two young scientists into our CFD team. We were also able to propose new, more ambitious, programmatic milestones that would not have been possible without the assistance from M&IC.

NUFT-C Project

William Glassley, John Nitao

1. **Scientific intent and scope**
   Water movement, contaminant transport, and chemical reactions between migrating fluids and minerals are common processes in the shallow levels of Earth's crust. These processes play a direct and controlling role in determining important aspects of many environmental, energy, and scientific concerns, including groundwater quality and quantity, subsurface pollutant containment and transport, petroleum basin evolution and petroleum migration, oil reservoir management, and sea water intrusion in coastal environments, to name a few. The NUFT-C code is designed to rigorously account for the coupled physical and chemical processes that occur as subsurface water migration takes place. This code has been applied to a number of diverse projects and programmatic areas, including Energy and Environment projects in carbon sequestration and high-level nuclear waste disposal (the Yucca Mountain Project), and multiple LDRD projects. It has also been applied to problems in other international nuclear waste disposal efforts, including the Japanese nuclear waste repository program.

   The capability to predict how fluid moves in Earth's crust, and what chemical interactions may occur, is critically important to present and future Laboratory interests. Whether

considering the impacts of climate change on soil chemistry and properties, contaminant transport and remediation activities, or carbon sequestration in subsurface environments, the processes modeled by the NUFT-C code must be accurately simulated. These and other related areas will play an important role in future Laboratory efforts in the Environment and Energy arena.

2. **Project funding source, size in FTEs, PI, and major contributors**
   Current and recent funding sources for Laboratory projects utilizing the NUFT-C code include DOE (Yucca Mountain Project, Carbon Sequestration program) and LDRD. Through bilateral agreements, work is also being funded to conduct simulation activities for the Japanese Nuclear Cycle Development Institute. Additional funding is anticipated in collaborative work involving the petroleum industry and DOE's NGOTP efforts.

   Roughly 5 to 10 FTEs work on Laboratory-sponsored projects utilizing NUFT-C. William Glassley and John Nitao acted as co-PIs for code development and application. Numerous individuals serve as PIs on independent projects utilizing the code.

3. **Code maturity level and history**
   NUFT-C has been under development at the Laboratory for approximately 4 years. The code is an outgrowth of the NUFT code for simulating thermal and hydrological processes for multicomponent fluid flow in fractured, porous media, for unsaturated and saturated, nonisothermal conditions. Coupling geochemical processes to thermal–hydrological effects was accomplished by explicitly and rigorously representing changes in the properties of the physical framework caused by chemical effects. Rigorous representation of these coupling effects is a significant advance in simulation capabilities. Institutional LDRD support was provided to accomplish this goal.

   The code consists of 250,000 lines written mainly in C and C++. Computations are conducted within a finite volume framework, with internal mesh generation. Cylindrical or rectangular coordinates allow fully 3D simulations. Time-stepping is automatically controlled. MPI library calls are used to decompose the problem over the multiple nodes of a massively parallel system.

   NUFT-C has been implemented on a number of existing and former computer systems. These include workstations (Sun, SGI, HP, PC/Linux), heterogeneous workstation networks, SMPs (Compaq Alpha clusters), and massively parallel machines (IBM SP2, Compaq MPP, PC/Linux cluster). Single CPU speeds of up to 25–40% peak are observed.

   NUFT-C has checkpoint restart capabilities that are fully functional. Independent scripts are used for run-time visualization. Post-processing visualization utilities are available, but have limited capabilities.

4. **Anticipated scientific advance at 2–8 teraFLOP/s**
   High-level nuclear waste disposal relies on the ability to simulate the thermal–hydrological–geochemical evolution of the potential repository site for thousands of years into the future. Currently, such simulations are simplified (limited number of chemical species considered at limited resolution) because of the heavy computational burden associated with such calculations. Consequently, there has never been a full simulation of a potential repository in which a simulated release of radionuclides from a leaking waste container has been represented.

Conducting such a calculation would be a breakthrough in the development of safety arguments for nuclear waste repositories. It would allow clear evaluation of the uncertainties associated with previous simplified and abstracted models, would allow unambiguous representation of the interplay between coupled processes that influence contaminant transport, and would provide a rigorous description of the rates and 3D flow pathways released radionuclides would follow.

This accomplishment would be an important breakthrough by virtue of addressing an important national issue (radioactive waste disposal) that impacts nuclear materials management. The currently most advanced simulation tool for this purpose (NUFT-C) has focused on simulating the evolution of the natural system, perturbed by the thermal impact of emplacing nuclear waste. These simulations have taken upwards of 80 hours on an IBM SP2, utilizing up to 256 processors. Consideration of radionuclide transport, which was not treated in these simulations, would increase the problem size by approximately two orders of magnitude.

5. **Anticipated scientific advance at 8–32 processor range**
   To date, there have been only a handful of NUFT-C runs that consider the repository behavior (thermal–hydrological–geochemical effects in the natural system without radionuclides). Many scenarios have yet to be evaluated in which different heat loadings and tunnel geometries are considered for their impacts on performance. Access to large banks of processors would allow a dramatic and systematic improvement in the ability to optimize a repository for overall safety and performance.

6. **Connection with experiment**
   Support of experimental science comes from dramatically increasing the ability to understand the consequences of interactions that have previously not been systematically studied in a coupled fashion. NUFT-C simulations provide clear insight into data limitations (permeability–porosity relationships, fracture–matrix interactions, reaction kinetics, etc.) that that can guide, and have guided, experimental studies. Similarly, experimental data on rock properties and on mineral and chemical thermodynamic, and kinetic parameters form the basis of the calculations we perform. Planned field and laboratory experiments form the only credible means for establishing initial and boundary conditions for the simulations. In addition, it is on the basis of experimental data that most code verification and validation activities are conducted.

7. **Requirements for scientific challenge at 2–8 teraFLOP/s**
   There are diverse suites of problems that could be addressed by this code, each with its own unique set of requirements. For most of these problems, we are run-time limited; it is common to have the need to simulate thousands of years of processes but be time-step constrained by the relatively short time constants affecting transport processes. For many of these problems, approximately 100 hours would be required on an 8 TF system with a terabyte of memory (assuming standard efficiencies we normally see when simulating complex chemical systems). The local/global disk space would be moderate (500 GB).

   Advanced visualization capabilities (high resolution and animation) are needed to thoroughly utilize the results of the computations. We have been pursuing this effort informally, and need to continue that effort in a dedicated way. Storage requirements would be on the order 10 TB.

Start dates would be shortly after computer time is provided, generally less than 1 month, if FTE support were available.

8. **Requirements for scientific challenge at intermediate capacity level**
   Such activities would represent a scaled-down version of the efforts described in the response to question 7. Focus would be on examination of specific chemical interactions and impacts. Run-time constraints would establish the level of chemical complexity that could be considered. But, without test runs, it is difficult to answer this question quantitatively.

   A 10–15 GB problem would require about 20 hours compute time on a 10–20 GF (observed) system. Required local/global disk space would be about 100 GB. Visualization needs would be similar.

9. **Machine use pattern at 2–8 teraFLOP/s**
   The algorithms and the applications that use the algorithms have been applied and tested in numerous applications to natural and synthetic systems and are well validated. Although less well evaluated, scaling issues have been explored and are reasonably well understood. Our primary constraint is linear solver technology. There has not been a large degree of effort committed to developing optimized solvers for these problems.

   Experimental and other data are available for testing some problems. However, the data set available is generally insufficient to allow testing of all components of the code simultaneously. Generally discrepancies result from conceptualization errors, as opposed to inaccuracies in the algorithms. Model refinement is a continuous and ongoing activity.

   The primary challenge in these simulations is constructing a conceptual model that realistically represents the system being considered but that also honors the limited data sets available that describe physical parameters that apply to the complex natural system being considered. Generally, these simulations are data starved. The general strategy is to construct a model based on the conceptualization and run numerous preliminary simulations to refine those initial and boundary condition estimates that had to be made in the absence of hard data. This is often the most time-consuming aspect of running these simulations, and can take weeks. Once simulation initial and boundary conditions are satisfied, the simulations are run and results compared to available data.

10. **Machine use pattern at intermediate capacity level**
    The answer here is the same as to question 9. In this instance, we would scale the problems and simulation times accordingly.

11. **Memory vs compute trade-off**
    This would not be useful for these simulations. We tend to be memory limited, and a decrease in memory for FLOP/s would generally be not beneficial.

12. **Comments on value of unclassified computing**
    In the area of nuclear waste management, there is a strong inertia to overcome that relates to the mind set of "We can't afford the time and effort to change horses in mid-stream" (in this case, horses being computational methods, i.e., changing from serial to high performance parallel machines). Nevertheless, our efforts are attracting increasing favorable attention nationally and internationally. It is now recognized, as a result of our efforts, that such simulation capabilities need to be adopted. In the foreseeable future, these simulation tools will be employed by the U.S. Yucca Mountain Project and by the Japanese Nuclear Cycle

Development Institute through a bilateral agreement. Increasing interest from regulatory elements is also being expressed (particularly by the Nuclear Regulatory Commission and the French Nuclear Energy Agency). It is anticipated that within the next 12–18 months there will be a significant shift by these entities toward reliance on such computational platforms.

Currently, LLNL is the only institution that has exploited this capability and applied it successfully in the nuclear repository arena. Without the resources the Lab has developed in this area, we would not be in the leadership position we currently occupy and intend to hold into the foreseeable future.

High-Resolution Global Climate Simulations

Philip B. Duffy, et al.

1. **Scientific intent and scope**
   Our project is to perform simulations of present and future global climates at higher spatial resolution than has ever been used in this type of simulation. Preliminary analyses of our initial high-resolution simulations of present climate indicate that the fine-scale (~100 km) detail in these simulations agrees well with observations. In addition, these simulations produce superior simulations of present climate even on scales that are resolved by coarse-resolution models (~1000 km). Our initial high-resolution simulations of climate changes for the next 100 years show very different results than coarse-resolution simulations in specific regions (e.g., western U.S., eastern Canada). Thus our simulations are providing (1) improved simulations of climate change on regional (~1000 km) spatial scales and (2) the first fine-scale (~100 km) predictions of global climate change. In short, our simulations appear to have both increased accuracy and increased detail compared to previous simulations. For these reasons, this work is generating strong impact among climate scientists. The increased accuracy and detail in our simulations also makes them uniquely well-suited for estimating the societal impacts of climate change. We therefore expect that when we make the results widely available, they will generate strong interest among scientists and social scientists working on this problem.

   Strategically, this work is important to the Atmospheric Science Division, the Energy & Environment Directorate, and to LLNL because it takes advantage of one of our strengths—high-end computing. We hope, therefore, to expand our funding and activities in this area. In addition, our strategy of making our results universally available for others to analyze is consistent with the role DOE likes us to play of performing a service for the climate-research community.

2. **Project funding source, size in FTEs, PI, and major contributors**
   We are currently funded by DOE Office of Biological and Environmental Research (within the Office of Science). Our total funding level is a little over one FTE. Team members are Philip Duffy (PI), Bala Govindasamy, Jose Milovich (CASC), and Starley Thompson.

3. **Code maturity level and history**
   The code we are using is the Community Climate Model (CCM) version 3.6.6, developed by the National Center for Atmospheric Research (NCAR). This is the leading U.S. climate model, in terms of both the number of users the fidelity of the results. The code is very widely used and has been heavily exercised on a wide range of machines, including several at

LLNL. The code achieves parallelism via both MPI message passing and OpenMP. Nonetheless, in the high-resolution configuration in which we use model, the parallel efficiency is poor. (We can use at most 200 or 300 processors efficiently on IBM systems.) This is due in part to a nonoptimal parallelization strategy (1D domain decomposition), and in part to a dynamics that requires global communication. (Both these problems could be eliminated by substitution of an alternative dynamical core, which we have recently "parallelized"). The code is checkpoint restartable.

4. **Anticipated scientific advance at 2–8 teraFLOP/s**
   The answer to this depends on how well our code scales on this machine. If, as expected, our code scales poorly, we could efficiently use large numbers of processors by running an ensemble of high-resolution calculations. Because of the chaotic nature of the atmosphere, an ensemble of calculations is needed to obtain the most accurate predictions. (We could, for example, run an ensemble of seasonal forecasts predicting the effects of the upcoming El Niño.) Of course, if our code did scale well on the new machine (which I don't expect) then we could push to even higher spatial resolutions.

5. **Anticipated scientific advance at 8–32 processor range**
   My group uses this level of computing to do coarse-resolution simulations, which constitute the bulk of the work we do. Increased access to this level of computing would allow us to perform our programmatic work more quickly and easily than we can now.

6. **Connection with experiment**
   The credibility of our predictions of future climate is based upon the ability of our models to simulate present climate. This ability is evaluated by comparing model results to a wide range of observations (satellite, surface oceanographic, geological, etc.). Thus, we are dependent upon observations to evaluate and improve our models.

7. **Requirements for scientific challenge at 2–8 teraFLOP/s**
   A number of these responses are approximate, based on our experience running on IBMs.

   - Ratio of bytes of memory/flop?
     Whatever value of this ratio the IBM systems (e.g., Frost) have works for us. Because our code does not scale well, what is more important than this ratio is for both the flops and memory per chip to be as high as possible.
   - Ratio of bytes of disk/flop?
     Disk space is not a big issue for us. Even at high resolution, our code creates output slowly enough that it is easy to periodically move it to mass storage. 100 GB of local disk is adequate.
   - Local disk needed?
     Of order 100 GB.
   - Visualization?
     We are using a product developed at LLNL known as VCS, and a successor known as VCDAT. These are only semi-satisfactory; a faster and more versatile visualization tool would be great.
   - Tertiary storage required?
     Maybe 5 TB at most.

- When ready?
  Now!

8. **Requirements for scientific challenge at intermediate capacity level**
   - Ratio of bytes of memory/flop?
     Our coarse resolution codes are not particularly demanding in terms of memory. Probably any system that LC would consider buying would work.
   - Disk space?
     Not a big issue. 100 GB is plenty.
   - Visualization?
     For this level of project, our present visualization capabilities are adequate.
   - Tertiary storage required?
     <1 TB.
   - When ready?
     Now!

9. **Machine use pattern at 2–8 teraFLOP/s**
   The preliminary high-resolution climate simulations we have already performed mean that we are ready to do large-scale production runs now. (Indeed, we are already doing such runs on Frost and other machines.) No evaluation of algorithms, etc. would be needed.

10. **Machine use pattern at intermediate capacity level**
    We would use this type of system to perform coarse-resolution simulations of the atmosphere, ocean, and other aspects of the climate system (e.g., carbon cycle). Again, we are ready to perform such simulations now (indeed, are already doing so on other systems). No evaluation of algorithms, etc. would be needed.

11. **Memory vs compute trade-off**
    No, 0.75 GB memory per processor is not enough for us.

12. **Comments on value of unclassified computing**
    Our access to M&IC computing has been directly responsible for much of our recent success. Our strategy for future success (indeed future survival) is closely tied to continued access to M&IC computing. Here's why. The general field we work in (climate modeling) is one in which universities etc. also participate. Since we are much more expensive than universities, we are forced to work on problems that universities cannot address. (If we compete head to head with them, we lose, because we are two to three times more expensive and not two to three times more productive.) In climate modeling, the single most important advantage we have over universities is our access to high-end computing.

    Thus our competitive strategy is to address problems (e.g., high-resolution modeling) that require big computing. Since it is clear that universities and even many other laboratories cannot attack these problems, this provides us with a *raison d'etre*. It is also noteworthy that emphasizing the use of high-end computing allows us to make the argument that there is important synergism between our work and the weapons programs. (And this is true, not just BS.) Obviously, for this strategy to succeed, we need continued access to high-end computing at LLNL. While we also have access to other, supposedly high-end facilities

through DOE (e.g., NERSC), in practice, these other facilities have been so oversubscribed that it has been difficult to do breakthrough science there.

To be specific about how M&IC computing has helped us: In the early days of TC2K, we used dedicated access to part of the machine to perform the highest-resolution near-equilibrium simulation ever performed with a global ocean model. We also performed the highest-resolution simulations of direct injection of CO_2 into the ocean (as a means of removing carbon from the atmosphere). We have used Frost to perform the highest-resolution simulations ever with a global climate model. All of this work has been widely recognized in the scientific community.

Global Atmospheric Chemistry

Doug Rotman, et al.

1. **Scientific intent and scope**
   The goal of our work is to understand the effects of natural and anthropogenic activities on the distribution of important atmospheric chemical species. These activities include natural events, such as volcanoes, biogenic surface emissions, lightning among many others, and anthropogenic activities such as energy use, solvents, biomass burning. Implementation of a National Energy Policy will strongly depend on knowing the impact of various energy technologies and policy scenarios on the environment and climate; hence, atmospheric chemistry bridges the issues of energy production and use and environmental quality. Our IMPACT model remains the only model capable of comprehensive tropospheric and stratospheric chemistry, so we provide a unique tool in this understanding. We are engaged in studies focused on regulatory analysis (e.g., do we understand the role of emissions of species A and should it be regulated) and also on improved prediction of climate change and its relationship to energy policy.

2. **Project funding source, size in FTEs, PI, and major contributors**
   We are currently funded by the DOE Office of Science and by NASA, with smaller projects at other agencies. Our team includes approximately 10–12 FTEs, including both atmospheric and computational scientists. PIs and major contributors include Doug Rotman, Cyndi Atherton, Peter Connell, Cathy Chuang, Jane Dignon, Dan Bergmann, John Tannahill, and Philip Cameron-Smith.

3. **Code maturity level and history**
   IMPACT has been in development and use at LLNL for nearly 8 years. It represents the primary large-scale atmospheric chemistry model in the DOE Office of Science and is the basis of the NASA Earth Science atmospheric chemistry activities. The model is approximately 100,000 lines of Fortran 90, with small areas of C programming and Perl scripting. Parallelization is accomplished through MPI in a 2D decomposition. (Over the next few months, OpenMP constructs will be implemented to take advantage of the shared memory on large parallel systems.) While the model is continually under development with additional or updated chemistry and physics, it is also used in production modes. We've had production runs ongoing for about 4–5 years. The model includes restart capability as well as user-specified output frequencies. Decomposition and level of parallelism, input and output characteristics, problem definition, and other parameters are all changeable at run time.

IMPACT has been implemented and run on many machines, including workstations (Sun, SGI), shared-memory machines (SGI Origin, Compaq cluster), vector machines (Cray C90, T90), and distributed-memory machines (Paragon, Meiko CS2, Cray T3D, Cray T3E, IBM SP, Compaq SC1, and others). IMPACT scales well on parallel machines, maintaining nearly 80% parallel efficiency on the SP at 100 processors. On the Compaq SC1 we have sustained between 200 and 250 MF per processor.

Our visualization techniques are primarily post-processing the archived NetCDF files using locally developed packages. These capabilities are primarily 2D, but some 3D graphics are coming on-line. Much of our high-end graphics work (and presentations on Computation's Power Walls) comes from collaborations with Computation's visualization teams.

4. **Anticipated scientific advance at 2–8 teraFLOP/s**
   Access to dedicated time at the 2–8 TF level would enable long-term "trend" runs that are currently not possible. By this I mean, currently we can do comprehensive chemistry and physics over relatively short time periods (where emissions and activities are held constant), but we cannot do the needed long-term simulations where trends in emissions are allowed to be integrated into the model simulation, allowing one to attempt the simulation of actual multidecades of atmospheric observations. Only then can one examine the interplay of energy-use emissions, infrequent volcanic eruptions, and related activities and how those interact to produce the historical species distributions.

5. **Anticipated scientific advance at 8–32 processor range**
   Such access would allow sensitivity studies of chemical mechanisms, physical parameterizations, input meteorological data, aerosol loading, and other chemistry and physics issues. These studies are essential to our understanding of how to simulate atmospheric chemistry. Also, simply having high-quality access to that processor count would allow us to nearly continuously have production runs in the hopper; this is essential to the production of science papers and analyses.

6. **Connection with experiment**
   We continually validate our model simulations to ground, aircraft, and satellite observations. Also, chemistry parameters and other physical parameters are directly input from laboratory experiments. In addition, our model is used in a coordinated fashion with ongoing observational campaigns to examine and understand data.

7. **Requirements for scientific challenge at 2–8 teraFLOP/s**

   - Local memory?
     In general, climate and chemistry models are not memory limited. The rule of thumb of 2 FLOP/s to 1 byte of memory is adequate (i.e., a 1 GF processor should have 512 MB memory). However, speed between memory and chip is essential. Memory bandwidth begins to compromise performance if the cross-sectional bandwidth is greater than one-eighth the desired sustained FLOP rate (i.e., a requirement of 8 FLOPs per byte bus volume).
   - Cache size?
     16 MB (or at least 8 MB)

- Inter-node bandwidth and latency?
  Inter-node bidirectional bandwidths need to be at least 5 GB/s and latency needs to be 5 μs or less.
- I/O bandwidth?
  For chemistry and climate models, this does not greatly affect throughput of simulations. More likely a driving force in this regard is any computer center need for checkpoint and restart.
- Disk capacity?
  Climate and chemistry models produce large output data sets. Attached storage needs to have 5–10 TB capacity.
- Software/library needs?
  Key software and libraries include NetCDF I/O library, TotalView debugger, parallel performance profiler/monitor, MPI and OpenMP, and standard mathematical libraries (BLAS, FFT, LAPACK, Spherepack).
- Batch queue?
  Climate and chemistry models need to simulate long-time-scale events, so climate and chemistry simulations run for long periods of time. Simulation throughput can be optimized with long batch queues.
- Tertiary storage?
  Storage requirements likely would not go beyond 5 TB.
- When ready?
  For many simulations we are ready now. However, for the very large trend run (as described above), we would need about 6 months to prepare the needed BCs, ICs and other parameters.

8. **Requirements for scientific challenge at intermediate capacity level**
   Many of the computational requirements above remain in effect. For example, disks, memory, etc., on TC2K is fine. Typically, our runs are not disk space nor memory limited (within reason), especially for most any machine LC is likely to purchase. Memory can become a problem if we use too few processors, so 1 GB per processors is likely needed. I am vague here because disks and memory are usually not an issue, so I do not have quality or determined values here. We are, however, time limited. We need long periods of time on these processors. Our runs can typically be hundreds of hours. We could start these runs immediately.

9. **Machine use pattern at 2–8 teraFLOP/s**
   Currently, we have 12 papers under development that are requiring simulations. Most would make use of IMPACT in its current form, simply applied in different ways. Each of those simulations could easily use 64–100 processors for weeks to a month. The IMPACT model is mature, relatively well understood, with quality (and analyzed) algorithms, chemistry, and physics. While the model is not perfect (what is?), its "characteristics" are known and assessed. Scaling and performance are, again, relatively well known and understood. Many of our studies involve the focus on the upper troposphere and lower stratosphere. This is an area of the atmosphere where few others can claim to analyze, but our combined troposphere/stratosphere chemistry is capable of doing that. The world knows of our model and its capabilities, but this series of papers could really put a stake in the ground. As a

whole, these studies involve gas phase chemistry, heterogeneous chemistry and aerosols. Taken together, they will provide insight into the role of energy use emissions in tropospheric ozone, the role of aerosols in radiative forcing, and the influence of tropopause ozone in climate change. With large allocation and access to a large machine, we could actually carry out the trend run (described above). This would be a major simulation result—basically the first ever using a model with comprehensive chemistry and physics. It would be a benchmark simulation in the atmospheric chemistry community. Historical emission databases would need to be created for the trend run.

10. **Machine use pattern at intermediate capacity level**
    The answer here is similar to number 9, although this would not allow the trend run. Moreover, the suite of simulations would have to be done in some order. Each run is approximately 3–6 weeks in length.

11. **Memory vs compute trade-off**
    Down to some level yes, but if the number of processors is less than 30–50, then reductions in memory are not useful. That said, in general, I would trade off memory for speed. As I mentioned above, we tend not to be memory limited, so speed is more important.

12. **Comments on value of unclassified computing**
    Access and use of the M&IC computational resources has been essential and critical to our success. The IMPACT model now stands as one of the world's most advanced 3D atmospheric chemistry models. We remain the only location able to carry out combined stratospheric and tropospheric chemistry simulations. Other locations (NCAR and NASA) are beginning the work in this area but are at least 2 to 3 years away from that capability. These successes come because of the quality staff in the Atmospheric Science Division (ASD) and also the state-of-the-art computational resources available to that staff at M&IC. We are the NASA core modeling site for 3D chemistry simulations because of our expertise in atmospheric chemistry/science and because of our ability to use large computers. We have been pushing the envelope in coupled chemistry climate modeling over the last few years, much of this done via LDRD and M&IC access. This has resulted in LLNL being selected as the primary institution for implementing interactive chemistry into the NCAR climate model that is the basis of a large SciDAC project, bringing together the skills of NCAR, DOE (LLNL, LANL, ANL, ORNL, PNNL, LBNL) and NASA Data Assimilation Office.

    Atmospheric chemistry is extremely CPU intensive. Many groups have had to cut back on modeling, run at reduced chemistry and physics, and use other methods to get simulation throughput. Access to M&IC has allowed us to move forward with quality models and carry out innovative and first-ever simulations.

    Indeed, the entire ASD is one of the foremost experts in applications of climate and chemistry models on large machines. Much of that expertise and experience comes from the investments made by LLNL into the M&IC and others over the last decade.

    Needs: The large calculations we carry out involve many processors over long periods of time. This sort of capability simulations requires high levels of availability and access. There is always a trade-off in many users and high access. Many users typically point to short batch queues and long periods of waiting in the queue for simulations to start. This is a highly efficient use of machines, and in many situations it represents optimal use. However, for large capability runs this is not feasible. Over the past years, M&IC has wisely used large

machines for relatively few users, thereby providing high-quality access. We would like to see that continue. Indeed, simulations obtained via that mechanism have been used in Computation's External Review presentations as well as presentations to Vic Reis and other high-level DOE officials. These simulations are also those that bring fame and (hopefully) fortune.

Computational Tools for Analysis of Earthquake Effects on Dams

Charles R. Noble

1. **Scientific intent and scope**
   The main objective of this study is to perform nonlinear dynamic earthquake time history analyses on Morrow Point Dam, which is located 263 km southwest of Denver, Colorado. The project poses many significant challenges, one of which is to model the entire Morrow Point Dam/Foundation Rock/Reservoir system, which includes accurate geology topography. Other challenges include appropriate modeling of the shear keys across the contraction joints and modeling the fluid–structure interaction. Dam safety is a very important issue, and if successful, this research will greatly enhance how the U.S. Bureau of Reclamation determines the risk posed to downstream populations. In addition, this research could save the U.S. Bureau of Reclamation dam modification costs. Moreover, the tools used to accurately simulate the geology, fluid, and structure could be used for countless other projects here at the Laboratory.
2. **Project funding source, size in FTEs, PI, and major contributors**
   The current funding is $80 K from the U.S. Bureau of Reclamation for this current year, $35 K Tech Base funds, and possibly $100 K for the following year. The number of FTEs is approximately one. The PI is Charles R. Noble.
3. **Code maturity level and history**
   The codes being used for this project are NIKE3D, DYNA3D, and Paradyn. These codes have been around for many years and are quite mature in their parallelism. These codes were developed at LLNL. NIKE3D, DYNA3D, and Paradyn can be run on pretty much any system. They can be used on SGIs, Suns, and ASCI platforms, for example. We can visualize the computational results using Griz, another LLNL-developed code.
4. **Anticipated scientific advance at 2–8 teraFLOP/s**
   We might be able to solve this large problem implicitly using NIKE3D, which would be a great achievement. Typically, these problems are computed using explicit codes because of the large size of the problem. However, we would eventually like to be able to solve these problems using an implicit code, because an implicit code has other capabilities that an explicit code does not have.
5. **Anticipated scientific advance at 8–32 processor range**
   See response to question 4.
6. **Connection with experiment**
   The U.S. Bureau of Reclamation is performing seismic tests on the large concrete blocks that typically make up a concrete arch dam. These experiments will help validate the new contact algorithms being developed for this project.

7. **Requirements for scientific challenge at 2–8 teraFLOP/s**
   I currently use an 8 GB memory computer to run the implicit part of the analysis. I would need at least this to run my current model and more if I gave the model more detail. I currently use approximately 20 GB of disk space to run the seismic portion of the analysis. I can be ready to start a large simulation in a very short period of time. All of the simulations are pretty much ready to be run.
8. **Requirements for scientific challenge at intermediate capacity level**
   See response to question 7.
9. **Machine use pattern at 2–8 teraFLOP/s**
   See response to question 7. No validation of the algorithms would need to be done on such a large system. We would only need a large system to run an explicit simulation much faster or run the simulation implicitly, which would be a significant achievement for this project.
10. **Machine use pattern at intermediate capacity level**
    See responses to questions 7 and 9.
11. **Memory vs compute trade-off**
    I would not trade away memory to get more nodes.
12. **Comments on value of unclassified computing**
    Without state-of-the-art computing that we have here at the Laboratory, I would not be able to do the work I currently do here, which include seismic, impact, and blast analyses of various structures. My work here at LLNL has also included studying terrorist scenarios, which has proven important after the events of September 11, 2001. In addition, the large simulations we perform challenge our current code capability and computers, making our codes here at LLNL more robust.