


NERSC Annual Report 2002



Science-of-Scale Projects

To promote the productivity of cutting-edge science, NERSC adapts its systems to provide the resources needed by data-intensive applications and provides "Red Carpet" support to large-scale strategic projects so that they can make rapid progress. These projects typically need support such as very large computing allocations, very large scratch disk, terabytes of usable memory and support for 64-bit computing, large-scale visualization, consulting support to make effective use of these resources, good bandwidth between the resources, and large archival storage.

The projects described below show how the unique resources of the IBM SP combined with excellent consulting support have produced scientific breakthroughs. They also demonstrate the size of the resources needed for rapidly growing, data-intensive projects. Three of these projects could have used all of NERSC's resources in FY 2002.

DOE Climate Change Prediction Program

Scientists in the DOE Climate Change Prediction Program recently completed a 1,000-year run of a powerful new climate system model on a supercomputer at NERSC. The millennium-long simulation of the new Community Climate System Model (CCSM2) ran for more than 200 uninterrupted days on the IBM SP supercomputer at NERSC. The lengthy run served as a kind of "shakedown cruise" for the new version of the climate model and demonstrated that its variability is stable, even when run for century-after-century simulations. The 1,000-year CCSM2 run had a total drift of just one-half of one degree Celsius, compared to older versions with two to three times as much variance.

A 1,000-year simulation demonstrates the ability of CCSM2 to produce a long-term, stable representation of the earth's climate (Figure 1). Few if any climate models in the world can

make this claim, since all previous simulations contained drifts too large to allow complete, uncorrected simulations to 1,000 years. In addition, the simulation provides scientists with a database to analyze the variability of weather and climate on time scales ranging from interannual to interdecadal to intercentennial. Few datasets exist which are as comprehensive as the one produced during this simulation.

CCSM2 tightly couples four complex models, including atmosphere and land modeling codes developed at the National Center for Atmospheric Research (NCAR) and ocean and sea ice models developed at Los Alamos National Laboratory. Computationally, the full CCSM2 code consists of five binaries which are organized to execute concurrently within a single job. The models exchange data at various frequencies appropriate to the physical, large-scale processes being simulated. CCSM2

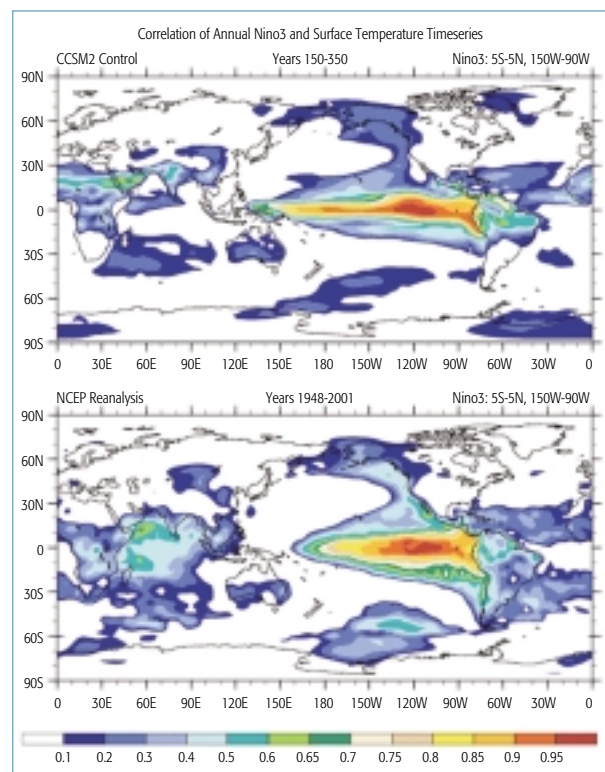


FIGURE 1 Two hundred years of modeling El Niño events and surface temperatures on the Community Climate System Model (CCSM2) closely correlate with 50 years of actual climate data.

requires 4.5 wall-clock hours on 144 1.5-Gflops CPUs of the NERSC IBM SP to complete one simulated year. NERSC gave CCSM2 special queue priority to complete this project in a timely fashion. Preliminary results of 800 model years were presented to 250 participants at the Seventh Annual CCSM Workshop held in Breckenridge, Colorado, on June 25–27, 2002.

INVESTIGATORS W. Washington, P. Gent, J. Hack, J. Kiehl, G. Meehl, and P. Rasch, National Center for Atmospheric Research; B. Semtner, Naval Postgraduate School; J. Weatherly, U.S. Army Cold Regions Research and Engineering Lab Laboratory.

PUBLICATION J. T. Kiehl and P. Gent, "The Control Climate Simulation from the Community Climate System Model (CCSM2)" (in preparation).

URL <http://www.cesm.ucar.edu/index.html>

High-Resolution Global Coupled Ocean/Sea Ice Modeling

The objective of this project is to couple a high-resolution ocean general circulation model with a high-resolution dynamic-thermodynamic sea ice model in a global context. Currently, such simulations are typically performed with a horizontal grid resolution of about 1 degree. At this resolution (about 30 to 50 km in the polar regions), the ocean model cannot resolve very narrow current systems (including fronts and turbulent eddies) that play a crucial role in the transport of heat and salt in the global ocean. Similarly, lower-resolution sea ice models cannot resolve important dynamics that occur in regions of complicated topography (such as the Canadian Archipelago).

This project is running a global ocean circulation model with horizontal resolution of approximately 1/10th degree—between 11 km and 2.5 km (Figure 2). This is the highest-resolution simulation even attempted with a such a realistic model. This configuration has dimensions of $3600 \times 2400 \times 40$, resulting in 177 million active ocean grid points (some grid points are on land). The code being used is the Parallel Ocean Program (POP), developed at LANL under the Department of Energy's CHAMMP program. At NERSC, 448 processors are used to run the model. One year can be simulated in about eight wall-clock days (86,000 processor hours), generating over 500 GB of output. Eight model years have been run to date,

with a goal of 30–50 years. After the ocean simulation has run for 10–15 model years, it will be coupled with a sea ice model to more accurately simulate the polar circulation.

The interaction of the ocean and overlying sea ice in global coupled numerical models is poorly understood, though very important. When ocean water freezes into sea ice, salt is released into the upper ocean, making it more dense. Conversely, when the ice melts, it creates a layer of fresh water that is less dense than the underlying ocean. This delicate balance between melting and freezing is very difficult to simulate with coarse grids. In particular, high vertical resolution is needed near the surface to simulate this salinity balance correctly. High horizontal resolution is required to properly simulate the current systems that advect these salinity anomalies into the open ocean. Inaccuracies in the surface ocean properties due to poor representation of ocean-ice interaction can have wide-ranging global consequences. Most notable is the possibility that too much fresh surface water can inhibit vertical convection in the northern seas (since it is less dense than the salty water beneath it), which then disrupts the entire global heat budget. Coarse-resolution simulations have found that the circulation and heat budget are extremely sensitive to the way sea ice is prescribed in ocean-only runs. The best tool for simulating the global circulation accurately is a high-resolution, fully coupled ocean-sea ice model.




FIGURE 2 High-resolution (1/10 degree) POP ocean model currents at 50m depth. Blue = 0; red > 150 cm/s.

INVESTIGATORS M. E. Maltrud and E. C. Hunke, Los Alamos National Laboratory; J. L. McClean, Naval Postgraduate School.

PUBLICATION In preparation.

URL <http://www.lanl.gov/orgs/t/t3/codes/pop.shtml>

Supernova Explosions and Cosmology

This collaboration brings together the SciDAC Supernova Science Center and the members of the PHOENIX/SYNPOL collaboration. The goal is a better understanding of supernovae of all types through simulation and model validation. Specific objectives are to clarify the physics of supernova explosions, to improve the reliability of such explosions as calibrated standard candles, and to measure fundamental cosmological parameters. Despite decades of research and modeling, no one understands in detail how supernovae work. The problem persists largely because, until recently, computer resources have been inadequate to carry out credible multi-dimensional calculations.

On June 4, 2002, at the American Astronomical Society meeting in Albuquerque, N.M., Michael Warren and Chris Fryer from Los Alamos National Laboratory presented the results of one of several projects in this collaboration, the first 3D supernova explosion simulation, based on computation at NERSC (Figure 3). This research eliminates some of the doubts about earlier 2D modeling and paves the way for rapid advances on other questions about supernovae.

Earlier one-dimensional simulations of core-collapse supernovae almost always failed to explode. Two-dimensional simulations were qualitatively different from 1D, leading to a robust explosion without fine-tuning of the star's physical properties. They showed that the explosion process is critically dependent on convection, the mixing of the matter surrounding the iron core of the collapsing star. It was believed that the results could again be changed radically by adding a third dimension, but the 3D simulations turned out to be similar to the 2D results. The explosion energy, explosion time scale, and remnant neutron star mass do not differ by more than 10 percent between the 2D and 3D models. With these 3D results, researchers are ready to attack more exotic problems that involve rotation and non-symmetric accretion.

The 3D simulation used a parallel smooth particle hydrodynamics (SPH) code coupled with a flux-limited diffusion radiation transport. Supernova calculations are computationally demanding because many processes, involving all

four fundamental forces of physics, must be modeled and followed for more than 100,000 time steps. Typical simulations (1 million particles) took about three months on the IBM SP at NERSC.

In the next five years, the Supernova Cosmology Project and the Nearby Supernova Factory experiments will increase both the quality and quantity of observational supernova data at low and high redshift by several orders of magnitude. The purpose of these experiments is to improve the use of supernovae as tools for cosmology by determining the underlying physics behind these catastrophic events and to utilize these tools to help us understand the dark energy that drives the acceleration of the universe. The only way to fully exploit the power of this amazing data set is to make a similar order-of-magnitude improvement in computational studies of supernovae, via spectrum synthesis and radiation hydrodynamics. The focus of the PHOENIX/SYNPOL collaboration's portion of this project is to start the process of creating 3D spectrum synthesis models of supernovae (Figure 4) in order to constrain the observations and place limits on the explosion models and progenitors of supernovae using the full-physics 1D models as a guide.

Currently two sets of spectrum synthesis codes, PHOENIX and SYNPOL, are used at NERSC to study the model atmospheres of supernovae. PHOENIX models astrophysical plasmas in one dimension under a variety of conditions, including differential expansion at relativistic velocities found in supernovae. The current version solves the fully relativistic radiative transport equation for a variety of spatial boundary conditions in both spherical and plane-parallel geometries for both continuum and line radiation simultaneously and self-consistently using an operator splitting technique. PHOENIX also solves the




FIGURE 3 Computer visualization shows (left to right) three stages of a simulated supernova explosion over a period of 50 milliseconds, starting about 400 milliseconds after the core begins to collapse. The surfaces show the material which is flowing outward at a speed of 1,000 kilometers/second. Left is the initial spherical implosion. Center, as in-falling gas approaches the core, it is exposed to a higher and higher influx of neutrinos that heat the gas and make it buoyant. Right, as more cold gas sinks in, it is heated and rises, resulting in enough convective energy transfer to create an explosion.

full multi-level non-local thermodynamic equilibrium (NLTE) transfer and rate equations for a large number of atomic species (with a total of more than 10,000 energy levels and more than 100,000 primary NLTE lines), including non-thermal processes. PHOENIX accurately solves the fully relativistic radiation transport equation along with the non-LTE rate equations (currently for ~150 ions) while ensuring radiative equilibrium (energy conservation).

SYNPOL is a 3D radiative transfer code to study the spectropolarimetry of supernovae. It is based on a Monte Carlo treatment of line formation via the Sobolev approximation and includes electron scattering. Because SYNPOL does not solve rate equations and does not do continuum transfer, it is not used for quantitative abundance determinations or for absolute flux calculations. Rather its value lies in establishing line identifications (the intervals of ejection velocity within which the presence of particular ions is detected) and in probing the geometry of the supernova and its ejecta. For a full 3D run, with signal-to-noise and resolution an order of magnitude greater than the observational data, approximately 10^{12} photons are generated within a Cartesian grid of 300 per side. Due to the size of the atomic data—over 42 million lines whose strengths can vary at each cube in the grid—the memory requirements and the time it takes to process the scattering of such a large number of photons are quite large: 1 million CPU hours for a 3D simulation with simplified physics, and 10 GB input and 1 GB

output per iteration, with 20 iterations per star model for 20 to 30 models.

NERSC provided a new 24-hour run queue to accommodate this simulation. Within the next two or three years, 100 times more CPUs will be needed to run 3D simulations with complex physics if there are no algorithmic improvements.

INVESTIGATORS P. Nugent, D. Kasen, and S. Perlmutter, Lawrence Berkeley National Laboratory; S. Woosley and G. Glatzmaier, University of California, Santa Cruz; P. Hauschildt, J. Aufdenberg, C. Shore, A. Schweitz, and T. Barman, University of Georgia; E. Baron, University of Oklahoma; T. Clune, Goddard Space Flight Center; A. Burrows, S. Hariri, A. Hungerford, P. Pinto, H. Sarjoughian, and B. Ziegler, University of Arizona; T. Evans, C. Fryer, M. Gray, W. Miller, and M. Warren, Los Alamos National Laboratory; F. Dietrich and R. Hoffman, Lawrence Livermore National Laboratory.

PUBLICATIONS C. L. Fryer and M. S. Warren, "Modeling core-collapse supernovae in three dimensions," *Astrophysical Journal* **574**, L65 (2002).
C. L. Fryer, D. E. Holz, and S. A. Hughes, "Gravitational wave emission from core collapse of massive stars," *Astrophysical Journal* **565**, 430 (2002).

D. N. Kasen, P. E. Nugent, L. Wang, and D. A. Howell, "Interpreting supernova polarization spectra," *Bulletin of the American Astronomical Society* **33**, 8409.

URLS <http://www.supersci.org/>
<http://www.lbl.gov/~nugent>

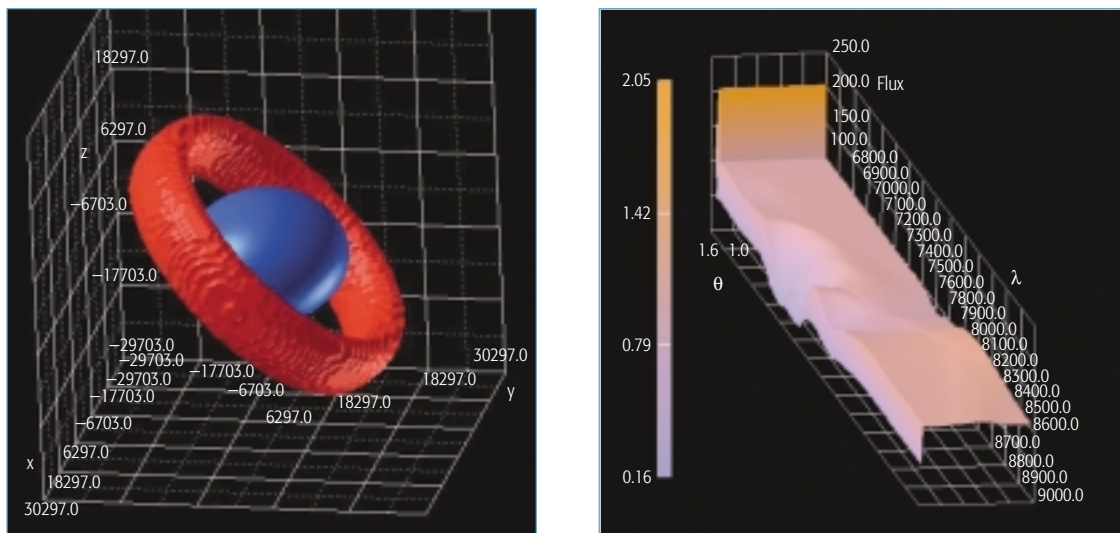


FIGURE 4 A spectrum synthesis calculation of a supernova atmosphere surrounded by a toroid. The layout of the atmosphere is presented on the left, while at the right is a graph of the flux vs. wavelength vs. viewing angle. As the viewing angle shifts towards the toroid, the strength of the absorption increases dramatically. Data that confirm such a model would for the first time put strong constraints on the progenitors of Type Ia supernovae. Such flux features are seen in the spectrum of SN 2001el.

Black Hole Merger Simulations

Physicists at the Max Planck Institute for Gravitational Physics are performing simulations of the spiraling coalescence of two black holes, a problem of particular importance for interpreting the gravitational wave signatures that will soon be seen by new laser interferometric detectors around the world. Detection of the first gravitational waves (or failure to do so) will strongly test Einstein's Theory of General Relativity, the results of which will have ramifications that extend throughout the world of physics. The Cactus simulation code is being used to perform the calculations. This is the first time ever that a spiraling merger of this type has been accurately simulated (Figure 5). The results so far indicate that the Meudon model for coalescence seems to match the simulation data more accurately than the competing Cook-Baumgarte model.

Collisions between black holes should theoretically create propagating gravitational waves, similar to the electromagnetic waves given off by distant stars. These ripples in space-time should be seen as subtle variations in the length of objects as they move through space. Recently built laser interferometric detectors such as LIGO and VIRGO are capable of measuring these subtle ripples in space. However, the gravitational wave signal that can be detected by these interferometers is so faint that it is very close to the level of noise in these devices. So simulations of the kinds of events that might produce gravitational waves can provide important insights into the gravitational wave signature produced by these events, potentially making the instruments more productive.

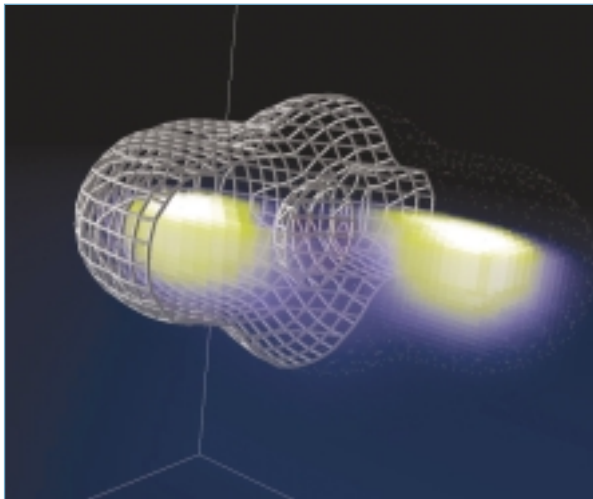


FIGURE 5 Visualization of binary black hole inspiral.

The Cactus code performs a direct evolution of Einstein's equations, which are a system of coupled nonlinear elliptic hyperbolic equations that contain thousands of terms if fully expanded. Consequently, the simulation resource requirements are enormous just to do the most basic of simulations. The simulations have been limited by both the memory and CPU performance of supercomputers as they attempt to move from calibrating against analytic black hole solutions to non-analytic astrophysically relevant cases in full 3D. The spiraling merger is just such a non-analytic case.

These simulations must use more than one-third of the NERSC IBM SP's available aggregate memory of 4.3 TB in order to achieve the resolution required to accurately simulate these phenomena. This simulation uses 1.5 TB of memory and more than 2 TB of disk space for each run. These runs typically consume 64 of the large-memory nodes of the SP (a total of 1,024 processors) for 48 wall-clock hours at a stretch. The simulation can use all 184 nodes, but this would only allow simulations that are fractionally larger than using the large-memory nodes due to memory/load-balancing issues.

NERSC provided access to a special queue to improve turnaround, opened ports to allow remote steering and Grid access, and provided consulting support for 64-bit integration and code debugging. In the space of two months, this simulation consumed 700,000 CPU hours, simulating three-fourths of a full orbit before coalescence. In the near future, this project could use 10 TB of disk for each run, 5 TB of uniform, user-available memory, and 15 million CPU hours.

INVESTIGATORS E. Seidel, M. Alcubierre, G. Allen, B. Brügmann, P. Diener, D. Pollney, T. Radke, and R. Takahashi, Max Planck Institute for Gravitational Physics; J. Shalf, Lawrence Berkeley National Laboratory.

PUBLICATIONS M. Alcubierre, B. Brügmann, P. Diener, M. Koppitz, D. Pollney, E. Seidel, and R. Takahashi, "Gauge conditions for long-term numerical black hole evolutions without excision," *Phys. Rev. D* (in press). M. Alcubierre, B. Brügmann, P. Diener, F. Guzman, S. Hawley, M. Koppitz, D. Pollney, and E. Seidel, "Dynamical evolution of binary black hole data," in preparation (2002).

URL <http://www.aei-potsdam.mpg.de/>