

Overview of NLCF FY 2006 Allocations

Arnold Tharrington*, **O. E. Bronson Messer**, and
Forrest Hoffman; Oak Ridge National Laboratory

ABSTRACT: The National Center of Computational Sciences is the National leader in capability computing and is designated as the National Leader Computing Facility (NLCF). Breakthrough computational scientific research is being achieved by world class researches in their respective fields. We present an overview of the FY 2006 projects allocated on NLCF. A more detailed overview is given for the projects with the 2 largest allocations.

KEYWORDS: NLCF, Allocations, and NCCS.

1. Introduction

The National Center of Computational Sciences (NCCS) was established in 1992 and in 2004 it was designated by the Secretary of Energy as the Leadership Computing Facility (LCF) program of the nation. The core computational resources consist of the Cray X1E (Phoenix) - 1024 multi streaming vector processor (MSP); the Cray XT3 (Jaguar) - 5294 nodes; and the SGI Altix (RAM) - 256 Intel Itanium2 processors.

In 2005, LCF was opened to scientist to perform scientific computational studies. Approximately 90% of the computer time is LCF allocated and the remaining 10% is INCITE allocated. The LCF computer allocations are open to scientific research consistent with the missions of the Department of Energy (DOE) Office of Science. INCITE computer allocations are open to scientific research for DOE, other government agencies, academia, or industry. The projects are in a broad range:

- Climate and Carbon Research (3)
 - Climate-Science Computational End Station Development and Grand Challenge Team - W. Washington, National Center for Atmospheric Research (LCF)

- The Role of Eddies in the Thermohaline Circulation - P. Cessi, Scripps Institution of Oceanography (LCF)
- Eulerian and Lagrangian Studies of Turbulent Transport in the Global Ocean - S. Peacock, University of Chicago (LCF)
- Astrophysics (3)
 - Multi-dimensional Simulations of Core-Collapse Supernovae (Mezzacappa) - A. Mezzacappa, Oak Ridge National Laboratory (LCF)
 - Multi-dimensional Simulations of Core-Collapse Supernovae A. Burrows - University of Arizona (LCF)
 - Ignition and Flame Propagation in Type Ia Supernovae - Stan Woosley; University of California, Santa Cruz (LCF)
- Combustion (1)
 - High-Fidelity Numerical Simulations of Turbulent Combustion - Fundamental Science Towards Predictive Models, Jackie Chen, Sandia National Laboratories (LCF)
- Nuclear Physics (1)
 - Ab initio Nuclear Structure Computations, David J. Dean - Oak Ridge National Laboratory (LCF)
- Computational Biology (2)
 - Next Generation Simulations in Biology: Investigating Biomolecular Structure, Dynamics

* Corresponding author

and Function through Multi-scale Modelling - P. K. Agarwal, Oak Ridge National Laboratory (LCF)

- Molecular dynamics simulations of molecular motors - M. Karplus, Harvard University (INCITE)

➤ Fusion (4)

- Gyrokinetic Plasma Simulation - W. W. Lee, Princeton Plasma Physics Laboratory (LCF)
- Exploring Advanced Tokamak Operating Regimes Using Comprehensive GYRO Gyrokinetic Simulations, - J. Candy, General Atomics (LCF)
- Simulation of Wave-Plasma Interaction and Extended MHD in Fusion Systems - D. B. Batchelor, Oak Ridge National Laboratory (LCF)
- Interaction of ETG and ITG/TEM gyrokinetic turbulence, - R. Waltz, General Atomics (INCITE)

➤ Theoretical Chemistry (1)

- An Integrated Approach to the Rational Design of Chemical Catalysts - R. Harrison, Oak Ridge National Laboratory & University of Tennessee, (LCF)

➤ Materials and nanomaterials theory (2)

- Predictive Simulations in Strongly Correlated Electron Systems and Functional Nanostructures, T. Schulthess - Oak Ridge National Laboratory (LCF)
- Direct Numerical Simulation of Fracture, Fragmentation and Localization in Brittle and Ductile Materials - M. Ortiz; California Institute of Technology (INCITE)

➤ High Energy Physics (2)

- Monte Carlo Simulation and Reconstruction of CompHEP-produced Hadronic Backgrounds to the Higgs Boson Diphoton Decay in Weak-Boson Fusion Production - H. Newmann, California Institute of Technology (LCF)
- Computational Design of the Low-loss Accelerating Cavity for the ILC - K. Ko, Stanford Linear Accelerator Center (LCF)

➤ Computer Science (1)

- Performance Evaluation and Analysis Consortium (PEAC) End Station - P. H. Worley, Oak Ridge National Laboratory (LCF)

➤ Industry (2)

- Real-Time Ray-Tracing - Evan Smyth, Dreamworks (INCITE)
- Development and Correlations of Large Scale Computational Tools for Flight Vehicles - M. Hong, The Boeing Company (INCITE)

Here we highlight two projects that benefit significantly from the computational resources of NLCF. Both projects involve simulations that are multi physics and multi scale making their realization on less capable platforms difficult. Not coincidentally, these two projects have been awarded the largest LCF allocations in FY 2006.

2. Climate-Science Computational End Station Development and Grand Challenge Team

Climate science is the study of long term weather, providing a statistical description of atmospheric and oceanic conditions over periods of weeks to millennia. The central challenge of modern climate science is the analysis and prediction of global change, planet-wide alterations of the Earth's environment, including the climate, vegetation productivity, ocean and water resources, atmospheric chemistry, and ecological systems. Global change is important nationally and internationally since such change has the potential to alter the capacity of Earth to sustain life.

Climate scientists employ complex models of processes affecting Earth's climate-including atmospheric and oceanic circulation, atmospheric radiation, land surface biophysics, land and ocean biochemistry, and atmospheric chemistry - in order to predict possible outcomes of natural and anthropogenically induced climate change. One important phenomenon affecting radiative forcing is the greenhouse effect. First proposed by the French mathematician Joseph Fourier in 1827, the greenhouse effect is the warming of the surface and atmosphere by trapped radiation (Figure 1).




Figure 1. Diagram of the greenhouse gas effect (Adapted from www.pewclimate.org).

The Earth receives energy from the sun in the form of shortwave radiation. Approximately 30% of this radiation is reflected back into space while 70% is absorbed by the oceans and land masses. This absorbed energy is then reradiated as long wave (thermal) radiation that is subsequently absorbed and reradiated by atmospheric gases, most notably water vapor, carbon dioxide, and methane.

The greenhouse effect is thought to maintain temperatures some 33 °C above what it would be otherwise. However, the temperature of the lower atmosphere has increased by about 0.6 °C averaged globally during the last century [1]. The consensus among the scientific community is that the warming trend is due, at least in part, to increased concentrations of CO₂ and other radiatively active gasses resulting from industrial activities and fossil fuel emissions. Applying the latest scientific knowledge and model results in the assessment of these effects is the job of the United Nation's Intergovernmental Panel on Climate Change (IPCC), which produces an internationally reviewed report every ten years or so describing their findings.

Since NCCS's inception, it has supported a large fraction of the climate system model runs supporting IPCC assessments in the United States. Significant computing resources are required by the complex parallel process models in order to simulate Earth's climate over decades to centuries, and ensembles of individual scenario runs are required to achieve statistically significant results. The Department of Energy (DOE) continues to support these activities by providing a sizable allocation of NCCS computer and personnel resources.

The Climate-Science Computational End Station (CCSE) project at the NCCS focuses on climate prediction based on various scenarios of anthropogenic emissions and energy policies. The CCSE Development and Grand Challenge Team is lead by Dr. Warren Washington, the principal investigator, along with co-principal investigators from a variety of DOE, NASA, and NSF laboratories and universities. The co-principal investigators are Drs. John Drake, Donald Anderson, Jr., David Bader, William Collins, Robert Dickinson, David Erickson, Peter Gent, Steven Ghan, Jim Hack, Philip Jones, Robert Malone, and William Schlesinger. The goals of the climate end station are 1) to simulate the dynamic ecological and chemical evolution of the climate system, 2) to deliver a next-generation climate model in three years, and 3) to develop and support the Community Climate System Model (CCSM) for use in climate simulation experiments [2,3].

The CCSM [3] consists of four individual component models--representing the atmosphere, the land surface, the oceans, and sea ice--connected by a coupler (Figure 2).

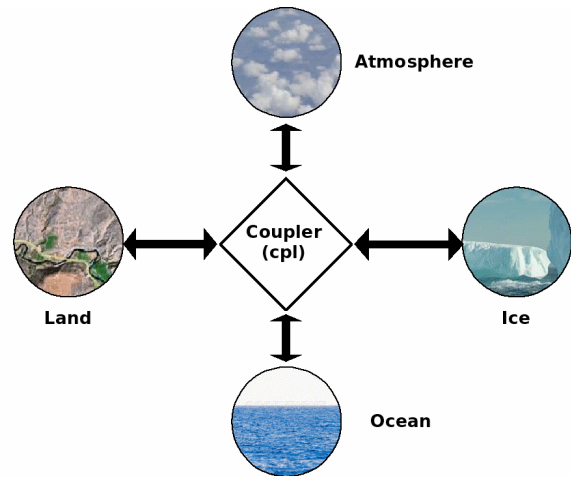


Figure 2. Schematic diagram of the CCSM framework.

When run together these models yield fully coupled simulations; however, the CCSM framework also provides for various configurations of both active and data component models and allows for easy interchange with alternate component models.

The atmospheric general circulation model (GCM) component in CCSM is the Community Atmosphere Model (CAM). CAM Version 3.0, the fifth generation GCM from the National Center for Atmospheric Research (NCAR), simulates atmospheric phenomena over disparate time and length scales. The computational approach is to divide the simulation into two phases: dynamics and physics. CAM consists of three alternative dynamics formulations (or dynamical cores), Eulerian, Semi-Lagrangian, and Finite Volume, which use different numerical techniques for simulating atmospheric diffusion and advection. A single physics package simulates convection, clouds and radiation, boundary layer processes, and chemistry.

The Parallel Ocean Program (POP) from Los Alamos National Laboratory serves as the ocean component model. This model solves the three-dimensional primitive equations for fluid motions on a sphere under hydrostatic and Boussinesq approximations. Spatial derivatives are computed using finite-difference discretizations which are formulated to handle any generalized orthogonal grid on a sphere. Various parameterizations account for subgrid scale processes.

The land surface is simulated by the Community Land Model (CLM). This model operates on the atmospheric grid, but represents various land surface biogeophysical processes using a nested subgrid hierarchy consisting of landunits, snow/soil columns, and plant functional types (PFTs) (Figure 3). The land surface interacts with the atmosphere model through the coupler providing albedo/energy feedbacks and the exchange of water, dust, carbon dioxide, and biogenic compounds.

The Community Sea Ice Model (CSIM) simulates sea ice thickness, thermal properties, advection, and other ice properties using an elasto-viscous plastic (EVP) approximation. The state variables of CSIM (Table 1) variables are evolved spatially and temporally by fundamental equations.

Symbol	Description
A_n	Sea ice area (fraction from 0 to 1)
V_{nl}	Sea ice volume per unit area(m)
E_{nl}	Sea ice internal energy per unit area ($J m^{-2}$)
V_{sn}	Snow volume per unit area (m)
T_{sn}	Surface Temperature of Snow/Ice ($^{\circ}C$)
\mathbf{u}	Sea ice velocity ($m s^{-1}$)
σ_{ij}	Stress Tensor components ($N m^{-1}$)

Table 1

The proposed computations are High Resolution Simulations and Improvements in the Atmosphere Land System (HRSAL), High Resolution Ocean Ice Simulations (HRSOI), and Biogeochemical Simulations with Carbon and Chemical Cycles (BSCC).

The HRSAL simulations will explore the benefits of running the atmosphere and land model at higher resolutions using the atmosphere component (CAM model). Changes in resolution affect the scales of motion available to the explicit solution of the governing equations - the governing equation will need reparameterization. The higher resolution makes the physics and dynamics of the simulations become more realistic and the physical processes may differ markedly from the lower resolution studies.

The HRSOI simulations will study thermohaline circulation over long time scales and small spatial scales using the ocean model coupled to an active ice model. Thermohaline circulation is a global ocean circulation and

it is driven by differences in the density of the sea water which is controlled by temperature (thermal) and salinity (haline). Thermohaline circulation is responsible for carrying significant amount of heat from the equator to the poles. Prior simulations have demonstrated that high resolutions are needed to adequately represent the ocean currents.

BSCC simulations will use the CCSM framework in order to understand how fossil fuel input, land coverage, and aerosol feedback will affect Earth's climate. These simulations will require several pre-calibration and equilibrium runs before the fully coupled model can be brought to bear upon the problem.

The primary end significance of CCSE project is the culmination of large-scale data archives (100 TB) of climate simulations that will be available for the larger scientific community. Secondly, the CCSM model will be strengthened with greater spatial resolutions, integration of biogeochemistry, dynamic simulations, atmospheric chemistry, etc. - the simulations will become more realistic. Finally, climate change is important to nation and world. The simulations and model development proposed in this end station will provide the best tools as to date in making rational, energy policy decisions that will benefit the sustention of planetary life.

3. Multi dimensional Simulations of Core-Collapse

This project focuses on multi-dimensional simulations of the core collapse during their supernova stage. The Principal Investigator is Dr. Anthony Mezzacappa. The co principal investigators are Drs. John Blondin, Stephen W. Bruenn, Christian Cardall, David J. Dean, John C. Hayes, W. Raphael Hix, Eric Myra, Jirina Stone, and Douglas Swesty [5].

Massive stars (*i.e.*, those with main sequence masses more than about 10 solar masses), live exceedingly short lives compared to our Sun (tens of millions vs. ten billion years). The timescales associated with the formation of the stratified "onion skin" structure developed in massive stars - where an inner iron core is surrounded by shells of lighter elements, including, silicon, oxygen, helium, and hydrogen - are remarkable. For example, a hypothetical twenty-five solar mass star born 11 million years ago exhausted the supply of hydrogen in its core about 700,000 B.C. The burning of the next available nuclear fuel, helium, continued from then until about 45,000 B.C., about the time Homo sapiens began to appear on Earth. The dawn of agriculture corresponds to core carbon ignition in our imaginary star (~ 10,000 B.C.). Neon

burning started approximately two decades ago; oxygen ignition occurred only half a decade ago; silicon burns for only one week and a half as the star approaches its spectacular death. It is important to note that the outer layers of the star, beyond the silicon and oxygen layers, will not directly participate in the birth of the supernova. The outer layers are thousands of kilometers distant from events that will take only milliseconds. The important events from this point on all take place within the iron core of the star.




Figure 4. Diagram of a star's life cycle. (Taken from Scientific American [8])

The iron core is inert to further nuclear burning because iron sits atop the curve of nuclear binding energy. The core does cool via neutrino emission and quasi-statically contracts as the evolution continues. The neutrinos emitted at the central densities before collapse escape freely carrying away energy and thermal pressure support from the core. Most of the pressure support, however, comes from the relativistic, degenerate sea of electrons in the core. This support is also being depleted by the neutrino emission: neutrinos are produced by electron captures on protons in the core, reducing the degeneracy of the electrons. More important, however, is the reduction in pressure caused by the photodissociation of iron as the temperature increases. Shortly after reaching a central density of 10^9 g/cm^3 , unable to support itself against its own gravity, the core collapses.

When the inner core reaches 2-3 times the density of nuclear matter, it rebounds as a unit, sending pressure waves outward. These pressure waves steepen to form a shock at a radius of some tens of kilometers. This shock begins to propagate outward, heating and dissociating material as it moves to larger radii.

In most simulations, the shock stalls at a radius of $\sim 200 \text{ km}$. The shock is robbed of pressure support by two

processes: (1) dissociation losses as nuclei pass through the shock and are converted to free nucleons and helium, and (2) electron capture on the resulting free protons, producing neutrinos that are able to escape the core. The inner core, which launched the shock, begins to settle into hydrostatic equilibrium over several milliseconds. This hot, distended object is the "proton-neutron star" (PNS). The PNS radiates neutrinos of all types as it cools and contracts, becoming a canonical neutron star if the explosion is successful. Without the further propagation of the shock, the PNS is doomed to be swallowed by a forming black hole as it accretes infalling matter and grows in mass.




Figure 5. Schematic of the supernova core-collapse phenomena.

Wilson [7] found that material behind the stalled shock can be heated by neutrinos radiated by the PNS, thereby reviving the shock and leading to an explosion. However, those early results have not been replicated by any other group, or by Wilson, without including the possibility of convection. Nevertheless, the notion of reenergizing the shock through neutrino heating has become the standard paradigm of modern supernova theory. This "delayed explosion mechanism" is almost universally accepted as the way supernovae explode, at least for stars of mass > 10 solar masses.

The typical delayed explosion configuration is shown in Figure 5. The neutrinosphere is analogous to the more familiar photosphere of stellar physics, and is the radius at which the neutrino optical depth equals $2/3$. The neutrinosphere and the shock bound a region within which neutrino emission and absorption, primarily on shock dissociated nucleons, profoundly affect the local matter. There is a region of net neutrino cooling near the neutrinosphere below a region of net heating nearer the shock. The radius where cooling is balanced by heating is called the gain radius. The efficiency with which matter is

heated between the gain radius and the shock determines the success or failure of the delayed explosion mechanism. If a parcel of matter can be heated sufficiently to reverse its infall, the shock can be reenergized.

The simulations undertaken by this LCF project are designed to ultimately answer how the supernova shock wave is revived. The source of energy in a core collapse supernova is the $\sim 10^{53}$ erg of gravitational binding energy associated with the formation of the neutron star. This gravitational binding energy is released after core bounce over ~ 10 s in the form of neutrinos. Core magnetic fields may also contribute to the revival of the supernova shock, and stellar rotation might also play a role.

The proposed computations will be in 3 stages with each stage differing primarily in the neutrino transport implementation planned.

The first stage will be hydrodynamics-only simulations to understand the role of a newly discovered supernova shock wave instability. The Stationary Accretion Shock Instability (SASI), discovered recently in two-dimensional simulations by Blondin, Mezzacappa, and DeMarino [8] has introduced a new and important ingredient in the mix that will ultimately determine the supernova explosion mechanism. The SASI comes about as streamlines of infalling matter are refracted by small "kinks" in the stalled shock. These refracted streamlines lead to vorticity production, producing eddies that propagate down to the surface of the PNS. The eddies are effectively reflected at the surface of the PNS, are directed back toward the shock, and produce further kinks in the shock, leading to a non-linear feedback loop. Through the SASI it is possible to generate a significant amount of angular momentum in the stellar core even in the case where the simulations are started with a spherically symmetric initial configuration. The SASI can grossly affect the explosion morphology, and therefore, the supernova observables associated with the "shape" of the explosion[†]. NLCF resources will be used to study the dynamics of the post-bounce shock in core-collapse supernovae in three dimensions and investigate the role of non-axisymmetric modes of the SASI in the spin-up of the neutron star left behind by the supernova. Simulations of the SASI are carried out with the Eulerian hydrodynamics code VH-1, primarily on the LCF Cray X1E. Each of these three-dimensional simulations

[†] For example, it is now known that supernova light is polarized at the 2% level; the SASI offers the first fundamental explanation for this observation given that such polarization could be produced by an explosion that is a prolate, which the SASI is capable of generating.

produces between 5 and 10 TB of data per run. Planned improvements include the implementation of a Yin-Yang grid [9].

The second stage includes simulations with "ray-by-ray" neutrino transport, in which transport is restricted to radial rays and lateral transport is suppressed. This will allow three-dimensional simulations to be performed in the near future with a significant level of realism (three-dimensional simulations including lateral transport will take years to develop and complete). These simulations will be carried out by combining a variant of the VH1 (PPM) hydrodynamics code with ray-by-ray multi-frequency neutrino transport, using the multi-frequency flux-limited diffusion code (MGFLD_TRAN) of Bruenn and the Boltzmann transport code (BOLTZTRAN) developed by Mezzacappa, Liebendörfer, & Messer. These ray-by-ray (RBR) models leverage the long development of these spherically symmetric 1-D neutrino transport codes by employing them to independently calculate neutrino transport along each radial direction on a 2D or 3D spherical polar grid. While this RBR approximation is limited in its treatment of neutrino transport within the proto-neutron star [which can only be explored using full (including lateral) 2D and 3D transport like the V2D and GENASIS codes], it should capture the essential physics above the neutrinosphere, it is relatively fast, easily extensible to 3D, and an excellent probe of the behavior of the supernova shock. The hydrodynamics code used, EVH1, is similar to VH1 described above, except that a realistic EOS is employed. In addition, nuclei in nuclear statistical equilibrium (NSE) at high densities and temperatures are interfaced with a multinuclear species EOS at lower densities and temperatures where NSE cannot be assumed. A reaction network is used to evolve the nuclei in the latter region. Comparison between the MGFLD/EVH1 and BOLTZTRAN/EVH1 RBR models will shed light on the accuracy of flux-limited diffusion (used in MGFLD_TRAN) when compared to more accurate but more costly Boltzmann transport (used in BOLTZTRAN) in the multi-dimensional context. The roadmap also includes the examination of a wide variety of progenitors, investigation of the impact of improved neutrino-matter interactions and equations of state, and the inclusion of thermonuclear reaction networks to provide greatly improved predictions of supernova elemental production.

And finally, there are two-dimensional simulations planned with two-dimensional multi frequency neutrino transport. The codes necessary to accomplish these goals are still under development, and will test even the capabilities of a petaflop machine.

Literature Cited

1. "Global Warming Basics", Pew Center on Global Climate Change, Wednesday, May 3, 2006 4:32:11 PM. Online available from: <http://www.pewclimate.org/global-warming-basics> (Accessed May 5, 2006.)
2. W. Washington, J. Drake, D. Anderson, Jr., D. Bader, W. Collins, R. Dickinson, D. Erickson, P. Gent, S. Ghan, J. Hack, P. Jones, R. Malone, and W. Schlesinger. "Climate Science Computational End Station Development and Grand Challenge Team", National Center for Computational Sciences Leadership Computing Facility 2006 Proposal (2006): Project No. CLI017
3. "Global Climate Change", National Science Foundation, Tuesday, May 11, 2004 10:53:26 AM. Online available from: <http://www.exploratorium.edu/climate> (Accessed May 5, 2006)
4. "CCSM 3.0 Public Release Home Page", Community Climate System Model, May 5, 2006 3:48:25 PM. Online available from: <http://www.cesm.ucar.edu/models/ccsm3.0/#docs> (Accessed May 5, 2006 15:52 PM)
5. A. Mezzacapa, J. Blondin, S. W. Bruenn, C. Cardall, D. J. Dean, J. C. Hayes, W. R. Hix, E. Myra, J. Stone, and D. Swesty, "Multi-dimensional Simulations of Core-Collapse Supernovae," National Center for Computational Sciences Leadership Computing Facility 2006 Proposal (2006): Project No. AST015
6. A. Mezzacappa, "ASCERTAINING THE CORE COLLAPSE SUPERNOVA MECHANISM: The State of the Art and the Road Ahead," *Annu. Rev. Nucl. Part. Sci.* 55 (2005): 467–515.
7. H. H. Bethe and J. R. Wilson, "Revival of a Stalled Supernova Shock By Neutrino Heating", *ApJ* 295 (1985): 14-23.
8. S. Woosley and T. Weaver, "The great supernova of 1987", *Scientific American*, 261 (1989): 24-32.
9. J. Blondin, A. Mezzacappa, & C. DeMarino, *Astrophys. J.* 584 (2003): 971
10. A. Kageyama, and T. Sato, *Geochemistry Geophysics Geosystems*, 5 (2004): 9005.

About the Authors

Arnold Tharrington is a member of the Scientific Computing Group of the National Center of Computational Science at Oak Ridge National Laboratory. He can be reached at: Oak Ridge National Laboratory; PO BOX 2008 MS 6008 Oak Ridge, TN 3783-6008; E-mail: arnoldt@ornl.gov

O. E. B. Messer is a member of the Scientific Computing Group of the National Center of Computational Science at Oak Ridge National Laboratory. He can be reached at: Oak Ridge National Laboratory; PO BOX 2008 MS 6008 Oak Ridge, TN 3783-6008; E-mail: bronson@ornl.gov

Forrest Hoffman is a member of the Climate Dynamics Group of the Computer Science and Mathematics division at Oak Ridge National Laboratory. He can be reached at: Oak Ridge National Laboratory; PO BOX 2008 MS 6016 Oak Ridge, TN 3783-6016; E-mail: bronson@ornl.gov