5.9 MODELING ENVIRONMENT FOR ATMOSPHERIC DISCOVERY

(http://www.ncsa.uiuc.edu/expeditions/MEAD)

Robert Wilhelmson[1], Kelvin Droegemeier[2], Sara Graves[3], Mohan Ramamurthy[4], Dale Haidvogel[5], Brian Jewett[4], Jay Alameda[1], Dennis Gannon[6]

[1]National Center for Supercomputing Applications (NCSA), University of Illinois at Urbana-Champaign

[2]Center for Analysis and Prediction of Storms, University of Oklahoma

[3]Information Technology and Information Center, University of Alabama at Huntsville

[4]Department of Atmospheric Sciences, University of Illinois

[5]Institute of Marine and Coastal Sciences, Rutgers University

[6]Computer Science Department, University of Indiana

1. INTRODUCTION[1]

Research enabled by execution of high resolution simulation suites is becoming increasingly possible with growth in computing, networking, and storage capabilities. For example, factors that influence low level mesocyclone development in severe storms after mergers are being investigated utilizing hundreds of simulations[2]. Ensembles represent another type of simulation suite that has become popular in forecasting and weather prediction research.

The opportunity to carry out hundreds of simulations within a suite brings with it new computational problems. These include execution of model simulations, management of the resulting large volumes of data, and subsequent data analysis, data mining, and visualization. Further, data mining and machine learning can be used to guide the selection of the set of parameters to be used in new simulations based on selected objectives within a research study. For example, if the objective is to determine what conditions (sets of parameters and environments) lead to tornadic formation, machine learning models can be constructed and modified from simulations as they are being carried out. These learning models can be, in turn, used to provide information on the set of parameters needed in new simulations to optimize the desired objectives (e.g, the strength of the low level mesocyclone in a storm).

Conducting research that involves launching hundreds of simulations coupled with analysis, datamining, probability determination, and visualization (both individually and collectively) is daunting when each job must be submitted and catalogued individually, with resulting data analyzed, merged, and visualized in many painstaking steps (even with scripting). Further, the advent of Grid computing brings with it new challenges.

2. CYBERINFRASTRUCTURE: THE TERAGRID

With the advent of the TeraGrid and the recently awarded Extensible Teragrid[3], it is now feasible to conduct simulation suites (with hundreds of simulations) and subsequent analysis/visualization at many sites. The TeraGrid is an outgrowth of the NSF PACI program which has two leading edge sites, NPACI in San Diego and NCSA in Champaign/Urbana.

The TeraGrid is a plan to advance the cyberinfrastructure for 21st century science and engineering. It is a response to the pressing need for greater computational power to enable experimentation, modeling, data analysis, and visualization activities that often involve large volumes of data and use of an expanding national computational grid[4]. It is a vision to develop and deploy a comprehensive computational, data management, and networking infrastructure of unprecedented scale and capability that couples distributed scientific instruments, terascale and petascale computing facilities, multiple petabyte data archives, and gigabit and beyond networks, all widely accessible by researchers including those in the atmospheric sciences, oceanography, and hydrology.

---

1 Corresponding author address:
Dr. Robert B. Wilhelmson
NCSA
152 Computer Applications Building
605 East Springfield Avenue
Champaign, IL 61820
email: bw@ncsa.uiuc.edu

2 See Jewett et al. (Applying Portal Technologies to Ensemble Modeling of Convection on the Grid) - P1.50 in this volume

3 http://www.teragrid.org/

4 http://access.ncsa.uiuc.edu/witg/ - multimedia addressing the question: What is a grid?

Quoting from a recent press release available at the TeraGrid web site, "The Extensible Terascale Facility (ETF) award expands the TeraGrid to five sites: the National Center for Supercomputing Applications (NCSA) at the University of Illinois, Urbana-Champaign; the San Diego Supercomputer Center (SDSC) at the University of California, San Diego; Argonne National Laboratory near Chicago; the Center for Advanced Computing Research (CACR) at the California Institute of Technology, Pasadena; and the Pittsburgh Supercomputing Center (PSC) at Carnegie Mellon University and the University of Pittsburgh.

This extended TeraGrid environment will provide the national research community with more than 20 teraflops of computing power distributed among the five sites and nearly 1 petabyte (1 quadrillion bytes) of storage capacity.

The award also ensures that the TeraGrid will be extensible and ready for expansion in the future. Additional sites will be able to connect to the TeraGrid, and the national research community will be able to take advantage of its high-performance resources. NCSA, SDSC, Argonne, and CACR were already part of the three-year, $53 million TeraGrid project, announced a year ago by the NSF. PSC had previously received a $45 million NSF award to build a terascale computing system, called TCS-1. The ETF award integrates these two efforts and their computing environments to create an extended terascale-level grid of data, computation, and visualization resources that will make possible new scientific discoveries. The five sites will be linked by the world's fastest dedicated optical research network[5], built in partnership with Qwest Communications and designed to accommodate additional connections."

3. CYBERINFRASTRUCTURE: SOFTWARE

The TeraGrid or any other Grid facility infrastructure requires "middleware[6]" for authentication of users, resource management, rapid movement of data between systems, and security. The National Middleware Initiative (NMI)[7,8] funded by NSF is an effort to deploy and evaluate middleware for use by the research community. The objective of the NMI is to ease the use and deployment of such middleware, making distributed, collaborative environments such as Grid computing and desktop video-conferencing more accessible. Grid packaging tools[9] are being used to support deployment that includes the Globus Toolkit[10], CondorG[11] and the Network Weather Service, along with security tools and best practices for enterprise computing such as eduPerson[12] and Shibboleth[13]. This software is in early release form and is being tested in a number of locations.

4. CYBERINFRASTRUCTURE: MEAD[14]

The scientific advances and societal benefits associated with a better understanding and more accurate prediction of hurricanes, severe storms, and other meso- and synoptic scale events are enormous; however, presently available computational and data management frameworks have not yet incorporated the full benefits of grid and web technologies for carrying out and analyzing/mining ensemble simulations. To help address existing limitations and to remove unnecessary human effort, we are developing a scalable framework -- known as Modeling Environment for Atmospheric Discovery (MEAD) -- for use in ensemble prediction and parameter studies. With MEAD it will be possible to launch hundreds of model simulations on a variety of grid resources and manage distributed model data stores and streams. Each simulation will not have to be submitted and catalogued individually in the traditional painstaking method of research. Metadata and the resulting large volumes of data (10's to 100's of terabytes) will then be made available through the MEAD portal for analysis, datamining, machine learning, and visualization.

Key modeling components of MEAD include the Weather Research and Forecasting Model (WRF)[15] and the Regional Ocean Modeling System (ROMS)[16]. The MEAD effort includes the coupling of these two models for studying hurricanes. Datamining and machine learning will be provided through D2K (Data to Knowledge)[17] and ADaM (Algorithm Development and Mining System)[18]. Other technology to be incorporated into MEAD includes ESML (Earth System Markup

[5] 30 – 40 gigabits per second
[6] http://www.educause.edu/ir/library/pdf/erm0241.pdf
[7] http://www.nsf-middleware.org/
[8] http://www.nsf.gov/od/lpa/news/02/pr0238.htm
[9] http://www.ncsa.uiuc.edu/Divisions/ACES/GPT/
[10] http://www.globus.org/
[11] http://www.cs.wisc.edu/condor/condorg/
[12] http://access.ncsa.uiuc.edu/Briefs/01Briefs/010227.eduPerson.html
[13] http://middleware.internet2.edu/shibboleth/
[14] http://www.ncsa.uiuc.edu/expeditions/MEAD
[15] http://www.wrf-model.org/
[16] http://marine.rutgers.edu/po/models/roms/index.php
[17] http://www.ncsa.uiuc.edu/Divisions/DMV/ALG/d2k
[18] http://datamining.itsc.uah.edu/adam/

Language)[19], HDF5[20], DODS[21] with GridFTP, wall visualization[22,23] and Chimera[24].

The MEAD research environment will include model and grid workflow management, metadata creation for model simulations and subsequently derived information including visualizations, data management of very large computed and derived data sets, and analysis/mining and visualization capabilities for large distributed datasets. Working groups in MEAD have been formed around the following themes.

Workflow and Portal Interface

Workflow infrastructure supports coordination of job submission on the Grid including authentication, scheduling, monitoring, and event handling. Researchers will submit jobs through a user portal adapted from the Web-EH toolkit developed at the Center for Analysis and Prediction of Storms (CAPS). The existing Web-EH environment supports account and user login management, experiment design and input parameter configuration, and file transfers between the client, web and compute servers. Via an extended Web-EH interface, users will be able to specify model parameters for suites of ROMS/WRF simulations.

Because MEAD will generate a large number of simultaneous jobs when it is executing, an adaptive "mapper" will be used to assign these jobs to the Grid resources available. This mapping will be based on technologies developed at Rice, which rely on performance models and system monitoring to ensure the distribution and scheduling of processes. In addition, XML descriptions of model, analysis, and visualization parameters will be developed for managing the workflow.

Metadata and Data Management

MEAD data management will leverage efforts within the Alliance petascale data quest and portal expeditions and will require coupling scientific and derived metadata with appropriate grid metadata for searching and acquiring data. ESML will be used for the interchange technology to handle the multiple data formats in MEAD. Based on XML, ESML consists of a schema and an associated library for use in data management.

The ESML interchange technology effort will focus on coupling ESML tools with the data Grid toolkit, including the GridFTP protocol extensions. The data structural and semantic information described by ESML complements the data content, location, and other information to be managed by the data Grid toolkit. This will provide the MEAD system not only with access to heterogeneous data sources, but will also deal with the distributed nature inherent in linking multiple models and multiple tiled data sets. Currently DODS and Chimera are being considered for providing access to model and derived data.

Model Coupling and I/O

The MEAD coupling and I/O effort is focused on two community models developed to study convective and regional flows: ROMS and WRF. ROMS can be run on either serial or parallel computer architectures. Features include high-order advection schemes, accurate pressure gradient algorithms, several subgrid-scale parameterizations, atmospheric, oceanic, and benthic boundary layers, biological modules, radiation boundary conditions, and data assimilation. WRF is also designed for use on parallel computers and for performance portability, maintainability, extensibility, readability, usability, run-time configurability, interoperability, and reuse in a limited area model with lateral boundary conditions and nesting. A modular, hierarchical software architecture facilitates multiple dynamic cores and plug-compatible physics in a code that operates efficiently over a broad range of computing architectures.

HDF5 will be incorporated into both ROMS and WRF to provide a new format choice for storing data. This will improve portability, facilitate data and software sharing, and provide for efficient parallel I/O and 64-bit addressing. Currently WRF and ROMS utilize the netCDF format, so the addition of HDF5 access will provide new I/O capabilities for running large simulations on parallel systems.

Coupling will be done using the WRF I/O API. The API is extensible and will be packaged for importation into ROMS. Design features will allow model "layers" for the possible inclusion of sea-ice models, wave-state models, etc. Components of the Model Coupling

---

[19] http://esml.itsc.uah.edu/
[20] http://hdf.ncsa.uiuc.edu/
[21] http://www.unidata.ucar.edu/packages/dods/
[22] http://www.ncsa.uiuc.edu/TechFocus/Deployment/DBox/index.html
[23] http://geowall.geo.lsa.umich.edu/indexmain.html
[24] http://www-unix.griphyn.org/chimera/ A Virtual Data System for Representing, Querying, and Automating Data Derivation

Toolkit (MCT)[25] will be used to do both time and space grid transformations between WRF and ROMS.

Data Mining and Machine Learning

MEAD will provide tools for analysis and machine learning utilizing ROMS and/or WRF produced data. This includes links to data mining and machine learning capabilities provided through D2K and ADaM, as well as to the Grid-enabled Parallel Problems Server (PPS)[26]. D2K is a rapid application development environment developed at NCSA that provides a standardized visual programming interface and supplies the user with modules and application templates used for a variety of data mining problems. These modules and templates may then be modified to accommodate the specific requirements of the user or developer. Any modified module or application may also be saved and re-used. D2K also supports distributed computing and parallel algorithm implementation,

ADaM was developed by the Information Technology and Systems Center, in response to the need to mine large scientific data sets specifically for geophysical phenomena detection and feature extraction. It provides a variety of processing tools, which allow easy integration of spatial and temporal variables of earth science data sets. This mining system consists of a series of interoperable data readers, preprocessing and analysis modules, and data writers. These modules can be linked together in many ways to create customized mining processes. New modules can easily be added to extend the functionality of the system.

PPS couples MATLAB on the front end with the power of high performance parallel computing on the backend. Currently many of the matrix manipulation (and indexing) functions in MATLAB can be carried out in this way.

Model and Performance Analysis

Model performance analysis will be provided through HPCView[27] and Prophesy[28]. HPCView is a toolkit for combining multiple sets of program profile data, correlating the data with source code, and generating a database that can be analyzed anywhere with a commodity Web browser. It is relatively language- and compiler-independent since it analyzes compiled code. Prophesy, on the other hand, is utilized with source code. It is an infrastructure for analyzing and modeling the performance of parallel and distributed applications. These applications are instrumented at the loop, function, or other user specified way. Performance data is then collected from many runs and then an analytical performance model is creased using optimization techniques.

Performance issues are also being studied using a finite difference model for exploration of ideas that could be incorporated into the two community models (NCOMMAS[29]). The difference methods are similar to those now being used in WRF. One effort at Rice is to seek performance improvement through source to source transformations including loop unroll and jam and loop fusion.

Another is to utilize the SHMOD (Shared Memory on Disk) paradigm as developed at the University of Minnesota to develop a fault tolerant code. SHMOD uses fast bulk data transfers over a cluster network to create the effect of a shared cluster memory that is used to store data contexts corresponding to model sub-domains. Reassignment of cluster nodes is possible if they fail and new nodes can be added at any time when they become available. Task management is handled through the use of a standard database. The database contains information on where the sub-domains are located, levels to which they have been updated, and the times and completion status information of individual task assignments. Global coordination in this way frees the network nodes from coordination tasks and is a significant step to realizing fault tolerance.

Visualization

Visualization components within MEAD will be through commonly used packages such as NCAR Graphics with NCL[30], Vis5D[31], VisAD[32], MATLAB, and the new IDV[33]. The IDV brings together the ability to display and work with satellite imagery, gridded data (primarily from model output), and surface, upper air, and radar data within a unified interface. It also provides 3-D views of the atmosphere and allows users to interactively slice, dice, and probe the data to create cross-sections, profiles, animations and value readouts of multi-dimensional data sets. Computation and

[25] http://www.mcs.anl.gov/acpi/mct
[26] http://www/mersc/gpv/~parry/text/ppserver/
[27] http://www.cs.rice.edu/~dsystem/hpcview
[28] http://prophesy.mcs.anl.gov/
[29] http://www.nssl.noaa.gov/~wicker/commas.html
[30] http://www.scd.ucar.edu/vets/Resources.htm
[31] http://www.ssec.wisc.edu/~billh/vis5d.html
[32] http://www.ssec.wisc.edu/~billh/visad.html
[33] http://www.unidata.ucar.edu/projects/metapps/idv_help/IDV_intro.html

display of built-in and user-supplied formula-based derived quantities is supported as well.

In addition, the Hierarchical Volume Renderer (HVR)[34] developed at the University of Minnesota is being extended for use with WRF and ROMS data. HVR runs on ordinary Windows PCs using Open GL software rendering. The hierarchal nature of the software enables display of subjects closer to the viewer's location at high resolution, whereas subjects farther away are shown at lower resolution. This feature greatly reduces the computer resource requirements and permits truly interactive behavior for even billion zone volumes.

Wall visualization software will be provided for the 40-tile (40 projectors and computer processors) scalable display wall at NCSA that provides a display surface of 8192 x 3840 pixels. It will be possible to display fields from multiple simulations and animate them in lock step. Each field is displayed in a window that can be sized as desired including across tiles. Single images can also be displayed across all the tiles.

Finally, a GeoWall[35] application for wall stereo display will be developed at the University of Illinois at Chicago in the Electronic Visualization Lab as part of the Scientific Workstations for the Future expedition.

Education

The educational effort associated with MEAD will include the development of educational tools and curricula centered around hurricanes, predictability, and scientific workspaces. Some of the learning modules will be inquiry-based and tailored to learning standards of the K-12 community. Students will be able to enter a simplified version of the MEAD portal interface and launch simulations on the Grid. For example, they will be able to follow the movement of balloons launched within a modeled wind field and investigate the predictability of forecasting where they will go and their spread – even when they start very close to one another.

Another effort will involve a hands-on experiment with spinning tops and the predictability of where they will end up on a table. This will depend on their translation and rotation speed as well as the surface on the table.

Finally, several full model data sets will be made available through VGEE[36,37,38]. VGEE aims to use learner-constructed visualizations as an anchor for student inquiry in the geosciences. It includes a set of tools specially designed to help learners connect visualization to physical principles and a model inquiry strategy to guide learners in using these additional tools. The inquiry strategy begins with identifying and relating patterns using a specially constructed visualization environment, explaining those patterns using idealized concept models, and integrating multiple explanations into a comprehensive understanding of the phenomena using probes. The visualization environment has now been ported to IDV.

5. UPDATED PREPRINT

This preprint will be updated periodically. Updates can be found at the MEAD website. This will include more detailed information on the MEAD environment and other hyperlinks.

6. ACKNOWLEDGEMENTS

MEAD is funded through the NCSA Alliance[39] expedition program, the PACI Education, Outreach, and Training Program (EOT), and the NCSA Faculty Fellows Program. The Alliance is led by Dan Reed. Co-leads of the MEAD expedition are Robert Wilhelmson and Sara Graves. Expedition funded PI's include Robert Wilhelmson (NCSA), Sara Graves (University of Alabama at Huntsville), Kelvin Droegemeier/Dan Weber (University of Oklahoma), Dale Haidvogel (Rutgers University), Paul Woodward (University of Minnesota), Bryant York (Portland State University), Charles Lee Isbell, Jr. (Georgia Tech), Parry Husbands (Lawrence Berkeley National Laboratory), and Rob Fowler (Rice University). Related participants include Mike Folk, Kent Yang, Al Rossi, Shawn Hampton, Mark Straka and John Towns (NCSA), Steve Tanner and Rajul Ramachandran (University of Alabama at Huntsville), Chris Moore (PMEL), Sarah Anderson and David Porter (University of Minnesota), and Anirban Mandal (Rice University).

Other expedition partners include Ian Foster and Veronika Nefedova (Argonne National Laboratory – co-lead and coordinator, respectively, of the Petascale Data Quest Expedition), Dennis Gannon (University of Indiana – co-lead of the Portal Expedition), Valerie Taylor (Northwestern University – co-lead of the Performance Engineering for Clusters and Grids Expedition), Rick Stevens (Argonne National Laboratory – co-lead of the Scientific Workspaces for

---

[34] http://www.lcse.umn.edu/hvr/hvr.html
[35] www.geowall.org
[36] http://www.unidata.ucar.edu/projects/metapps/webstartTest/VGEE/
[37] http://www.dlese.org/vgee/
[38] http://ww2010.atmos.uiuc.edu/(Gh)/abt/pubs.rxml
[39] http://www.ncsa.uiuc.edu/About/Alliance/

the Future Expedition) and Jason Leigh and Tom Defanti (University of Illinois at Chicago). EOT involvement includes Scott Lathrop (NCSA), Bob Panoff (Shodor Foundation), and Steve Gordon (Ohio State University). Mohan Ramamurthy and Brian Jewett are NCSA Faculty Fellows. Other collaborators include Joe Klemp, John Michalakes, and Don Middleton (NCAR), Dan Schaeffer and Tom Henderson (FSL/NOAA), Rob Jacob (Argonne National Laboratory), Lou Wicker (National Severe Storms Laboratory), Bill Hibbard (University of Wisconsin), and Michael Welge, David Tcheng and Dave Semeraro (NCSA).

MEAD's will coordinate with the DOE funded Earth System Grid (ESG)[40] effort lead by Ian Foster (ANL), Don Middleton (NCAR), and Dean Williams (LLNL). The primary goal of ESG is to address the formidable challenges associated with enabling analysis of and knowledge development from global Earth System models. While ESG is focusing on global coupled models, MEAD will focus on mesoscale coupled systems.

[40] http://www.earthsystemgrid.org/