Long Term Climate Simulations Using an Adaptable Virtual Machine.

J.B. Drake G.A. Geist

October 28, 1997

1 The scientific problem and DOE's interest.

Climate change research has become increasingly dependent on large computations as the demand for realism in climate models has grown. The latest generation of models, as typified by the NCAR Climate System Model (CSM), no longer requires correction of fluxes between the atmosphere and the ocean to avoid climate drift. This was achieved by increasing the resolution and improving the physical models thus increasing the computational complexity.

The DOE is primarily interested in predictive capability for climate in the decadal and century time scales. Simulations involving coupled ocean, atmosphere, land, carbon and ice models are needed which integrate for hundreds of years. This capability has a direct relationship to our ability to scientifically assess the effects of greenhouse gases on changing global and regional climates as well as our scientific understanding of climate feedbacks and fundamental processes which make up the climate system.

The present computing situation as an obstacle to the science.

With the dramatic increase in computational speed of both supercomputers and PC's it is rather odd that the climate research community is lacking in computational resources. While it is true that climate research can probably absorb any foreseeable increase in computational capability, the present situation is that existing models are ahead of available resources. Computer systems capable of hundreds of Gigaflop/sec performance are not readily available. This is primarily due to lack of high performance computer vendors and a relatively small market for such machines in comparison with the business and home PCs. The vendors have essentially instructed the scientific community to build their own machines. The powerful machines in existence now are all based on multiple processors coordinated over a fast network. Clusters of shared memory machines are connected to each other with message passing systems. The emergence of standards for message passing protocols and the existence of compiler technology able to exploit shared memory parallelism are the bright points in this otherwise untenable situation.

2 The rise of middle ware and the network computer

To meet the challenge of utilizing multiple heterogeneous processors a class of software has arisen which attempts to hide the network fabric from the scientific user. In the coupled climate models, such as the CSM, a software layer has been developed that handles the communication of control information and flux data between the component models. In the CSM this is called the MCL- the Model Coupling Library. The MCL uses Ethernet, hippi channels or other specialized communication The default message passing system in the MCL is PVM (Parallel Virtual Machine) developed by on of the PI's to this proposal. The CHAMMP sponsored projects at Oak Ridge, at the National Center for Atmospheric Research and Los Alamos have successfully used the MCL adapted to Cray's, SGI clusters, networked workstations and Intel Paragons to perform coupled climate simulations. Increasingly, the resources of several computers, and possibly several computing centers, are required to perform the critical simulations needed in the development of new models and in assessment activities.

Problems for long running simulations

With this mix of machines and centers together with the complexity of running ocean and atmospheric models coupled in parallel has come a coordination nightmare. Even building the models on the different machines and linking with libraries which are not in a standard location requires careful attention. Accommodating the usage policies and machine schedules for coupled runs is also difficult. For a long running simulation, checkpoint and restart capabilities are absolutely essential because the reliability of a networked computer is very low. A node can go down, particularly if it is a multiuser system, or the network can become congested and fail. Currently there are no recovery mechanisms except to restart from the latest check point. If this were not the only way to achieve throughput, it would have been abandoned. In comparison to the single machine shops of the last decade, the current computing environment is hostile to long running production jobs.

Fault tolerance and adaptability

At a minimum the environment needs to have a measure of fault tolerance and adaptability. For a simulation that may take several months to complete repeated restarts are a major distraction to the scientist. Over such a long period, the computing environment is likely to change as well. New machines may come online which offer significant performance advantages. New network fabrics may be enabled or network performance altered. To effectively exploit new hardware currently requires re-engineering the the code. In particular, the MCL configuration must be modified to let the simulation use new, fast channels. The ability to perform long running production climate simulations in this rapidly changing computing environment requires software systems that are both fault tolerant and adaptable. Such a system does not at present exist.

Design to the hardware replaced by design to the middle ware.

What we see as the next generation scientific computing paradigm is a shift from machine specific coding to design around a virtual machine. The virtual machine is a software layer that handles the specifics of machine and network. It is not a seamless environment; that is too general. Rather, the virtual machine concept allows specific local capabilities to be added to an individuals virtual machine in a coordinated and user controlled manner. The distributed computing is particularly vulnerable to network problems. Beyond simple availability there are real time performance issues that must be considered for long running jobs. When a shared network becomes conjested the climate simulation will be impacted. In a rich network, one with sufficient connections, the conjestion is often a result of a static routing. The virtual machine should provide the capability to dynamically add links or exploit faster connections when required. The complexity of the network, its performance and its reliability, cannot be effectively handled within the application. Nor can the scientist continually interface with computer operators and network operators at remote sites to ensure that the climate simulation is progressing. Another method of supporting the production climate runs must be found.

A software layer should control the definition of the virtual machine which could then adapt to exploit network and machine capabilities. Scientific application software would then be developed to meet the interface standards for this virtual machine API. In this way, for example, the climate models would become machine/vendor independent and portable allowing for single source development.

The HARNESS project

The HARNESS project offers a new conceptual design and machine paradigm that will make possible a stable and dynamic computing environment for the execution of long running jobs. HARNESS stands for Heterogeneous Adaptable Reconfigurable NEtworked SystemS. At present this project is unfunded and in proposal form (included) to the DOE MICS office. But it is the first distributed computing and networking proposal for middle ware that incorporates the right computer science approaches with an application-centric view. It is proposed by the developers of PVM as the next generation tool for distributed computing. Since the PVM software proved so useful to climate modeling there good reason to expect that HARNESS will become a central building block of the next generation models. What we propose is to collaborate with the HARNESS project in its early phases so that the goals of the Climate Change Prediction Program (CCPP formerly CHAMMP) will be complemented. In short, to ensure that HARNESS is immediately usable for long running climate simulations.

3 A case study: Coupled climate model

As a result of a CHAMMP Model Development Team project, the NCAR Climate System Model, the NCAR Community Climate Model (CCM3) and the Los ALamos Parallel Ocean Program (POP) were implemented in a general parallel heterogenous computing environment. These codes (and the new emerging models) will be adapted to the HARNESS framework without much change. *The goal of the project will be to perform a coupled climate simulation with the CSM which will run continuously (without restart) for 1 month adapting to changing computer and network hardware.* We hope to demonstrate simple execution with the HARNESS system within the first nine months and gradually increase the stress on fault tolerance, adaptability and performance.

Decadal to Century long simulations resource requirements

To perform a meaningful simulation with the CSM, such as the 125 year increasing CO2 simulation recently run on a 32 processor NEC SX-4, requires significant computational resources. The simulation mentioned required approximately 8 days of dedicated time. The 32 processor SX-4 vector supercomputer is rated at a peak performance of 64 Gflops/sec. The data generated from this run occupies 170 Gbytes. (for further information about this interesting simulation see (http://www.ucar.edu/staffnotes/9705/scibrief.html)

The CSM Model Coupling Layer

The critical element of the CSM which must be rewritten is the MCL. The capabilities of HARNESS will be used to define the network functions in the MCL. The control network for coordination of tasks must be adapted to fault tolerance and a collective view of the status of the simulation. If one processes fails, this cannot be allowed to destroy the continuity of the simulation, even if this process is the driver which functions as the flux coupler. The data network links must also be modifiable and specifiable through the HARNESS virtual machine. This will allow adding high speed links even after the simulation has been started. We propose to treat the MCL as a special, application specific, *plug in* to HARNESS. This means that coupled models use a specialized communication fabric that encapsulates how model components need to communicate but under the management of the HARNESS system. Thus the fault tolerance and adaptability features would not be coded inside the MCL but rather in the HARNESS framework.

Implications for climate simulations.

The implications for how climate simulations are managed in this next generation model are very broad. Once a virtual machine is defined for a given simulation, it could persist indefinitely. This offers even more stable computing environments than are available at a single machine computing center. Performance tuning of application code could be somewhat replaced by performance tuning of the virtual machine in a real time setting. The performance of the machine could be optimized to achieve optimal performance for the particular climate simulation rather than for a mix of users running a variety of applications.

The ability to add new processors to the simulation for increased throughput is also interesting. This could allow effective use of lightly loaded computers in a center that sees peaks of usage during certain hours of the day. From an application point of view, HARNESS offers this capability under the control of the user instead of at the discretion of the centers system administrators.

Collaborative, adaptable computing centers.

A further implication may be in the way that computer centers could collaborate for large science runs. Projects linking national computer centers have heretofore been largely demonstration projects and publicity stunts rather than offering a production capability for long science runs. Experimental networks are often used which push the envelope of bandwidth and are relatively unstable. One doesn't expect such systems to perform well and give quality of service unless there is an adaptive, fault tolerance built into the managing software.

When this software exits, and HARNESS is an instantiation of such a system, it may become possible for remote centers to develop collaborative agreements for the sharing of cycles. For example, rather than being isolated for a run at a single center, it should be possible for a simulation to be started on a local machine. If the user has available time at a center like the DOE NERSC, a remote machine could be added to the virtual machine until the NERSC account is depleted. The next months allocation might be used in a similar fashion to contribute to the same continuously running climate simulation. These capabilities no longer seem like unattainable possibilities. The conceptual design is clear. Given the past record of the PVM developers and the number of scientific applications that have been adapted for the message passing distributed computing, it may not be a big stretch to take the next step.

Uninteruppted science.

The thing that we hope will be accomplished as a result of this work is the means to support uninteruppted science. In the case of climate research this could be accomplished by a stable, high performance computing environment. The proposed solution is to define a virtual machine environment which offers the stability while at the same time adapting to the latest set of hardware. The scientist should be able to concentrate on what is being computed rather than how and where it is being computed.

4 Facilities for proposed research

Development test bed

ORNL and its collaborators have a rich variety of experimental networks operating on a departmental level. These include ATM, FDDI, HIPPI, Myrinet. For development purposes these networks will be adequate to test software. The availability of small SGI cluster machines, networked workstations and the ORNL CCS computers will also provide adequate compute power to test our work with short coupled climate simulations. In particular, the fault tolerance procedures and the adaptability requirements can be developed in this setting which is locally controlled and readily available.

Demonstration test bed

To make a reasonable demonstration for a coupled climate simulation in the decadal to century time scale, however, a larger scale must be considered. Machines with multi-Gigaflop performance must be available. Proposals have been made for the ORNL CCS to provide a SRC 6 computer as a test bed. The power of this system would fall in the multi-Gigaflop range with possible follow on machines in the teraflop range. So this is one possibility as part of a demonstration test bed. Another is to include the NERSC machines at LBL, the Cray C90 and the Cray T3E.

If the research project is successful, then plans can be made to install these capabilities at the LANL Advanced Computing Laboratory and at the National Center for Atmospheric Research. Clusters of SGI Origin machines could be supported as part of the virtual machine. Plans for this later stage are not included in this proposal but would form a logical extension of the work.

Network Issues: Quality of Service, Latency, and Bandwidth

The national network required to support the linking of centers must provide a high level of quality of service and high bandwidth for this to make sense. The rich ATM network supported by ESnet connections, particularly into NERSC, would be important for the success of this project. The latency requirements for the coupled simulations are relatively low as component models communicate surface states and fluxes across the network. High bandwidth minimizes the wait time of the models so bandwidth in the 100Mb/s range is a minimum requirement.

Last but not least: Archival storage

It is a recognized fact that climate models stress all the components of a computing environment. Thus some plan must be given for the development of an archival storage system that would work with HARNESS. The ORNL CCS has developed a fine storage system capability using the Unitree system and now HPSS. The CCS could participate in this proposal by providing a HPSS plug-in to the HARNESS system. This would allow for a distributed climate simulation to use local (or remote) storage systems as a service of the virtual machine. We expect that this is not a small task, but one that would be very useful and complete the implementation of a virtual machine useful for long climate simulations. We propose to consider the extension of the system to include storage only after successful completion of the demonstration project. At that point we need to move beyond a demonstaration to a significant, new coupled run where archival storage will be required. The CCS is the logical partner for this work but we are cognizant of the need to provide interface to the NCAR archive as well as the LLNL PCMDI archive. These are extensions of the work proposed here.

5 Budget Outline

Two groups are involved with this proposal. First, the application group working with the CSM led by John Drake of ORNL. Second is the research consortium developing HARNESS consisting of ORNL, Emory University and the University of Tennessee. The ORNL leader for this effort is Al Geist and will also be the point of contact and co-PI for this proposal. A third group that could become involved is the ORNL CCS who could provide the expertise in storage systems and (contingent on computer resources) provide cycles for the demonstration test bed.

We would propose that OBER/CCPP provide 1.5 FTE for the CSM work. OCTR/MICS would then match with 1.5 FTE for HARNESS work. We take a nominal dollar amount on an FTE to be $200K per year. So the total commitment from OBER/CCPP would be $300K and similarly from OCTR/MICS. The total proposal is for $600K per year.

The duration of the proposal is 3 years. The first year would be dedicated to design and development of the HARNESS software and MCL plug-in. Second year would be for the demonstration projects and the third year would be to reengineer the software based on what has been learned, extending to other centers preparing support for production climate simulations.