Climate Modeling with a Linux Cluster

PAGE 290

Until recently, computationally intensive calculations in many scientific disciplines have been limited to institutions which have access to supercomputing centers. Today, the computing power of PC processors permits the assembly of inexpensive PC clusters that nearly approach the power of supercomputers. Moreover, the combination of inexpensive network cards and Open Source software provides an easy linking of standard computer equipment to enlarge such clusters. Universities and other institutions have taken this opportunity and built their own mini-supercomputers on site.

Computing power is a particular issue for climate modeling and impacts the community. The purpose of this article is to make available a Linux cluster version of the Community Climate System Model developed by the National Center for Atmospheric Research (NCAR; http://www.cgd.ucar.edu/csm).

The Community Climate System Model (version CCSM2.0.1) is a comprehensive, atmosphere-ocean general circulation model which is used for simulations of the past, present, and future climate *[Blackmon et al.,* 2001]. The releases by NCAR are tailored to IBM SP3/4 and SGI machines. Within the framework of the Swiss National Centers of Competence in Research (NCCR Climate), the current model and environment manager of CCSM2.0.1 was adapted to run on Linux clusters. The code is made available to the climate model community in order to enable wider use of this model by research groups with access to PC clusters. A migrated version with a brief installation guide can be downloaded at http://www.climate.unibe.ch/CCSM.

To demonstrate the scalability and performance of the CCSM2.0.1 for Linux, a series of short runs was carried out on a PC-Linux cluster. To compare these results with a supercomputer, additional tests were made with the climate model on an IBM SP4. The SP4 consists of several pSeries 690 Regattas and is equipped with 1.3-GHz Power 4 processors, while the Linux cluster is made up of 20 2-Ghz AMD processors and 34 1.3-GHz processors. The faster cluster PCs are connected by a 1-Gbit network and the slower ones use a 100-Mbit network (see Technical Notes and Figure 1).

Note that supercomputer processors cost approximately 10 times more than PC processors. Comparing the performance of the climate model on the two machines, the price performance benefit of the Linux cluster is more than evident (Figure 2). Running the climate model on 12 CPUs, the Linux version reaches nearly 80% of the SP4 performance. By further increasing the number of processors, the Linux cluster is not able to follow the scalability gain of the supercomputer, but with 48 processors, it still reaches approximately 20 to 25% of the SP4 performance. One of the main reasons why the SP4 is still faster is its superior network and the sophisticated memory connections.

The limitations of the Linux cluster network can be shown using different network speed.

Fig. 1. Part of the Linux cluster installed at the Division of Climate and Environmental Physics of the Physics Institute at the University of Bern, Switzerland. The PCs are connected by a 1-Gbit network, and the slower nodes by a 100-Mbit network.

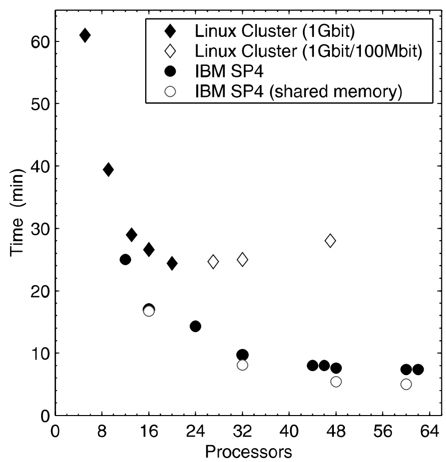

Fig. 2. The wall-clock time to calculate one model month of the NCAR CCSM2.0.1 climate model running on an IBM SP4 (dots), and on a Linux cluster (diamonds), as a function of the number of processors. All tests have been carried out with a resolution of approximately 3.8° for the atmosphere and a ~ 3.6 x 1.8° resolution for the ocean component. The maximum speed of the Linux cluster is reached with 20 processors.

Using 1-Gbit connections only throughout, the performance scales well, but combined with the slower 100-Mbit connections, the trend breaks beyond 20 processors, mainly due to the slower network (diamonds in Figure 2). An additional factor for this cluster has to be taken in account regarding the processors with the 100-Mbit network. The inhomogeneity of the PC clock speed could possibly limit the overall performance of the computation on our PC cluster. An attempt was made to minimize this effect by assigning the computationally less-intensive model components to slower machines, but this yielded no significant improvement. During these tests, the load of the machines never reached 100%. This implies that the performance is practically limited by the network bandwidth and latency time, rather than the speed of the individual processors. With the given network hardware on the cluster, the maximum performance is reached with 20 PCs. However, by equipping all machines with a 1-Gbit network or with high-speed interconnections, such as a Myrinet network, a further performance increase is expected.

There are additional advantages of running the CCSM2.0.1 on a Linux cluster. First, during the performance test, the possibility of improving the overall cluster performance, by superimposing two models which run simultaneously on the same group of PCs, was identified. The loss of individual model performance was less than 50% on a homogeneous 100-Mbit network. This method ensures better processor utilization in the case of a network with high latency times, and can be used to perform simultaneous ensemble simulations. Second, PC clusters are also attractive because of the possibility of performing numerous ensemble simulations using large farms of PCs at reasonable cost.

Multi-centennial climate simulations with comprehensive general circulation models can now be performed on PC clusters. Ensemble simulations are also within the range of PC cluster applications. Over 500 years of a control simulation and several hundred years of global warming scenarios were conducted in a couple of months with our Linux cluster.

The future of PC clusters seems promising. High-speed interconnect initiatives such as the Myrinet 2000 network or Infiniband technology, among others, will continue to provide

Technical Notes

- The IBM SP4 of the Swiss Center for Scientific Computing (CSCS; http://www.cscs.ch/) consists of a farm of eight pSeries 690 Regatta SMP with 32 CPUs each (Power 4 at 1.3 GHz, 5.2 GFlops).
- The Linux cluster consists of 54 state-of-the-art PCs, a part of which is shown in Figure 1. The climate model output data are saved with a Network File System (NFS) exported file system to a 6-terabyte IDE-Raid storage system.
- On the Linux cluster, the CCSM2.0.1 was compiled with the Portland pgf90 compiler using the embedded MPI library. On the IBM used, the AIX mpxlf90 compiler was used. Note that the PCs of the Linux cluster consist of single processor with single memory, while the SP4 is equipped with multiprocessor shared memory system. Therefore, the shared memory code at the Linux system was not activated. However, the improvement using shared memory on the IBM SP4 is relatively small below 24 processors (circles in Figure 2). Using more processors, the SP4 performance improvement is significant.

improvements in cluster performance. We are now working to target Linux clusters in future versions of CCSM. The next version CCSM3.0 (released June 2004) is also available for Linux. We believe that Linux clusters are ready to be used in climate modeling, and because performance tests have been carried out with version CCSM2.0.1.

Acknowledgment

This work is supported by the National Center of Competence in Research (NCCR Climate), Switzerland, (http://www.nccr-climate.unibe.ch/) in close collaboration with National Center for Atmospheric Research (NCAR; http://www.cgd.ucar.edu/csm) and the Swiss National Supercomputing Centre (CSCS; http://www.cscs.ch/).

Reference

Blackmon, M., et al. (2001), The community climate system model, *Bull. Am. Meteorol. Soc.*, *82(11)*, 2357–2388.

—M. Renold, U. Beyerle, C. C. Raible, R. Knutti, and T. F. Stocker, University of Berne, Switzerland; and T. Craig, National Center for Atmospheric Research, Boulder, Colo.

For more information, contact Manuel Renold, Climate and Environmental Physics, Physics Institute, University of Berne, Sidlerstrasse 5, CH-3012 Bern, Switzerland; E-mail: renold@climate.unibe.ch.

MEETINGS

Thermal Processes in the Context of EarthScope

PAGE 292

The EarthScope project is designed to elucidate the dynamics and evolution of the North American continent. To reach that goal, focused studies of the continental seismic structure, the strength of the lithosphere and modes of deformation, the patterns of stresses and strain, and the manner in which the lithosphere responds to excitation during earthquakes are all a high priority.

The common factor linking these EarthScope centric issues is the thermal state of the lithosphere. Although not an explicit component of EarthScope in the sense of USArray, PBO or SAFOD, an understanding of the thermal structure and associated thermal processes plays a key role in interpreting the observations generated by the EarthScope infrastructure. In order to help guide EarthScope science to exploit most effectively the links among these observations, a workshop was held to bring together practitioners from the fields of seismology, geodesy, thermochronology, and crustal deformation with researchers who focus on thermal processes. The goal was to define the high priority issues to be addressed through integration of thermal geophysics with EarthScope science.

Forty-two researchers and educators gathered in March 2004 at the University of Utah to discuss thermal processes within the context of EarthScope. Workshop participants included individuals from the U.S., Canada, Germany, and The Netherlands. Invited talks focused on the general theme of combining geodynamic, geodetic, seismic, and other lithospheric data sets with thermal data to enhance our understanding of Earth processes; posters displayed specific studies covering similar themes. Working groups were formed to relate represented subdisciplines to EarthScope goals. Breakout groups were largely devoted to discussions of thermal processes and EarthScope goals. Plenary sessions summarized breakout group discussions and addressed questions posed during the meeting.

All of the sessions explored integrating thermal data with other geophysical and geological data acquired through EarthScope. Within the context of EarthScope, knowledge of the thermal state of North America is important to understanding: 1) continental dynamics and evolution; 2) the relationship between mantle dynamics and crustal tectonics; 3) the behavior of active fault systems; 4) earthquake nucleation and rupture processes; 5) natural hazard reduction; and 6) magmatic systems. Summaries of discussions covering these topics are discussed here.

Temperature and Thermal Processes

Quantifying the flow of thermal energy is part of understanding Earth dynamics and compels integrating thermal data with geodetic, seismic, and other data collected through EarthScope. Better models of heat transfer and estimates of thermophysical rock properties are critical to understanding the thermal regime in the crust and upper mantle. Improvements to thermal models will likely come from integrating data from such sources as geochemistry, xenoliths, and seismology.

A long-standing uncertainty in estimating the thermal state of the lithosphere is the distribution of heat-producing elements (HPE) as a function of depth. Heat production is a function of geochemical environment and geologic history, and although it is clear that HPEs generally decrease with depth, the exact nature of this decrease is less certain. The linear relationship between surface heat flow and heat production that provided a basis for defining a depth distribution of HPEs has become less clear as new data have been acquired.

The Evolution and Dynamics of North America

Heat transport in the Earth drives deformation, magmatism, and continental evolution. Surface heat flow data provide primary observations shaping our understanding of the thermal, petrologic, and tectonic evolution of the continents. Outstanding problems for understanding continental evolution include: 1) stabilization of continental lithosphere; 2) the rate of continental growth through time; and 3) an understanding of the thermal and geodynamic characteristics of mobile belts. For example, what are the respective roles of fluids and heat transfer in keeping mobile belts weak and long-lived? How do mobile belts evolve toward stable cratons and platforms? Answers to these questions will likely come through combining EarthScope studies of active tectonics and exhumed cratons and platforms with, among other endeavors, thermal modeling of tectonic processes.

Mantle Structure and Dynamics of North America

A critical component to the success of understanding mantle structure and dynamics is the ability to discern thermal and compositional variations in the mantle. The seismic properties of the mantle will be mapped at an unprecedented spatial resolution through USArray. To exploit this increased resolution of mantle seismic structure, an improved ability is needed to distinguish among seismic velocity anomalies produced by the competing influences of temperature, composition, rheology (anelasticity), and fluids. Better estimates of seismic properties of mantle materials and their variation with respect to temperature, composition, and fluids is critical.

Active Faulting and Earthquake Processes

EarthScope presents the opportunity to increase greatly the quality and quantity of data that can be used to unravel the processes associated with active faulting and earthquake processes. The themo-mechanical state of earthquake source regions and along active faults play an important role in seismogenesis, but are incompletely known. For example, the overall strength of the lithosphere and the partitioning between brittle and ductile deformation styles depend, to a great extent, on the thermal state of the lithosphere.

Within zones of active faulting, the base of the seismogenic zone, whether defined as a transition from unstable (velocity-weakening) to stable (velocity-strengthening) sliding or the transition from brittle failure to ductile flow, appears to be determined in part by variations in temperature. Both heat flow and fission-track annealing studies have been applied in