DOE/ER/61215--1

DE92 019457

CHAMMP PROGRESS REPORT

March 10, 1992

PROGRESS TO DATE

The first phase of the proposed work is largely completed on schedule. Scientists at the San Diego Supercomputer Center (SDSC) succeeded in putting a version of the Hamburg isopycnal coordinate ocean model (OPYC) onto the INTEL parallel computer. The role of the SIO Principal Investigators in this work phase was advisory in nature and thus consumed only a small amount of time (and funds). To further assist the SDSC people in their effort we arranged for the author of the code (J. Oberhuber) to spend two weeks working with them in the conversion process. A brief summary of the SDSC effort is attached as Appendix A.

In summary, it does not appear that the code can be speeded up enough in its present form to be practically useful for 1000+ year integrations. The same conclusion has been reached by a different group working on the OPYC code at Los Alamos. That group is planning a serious restructuring of the code in the coming year so that it will be more amenable to parallel computing.

YEAR 2 ACTIVITY

Ocean model

Due to the slow run speeds of the OPYC on the parallel machine we have decided to use another ocean model during the first part of phase 2. The model we have chosen is the Large Scale Geostrophic (LSG) model from the Max Planck Institute. The LSG model has been successfully integrated for 10,000 years of simulation and demonstrated large and interesting natural variability. Best of all, it is very fast...about 10 seconds of CPU time on a Y-MP for 1 year of integration. In fact, it is so fast that we have decided it is not worth the effort to put it on a parallel machine. At this point, we are hopeful of efficiently running it on a dedicated work station...less than 5 minutes of CPU time per year of integration or 4 days per 1000 year simulation.

In retrospect, we should have planned to use the LSG model in the project from its very outset. Its speed will allow us many cheap runs to try out ideas, make mistakes and learn how to handle the frightening data volumes associated with the long integrations. The area

MASTER

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

of handling the data volume generated by the OGCM is a particular concern for we need to substantially upgrade our visualization capabilities to meet the data volume/display/analysis problem. Underlying these problem and model shift is the fact that the LSG is a relatively good model for looking at large scale, low frequency ocean physics.

In the latter half of Phase 2, after obtaining considerable experience with the LSG, we hope to insert a more complex ocean model into the program. Computer scientists at Livermore are hard at work converting the GFDL OGCM to the CM5. It is anticipated that an operational code may be ready by fall. In the meantime, the Semptner-Chervin code is anticipated to be ready for general use on the CM5, perhaps as early as this summer. Finally, the recoding of the OPYC model may be completed before this phase is finished and, if so, offers another possible ocean model. In short, it appears that within the year there will be one or more ocean models that can be substituted for the LSG. It is doubtful, however, that any of them will be as fast and cheap as the LSG, so again it seems sensible to get started with this efficient model.

Statistical Atmosphere Model

Near the beginning of the new fiscal year we will start construction of the atmospheric anomaly model that will initially be coupled to the LSG. By that time we should have complete results from a 20 year integration with the MPI T42 AGCM. The forcing fields derived from this run (heat, moisture and momentum fluxes) will first be compared with observationally derived fluxes. If the agreement is reasonably good then we will parameterize the fluxes in terms of LSG SST, phase of the seasonal cycle,etc. It appears quite likely that much of the parameterization will consist of representing the flux data as a quasi-random process.

An effort will also be made to carry out the same type of parameterization only using observed quantities as opposed to those derived from GCMs. This effort is more problematical due to the well known data sparsity in many regions of the oceans.

Interactions

The SIO investigators expect to cooperate with several other segments of the program. Specifically, we are interested in making long runs with the GFDL and S-C OGCMs to see how well they represent the seasonal cycle in the ocean and key aspects of deep ocean circulation. If available we would like to do the same thing with the OPYC model. The next step would be to couple our statistical atmosphere to the other ocean models and see if the variability produced is similar to that expected from the coupled LSG runs. All of this work will require close cooperation with the computer scientists at LLNL and LANL, as well as the oceanographers involved in the modeling effort.

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

SDSC Summary:

The goal of the initial phase of the SDSC effort is to

a) implement the OPYC model on the 64-node iPSC/860 parallel computer at SDSC

b) evaluate its performance on that platform (and its near term successor, the Intel Paragon) relative to the fastest available current generation vector supercomputers (i.e., the CRAY Y-MP and its successor, the CRAY C90) for long term (100 year) global integrations. Such integrations may take as much as 2000 single processor CRAY Y-MP hours for grid sizes of interest.

The iPSC/860 has a theoretical peak speed of 40 Mflops for each of its i860 chip nodes or 2560 Mflops for the full 64-node configuration, while the CRAY Y-MP has a single processor peak speed of 333 Mflops. Thus large potential speedups relative to the CRAY in principle are possible if a) good performance on individual nodes is obtained and b) effective parallelization is achieved.

Initial efforts have concentrated on evaluating and maximizing performance of the most computationally intensive subroutines on individual i860 nodes and comparing results to those on the CRAY Y-MP. The base case for overall evaluation relative to the CRAY was taken to be the "T42" x-y-z grid of dimension 130 x 62 x 9. The CRAY version of the model is well vectorized and runs quite fast (approximately 150 Mflops on a single processor). The model author, J. Oberhuber, has recently made several performance enhancements to the CRAY version that are claimed to increase the speed by 20%. We are in the process of implementing these changes now and have not yet verified the improvements, but believe them to be valid based on some similar improvements obtained in our own work.

Although individual subroutines of the of the T42 grid OPYC model can be tested in isolation on single i860 nodes, the entire model is too large to run on one node. With the assistance of J. Oberhuber, who spent a week working with us in the early part of the project, a simplified "box" model with a flat bottom topography on a 22 x 22 x 3 grid, cabable of being run on a single node, was constructed to aid in evaluating overall node performance. Also, a single node self-contained version of those components of the model associated with the solution of the wave equations was constructed. Additionally, test routines for performance evaluation were written for the most computationally intensive subroutines on the full grid.

Generally, performance on individual nodes has been quite disappointing. The box model ran at approximately 1.4 Mflops on a single node in its original form. After an extensive optimization effort, this was increased to 3.2 Mflops. While somewhat faster speeds may be obtained on the full grid, we believe the increment will be small. Similarly, the computationally most intensive subroutines, which consist of direct (Gaussian elimination) solvers for block tridiagonal systems arising from the wave equations and tridiagonal and 2 by 2 block tridiagonal systems arising from advection and diffusion equations, all ran at less than 1.5 Mflops when dimensioned for the full grid. Again, optimization efforts resulted in speed increases to approximately 3 Mflops. Thus overall, we believe individual node performance, before overhead for parallelization and internode communication, will be limited to slightly over 3 Mflops. The technical reasons for this relatively low performance are related to the internal architecture of the i860 chip, which has a limited memory bandwidth and very poor floating point divide performance. The consequent performance penalties are particularly severe for Gaussian elimination applied to the types of very limited bandwidth linear systems which dominate the OPYC model.

Furthermore, while the tridiagonal and 2 x 2 block tridiagonal systems occur in sufficiently many parallel instantiations to allow effective parallelization over 64 (or even 128) processors for the base case, the block tridiagonal systems that arise from the wave equations do not. The number of such parallel systems that arise is half the number of y grid points (31 in the base case). Thus on a 64 node system, approximately half the processors would be idle at any given time when solving the wave equations. Unfortunately, this is the dominant portion of the entire model calculation.

When internode communications overhead is also considered, we expect that the overall performance of the 64-node iPSC/860 will be less than 75 Mflops on the OPYC model, and that a larger system will not help appreciably. Increased node performance will, of course, increase total system performance. However, the next generation Intel machine, the Paragon, is expected to have only a factor of two improvement in node performance. A parallel implementation thus will still be considerably slower than a single CRAY Y-MP processor version. Also, the CRAY C90 is now becoming available and has a single processor peak speed of 1000 Mflops. The OPYC model is currently being benchmarked on the C90, and is expected to perform very well due to the long vector lengths in most of the computations. Thus for the immediate future, we believe the OPYC model, in its current algorithmic form, will be far better suited to state-of-the-art vector supercomputers than parallel supercomputers. A major algorithmic reformulation of the model, particularly to rely on iterative rather than direct linear system solvers, appears to be the only possibility to allow a parallel implementation to become competitive.

As an alternative to the OPYC model, the Hamburg Large Scale Geostrophic Ocean General Circulation Model (LSG-OGCM) was obtained from the Max Planck Institute for Meteorology. With the help of M. Lautenschlager from MPI, the original Convex version was ported to the CRAY Y-MP. Performance on the CRAY is excellent, with computational speeds of approximately 150 Mflops on a single processor. The model physics are considerably simpler than those of the OPYC model, and consequently ruuning times for a given simulation period are much faster. In particular, the basic time step of 30 days for a 72 x 76 x 11 grid takes approximately 0.75 seconds, allowing the completion of a 100 year integration in less than 15 CPU minutes. The code is somewhat memory intensive (approximately 48 Mbytes for this grid), as the matrix factors of a large system of linear equations are saved and used at each time step. However, it appears to be feasible to port the code to a high performance workstation and obtain reasonable performance (e.g. 100 year integrations within several wall-clock hours).

END

DATE
FILMED
9/28/92