Progress Report Project Number: 3211-2058

Scaling Climate Models for Future Computer Architectures

A.S. Bland [1], J.B. Drake [1], P.H. Worley [1], and J. B. White, III [1]
[1]Computer Science and Mathematics Division

Abstract

The 40 TFLOP/s Japanese Earth Simulator (ES40) has challenged the U.S. climate-modeling community and computer industry to accelerate the development of both the models and the computers required to run them, to keep our country competitive in both climate science and policy decisions. This project will characterize the performance of current U.S. models and suggest research areas for improvements. Researchers will use representative kernel codes to understand the factors that limit performance and scalability and then work to mitigate these factors. The second phase will be to instrument and evaluate the latest full climate model, CCSM2, to improve its performance. The objective is to prepare ORNL to take advantage of future funding opportunities from DOE to develop and deploy an ultra-scale computer system for the U.S. science community.

Body of Progress Report

The objective of this project is to prepare ORNL to take advantage of future funding opportunities from DOE to develop and deploy an ultra-scale computer system for the U.S. science community. The approach to this is to begin scaling the components of the Community Climate Systems Model (CCSM) up to the resolution and physical fidelity that is expected to be used on this new class of computer systems and to port these codes to a new generation of high-end computer architectures.

This project has three major technical thrusts:

1. Benchmarking performance of a dynamical kernel on a high-resolution model to identify hardware and software issues with current and planned architectures.
2. Evaluation and implementation of performance enhancements for global spectral atmospheric models.
3. In-depth scientific performance study of the CCSM2 at high resolution on IBM p690 and Cray X1 systems.

The performance benchmark targets a spectral model of atmospheric dynamics. We will match the resolution of T1279L96 used in the ES40 benchmark that achieved 26 TFLOP/s. As currently configured, both the IBM p690 and Cray X1 at ORNL are currently insufficient to run this problem resolution. However, we have run experiments using problem sizes with large horizontal resolutions (up to T680 on the Cray and T1279 on the IBM) with small numbers of vertical levels, and problem sizes with smaller horizontal resolutions and up to 92 vertical levels, determining the processor and memory requirements necessary to run the target problem size.

The experiments on the Cray were preceded by extensive performance tuning of the kernel code PSTSWM to improve its vectorization. We have improved the performance by a factor of 10 over the original version of the code on a single processor, achieving approximately 50% of peak for the largest horizontal resolution. Experiments indicate that maximum serial performance is achieved by computing on a few vertical levels at a time. This is not the approach used in the current production spectral model (nor in the Earth Simulator code), and has motivated investigating alternative parallelization strategies in the SciDAC project concerned with development of the production atmospheric model. The next step will be to run the full T1279L96 problem resolution on the upgraded Cray X1. Preliminary optimizations of the parallel algorithms for this large run have already been completed using the current system.

A high resolution ocean configuration has been targeted using the POP code. The ES40 is showing performance of hundreds of Gflops in a 60 node configuration for POP at 1/10-degree resolution. The Cray X1 exceeds the ES40 performance for the low-resolution case, and we expect to find competitive performance with the high resolution case as well.

Results to date show that the Cray X1 architecture is superior to the IBM p690 in both performance and scalability for the CCSM. Future work will focus on completing scaling to T1279L96 resolution for benchmarking and maximizing performance of the full model for high-resolution configurations.