Geophysical Research Abstracts
Vol. 16, EGU2014-16712-1, 2014
EGU General Assembly 2014

Toward an efficient and scalable IO mechanism in the US DOE Earth System Model

Valentine Anantharaj (1,2)

(1) National Center for Computational Sciences, Oak Ridge National Laboratory, USA (anantharajvg@ornl.gov), (2) Climate Change Science Institute, Oak Ridge National Laboratory, USA

The number of cores (both CPU as well as accelerator) in large scale systems has been increasing rapidly over the past several years. In 2008, there were only 5 systems in the Top500 list that had over 100,000 cores whereas the number of system with such capability has jumped to 26 in Nov 2013. This growth however has also increased the risk of hardware failure rates, necessitating the implementation of fault tolerance mechanism in applications. The checkpoint and restart (CR) approach is commonly used to save the state of the application and restart at a later time either after failure or to continue execution of experiments.

The implementation of an efficient CR mechanism will make it more affordable to output the necessary CR files more frequently. The availability of larger systems (more nodes, memory and cores) has also facilitated the scaling of applications. Nowadays, it is more common to conduct global climate simulation experiments at 1 deg horizontal resolution (atmosphere), often requiring about 10^3 cores. At the same time, a few climate modeling teams that have access to a dedicated cluster and/or large scale systems are involved in modeling experiments at 0.25 deg horizontal resolution (atmosphere) and 0.1 deg resolution for the ocean. These ultrascale configurations require the order of 10^4 to 10^5 cores. It is not only necessary for the numerical algorithms to scale efficiently but the input/output (IO) operations must also scale accordingly.

An ongoing series of ultrascale climate simulations, using the Titan supercomputer at the Oak Ridge Leadership Computing Facility (ORNL), is based on the spectral element dynamical core of the Community Atmosphere Model (CAM-SE) on the U.S. Department of Energy (DOE) Earth System Model, which is a branch of the Community Earth System Model. The CAM-SE dynamical core for a 0.25 deg configuration has been shown to scale efficiently across 100,000 cores. Under certain circumstances on Titan, the output of restart files could require as much as 20 minutes of wallclock time that typically translates to nearly 1.4% of the 24 hour wallclock time using 60,000 cores at a cost of 36,000 titan core hours for the IO. At this scale, there is an increased risk that the simulation could be terminated due to hardware failures, resulting in a loss that could be as high as 2.6 million titan core hours. Increasing the frequency of the output of CR files could mitigate this loss but at the cost of additional CR overhead. Increasing the CR frequency to every hour would result in 33% overhead for restart IO. We are implementing a more efficient CR mechanism in CESM. Our early implementation has demonstrated a nearly 10X performance improvement for a 1 deg CAM-SE configuration using nearly 10^3 cores. We are in the process of scaling our implementation to 10^5 cores.