IRS2012-539
International Radiation Symposium 2012
Dahlem Cube, Berlin, Germany, 06 – 10 August 2012

Development of a GPU compatible version of the fast radiation code RRTMG

E. Mlawer (1), D. Berthiaume (1), M. Iacono (1), K. Cady-Pereira (1), M. Suarez (2), L. Oreopoulos (2), and D. Lee (3)

(1) Atmospheric and Environmental Research (AER), Lexington, MA, United States (mlawer@aer.com), (2) NASA-GSFC, Greenbelt, MD, USA , (3) GESTAR/USRA, Columbia, MS, USA

The absorption of solar radiation and emission/absorption of thermal radiation are crucial components of the physics that drive Earth's climate and weather. Therefore, accurate radiative transfer calculations are necessary for realistic climate and weather simulations. Efficient radiation codes have been developed for this purpose, but their accuracy requirements result still necessitate that as much as 30% of the computational time of a GCM is spent computing radiative fluxes and heating rates. The overall computational expense constitutes a limitation on a GCM's predictive ability if it becomes an impediment to adding new physics to or increasing the spatial and/or vertical resolution of the model. The emergence of Graphics Processing Unit (GPU) technology, which will allow the parallel computation of multiple independent radiative calculations in a GCM, will lead to a fundamental change in the competition between accuracy and speed.

The fast radiation code RRTMG, developed at Atmospheric and Environmental Research (AER), is utilized operationally by many dynamical models throughout the world. We will present the results from the first stage of an effort to create a version of the RRTMG radiation code designed to run efficiently in a GPU environment. This effort will focus on the RRTMG implementation in GEOS-5. RRTMG has an internal pseudo-spectral vector of length ∼100 that, when combined with the much greater length of the global horizontal grid vector from which the radiation code is called in GEOS-5, makes RRTMG/GEOS-5 particularly suited to achieving a significant speed improvement through GPU technology. This large number of independent cases will allow us to take full advantage of the computational power of the latest GPUs, ensuring that all thread cores in the GPU remain active, a key criterion for obtaining significant speedup. The CUDA (i.e. Compute Unified Device Architecture) Fortran compiler developed by PGI and Nvidia will allow us to construct this parallel implementation on the GPU while remaining in the Fortran language. This implementation will scale very well across various CUDA supported GPUs such as the recently released Fermi Nvidia cards. We will present the computational speed improvements of the GPU-compatible code relative to the standard CPU-based RRTMG with respect to a very large and diverse suite of atmospheric profiles. This suite will also be utilized to demonstrate the minimal impact of the code redevelopment on the accuracy of radiation calculations.

The GPU-compatible version of RRTMG will be directly applicable to future versions of GEOS-5, but is also likely to lead to significant associated benefits for other GCMs that employ RRTMG.