High Rayleigh Number Mantle Convectionon GPU

© *D. Sanchez*[1], *D. Yuen*[1], *G. Wright*[2], *G. Barnett Jr.*[3], *2010*

[1]Minnesota Supercomputing Institute, University of Minnesota, USA
[2]Department of Mathematics, Boise State University, USA
[3]Department of Applied Mathematics, University of Colorado, Boulder, USA

The performance, potential growth, availability, and affordability of GPUs has made them attractive to scientists for many years. Although historically cumbersome and difficult to use for scientific software, the introduction and refinement of high-level development tools, such as CUDA, have made GPU computing accessible. With the advent of architectures, such as NVIDIA's Fermi, which explicitly cater to scientists by enabling more memory, faster access to that memory, and better double-precision support than ever before, members of the computational world are finding GPU difficult to ignore.

Using finite-difference methods with second-order accuracy in space and third-order accuracy in time, we investigate 2D and 3D thermal convection at the infinite Prandtl number limit, at resolutions on the order of 1000×2000 2D and 400×400×200 3D grid points. Our CUDA code makes extensive use of

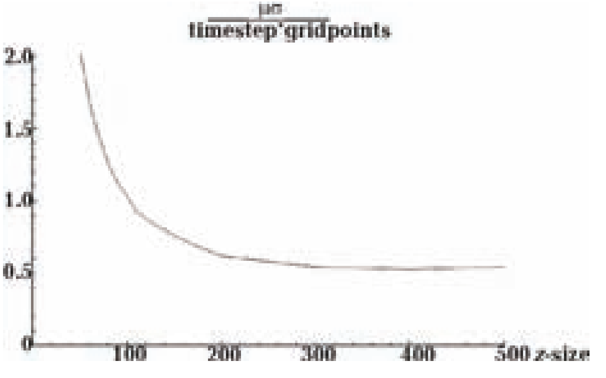

Fig. 1. Performance scaling with grid size on Tesla C1060.

Fig. 2. 2D mantle convection at a Rayleigh number of 10^9.

highly-optimized CUBLAS routines, allowing us to unlock a significant fraction of GPU's performance. This performance has enabled us to study the behavior of high Rayleigh number simulations, on the order of 10^9, in 2D and 10^7 in 3D over sufficient time scales to see evidence of flow-reversal (Fig. 1).

We compare our CUDA code's performance with Jacket-accelerated Matlab code and CPU-only Matlab code across the Tesla C1060, Tesla C2050, and GTX 480 (Fig. 2).