GPU Accelerated Reduced MHD Simulations of Coronal Loops

LiWei Lin[1], Chung-Sang Ng[2], Amitava Bhattacharjee[1]

[1]*Space Science Center, University of New Hampshire, Durham, NH 03824*

[2]*Geophysical Institute, University of Alaska, Fairbanks, AK 99775*

We present a comprehensive re-programming of a 3D reduced MHD code for hardware acceleration using general purpose graphics processing units (GPGPUs). It is built on a conventional numerical scheme (pseudo-spectral semi-implicit), parallelized using MPI, and GPU accelerated with Nvidia's Compute Unified Device Architecture. The code is currently used to study a 3D model of coronal loop heating [Arxiv:1106.0515]. We will discuss our general porting strategy and report code performance and detailed code tracing on GPU accelerated supercomputers (NCSA/Lincoln, NERSC/Dirac, and NICS/Keeneland). For the highest resolution tested (2048^2x256), the chip-to-chip speedup is 18x comparing Xeon Nehalem QC CPUs with Nvidia Fermi GPUs. The code scales well up to 256 GPUs, effectively giving a speedup of 46x compared with our previous best performance on conventional CPU clusters. We also demonstrate that while the GPU accelerators now undertake the bulk of the computational burden, otherwise idle cores on the host CPUs can be employed for run-time post processing. A test case will be presented in which magnetic island coalescence is examined in a 3D line-tied geometry. The system is prescribed with a very large Lundquist number in order to induce magnetic flux-tube sloshing, necessitating both high spatial and temporal resolution. This results in storage requirements that practically inhibit the downloading of the full data set for off-site analysis. We show that several post processing tasks (O-point tracking, current sheet identification and characterization, 3-D visual renderings etc.) can be run concurrently on CPU cores alongside the CPU-GPU numerical simulation itself, thus precluding the need for auxiliary post processing runs and excessive off-site data movement. This work is supported by NASA grants NNX08BA71G, NNX06AC19G, DOE grant DE-FG02-07ER54832, NSF grant AGS-0962477, and NSF TeraGrid grants from NCSA (TG-PHY100057) and NICS (UT-NTNL0092).