Abstract Submitted
for the 4CF11 Meeting of
The American Physical Society

Cell-based Adaptive Mesh Refinement on the GPU with Applications to Exascale Supercomputing DENNIS TRUJILLO, New Mexico State University, ROBERT ROBEY, NEAL DAVIS, DAVID NICHOLAEFF, Los Alamos National Lab XCP-2 — We present an OpenCL implementation of a cell-based adaptive mesh refinement (AMR) scheme for the shallow water equations. The challenges associated with ensuring the locality of algorithm architecture to fully exploit the massive number of parallel threads on the GPU is discussed. This includes a proof of concept that a cell-based AMR code can be effectively implemented, even on a small scale, in the memory and threading model provided by OpenCL. Additionally, the program requires dynamic memory in order to properly implement the mesh; as this is not supported in the OpenCL 1.1 standard, a combination of CPU memory management and GPU computation effectively implements a dynamic memory allocation scheme. Load balancing is achieved through a new stencil-based implementation of a space-filling curve, eliminating the need for a complete recalculation of the indexing on the mesh. A cartesian grid hash table scheme to allow fast parallel neighbor accesses is also discussed. Finally, the relative speedup of the GPU-enabled AMR code is compared to the original serial version. We conclude that parallelization using the GPU provides significant speedup for typical numerical applications and is feasible for scientific applications in the next generation of supercomputing.

Dennis Trujillo
New Mexico State University

Date submitted: 16 Sep 2011 Electronic form version 1.4