

Swept time-space domain decomposition on GPUs and heterogeneous computing systems

By Daniel J. Magee

Candidate for Master of Science in Mechanical Engineering

Major Professor: Dr. Kyle E. Niemeyer

Abstract

Modern scientific and engineering problems often require simulations with a level of resolution difficult to achieve in reasonable amounts of time—even in effectively parallelized programs. Therefore, applications that exploit high performance computing (HPC) systems have become invaluable in academia and industry over the past two decades. Addressing the questions that arise from continual scientific advancement requires solutions from hardware and software are required to supply the necessary throughput for demand across scientific disciplines.

The most important development on the hardware side has been the General Purpose Graphics Processing Unit (GPGPU), a class of massively parallel device that now composes a substantial portion of the computational power of the top 500 supercomputers. As these systems grow, barriers to increased performance arise from small costs accumulated over innumerable iterations such as latency, the fixed cost of memory accesses, which becomes significantly larger when access requires communication between two distant CPU processes. This thesis implements and analyzes swept time-space domain decomposition, a communication avoiding scheme for time-stepping stencil codes, for GPGPU and heterogeneous (CPU/GPU) architectures.

The GPGPU program significantly improves the execution time of finite-difference solvers for relatively simple one-dimensional time-stepping partial differential equations (PDEs). The swept decomposition code showed speedups of 2-9x compared with simple GPU domain decompositions and 7-300x compared with parallel CPU versions over a range of problem sizes: $10^3 - 10^6$ spatial points. However, for a more sophisticated one-dimensional system of equations discretized with a second-order finite-volume scheme, the swept rule performs 1.2-1.9x than a standard implementation for all problem sizes. The program targeting heterogeneous systems with distributed memory patterns performs significantly better on both simple problems, speedup 4-18x, and more complex equation systems, speedup 1.5-3x, over the range of problem sizes: 10^5-10^7 spatial points. This demonstrates the benefit of GPU architecture and the contingent effectiveness of swept time-space decomposition for accelerating explicit PDE solvers on current computational architectures.

Friday, June 8, 2018

9:00 am, Rog 226

School of Mechanical, Industrial
and Manufacturing Engineering



Oregon State
University