REV - [5-19-10]

Benchmark Evaluation of Radar Processing Algorithms on Graphics Processor Units (GPUs)

Scott Sawyer, Rick Pancoast, Mike Iaquinto, Rathin Putatunda, Rex Bennett, John Broadbent, Scott Harrington, Edward Dunne

Lockheed Martin MS2
199 Borton Landing Road, P.O. Box 1027
Moorestown, NJ 08057-0927
{scott.m.sawyer, rick.pancoast, michael.a.iaquinto, rathindra.n.putatunda, rex.bennett, john.d.broadbent, scott.harrington, edward.dunne}@lmco.com

Abstract

Graphics Processor Units (GPUs) are known to offer high performance and power efficiency for processing algorithms that map well to their highly parallel architecture. Due to their lower cost, better power efficiency, and relative ease of development, GPUs show potential to replace Field-Programmable Gate Arrays (FPGAs) and general purpose processors (GPPs) in many defense applications, including radar signal processing. This briefing describes the effort to benchmark core radar signal processing algorithms on commercial-off-the-shelf (COTS) GPU processor architectures, and to evaluate relative GPU performance and the effectiveness of available software environments and Application Programming Interfaces (APIs) for implementing DoD applications.

Objectives

The objective of this research was to implement commonly used radar processing algorithms on GPU platforms to benchmark processor utilization and to evaluate the software environment. Benchmarks were developed using NVIDIA's CUDA library, and we examined the potential benefits of a VSIPL-based implementation. We also investigated the effects of the inherent I/O constraints of commercial GPUs on data-intensive algorithms. A discussion follows on techniques for mapping algorithms to the GPU architecture, opportunities to improve I/O and computational performance, and areas for future research.

Several key radar benchmark algorithms were selected for implementation by an inexperienced team of implementers in order to gauge the difficulty level of porting algorithms from a heterogeneous architecture, consisting of general purpose PowerPC processors and high density FPGAs, to a GPU-based architecture.

Performance metrics (including the programming environment, throughput, latency, size, weight, power and cost) of the GPU-based architecture are contrasted and compared with the more conventional heterogeneous GPP / FPGA architectures.

Conclusion

Comparisons of a GPU-based radar signal processor, contrasted with a more conventional heterogeneous mix of PowerPCs and FPGAs, show that the GPU-based architectures demonstrate a significant potential for higher achievable performance with a significant reduction in size, weight, power and cost. The software development environment for GPU architectures, while not quite as friendly as a GPP environment, has demonstrated a much simpler development environment and software life-cycle as compared with an algorithmic FPGA environment.