


Dr. Lars Nyland, UNC Chapel Hill and nVidia
April 29th, 2009
Room 213, Petty Science Building.
11:00am - 12:00pm

“Ubiquitous Compute Acceleration”

Abstract


In this talk, I will cover three topics: 1) The NVIDIA Tesla Architecture, 2) The CUDA programming language, and 3) and recent work on N-Body simulation, concluding with examples of similarly accelerated computing for broader audiences. The Tesla Architecture supports both graphics and non-graphics computation, using an array of custom processors on a single chip. The programming model is neither SIMD nor MIMD, but somewhere in between, where we can exploit the advantages of each. The current performance part has 240 processors running at 1.3 - 1.6 GHz. With dual-issue capabilities, this places the peak performance over 1 TFLOP. CUDA is the C programming language with a few extensions for programming the Tesla chips. These include thread launch/terminate, synchronization, sharing, and atomic operations.



In a collaborative effort with Jan Prins (UNC CS) and Mark Harris (NVIDIA), we have written an N-Body simulator using CUDA that runs on Tesla hardware. We achieve a sustained computational rate over 400 GFLOPS, or 30k bodies interacting at nearly 30 steps/second. This is substantially faster than any conventional CPU, as the core of the computation relies on $1/\sqrt{x}$, which is implemented as an optimized instruction, because it is required in graphics (and physics) for normalizing vectors. I'll summarize with thoughts about the availability of accelerated computing.

Biography

Lars Nyland is a senior architect in the compute group at NVIDIA, where he designs, develops and tests architectural features to support non-traditional uses of graphics processors. Prior to joining NVIDIA, Lars was an associate professor of computer science at the Colorado School of Mines in Golden, Colorado. He ran the Thunder Graphics Lab, where demanding computational applications were coupled with immersive, 3D graphics. Prior to that, he was a member of the research faculty at UNC, Chapel Hill,



where he was a member of the high-performance computing and image-based rendering groups. Some notable achievements were the development of the DeltaSphere scene digitizer and its use at Monticello to provide an immersive experience for visitors to the New Orleans Museum of Art's Jefferson and Napoleon exhibit. He also spent considerable time studying N-Body algorithms, parallelizing N-Body algorithms for Molecular Dynamics, and parallel programming languages. Lars earned his PhD at Duke University in 1991 under the direction of John Reif, exploring high-level parallel programming languages.