

Advanced Solver Components on High-End Graphics Hardware

Manfred Liebmann

Institute for Mathematics and Scientific Computing

University of Graz

manfred.liebmann@uni-graz.at

ABSTRACT

With the introduction of the CUDA Toolkit for high performance graphics cards, Nvidia Corporation significantly lowered the barrier for the development of computational kernels that execute directly on the graphics processing units (GPUs).

The presentation will outline the porting issues that arise from the extremely high parallelism of the graphics hardware with hundreds of processing units and thousands of simultaneous threads. To really use the raw processing power of the graphics card it is absolutely mandatory to design data structures and algorithms with this high degree of parallelism in mind.

Although a whole new class of algorithmic difficulties arises on the graphics hardware, first results from a highly tuned sparse matrix-vector kernel give a 10 – 20 fold performance advantage over high-end Intel and AMD processors and thus promise high performance solver components on inexpensive graphics boards.