

## **CUDA Techniques in Computational Mechanics**

**Peng Wang**

### **Summary**

Current trends in high performance computing (HPC) are moving towards the availability of several cores on the same chip of contemporary processors in order to achieve speed-up through the extraction of potential fine-grain parallelism of applications. The trend is led by GPUs, which have been developed exclusively for computational tasks as massively-parallel co-processors to the CPU. During 2010 an extensive set of new HPC architectural features were developed in the third generation of NVIDIA GPUs (Fermi), giving computational mechanics an opportunity to expand use of GPU modelling and simulation.

This presentation will examine examples relevant to industry-scale HPC practice of GPU-accelerated computational structural mechanics and computational fluid dynamics software that support product design in manufacturing industries. For each example, the key CUDA porting strategy will be described. Performance results compare use of conventional CPUs with and without GPU acceleration.

