

GPU Technology Conference, May 14-17, 2012  
McEnery Convention Center, San Jose, California  
[www.gputechconf.com](http://www.gputechconf.com)

---

## Sessions on **Molecular Dynamics** (subject to change)

**IMPORTANT:** Visit <http://www.gputechconf.com/page/sessions.html> for the most up-to-date schedule.

---

### S0268 - Virtual Process Engineering - Realtime Simulation of Multiphase Systems

Wei Ge (Institute of Process Engineering, Chinese Academy of Sciences)

Day: Tuesday, 05/15 | Time: 9:00 am - 9:50 am

Topic Areas: Computational Fluid Dynamics; Molecular Dynamics; Computational Physics; Algorithms & Numerical Techniques

Session Level: Advanced

Realtime simulation and virtual reality with quantitatively correct physics for industrial processes with multi-scale and multiphase system is once a remote dream for process engineering, but is becoming true now with CPU-GPU hybrid supercomputing. Numerical and visualization methods for such simulations on thousands of GPUs will be reported with applications in chemical and energy industries.

### S0351 - Strong Scaling for Molecular Dynamics Applications

Sarah Tariq (NVIDIA)

Day: Tuesday, 05/15 | Time: 2:30 pm - 2:55 pm

Topic Areas: Molecular Dynamics; Cluster Management; Life Sciences

Session Level: Intermediate

In this session we will talk about how to improve strong scaling for molecular dynamics applications. Using the NAMD molecular dynamics code as our primary case study, we will discuss the types of issues that can impede scaling, how to use already available and custom tools to discover such issues, and how to build a model to help analyze and predict scaling performance. Although this session is primarily focused on molecular dynamics applications, most of the lessons can be applied equally well to many other areas and applications.

### S0036 - Multiparticle Collision Dynamics on GPUs

Elmar Westphal (Forschungszentrum Juelich)

Day: Tuesday, 05/15 | Time: 3:00 pm - 3:50 pm

Topic Areas: Computational Physics; Computational Fluid Dynamics; Molecular Dynamics

Session Level: Intermediate

See how we employ GPUs to simulate the interaction of millions of solvent and solute particles of a fluid system. Often the domain of large cluster system, the most time consuming part of our simulations can now be done on desktop PCs in reasonable time. This contribution shows how GPUs can effectively be used to accelerate existing programs and how techniques like streaming and increased data locality significantly enhance calculation throughput. It also shows how a GPU-optimized program structure yields usually expensive additional functionality "almost free". Furthermore, a well-scaling single-node/multi-GPU implementation of the program is presented.

**S0108 - An Innovative Massively Parallelized Molecular Dynamic Software**

Thomas Guignon (IFPEN), Ani Anciaux Sedrakian (IFP Energie Nouvelles)

Day: Tuesday, 05/15 | Time: 4:00 pm - 4:25 pm

Topic Areas: Molecular Dynamics; Supercomputing; Application Design &amp; Porting Techniques

Session Level: Intermediate

In this paper, we present how we improved the speedup of the electronic structure calculator VASP by more than an order of magnitude. Recently, the research works done (at IFP Energies Nouvelles) have shown that by coupling traditional clusters or High Performance Computing (HPC) machines with accelerators based on graphical processor units (GPUs), by recording the most time consuming parts of the codes (with programming languages like CUDA, OpenCL) and offloading them on the graphic chips, it is possible to reduce the computing time to ensure a speedup of a factor of 5 to 15.

**S0171 - Numerical Modeling of 3D Anisotropic Seismic Wave Propagation On MultiGPU Platforms**

Denis Sabitov (Schlumberger)

Day: Wednesday, 05/16 | Time: 9:00 am - 9:50 am

Topic Areas: Energy Exploration; Algorithms &amp; Numerical Techniques; Supercomputing; Molecular Dynamics

Session Level: Intermediate

We present an efficient and accurate numerical algorithm for the simulation of seismic experiments. The basis of the approach is a heterogeneous spectral element method implemented on MultiGPU applied to anisotropic elastic wave equation. The approach was designed to simulate wave propagation in 3D arbitrary anisotropic elastic media. Due to the use of an unstructured grid, the spectral element algorithm enables handling complicate geometries of the layers. We discuss results and computational efforts of simulation on MultiGPU platform. Several aspects of the code implementation are considered: optimal domain decomposition, data transfers between GPU by means of P2P and UVA, etc.

**S0010 - Towards Routine Microsecond Molecular Dynamics Simulations on Commodity Hardware**

Ross Walker (University of California San Diego)

Day: Wednesday, 05/16 | Time: 9:00 am - 9:50 am

Topic Areas: Molecular Dynamics; Life Sciences

Session Level: Advanced

The original AMBER 11 provided performance on one GPU equivalent to an 8 node cluster and almost 60ns/day for 8 GPUs running the JAC production benchmark without additional approximations outstripping the performance of all conventional supercomputers. Here we describe further optimization of the code, coupled with hardware and software advances on the part of NVIDIA, that provides performance of >50ns/day on a single GPU with multiple GPUs providing simulation rates on systems the size of DHFR approaching a microsecond per day. This brings performance levels on desktops and commodity hybrid clusters to levels previously only considered possible using custom silicon.

**S0058 - Advancing GPU Molecular Dynamics: Rigid Bodies in HOOMD-blue**

Joshua Anderson (University of Michigan), Trung Dac Nguyen (University of Michigan)

Day: Wednesday, 05/16 | Time: 10:00 am - 10:50 am

Topic Areas: Molecular Dynamics; Computational Physics

Session Level: Intermediate

Learn how rigid body dynamics are implemented in HOOMD-blue. Previous releases were capable of executing classical molecular dynamics -- where free particles interact via smooth potentials and their motion through time is computed using Newton's laws. The latest version allows particles to be grouped into bodies that move as rigid units. Users can now simulate materials made of cubes, rods, bent rods, jacks, plates, patchy particles, bucky

balls, or any other arbitrary shapes. This talk covers how these algorithms are implemented on the GPU, tuned to perform well for bodies of any size, and discusses several use-cases relevant to research.

#### **S0142 - VMD: High Performance Molecular Visualization and Analysis on GPUs**

**John Stone (University of Illinois at Urbana-Champaign)**

**Day:** Wednesday, 05/16 | **Time:** 2:00 pm - 2:50 pm

**Topic Areas:** Molecular Dynamics; Algorithms & Numerical Techniques; Computer Graphics

**Session Level:** Intermediate

This talk will present recent successes in the use of GPUs to accelerate interactive molecular visualization and analysis tasks on desktop computers, and batch-mode simulation and analysis jobs on GPU-accelerated HPC clusters. We'll present Fermi-specific algorithms and optimizations and compare with those for other devices. We'll also present performance and performance/watt results for VMD analysis calculations on GPU clusters, and conclude with a discussion of ongoing work and future opportunities for GPU acceleration, particularly as applied to the analysis of petascale simulations of large biomolecular complexes and long simulation timescales.

#### **S0127 - Petascale Molecular Dynamics Simulations on GPU-Accelerated Supercomputers**

**James Phillips (University of Illinois)**

**Day:** Wednesday, 05/16 | **Time:** 3:00 pm - 3:25 pm

**Topic Areas:** Molecular Dynamics; Application Design & Porting Techniques; Parallel Programming Languages & Compilers; Supercomputing

**Session Level:** Intermediate

The highly parallel molecular dynamics code NAMD was chosen in 2006 as a target application for the NSF petascale supercomputer now known as Blue Waters. NAMD was also one of the first codes to run on a GPU cluster when G80 and CUDA were introduced in 2007. How do the Cray XK6 and modern GPU clusters compare to 300,000 CPU cores for a hundred-million-atom Blue Waters acceptance test? Come learn the opportunities and pitfalls of taking GPU computing to the petascale and the importance of CUDA 4.0 features in combining multicore host processors and GPUs in a legacy message-driven application.

#### **S0207 - GPU Enabled Macromolecular Simulation: Challenges and Opportunities**

**Michela Taufer (University of Delaware), Sandeep Patel (University of Delaware)**

**Day:** Wednesday, 05/16 | **Time:** 3:30 pm - 3:55 pm

**Topic Areas:** Molecular Dynamics; Algorithms & Numerical Techniques

**Session Level:** Advanced

GPU enabled simulation of fully atomistic macromolecular simulation is rapidly gaining momentum, enabled by the massive parallelism and due to parallelizability of various components of the underlying algorithms and methodologies. The massive parallelism in the order of several hundreds to few thousands of cores, presents opportunities as well poses implementation challenges. In this talk dive deep into the various key aspects of simulation methodologies of macro molecular systems specifically adapted to GPUs. Learn some of the underlying challenges and get the latest solutions devised to tackle them in the FENZI code for fully atomistic macromolecular simulations.

#### **S0363 - Efficient Molecular Dynamics on Heterogeneous GPU Architectures in GROMACS**

**Szilárd Páll (KTH Royal Institute of Technology), Berk Hess (KTH Royal Institute of Technology)**

**Day:** Wednesday, 05/16 | **Time:** 4:00 pm - 4:25 pm

**Topic Areas:** Molecular Dynamics; Computational Physics; Life Sciences

**Session Level:** Intermediate

Molecular Dynamics is an important application for GPU acceleration, but many algorithmic optimizations and features still rely on code that prefers traditional CPUs. It is only with the latest hardware and software we have been able to realize a heterogeneous GPU/CPU implementation and reach performance significantly beyond the state-of-the-art of hand-tuned CPU code in our GROMACS program. The sub-millisecond iteration time poses challenges on all levels of parallelization. Come and learn about our new atom-cluster pair interaction approach for non-bonded force evaluation that achieves 60% work-efficiency and other innovative solutions for heterogeneous GPU systems.

### **S0139 - GPU-Based Molecular Dynamics Simulations of Protein and RNA Assembly**

**Samuel Cho (Wake Forest University)**

**Day:** Wednesday, 05/16 | **Time:** 5:00 pm - 5:25 pm

**Topic Areas:** Molecular Dynamics; Computational Physics

**Session Level:** Intermediate

Protein and RNA biomolecular folding and assembly problems have important applications because misfolding is associated with diseases like Alzheimer's and Parkinson's. However, simulating complex biomolecules on the same timescales as experiments is an extraordinary challenge due to a bottleneck in the force calculations. To overcome these hurdles, we perform coarse-grained molecular dynamics simulations where biomolecules are reduced into simpler components. Furthermore, our GPU-based simulations have a significant performance improvement over CPU-based simulations, which is limited to systems of 50-150 residues/nucleotides. The GPU-based code can simulate protein/RNA systems of 400-10,000+ residues/nucleotides, and we present ribosome assembly simulations.

### **S0057 - GPU-Accelerated Molecular Dynamics Simulation of Solid Covalent Crystals**

**Chaofeng Hou (Institute of Process Engineering, Chinese Academy of Sciences)**

**Day:** Thursday, 05/17 | **Time:** 9:00 am - 9:25 am

**Topic Areas:** Molecular Dynamics; Algorithms & Numerical Techniques; Supercomputing

**Session Level:** Intermediate

An efficient and highly scalable algorithm for molecular dynamics (MD) simulation (using sophisticated many-body potentials) of solid covalent crystals is presented. Its effective memory throughput on a single C2050 GPU board reached 102 GB/s (81% of the peak), the instruction throughput reached 412 Ginstr/s (80% of the peak), and 27% of the peak flops of a single GPU was obtained. Parallel efficiency of the algorithm can be as high as 95% on all 7168 GPUs of Tianhe-1A, reaching possibly a record in high performance of MD simulations, 1.87Pflops in single precision.

### **S0315 - Probing Bio-Nano Interface Structure from Microsecond Molecular Dynamics on GPUs**

**Olexandr Isayev (Case Western Reserve University)**

**Day:** Thursday, 05/17 | **Time:** 10:00 am - 10:25 am

**Topic Areas:** Molecular Dynamics; Life Sciences

**Session Level:** Advanced

Using the latest algorithmic development in molecular dynamics on multiple GPUs over MPI, and technologies like GPUDirect it is now possible to address problems of interaction at bio-nano interface via large scale atomistic simulations. This talk will discuss the aspects of DNA-nanotube interactions and SWCNT induced conformational changes in DNA nucleosome structure. We will also address technical challenges upon porting and tuning AMBER 11 code on Condor GPU cluster at AFRL.

**S0122 - Computational Screening of Novel Carbon Capture Materials**

Jihan Kim (Berkeley Lab), Berend Smit (UC Berkeley/Berkeley Lab)

Day: Thursday, 05/17 | Time: 10:30 am - 10:55 am

Topic Areas: Molecular Dynamics

Session Level: Intermediate

Discover how GPUs are used to identify optimal framework structures for carbon dioxide separation with the goal of reducing carbon emission. We describe the algorithm behind our GPU software tool that iterates through a database of hypothetical zeolites and computes the selectivity of each of the structures. The code can be easily extended to simulate other adsorbent structures such as ZIFs (zeolitic imidazolate frameworks) and provide valuable insights to both theorists and experimentalists who have interest in carbon capture research.

**S0334 - The Fast Multipole Method on CPU and GPU Processors**

Eric Darve (Stanford)

Day: Thursday, 05/17 | Time: 3:00 pm - 3:25 pm

Topic Areas: Computational Physics; Molecular Dynamics; Algorithms &amp; Numerical Techniques

Session Level: Advanced

The fast multipole method (FMM) is a widely used numerical algorithm in computational engineering. Accelerating the FMM on CUDA-enabled GPUs is challenging because the FMM has a complicated data access pattern, mostly during the so-called multipole-to-local (M2L) operation. We have created several schemes to optimize the M2L and have attained a performance of over 350 (resp. 160) Gflop/s for single (double) precision arithmetic. The optimal algorithm was incorporated into a complete FMM code, which can accept any smooth kernel as specified by the user, making it very flexible. We have also developed a highly efficient CPU version.

**S0111 - An Efficient CUDA Implementation of a Tree-Based N-Body Algorithm**

Martin Burtscher (Texas State University)

Day: Thursday, 05/17 | Time: 3:30 pm - 4:20 pm

Topic Areas: Application Design &amp; Porting Techniques; Astronomy &amp; Astrophysics; Molecular Dynamics; Supercomputing

Session Level: Advanced

This session presents a complete CUDA implementation of the irregular Barnes-Hut n-body algorithm. This algorithm repeatedly builds and traverses unbalanced trees, making it difficult to map to GPUs. We explain in detail how our code exploits the architectural features of GPUs, including lockstep operation and thread divergence, both of which are commonly viewed as hurdles to achieving high performance, especially for irregular codes. On a five million body simulation running on a Tesla C2050, our CUDA implementation is 30 times faster than a parallel pthreads version running on a high-end 6-core Xeon.