UNIVERSITY of DELAWARE

GPU-enabled Macromolecular Simulations: Challenges and Opportunities

Michela Taufer[1]
with Sandeep Patel[1]
and Narayan Ganesan[2]

[1] University of Delaware
[2] Stevens Institute of Technology

MD Simulations

- MD simulations study the dynamics of individual molecules (molecular positions) as in a motion picture[1]
  - MD simulations are iterative executions of MD steps
  - Given initial atomic positions and velocities, obtain those at later times by integrating ordinary differential equations
- MD simulations complement experiments
  - Critical for atomic-level insights
- Limitations in MD simulations:
  - Length and time scales restricted at the fully atomistic level
- GPUs can provide us with the computing power to cope with large length and time scales

[1] J.M. Haile, Molecular Dynamics Simulation, John Wiley and Sons, Inc. (1992)

MD on GPUs

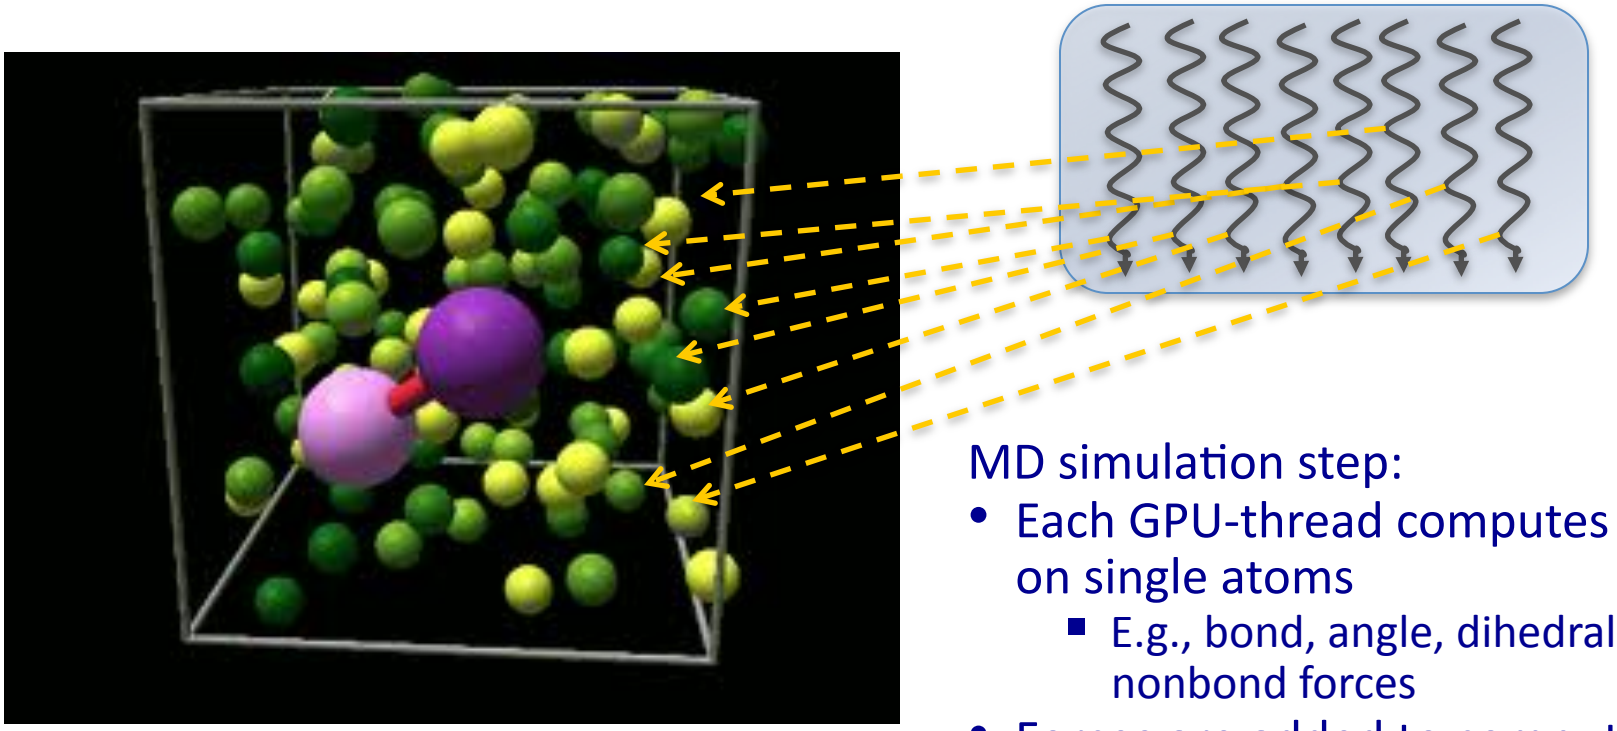

Force -> Acceleration -> Velocity -> Position

MD simulation step:

- Each GPU-thread computes forces on single atoms
  - E.g., bond, angle, dihedrals and, nonbond forces
- Forces are added to compute acceleration
- Acceleration is used to update velocities
- Velocities are used to update the positions

Aspects of Realistic MD Simulation

- Realism of model: the mathematical model reproduces the behavior of the real physical system[1]
- Validity and accuracy of simulation: simulations may suffer from uncertainties in reaching an equilibrium or errors[1]
  - Statistical, numerical, and round-off errors

[1] J.M. Haile, Molecular Dynamics Simulation, John Wiley and Sons, Inc. (1992)

FLOPS

NS/DAY

PERFORMANCE

NVE

MD OF ENSEMBLES

NPT

NVT

SCALABILITY

MATHEMATICAL MODEL

INTERACTION POTENTIAL

ASPECTS OF REALISTIC MD SIMULATION

ENERGY FLUCTUATIONS

FORCE FIELD

ELECTROSTATIC INTERACTIONS

VALIDATION AND ACCURACY

PME

FREEZING FAST DEGREES OF MOTION

EMPIRICAL VALIDATION

IMPOSING CONSTRAINTS ON INTERATOMIC DISTANCES

SINGLE PEPTIDE DIFFUSION

SHAKE / RATTLE

STUDY PHYSICAL-CHEMICAL SYSTEMS

RESTRAINTS

PROTEIN-MEMBRANE INTERACTION

FLOPS

NS/DAY

PERFORMANCE

NVE

MD OF ENSEMBLES

NPT

NVT

SCALABILITY

MATHEMATICAL MODEL

INTERACTION POTENTIAL

ASPECTS OF REALISTIC MD SIMULATION

ENERGY FLUCTUATIONS

FORCE FIELD

ELECTROSTATIC INTERACTIONS

VALIDATION AND ACCURACY

PME

FREEZING FAST DEGREES OF MOTION

EMPIRICAL VALIDATION

IMPOSING CONSTRAINTS ON INTERATOMIC DISTANCES

SINGLE PEPTIDE DIFFUSION

SHAKE / RATTLE

STUDY PHYSICAL-CHEMICAL SYSTEMS

RESTRAINTS

PROTEIN-MEMBRANE INTERACTION

FLOPS

NS/DAY

PERFORMANCE

NVE

MD OF ENSEMBLES

NPT

NVT

SCALABILITY

MATHEMATICAL MODEL

INTERACTION POTENTIAL

ASPECTS OF REALISTIC MD SIMULATION

ENERGY FLUCTUATIONS

FORCE FIELD

ELECTROSTATIC INTERACTIONS

VALIDATION AND ACCURACY

PME

FREEZING FAST DEGREES OF MOTION

EMPIRICAL VALIDATION

IMPOSING CONSTRAINTS ON INTERATOMIC DISTANCES

SINGLE PEPTIDE DIFFUSION

SHAKE / RATTLE

STUDY PHYSICAL-CHEMICAL SYSTEMS

RESTRAINTS

PROTEIN-MEMBRANE INTERACTION

FLOPS

NS/DAY

NVE

PERFORMANCE

MD OF ENSEMBLES

NPT

SCALABILITY

NVT

MATHEMATICAL MODEL

INTERACTION POTENTIAL

ASPECTS OF REALISTIC MD SIMULATION

ENERGY FLUCTUATIONS

FORCE FIELD

ELECTROSTATIC INTERACTIONS

VALIDATION AND ACCURACY

PME

FREEZING FAST DEGREES OF MOTION

EMPIRICAL VALIDATION

IMPOSING CONSTRAINTS ON INTERATOMIC DISTANCES

SINGLE PEPTIDE DIFFUSION

SHAKE / RATTLE

STUDY PHYSICAL-CHEMICAL SYSTEMS

RESTRAINTS

PROTEIN-MEMBRANE INTERACTION

FLOPS

NSEC/DAY

PERFORMANCE

SCALABILITY

NVE

MD OF ENSEMBLES

NPT

NVT

MATHEMATICAL MODEL

INTERACTION POTENTIAL

ASPECTS OF REALISTIC MD SIMULATION

ENERGY FLUCTUATIONS

FORCE FIELD

ELECTROSTATIC INTERACTIONS

VALIDATION AND ACCURACY

PME

FREEZING FAST DEGREES OF MOTION

EMPIRICAL VALIDATION

IMPOSING CONSTRAINTS ON INTERATOMIC DISTANCES

SINGLE PEPTIDE DIFFUSION

SHAKE / RATTLE

STUDY PHYSICAL-CHEMICAL SYSTEMS

RESTRAINTS

PROTEIN-MEMBRANE INTERACTION

FLOPS

NS/DAY

PERFORMANCE

NVE

MD OF ENSEMBLES

NPT

NVT

SCALABILITY

MATHEMATICAL MODEL

INTERACTION POTENTIAL

ASPECTS OF REALISTIC MD SIMULATION

ENERGY FLUCTUATIONS

FORCE FIELD

ELECTROSTATIC INTERACTIONS

VALIDATION AND ACCURACY

PME

FREEZING FAST DEGREES OF MOTION

EMPIRICAL VALIDATION

IMPOSING CONSTRAINTS ON INTERATOMIC DISTANCES

SHAKE / RATTLE

RESTRAINTS

SINGLE PEPTIDE DIFFUSION

STUDY PHYSICAL-CHEMICAL SYSTEMS

PROTEIN-MEMBRANE INTERACTION

FEN ZI

Yun Dong de FEN ZI = Moving MOLECULES

- FEN ZI enables GPU-based MD simulations in NVT, NVE, and NPT[1] ensembles and energy minimization[2]
  - MD forces are all computed on GPU
- Force field used: CHARMM force field[3]
- Lennard-Jones interactions:
  - Switching or shifting
- Long distance electrostatic interactions:
  - Ewald summation method[5]
  - Reaction field[6]
- Solvent:
  - Explicit or implicit model
  - TIP3 water model
  - Flexible SPC/Fw water model[4]

[1] H. C. Andersen, J. Chem. Phys., 72 (1980) 2384-2393
[2] M. C. Payne, et al., Rev. Mod. Phys., 64 (1992) 1045-1097
[3] B. R. Brooks, et al., J. Comp. Chem., 4 (1983) 187{217
[4] Y. Wu, et al., J. Chem. Phys., 124, 024503, 2006
[5] U. Essmann, et al., J. Chem. Phys., 103 (1995) 8577
[6] G. Hummer, et al., J. Phys. Condens. Matter (1994)

FLOPS

NS/DAY

NVE

PERFORMANCE

MD OF ENSEMBLES

NPT

SCALABILITY

NVT

MATHEMATICAL MODEL

INTERACTION POTENTIAL

ASPECTS OF REALISTIC MD SIMULATION

ENERGY FLUCTUATIONS

FORCE FIELD

VALIDATION AND ACCURACY

ELECTROSTATIC INTERACTIONS

PME

FREEZING FAST DEGREES OF MOTION

EMPIRICAL VALIDATION

IMPOSING CONSTRAINTS ON INTERATOMIC DISTANCES

SINGLE PEPTIDE DIFFUSION

STUDY PHYSICAL-CHEMICAL SYSTEMS

SHAKE / RATTLE

PROTEIN-MEMBRANE INTERACTION

RESTRAINTS

Short and Long Range Interactions

- Each iteration computes forces on each particle due to:
  - Bonded interactions [1]
    - Bonds
    - Angles (ANGLes, UREY-b)
    - Dihedrals
    - Improper
  - Non-bonded interactions [1]
    - Van der Waals
    - Electrostatic with PME
      - Direct space energy
      - Reciprocal space energy
      - Self energy

$$V_{LJ} = \sum_{i,j}^{\text{pairs}} \left(4\epsilon_{ij} \left[\left(\frac{\sigma_{ij}}{r_{ij}} \right)^{12} - \left(\frac{\sigma_{ij}}{r_{ij}} \right)^{6} \right] \right)$$

$$E_{\text{dir}} = \sum_{i=1}^{N-1} \sum_{j>i}^{N} \frac{q_i q_j \operatorname{erfc}(\beta r_{ij})}{r_{ij}}$$

$$E_{\text{rec}} = \frac{1}{2\pi V} \sum_{\vec{m} \neq 0} \frac{\exp(-\pi^2 \vec{m}^2 / \beta^2)}{\vec{m}^2} S(\vec{m}) S(-\vec{m})$$

$$E_{\text{corr}} = -\frac{1}{2} \sum_{(i,j) \in \text{Excl}} \frac{q_i q_j \operatorname{erf}(\beta |r_i - r_j|)}{|r_i - r_j|} - \frac{\beta}{\sqrt{\pi}} \sum_{i=1} N q_i^2$$

[1] J. Phys. Chem. B, 1998, 102, 3586; J. Comput. Chem., 2004, 25, 1400; J. Comput. Chem., 2000, 21, 86, ibid. 105ff

Interaction Implementation on GPUs

- Bond-, angle-, and dihedral interactions handled by a kernel
- Bond-, angle-, and dihedral lists never require updating
  - Constructed once on CPU and copied to GPU
- Non-bonded interactions (i.e., Lennard-Jones and direct space electrostatic) are handled by two kernels
  - One for building non-bond neighbor list and one for computing interactions

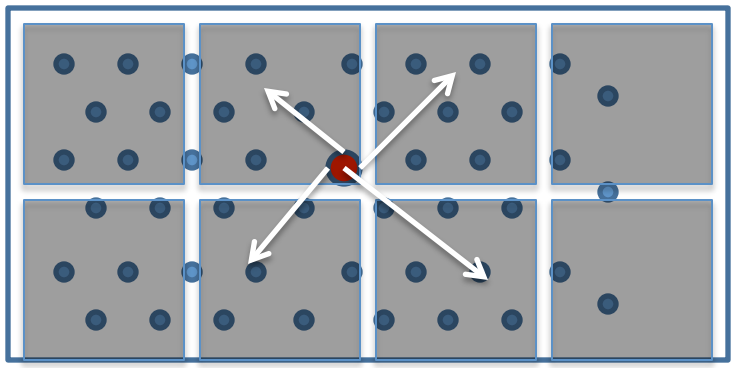

- Cell-based neighbor list
  - Divide the domain into equal cells of size = cutoff r_{cut}
  - Search only in current cell and 26 adjacent cells for neighboring atoms
  - List only needs to be updated when atoms move more than buffer cutoff $r_{list} > r_{cut}$

Electrostatic Interactions

- Divide interactions into short range (Direct Space) and long range (Reciprocal Space)

Direct space using non-bond list (Short Range)

$$E_{dir} = \frac{q_i q_j \mathrm{erfc}(\beta \mid r_i - r_j \mid)}{4\pi\varepsilon_0\varepsilon_1 \mid r_i - r_j \mid}$$

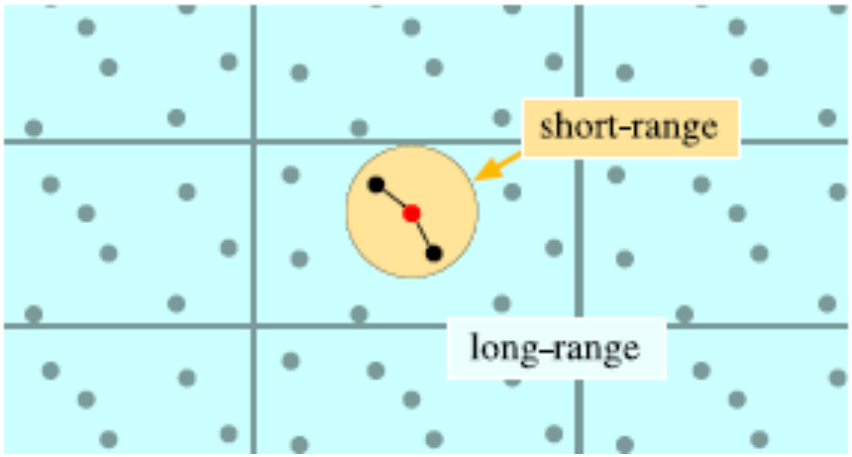

Reciprocal space using FTTs (Long Range)

$$E_{rec} = \frac{1}{2\pi V} \sum_{m \neq 0} \frac{\exp(-\pi^2 m^2 / \beta^2)}{m^2} S(m) S(-m)$$

Steps in SPME

Smooth Particle Mesh Ewald (SPME)

1 Put charges on grids

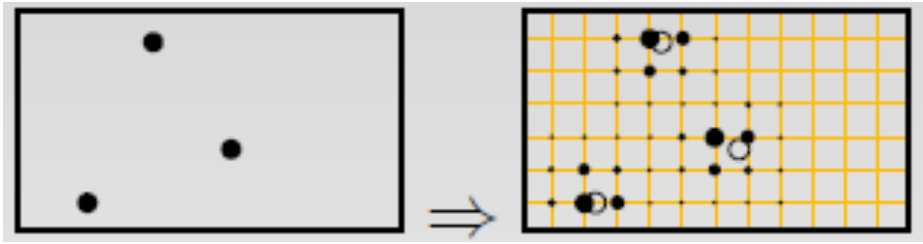

2 FFT of charge grid

3 Multiply charges by structure constants

4 FFT back

Convolution yields potential at grid points which have to be summed

5 Compute force on atom *i* by calculating $\frac{\partial U}{\partial r_i}$

Charge Spreading

- Each charge spread on a 4x4x4 grid points in 3-D[1]
  - Grid spacing 1 A by a cardinal B-Spline of order 4
  - Create a 3 dimensional Charge Matrix "Q"
- Mesh-based charge density
  - Approximation by sum of charges at each grid point
  - Multiple charges can influence a single lattice point

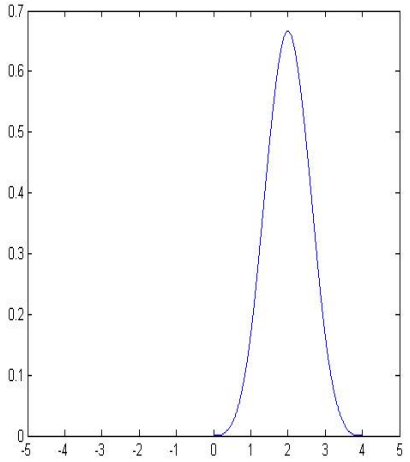

Charges

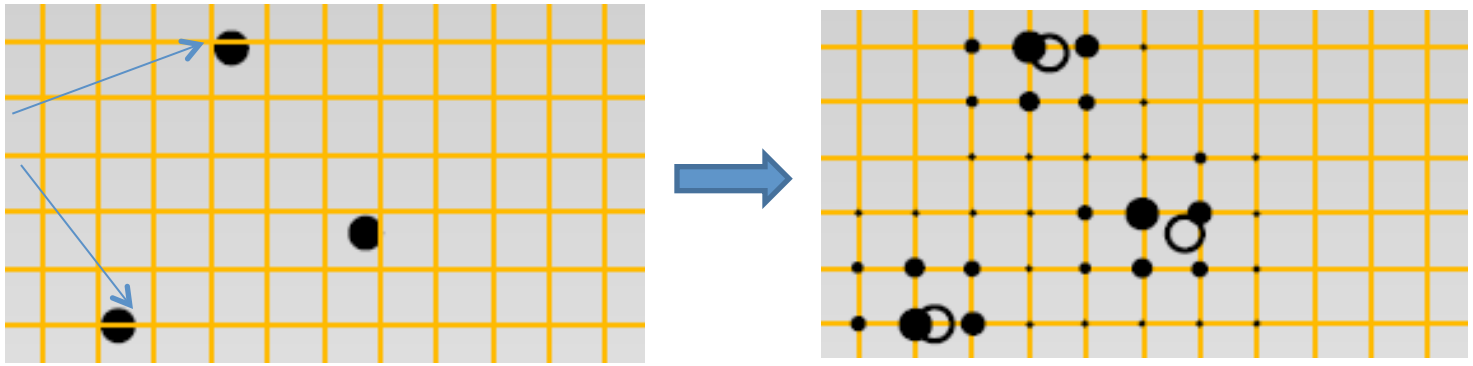

$$Q(k_1,k_2,k_3) = \sum_{i=1..n} q_i M_4(x_i - k_1) M_4(y_i - k_2) M_4(z_i - k_3)$$

x_i y_i z_i: position of the i^{th} charge; k_1 k_2 k_3: index of the lattice point

[1] Essman *et al.*, J. Chem. Phys. 1995

CPU vs. GPU Charge Spreading

CPU spreading of charges:

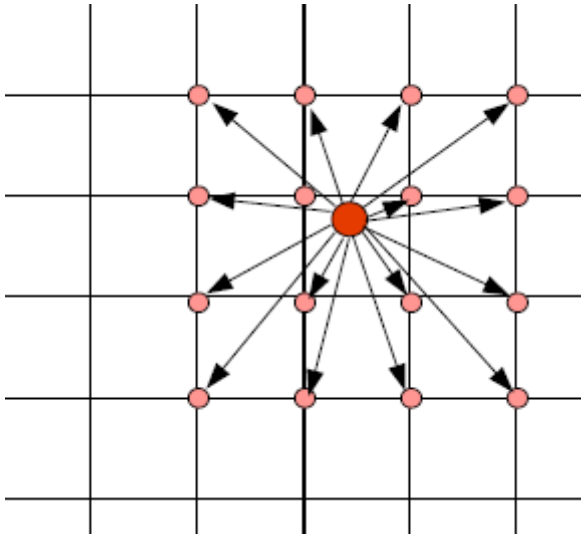

GPU gathering of charges by a cardinal B-Spline of order 4:

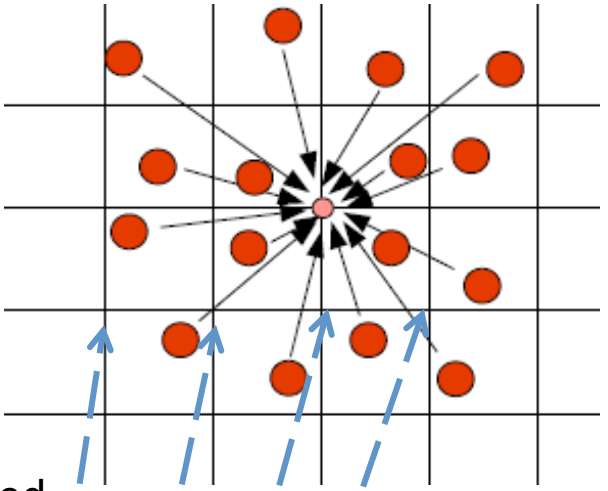

Each thread is assigned to a lattice point

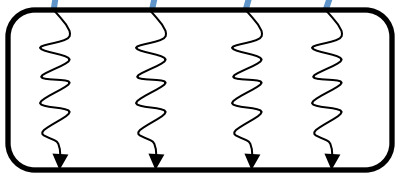

- Charge spreading on GPU can be parallelized easily by the grid points instead of the atoms
- Each thread works on a single or a set of grid points

Efficient and Scalable Charge Spreading

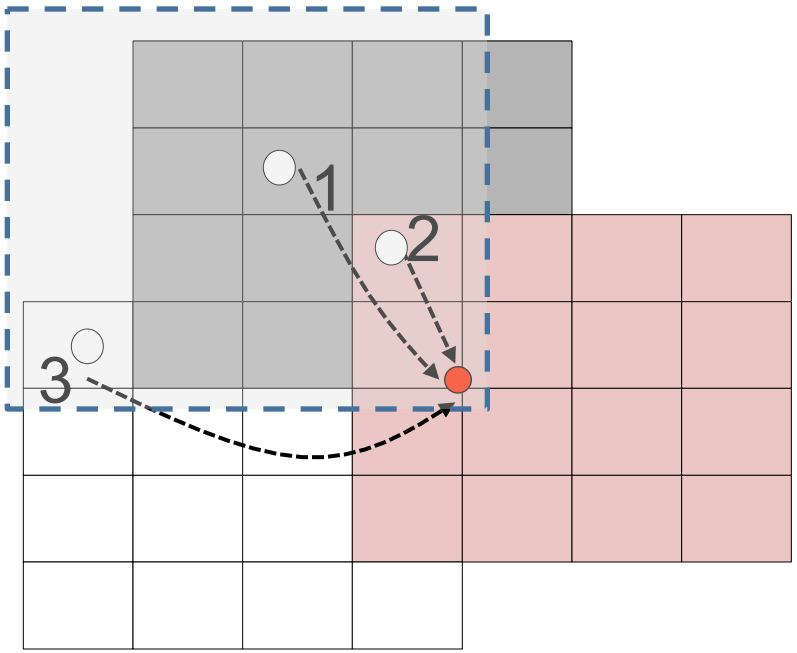

- Our charge spreading on GPUs is done by maintaining a list of charges within a 4x4x4 neighborhood of each lattice point
- The lattice neighbor list is built only once and efficiently updated throughout the rest of the simulation

Efficient and Scalable Charge Spreading

- When a charge moves from location 2 to 2', the neighborhood list of associated lattice points needs to be updated
- We update the neighborhood lists in parallel
  - A 1-to-1 mapping between lattice points gaining the charge and lattice points losing the charge is used
  - Threads of points losing the charge update list of points gaining the charge

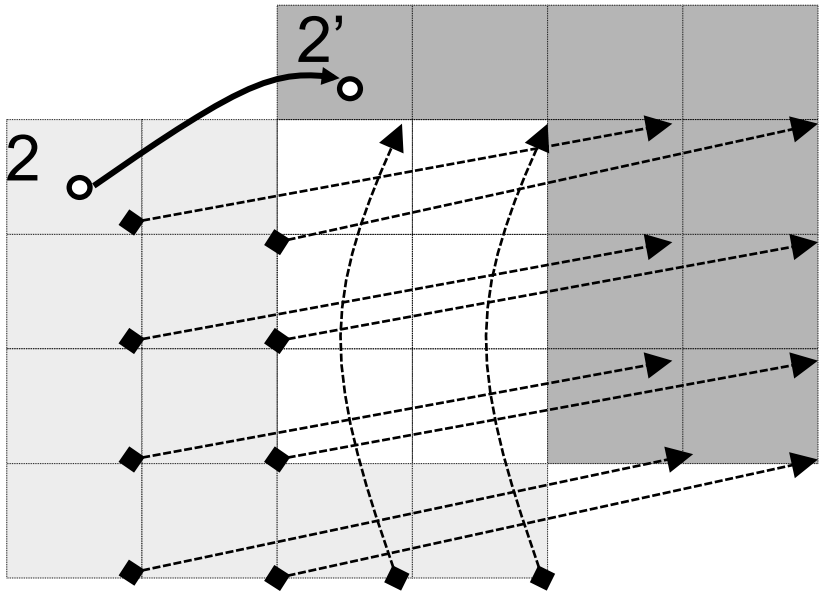

FLOPS

NS/DAY

PERFORMANCE

NVE

MD OF ENSEMBLES

NPT

NVT

SCALABILITY

MATHEMATICAL MODEL

INTERACTION POTENTIAL

ASPECTS OF REALISTIC MD SIMULATION

ENERGY FLUCTUATIONS

FORCE FIELD

ELECTROSTATIC INTERACTIONS

VALIDATION AND ACCURACY

PME

FREEZING FAST DEGREES OF MOTION

EMPIRICAL VALIDATION

IMPOSING CONSTRAINTS ON INTERATOMIC DISTANCES

SINGLE PEPTIDE DIFFUSION

SHAKE / RATTLE

STUDY PHYSICAL-CHEMICAL SYSTEMS

RESTRAINTS

PROTEIN-MEMBRANE INTERACTION

MD of Ensembles

- MD can simulate in the NVE, NVT, or NPT ensembles
- Microcanonical or constant NVE
  - Conventional MD simulation conserving total energy
- There exists different algorithms to implement temperature and pressure control mechanisms
- Canonical or constant NVT
  - Scaling velocities:
    - temperature depends on velocities
    - correct the velocities every 20,000 steps to keep desired temperature
- Isothermal-isobaric or constant NPT
  - Using a simple Nosé-Hoover style isothermal-isobaric molecular dynamics[1]

[1] Kalibaeva G. et al., *Molecular Physics* 101(6): 765-778, 2002

FLOPS

NS/DAY

NVE

PERFORMANCE

MD OF ENSEMBLES

NPT

SCALABILITY

NVT

MATHEMATICAL MODEL

INTERACTION POTENTIAL

ASPECTS OF REALISTIC MD SIMULATION

ENERGY FLUCTUATIONS

FORCE FIELD

ELECTROSTATIC INTERACTIONS

VALIDATION AND ACCURACY

PME

FREEZING FAST DEGREES OF MOTION

EMPIRICAL VALIDATION

IMPOSING CONSTRAINTS ON INTERATOMIC DISTANCES

SINGLE PEPTIDE DIFFUSION

SHAKE / RATTLE

STUDY PHYSICAL-CHEMICAL SYSTEMS

RESTRAINTS

PROTEIN-MEMBRANE INTERACTION

Cell Membranes

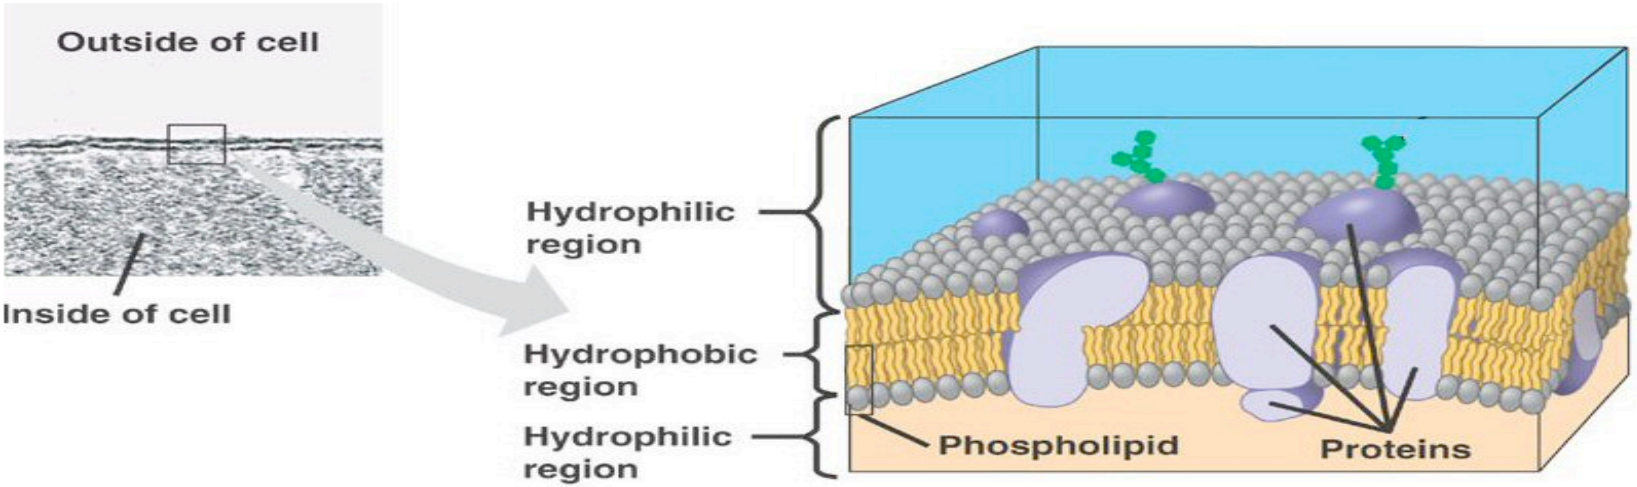

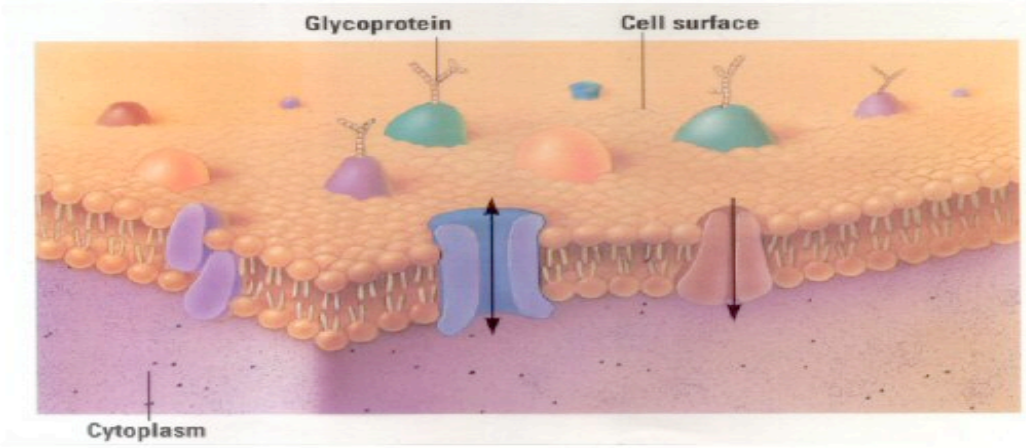

Lipid membranes are responsible for physiological functions and dysfunctions

DMPC Lipid Bilayers

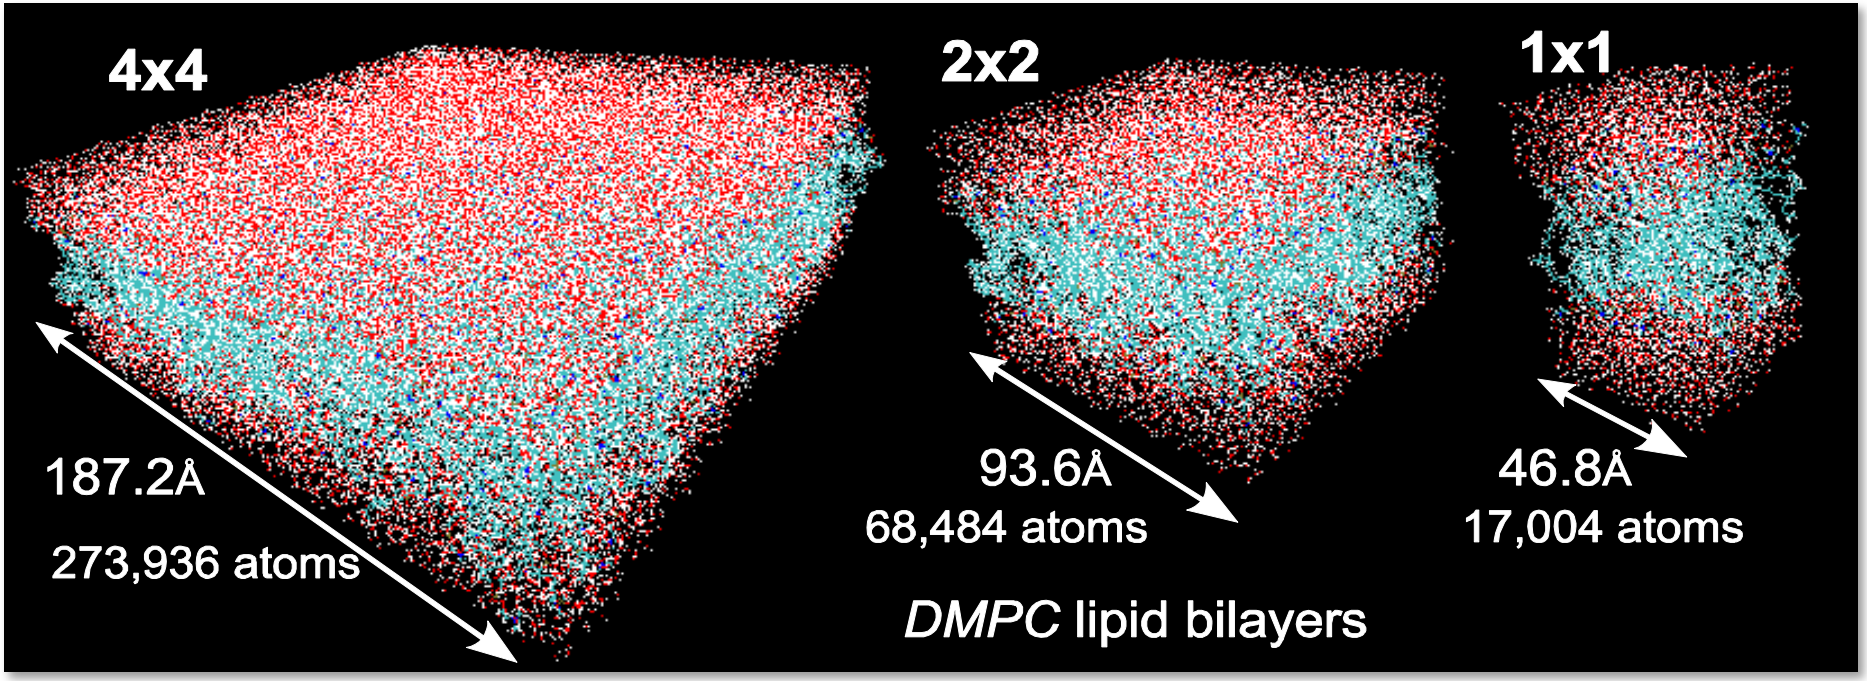

Large system Medium system Small system

DMPC Lipid Bilayers: Simulations (I)

Small system:

- Number of atoms: 17,004
  - 14,096 bonds
  - 19,108 angles
  - 22,536 dihedrals
- Size: 46.8A x 46.8A x 76.0A
- Water molecules: 2,836
- Temperature: 298 K

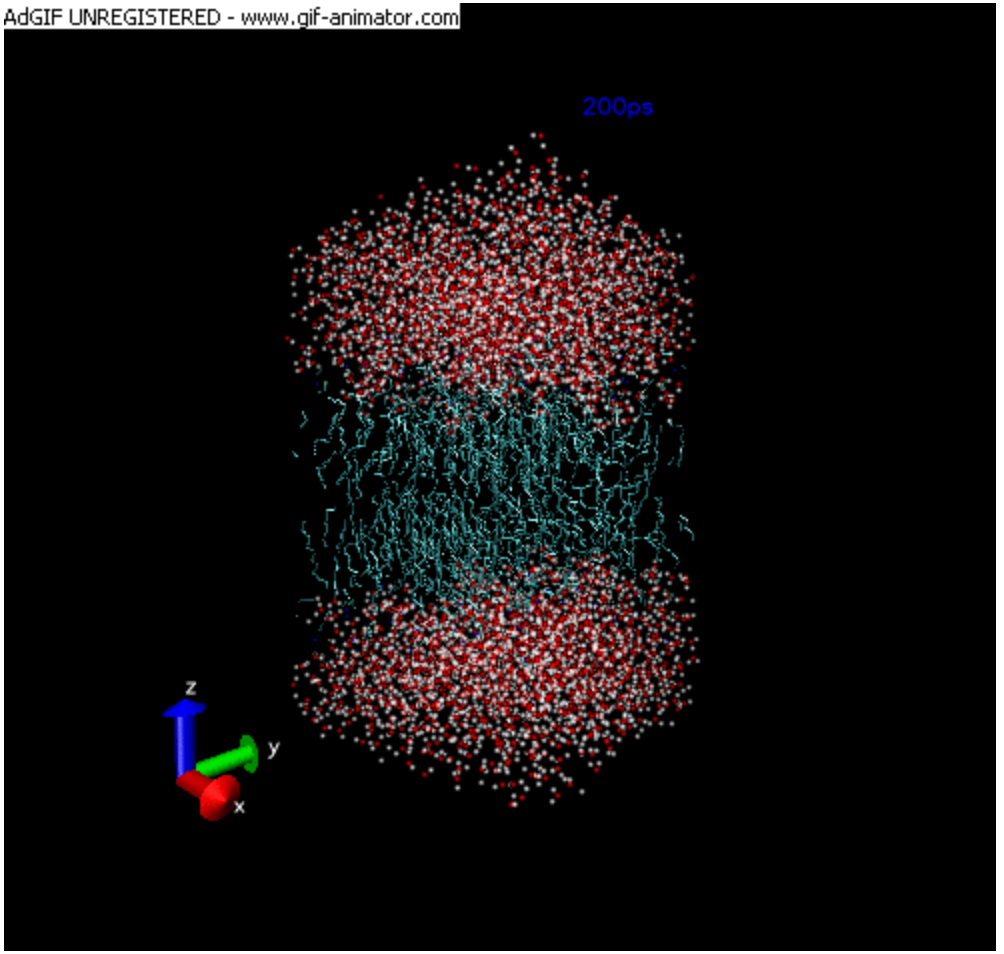

DMPC Lipid Bilayers: Simulations (II)

Medium system:

- Number of atoms: 68,484
  - 56,696 bonds
  - 76,588 angles
  - 360,576 dihedrals
- Size: 93.6 A x 93.6A x 76.0A
- Water molecules: 11,500
- Temperature: 298 K

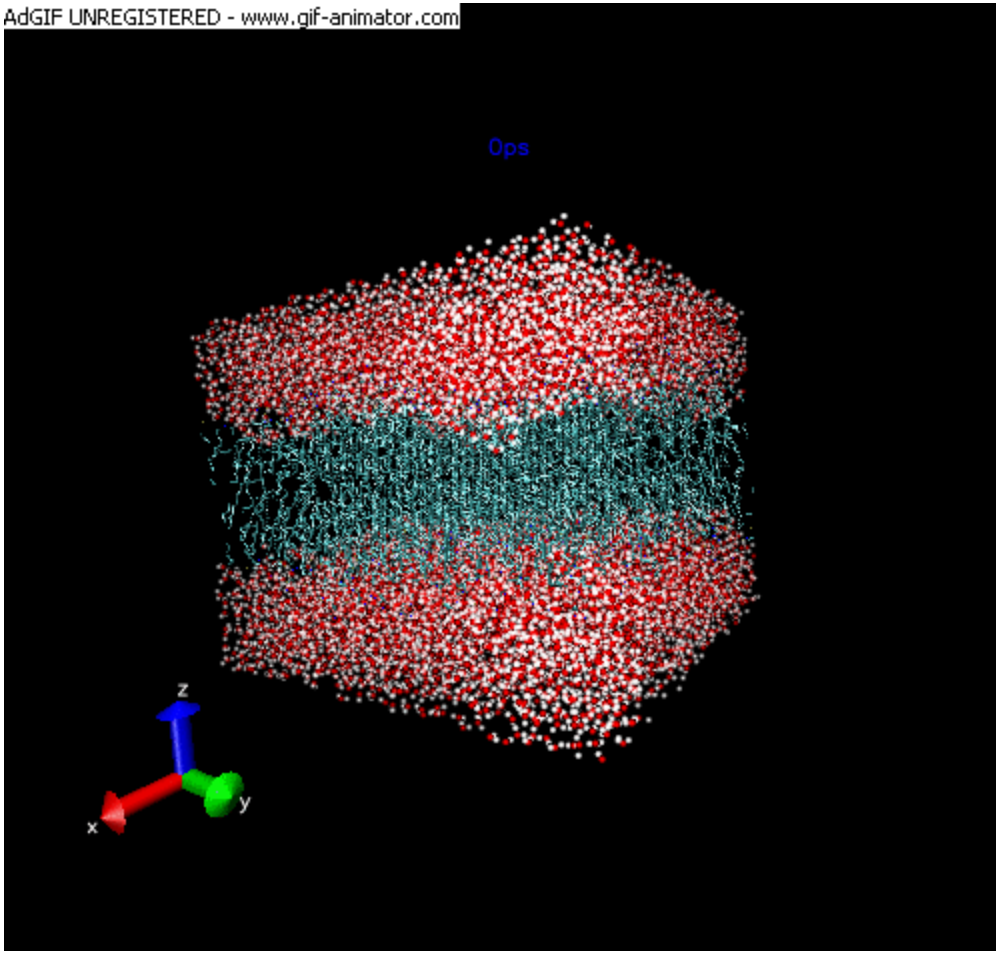

DMPC Lipid Bilayers: Simulations (III)

Large system:

- Number of atoms: 273,936
  - 226,784 bonds
  - 306,352 angles
  - 360,576 dihedrals
- Size: 187.2A X 187.2A X 76.0A
- Water molecules: 46,863
- Temperature: 298 K

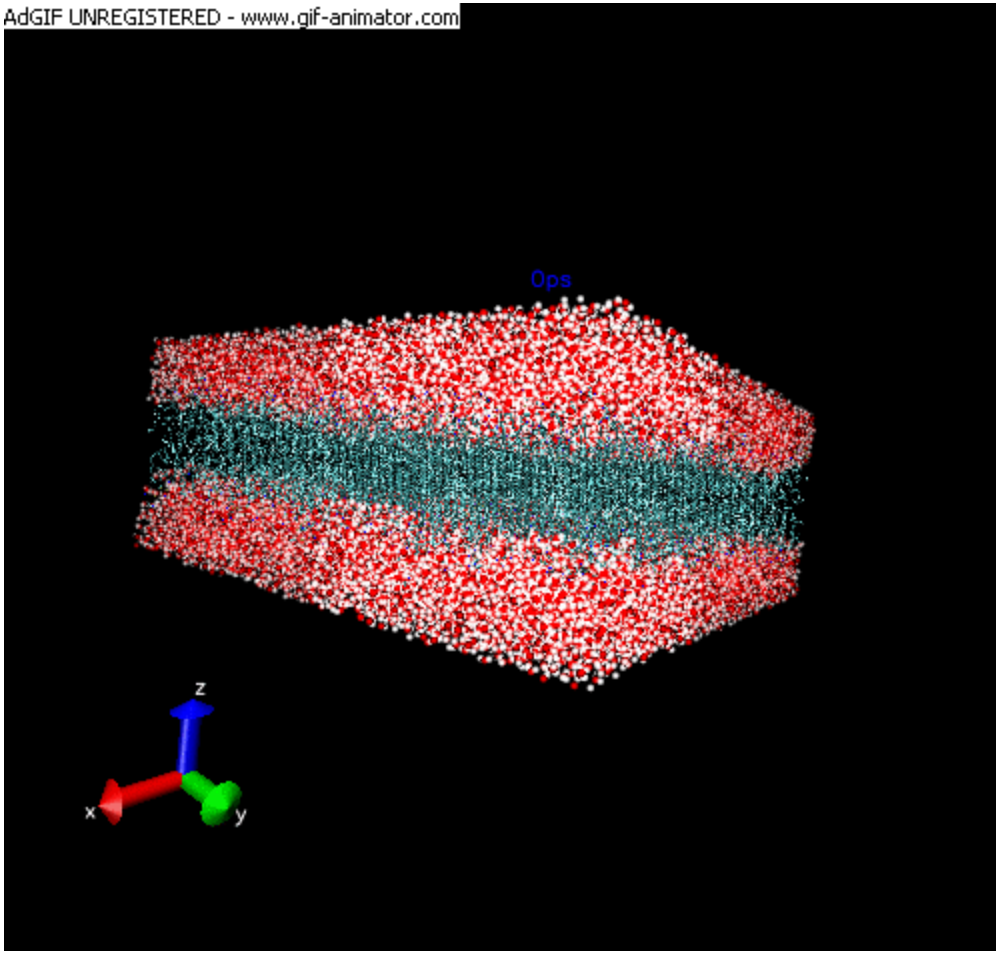

Accuracy: Comparison with Other Codes

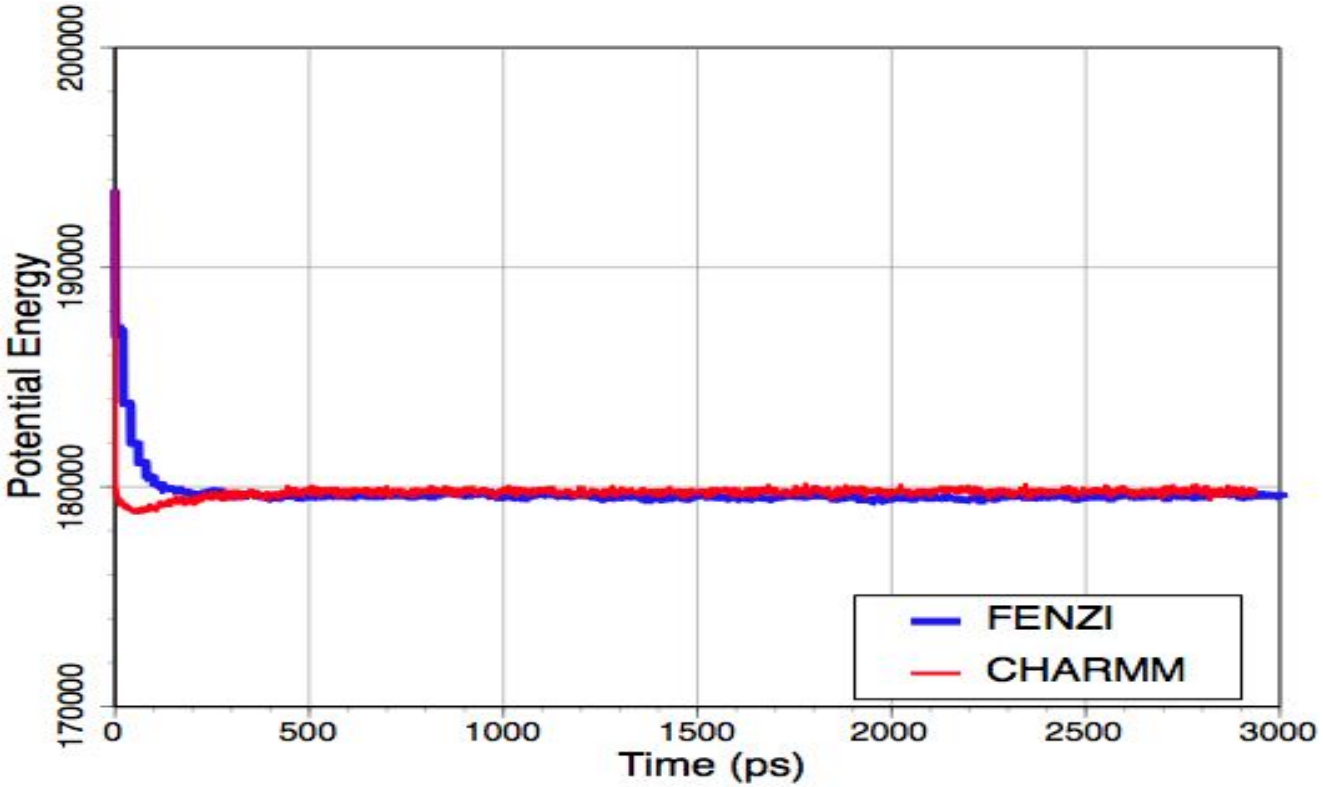

- Several energies fluctuate around same average values
- Before equilibrium is reached, energy drifting is due to different thermostats, i.e., Langevin vs. velocity reassignment

Accuracy: Comparison with Other Codes

3ns of NVT MD simulation with 1fs step size: CHARMM on single core, 64 bits; FENZI on GTX 480, 32 bits

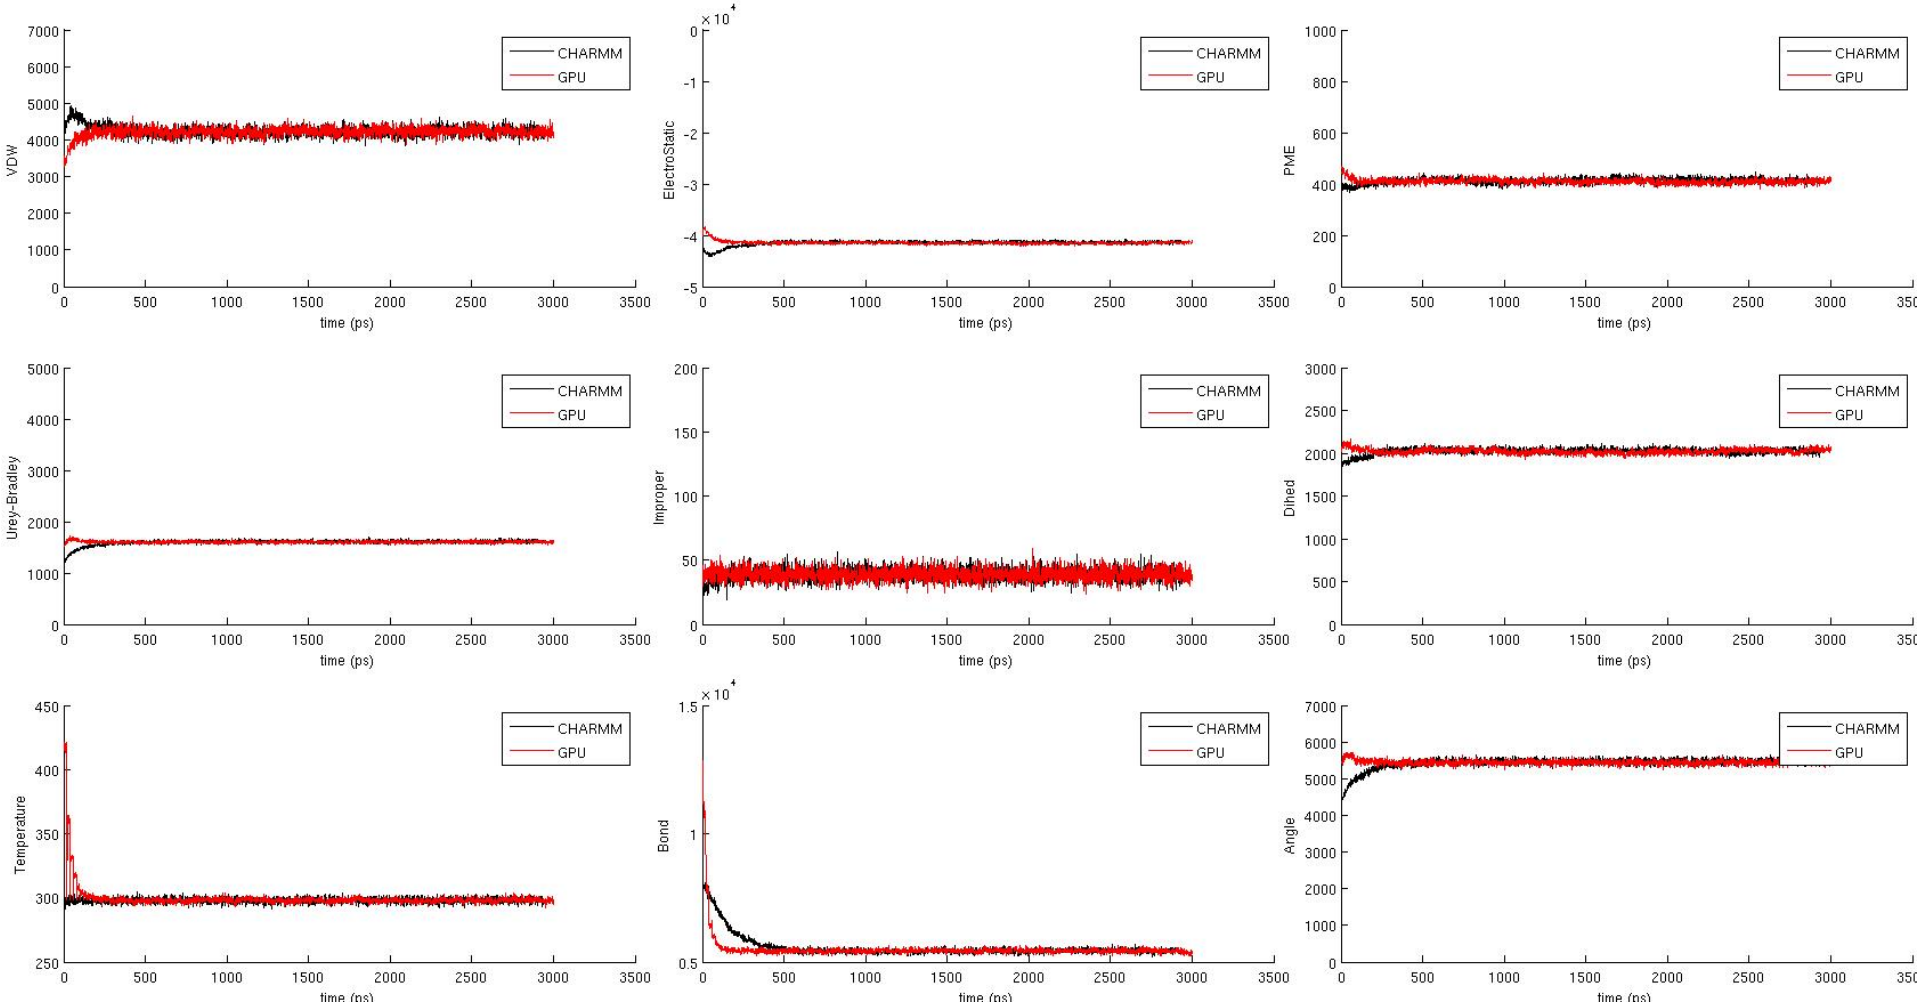

Accuracy: Energy Fluctuations

- A plot of the energy fluctuations versus time step size should follow an approximately logarithmic trend [1]
- FEN ZI fluctuations are proportional to time step size for large time step size
  - Larger than 0.5 fs
- A different behavior for step size less than 0.5 fs is consistent with results previously presented and discussed in other work [2]

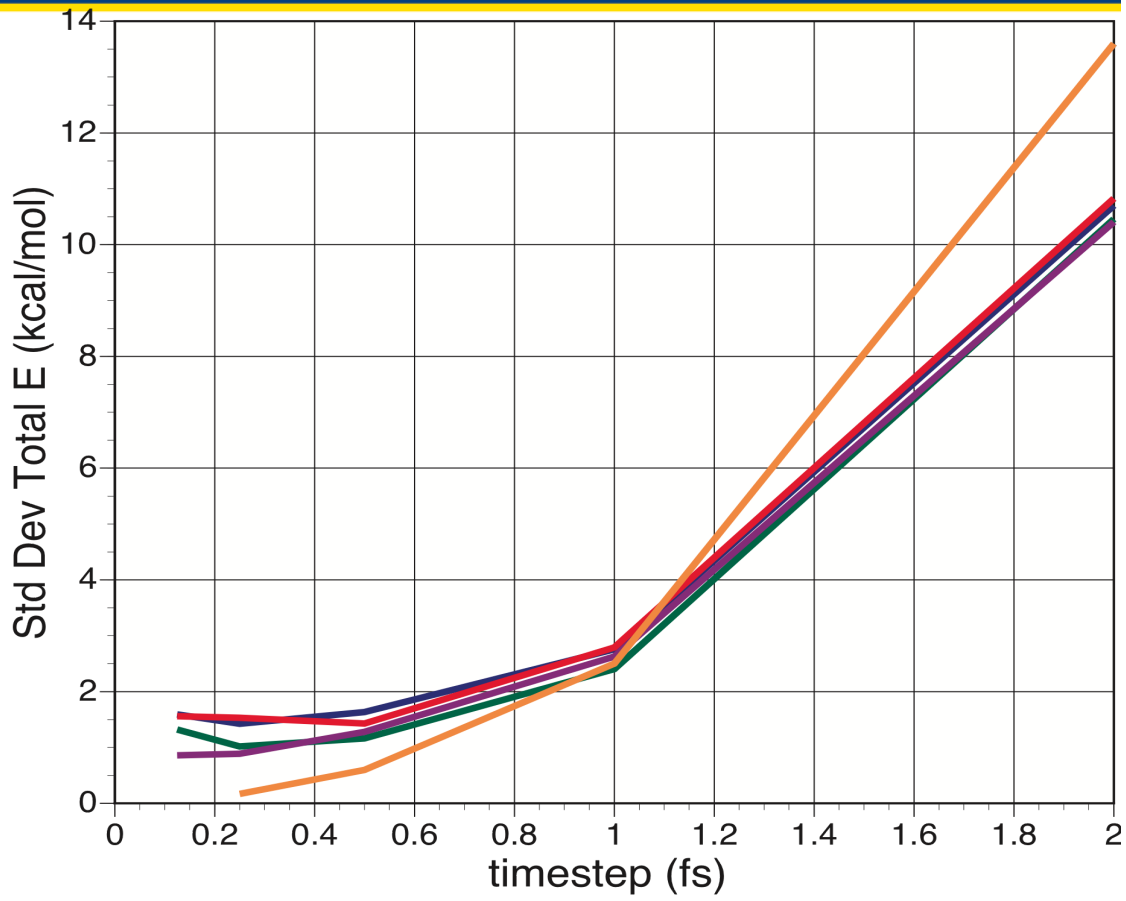

[1] Allen and Tildesley, Oxford: Clarendon Press, (1987)
[2] Bauer et al., J. Comput. Chem. 32(3): 375 – 385, 2011

Empirical Analysis

- We study the structural and electrostatic properties:
  - Mass density profiles of various chemical groups within the membrane
  - Mass density of water along the membrane
  - Electron and charge density profiles along the membrane
  - Surface potential due to water and lipid
  - Order parameters for the lipid tails
- We find that simulation results match experiment observations across the various membrane sizes[1]
- We surprisingly find that the structural properties are robust across the various membrane sizes[1]
  - Atomic number density, electron density, and electrostatic potentials remain consistently equivalent across the small, medium, and very large system

[1] Ganesan N. et al., J. Comp. Chem 32(14): 2958 – 2973, 2011

FEN ZI Analysis: Small Membrane

100 ns of FENZI simulation

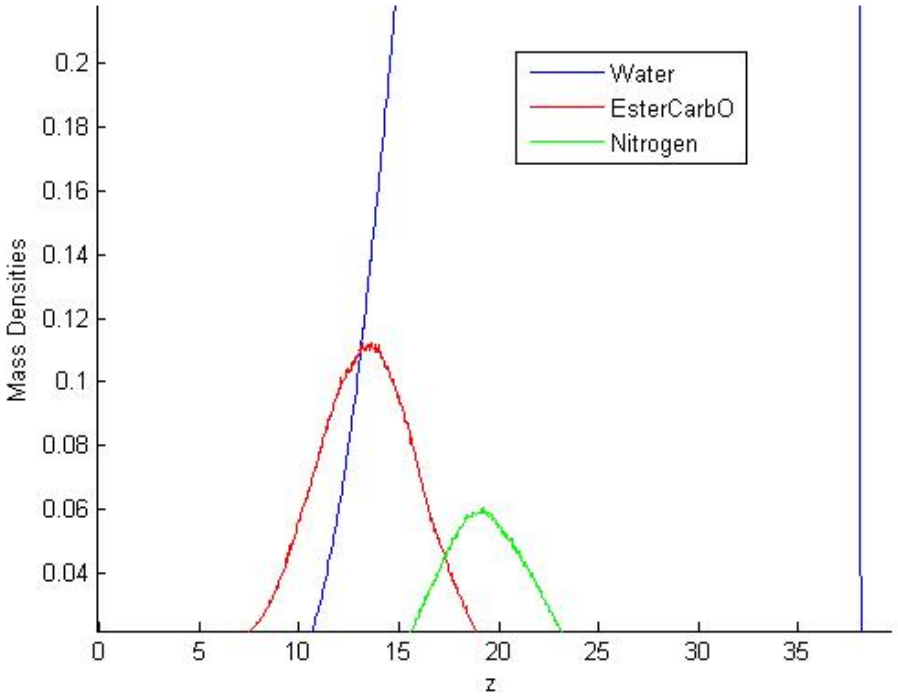

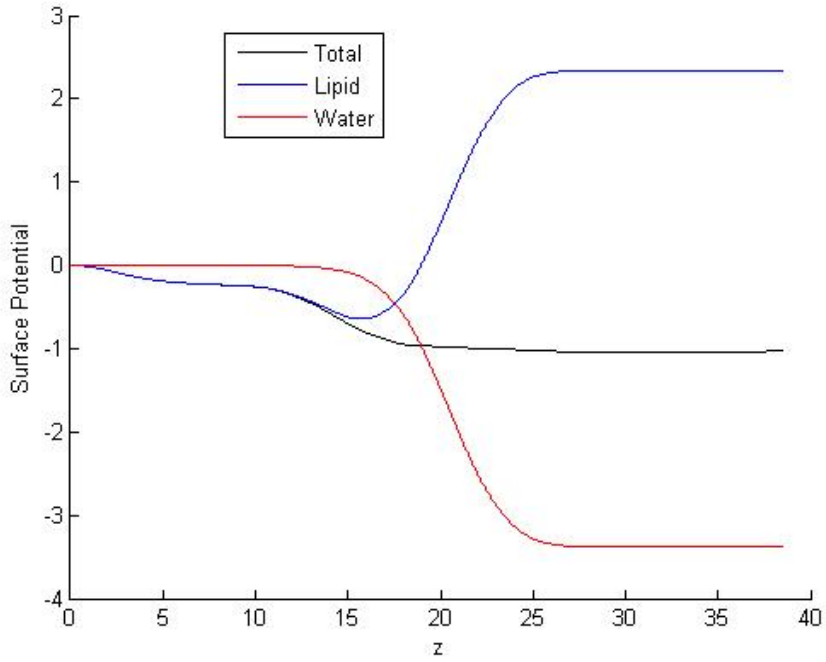

- Mass density histograms matches the membrane position
- Surface potential is 0.9 V as expected in a lipid membrane[1]

[1] Ganesan N. et al., J. Comp. Chem 32(14): 2958 – 2973, 2011

FEN ZI Analysis: Large Membrane

20 ns of FENZI simulation

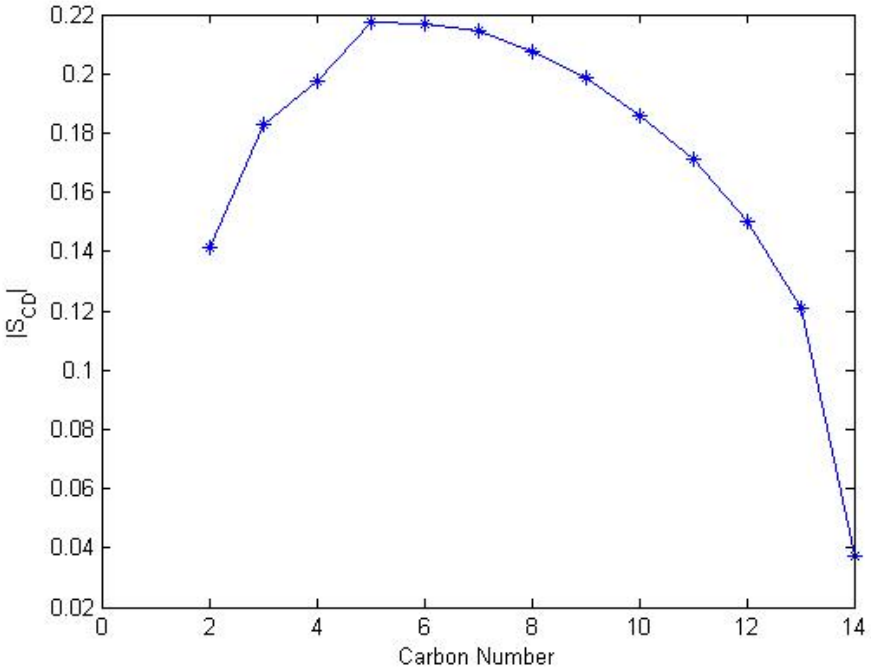

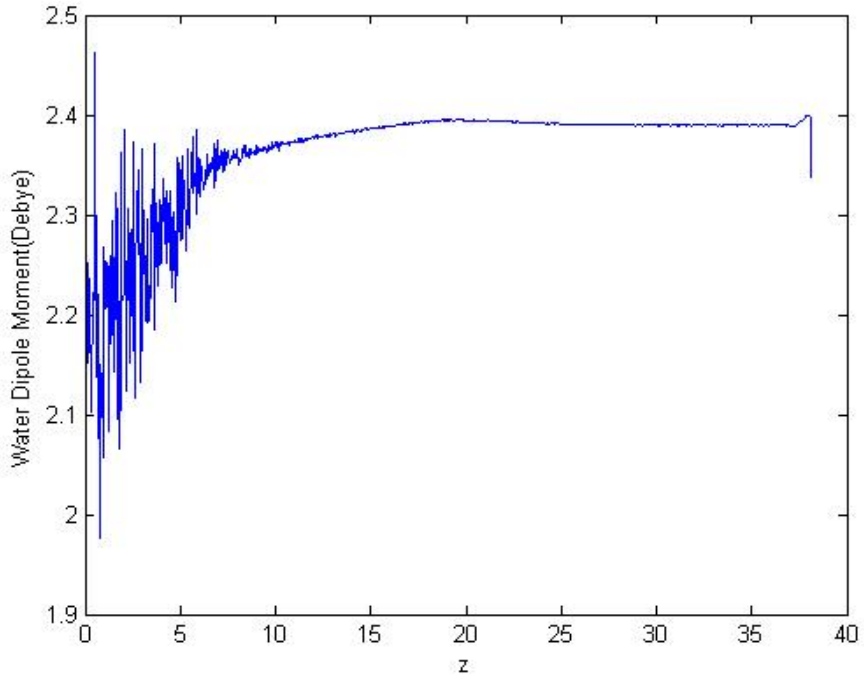

- Both order parameters and water dipole moment matched expected experiment results in a lipid membrane[1]

[1] Ganesan N. et al., J. Comp. Chem 32(14): 2958 – 2973, 2011

FLOPS

NSEC/DAY

PERFORMANCE

SCALABILITY

NVE

MD OF ENSEMBLES

NPT

NVT

MATHEMATICAL MODEL

INTERACTION POTENTIAL

ASPECTS OF REALISTIC MD SIMULATION

ENERGY FLUCTUATIONS

FORCE FIELD

ELECTROSTATIC INTERACTIONS

VALIDATION AND ACCURACY

PME

FREEZING FAST DEGREES OF MOTION

EMPIRICAL VALIDATION

IMPOSING CONSTRAINTS ON INTERATOMIC DISTANCES

SINGLE PEPTIDE DIFFUSION

SHAKE / RATTLE

STUDY PHYSICAL-CHEMICAL SYSTEMS

RESTRAINTS

PROTEIN-MEMBRANE INTERACTION

Performance: FEN ZI vs. CHARMM

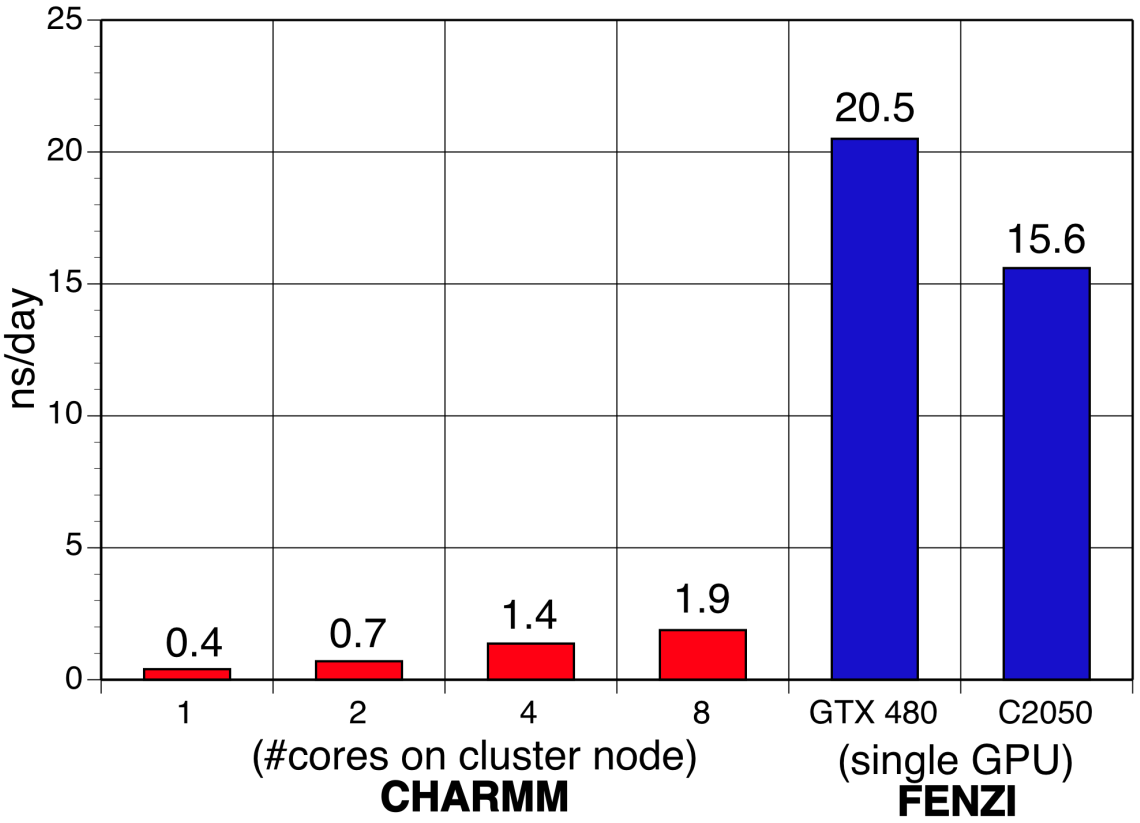

- Parallel MPI CHARMM optimized for and executed on the cores of a single node (i.e., 2.6 GHz, dual quad-core, 8GB memory Intel Xeon)
- FEN ZI optimized for and executed on one GTX 480 and C2050 GPUs (Fermi) using single precision and CUDA 3.1

FEN ZI speedup for a single-precision MD simulation on one GPU is up to 10X the same simulation on one 8-core, double-precision node[1]

[1] Ganesan N. et al., J. Comp. Chem 32(14): 2958 – 2973, 2011

FEN ZI Scalability

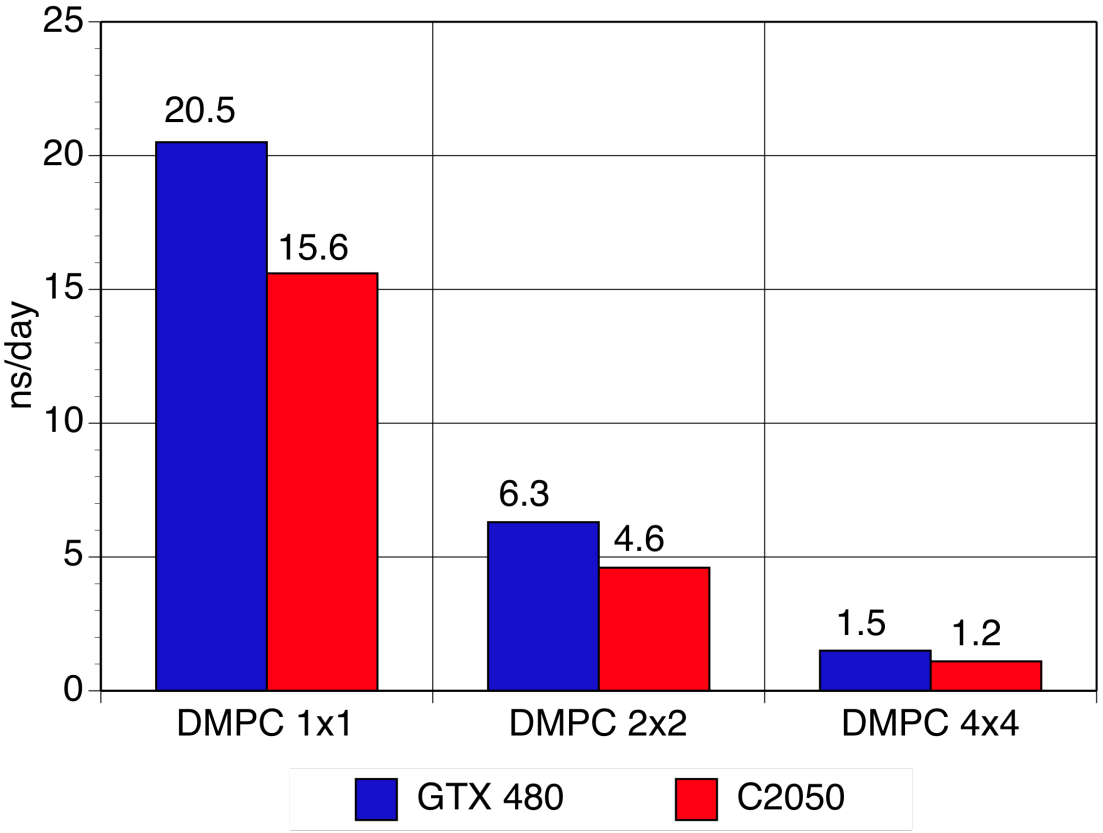

- Simulations of three lipid bilayer membranes (DMPC) with three different sizes, each four time larger than the previous
- FEN ZI simulations were performed on a GTX 480 GPU and a C2050 GPU (Fermi)

FEN ZI allows us to simulate larger membranes, larger than simple regions, over longer time interval in significant turnaround times[1]

[1] Ganesan N. et al., J. Comp. Chem 32(14): 2958 – 2973, 2011

FLOPS

NS/DAY

PERFORMANCE

NVE

MD OF ENSEMBLES

NPT

NVT

SCALABILITY

MATHEMATICAL MODEL

INTERACTION POTENTIAL

ASPECTS OF REALISTIC MD SIMULATION

ENERGY FLUCTUATIONS

FORCE FIELD

ELECTROSTATIC INTERACTIONS

VALIDATION AND ACCURACY

PME

FREEZING FAST DEGREES OF MOTION

EMPIRICAL VALIDATION

IMPOSING CONSTRAINTS ON INTERATOMIC DISTANCES

SHAKE / RATTLE

RESTRAINTS

SINGLE PEPTIDE DIFFUSION

STUDY PHYSICAL-CHEMICAL SYSTEMS

PROTEIN-MEMBRANE INTERACTION

Protein-Membrane Interaction

- Protein molecules are absorbed by cells walls constituted by membranes
- Interaction of proteins with membranes is biologically relevant
- Inject a protein molecule into the membrane system
- Identify types of proteins that are easily absorbed
- Study pathological conditions and behaviors

WALP16 peptide penetrating DMPC 2x2

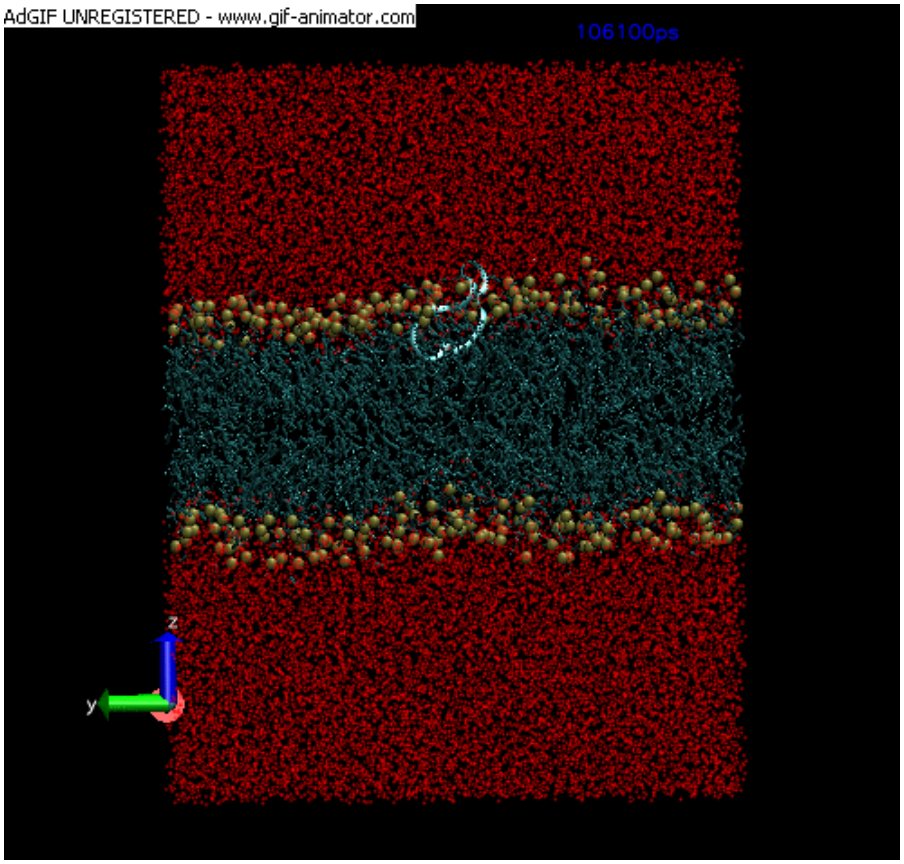

[1] Taufer M. et al., CiSE (2012) – In preparation

Multiple CPUs vs. Single GPU

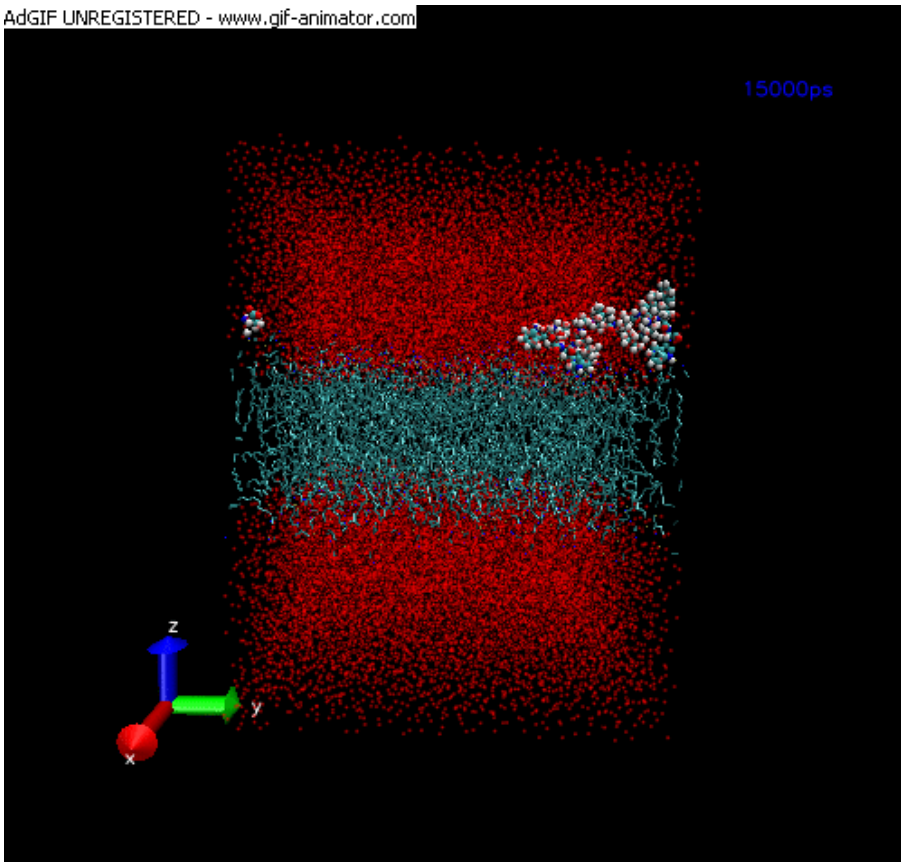

Projected performance of an 8-core node of a cluster running CHARMM

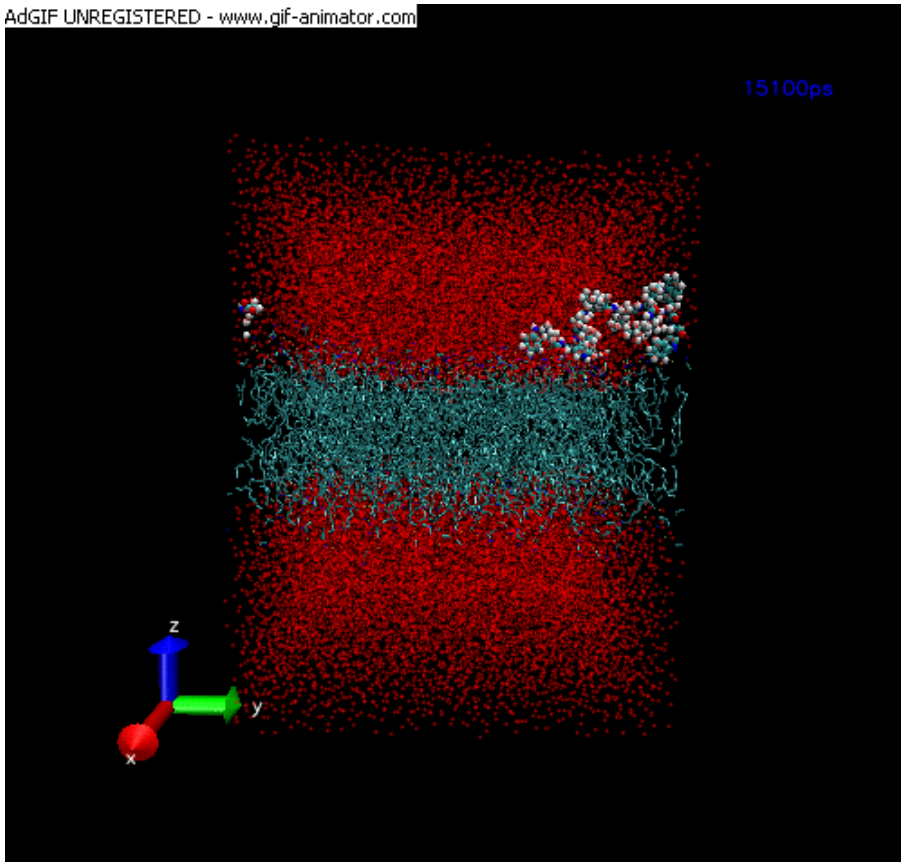

Performance of FENZI on one GTX 480

Single Peptide Diffusion

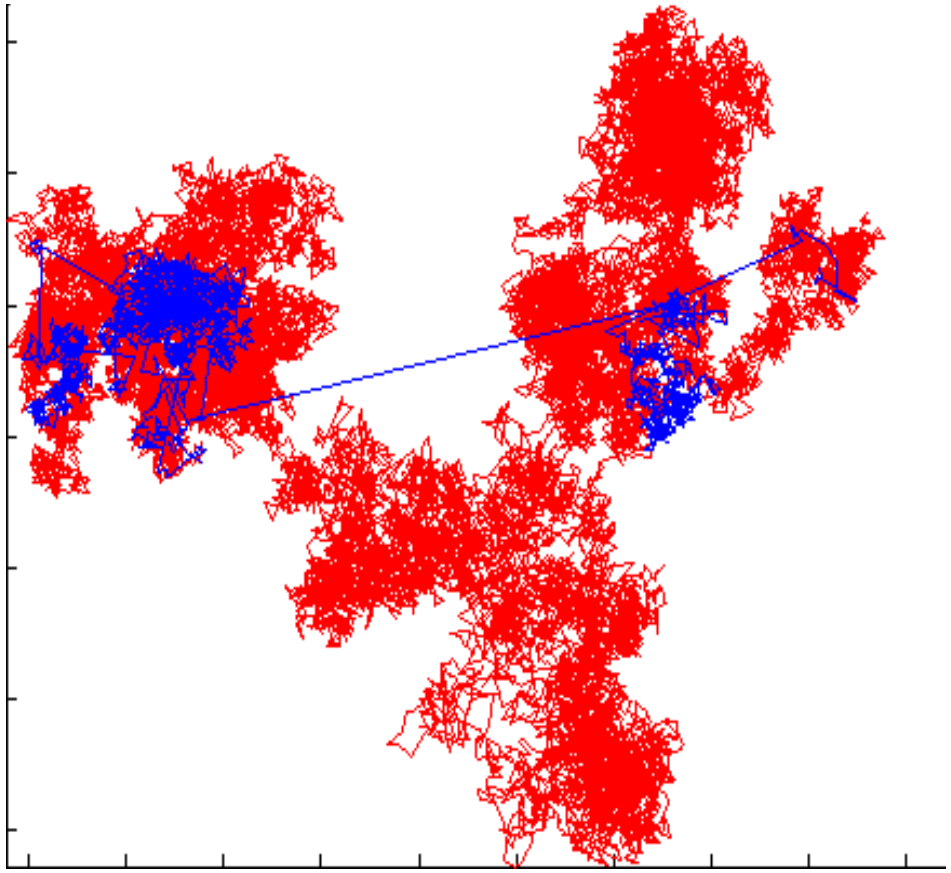

- Study the kinematic behavior of the WALP16 peptide on the surface of the membrane
- Larger membrane size and timescale on GPU enables us to observe the scooting behavior of the system:
  - Red regions indicate scooting while outside the membrane
  - Blue regions indicate scooting while penetrated

[1] Taufer M. et al., CiSE (2012) – In preparation

Summary

- GPUs enable unprecedented levels of system size being simulated in unprecedented turnaround times
- Realistic simulations of molecular systems can be implemented on GPUs while preserving model realism and simulation validity
- MD simulations with FEN ZI outline:
  - The structural properties in membranes are robust across the various simulation sizes
  - The diffusion properties of a single peptide when scooting above the membrane surface

GPU computing with its challenges and opportunities occupies a solid place in the molecular computational science community

Acknowledgments

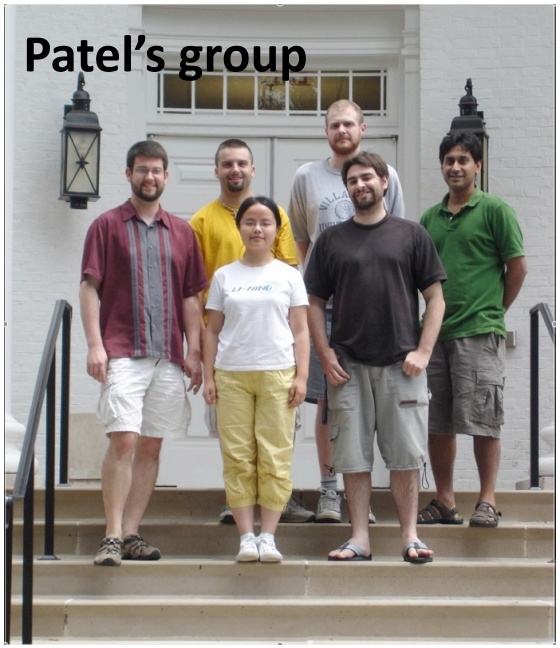

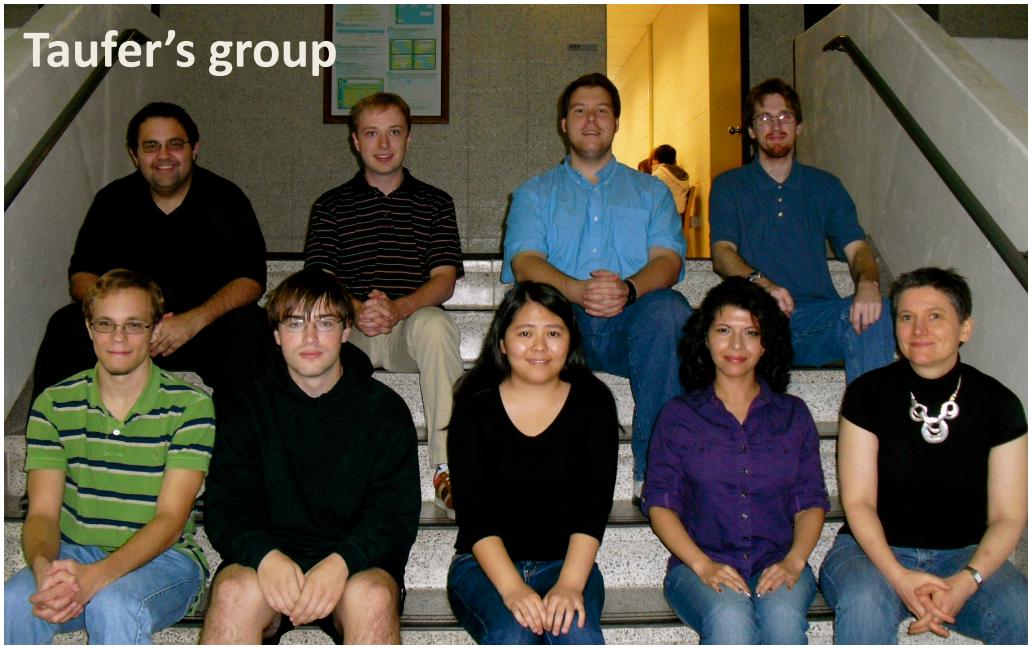

Related work:

Ganesan et al., JCC 2011

Bauer et al., JCC 2011

Davis et al., BICoB 2009

Contact:

taufer@udel.edu , sapatel@udel.edu

Sponsors:

GPU TECHNOLOGY CONFERENCE

Explore GTC On-Demand

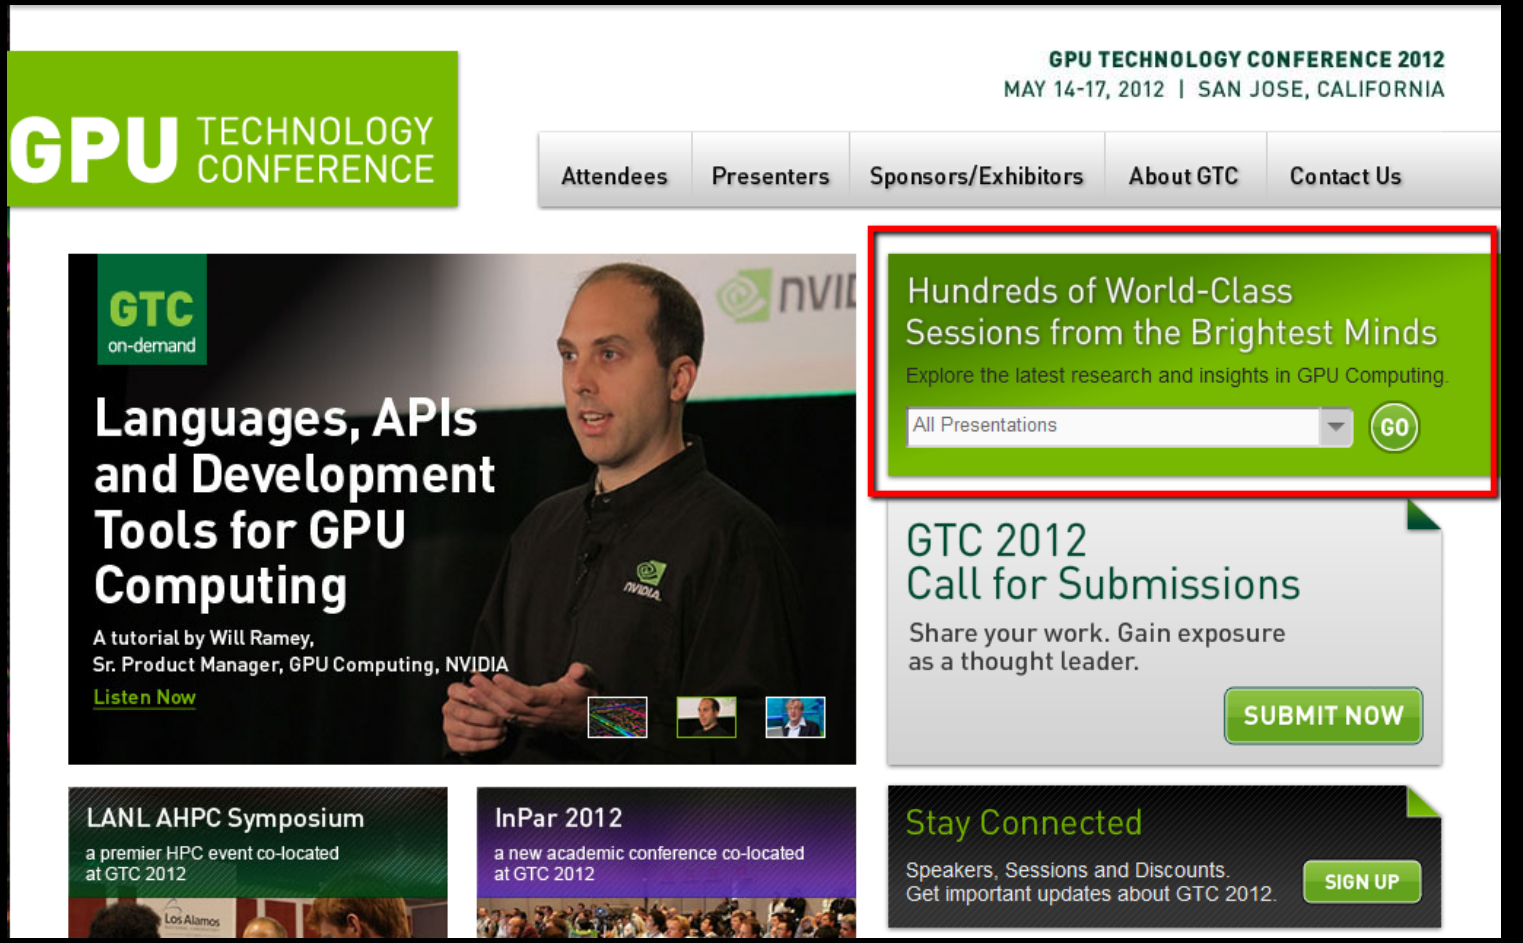

www.gputechconf.com