

# GPU-Accelerated Visualization and Analysis of Petascale Molecular Dynamics Simulations

John E. Stone

Theoretical and Computational Biophysics Group

Beckman Institute for Advanced Science and Technology

University of Illinois at Urbana-Champaign

**<http://www.ks.uiuc.edu/Research/vmd/>**


**<http://www.ks.uiuc.edu/Research/gpu/>**

First Imaging Initiative Workshop: Tomography and Ptychography,  
Argonne National Laboratory, September 30, 2014



# VMD – “Visual Molecular Dynamics”

- Visualization and analysis of:
  - molecular dynamics simulations
  - quantum chemistry calculations
  - particle systems and whole cells
  - sequence data
- User extensible w/ scripting and plugins
- <http://www.ks.uiuc.edu/Research/vmd/>




Poliovirus

Structural Similarity	
tho-a	cccc
foor-a	cccc
tyea-a	cccc
scyl-a	cccc
foyl-a	cccc
tho-a	cccc


  

Sequence Similarity	
tho-a	cccc
foor-a	cccc
tyea-a	cccc
scyl-a	cccc
foyl-a	cccc
tho-a	cccc


Ribosome Sequences



Electrons in  
Vibrating Buckyball



Cellular Tomography  
Cryo-electron Microscopy




Whole Cell Simulations


# Goal: A Computational Microscope

Study the molecular machines in living cells

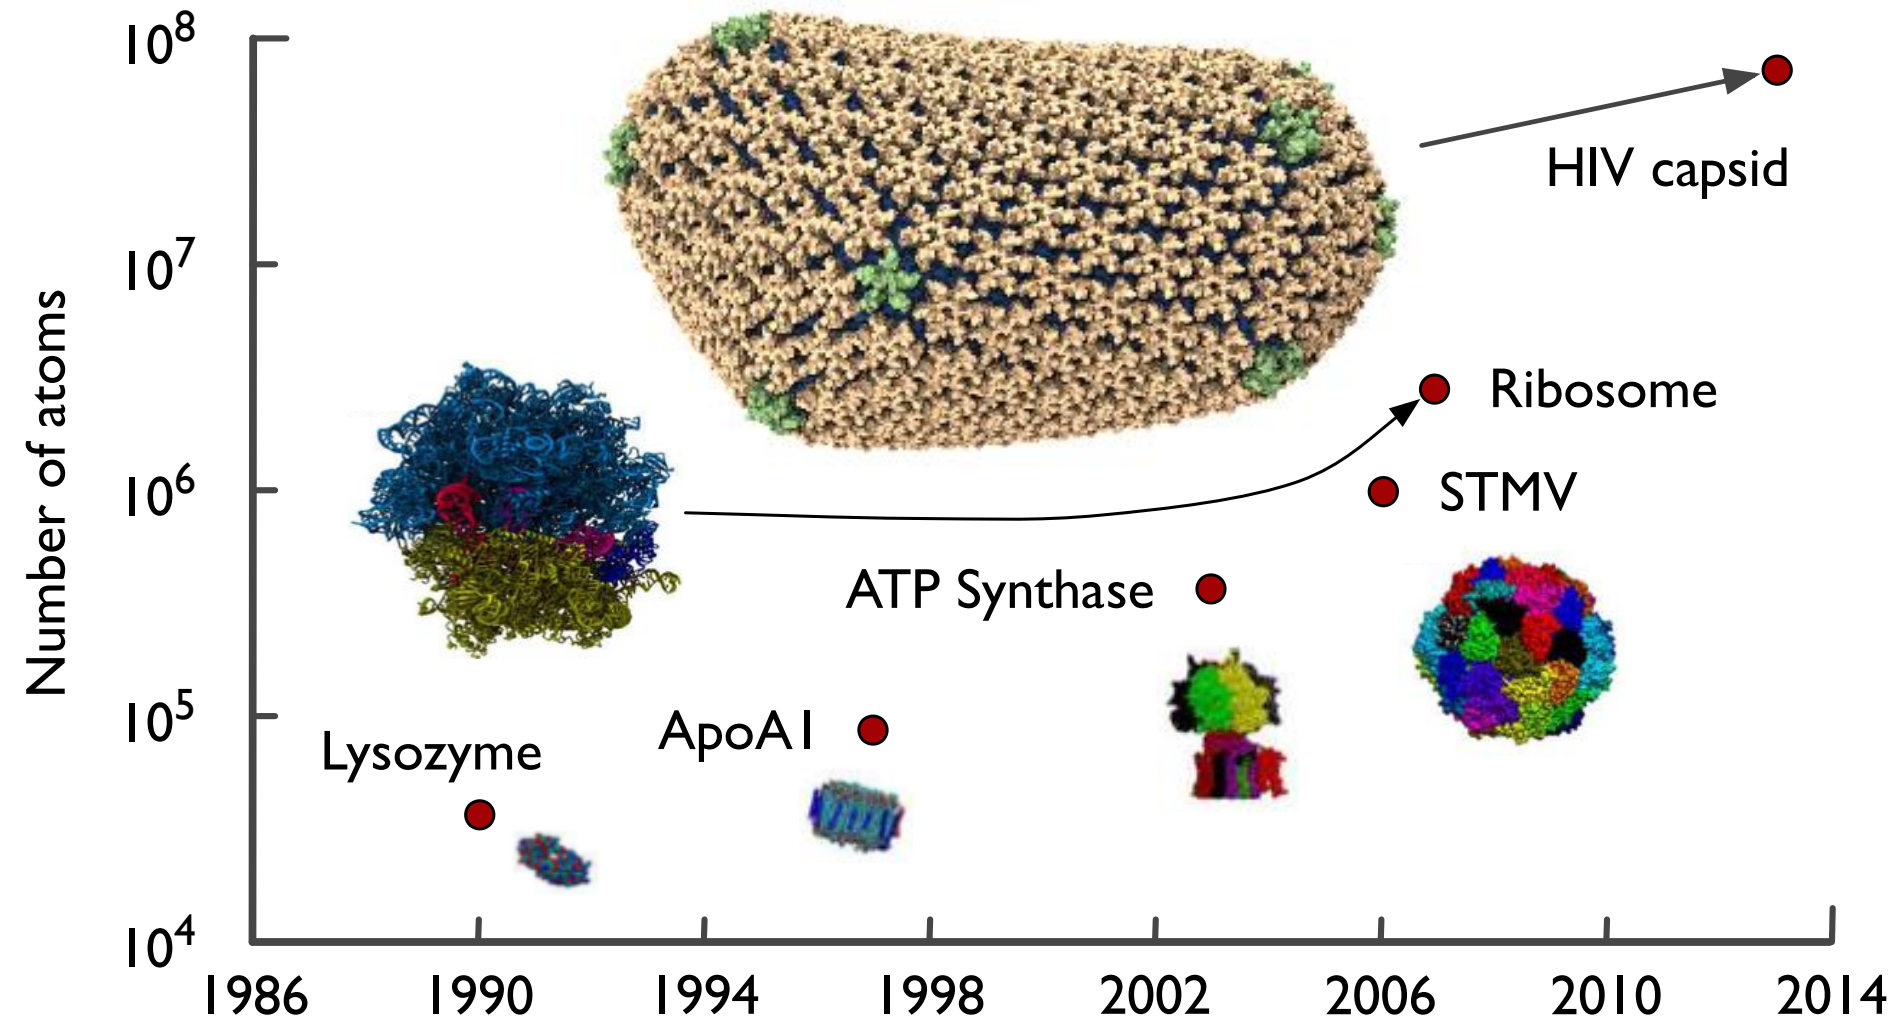
Ribosome: target for antibiotics



Poliovirus




# NAMD and VMD Use GPUs & Petascale Computing to Meet Computational Biology's Insatiable Demand for Processing Power



# First Simulation of a Virus Capsid (2006)

## Satellite Tobacco Mosaic Virus (STMV)



First MD simulation of a complete virus capsid


STMV smallest available capsid structure

**STMV simulation, visualization, and analysis pushed us toward GPU computing!**


MD showed that STMV capsid collapses without its RNA core

**1 million atoms  
A huge system for 2006**


t = 0




t = 5ns



t = 10ns




# Peak Arithmetic Performance Trend





# Peak Memory Bandwidth Trend

Theoretical GB/s



# CUDA GPU-Accelerated Trajectory Analysis and Visualization in VMD


VMD GPU-Accelerated Feature or Kernel	Typical speedup vs. multi-core CPU (e.g. 4-core CPU)
Molecular orbital display	30x
Radial distribution function	23x
Molecular surface display	15x
Electrostatic field calculation	11x
Ray tracing w/ shadows, AO lighting	7x
Ion placement	6x
MDFFF density map synthesis	6x
Implicit ligand sampling	6x
Root mean squared fluctuation	6x
Radius of gyration	5x
Close contact determination	5x
Dipole moment calculation	4x





# Time-Averaged Electrostatics Analysis on Energy-Efficient GPU Cluster


- **1.5 hour** job (CPUs) reduced to **3 min** (CPUs+GPU)
- Electrostatics of thousands of trajectory frames averaged
- Per-node power consumption on NCSA “AC” GPU cluster:
  - CPUs-only: 448 Watt-hours
  - CPUs+GPUs: 43 Watt-hours
- GPU Speedup: **25.5x**
- Power efficiency gain: **10.5x**



**Quantifying the Impact of GPUs on Performance and Energy Efficiency in HPC Clusters.** J. Enos, C. Steffen, J. Fullop, M. Showerman, G. Shi, K. Esler, V. Kindratenko, J. Stone, J. Phillips. *The Work in Progress in Green Computing*, pp. 317-324, 2010.

# NAMD Titan XK7 Performance August 2013

NAMD on Titan Cray XK7 (2fs timestep with PME)



# VMD Petascale Visualization and Analysis

- Analyze/visualize large trajectories too large to transfer off-site:
  - Compute time-averaged electrostatic fields, MDFFF quality-of-fit, etc.
  - User-defined parallel analysis operations, data types
  - Parallel ray tracing, movie making
- Parallel I/O rates up to **275 GB/sec** on 8192 Cray XE6 nodes – can read in **231 TB in 15 minutes!**
- Multi-level dynamic load balancing tested with up to 262,144 CPU cores
- **Supports GPU-accelerated Cray XK7 nodes for both visualization and analysis usage**



NCSA Blue Waters Hybrid  
Cray XE6 / XK7 Supercomputer

22,640 XE6 CPU nodes

4,224 XK7 nodes w/ GPUs support  
fast VMD OpenGL movie  
rendering and visualization

# Molecular Dynamics Flexible Fitting (MDFF)

X-ray crystallography




APS at Argonne

MDFF



Electron microscopy



FEI microscope



ORNL Titan



Acetyl - CoA Synthase



Flexible fitting of atomic structures into electron microscopy maps using molecular dynamics.

L. Trabuco, E. Villa, K. Mitra, J. Frank, and K. Schulten. *Structure*, 16:673-683, 2008.

# Molecular Dynamics Flexible Fitting - Theory

Two terms are added to the MD potential

$$U_{total} = U_{MD} + U_{EM} + U_{SS}$$




An external potential derived from the EM map is defined on a grid as

$$U_{EM}(\mathbf{R}) = \sum_j w_j V_{EM}(\mathbf{r}_j)$$

$$V_{EM}(\mathbf{r}) = \begin{cases} \xi \left( 1 - \frac{\Phi(\mathbf{r}) - \Phi_{thr}}{\Phi_{max} - \Phi_{thr}} \right) & \text{if } \Phi(\mathbf{r}) \geq \Phi_{thr}, \\ \xi & \text{if } \Phi(\mathbf{r}) < \Phi_{thr}. \end{cases}$$


A mass-weighted force is then applied to each atom

$$\mathbf{f}_i^{EM} = -\nabla U_{EM}(\mathbf{R}) = -w_i \partial V_{EM}(\mathbf{r}_i) / \partial r_i$$




# Structural Route to the all-atom HIV-1 Capsid

1st TEM (1999)




Ganser et al. *Science*, 1999  
 Briggs et al. *EMBO J*, 2003  
 Briggs et al. *Structure*, 2006

1st tomography (2003)



cryo-ET (2006)


Crystal structures of separated hexamer and pentamer




Pornillos et al. , *Cell* 2009, *Nature* 2011

High res. EM of hexameric tubule, tomography of capsid,  
**all-atom model of capsid by MDFF w/ NAMD & VMD,**  
**NSF/NCSA Blue Waters computer at Illinois**



hexameric tubule



Li et al., *Nature*, 2000



Byeon et al., *Cell* 2009




**Zhao et al. , *Nature* 497: 643-646 (2013)**



# Evaluating Quality-of-Fit for Structures Solved by Hybrid Fitting Methods


Compute Pearson correlation to evaluate the fit of a reference cryo-EM density map with a **simulated density map** produced from an **all-atom structure**.



# GPUs Can Reduce Trajectory Analysis Runtimes from Hours to Minutes

GPUs enable laptops and desktop workstations to handle tasks that would have previously required a cluster, or a very long wait...

GPU-accelerated petascale supercomputers enable analyses that were previously impractical, allowing detailed study of very large structures such as viruses




**GPU-accelerated MDFF Cross Correlation Timeline**

**Regions with poor fit**


**Regions with good fit**




(a)




(b)



(c)



(d)




# MDFFF Density Map Algorithm

- Build spatial acceleration data structures, optimize data for GPU
- Compute 3-D density map:

$$\rho(\vec{r}; \vec{r}_1, \vec{r}_2, \dots, \vec{r}_N) = \sum_{i=1}^N e^{-\frac{|\vec{r}-\vec{r}_i|^2}{2\alpha^2}}$$


- Truncated Gaussian and spatial acceleration grid ensure linear time-complexity



**3-D density map  
lattice point and  
the neighboring  
spatial acceleration  
cells it references**

# Single-Pass MDFF GPU Cross-Correlation



3-D density map decomposes into 3-D grid of 8x8x8 tiles containing CC partial sums and local CC values



**Spatial CC map and overall CC value computed in a single pass**

Small 8x8x2 CUDA thread blocks afford large per-thread register count, shared memory

Each thread computes 4 z-axis density map lattice points and associated CC partial sums



**Threads producing results that are used**

**Inactive threads, region of discarded output**

Padding optimizes global memory performance, guaranteeing coalesced global memory accesses


**Grid of thread blocks**

# VMD GPU Cross Correlation Performance

	<b>RHDV</b>	<b>Mm-cpn open</b>	<b>GroEL</b>	<b>Aquaporin</b>
<b>Resolution (Å)</b>	<b>6.5</b>	<b>8</b>	<b>4</b>	<b>3</b>
<b>Atoms</b>	<b>702K</b>	<b>61K</b>	<b>54K</b>	<b>1.6K</b>
<b>VMD-CUDA Quadro K6000</b>	<b>0.458s</b> <b>34.6x</b>	<b>0.06s</b> <b>25.7x</b>	<b>0.034s</b> <b>36.8x</b>	<b>0.007s</b> <b>55.7x</b>
VMD-CPU-SSE 32-threads, 2x Xeon E5-2687W	0.779s 20.3x	0.085s 18.1x	0.159s 7.9x	0.033s 11.8x
<b>Chimera 1-thread Xeon E5-2687W</b>	<b>15.86s</b> <b>1.0x</b>	<b>1.54s</b> <b>1.0x</b>	<b>1.25s</b> <b>1.0x</b>	<b>0.39s</b> <b>1.0x</b>
<b>VMD CPU-SEQ (plugin)</b> 1-thread Xeon E5-2687W	62.89s 0.25x	2.9s 0.53x	1.57s 0.79x	0.04s 9.7x

**GPU-Accelerated Analysis and Visualization of Large Structures Solved by Molecular Dynamics Flexible Fitting.** J. E. Stone, R. McGreevy, B. Isralewitz, and K. Schulten. Faraday Discussion 169, 2014. **(In press)**


# VMD RHDV Cross Correlation Timeline on Cray XK7



	<b>RHDV</b>
<b>Atoms</b>	<b>702K</b>
<b>Traj. Frames</b>	<b>10,000</b>
<b>Component Selections</b>	<b>720</b>
<b>Single-node XK7 (projected)</b>	<b>336 hours (14 days)</b>
<b>128-node XK7</b>	<b>3.2 hours 105x speedup</b>
<b>2048-node XK7</b>	<b>19.5 minutes 1035x speedup</b>


Calculation would take **5 years** using original serial VMD CC plugin on a workstation!

## RHDV CC Timeline




# Ray Tracing Molecular Graphics

- Ambient occlusion lighting, shadows, reflections, transparency, and more...
- Satellite tobacco mosaic virus capsid w/  $\sim 75\text{K}$  atoms



Standard OpenGL  
rasterization






VMD w/ new GPU ray tracing engine  
based on CUDA + OptiX

# Lighting Comparison



**Two lights, no shadows**

**Two lights, hard shadows, 1 shadow ray per light**

**Ambient occlusion + two lights, 144 AO rays/hit**




# VMD “QuickSurf” Representation, Ray Tracing




**All-atom HIV capsid simulations w/ up to 64M atoms on Blue Waters**



# BW VMD/Tachyon Movie Generation




# BW VMD/Tachyon Movie Generation




20 M atom chromatophore patch

# GPU Ray Tracing of HIV-1 on Blue Waters


- **Ambient occlusion lighting,** shadows, transparency, antialiasing, depth cueing, **144 rays/pixel minimum**
- 64 million atom virus simulation
- 1000+ movie frames
- Surface generation and ray tracing stages each use  $\geq$  75% of GPU memory



# VMD GPU Ray Tracing of HIV-1 Capsid



# HIV-1 Parallel HD Movie Rendering on Blue Waters Cray XE6/XK7



New “TachyonL-OptiX” on XK7 vs. Tachyon on XE6:

K20X GPUs yield **up to eight times** geom+ray tracing speedup

Cray XE6: 2x Opteron 62xx CPUs (32-cores)

Cray XK7: 1x Opteron 62xx CPU (16-cores) + NVIDIA Tesla K20X

Node Type and Count	Script Load Time	State Load Time	Geometry + Ray Tracing	Total Time
<b>256 XE6 CPU nodes</b>	7 s	160 s	<b>1,374 s</b>	<b>1,541 s</b>
512 XE6 CPU nodes	13 s	211 s	808 s	1,032 s
64 XK7 Tesla K20X GPUs	2 s	38 s	655 s	695 s
128 XK7 Tesla K20X GPUs	4 s	74 s	331 s	410 s
<b>256 XK7 Tesla K20X GPUs</b>	7 s	110 s	<b>171 s</b>	<b>288 s</b>

**GPU-Accelerated Molecular Visualization on Petascale Supercomputing Platforms.**

Stone et al. In UltraVis'13: Eighth Workshop on Ultrascale Visualization Proceedings, 2013.

# Acknowledgements

- Theoretical and Computational Biophysics Group, University of Illinois at Urbana-Champaign
- NCSA Blue Waters Team
- NVIDIA CUDA Center of Excellence, University of Illinois at Urbana-Champaign
- Many of the staff at NVIDIA and Cray
- Funding:
  - NSF OCI 07-25070
  - NSF PRAC “The Computational Microscope”
  - NIH support: 9P41GM104601, 5R01GM098243-02



# NIH BTRC for Macromolecular Modeling and Bioinformatics

1990-2017

**Beckman Institute  
University of Illinois at  
Urbana-Champaign**



# GPU Computing Publications

<http://www.ks.uiuc.edu/Research/gpu/>

- **Runtime and Architecture Support for Efficient Data Exchange in Multi-Accelerator Applications** Javier Cabezas, Isaac Gelado, John E. Stone, Nacho Navarro, David B. Kirk, and Wen-mei Hwu. IEEE Transactions on Parallel and Distributed Systems, 2014. (Accepted)
- **Unlocking the Full Potential of the Cray XK7 Accelerator** Mark Klein and John E. Stone. Cray Users Group, 2014. (In press)
- **Simulation of reaction diffusion processes over biologically relevant size and time scales using multi-GPU workstations** Michael J. Hallock, John E. Stone, Elijah Roberts, Corey Fry, and Zaida Luthey-Schulten. Journal of Parallel Computing, 2014. (In press)
- **GPU-Accelerated Analysis and Visualization of Large Structures Solved by Molecular Dynamics Flexible Fitting** John E. Stone, Ryan McGreevy, Barry Isralewitz, and Klaus Schulten. Faraday Discussion 169, 2014. (In press)
- **GPU-Accelerated Molecular Visualization on Petascale Supercomputing Platforms.** J. Stone, K. L. Vandivort, and K. Schulten. UltraVis'13: Proceedings of the 8th International Workshop on Ultrascale Visualization, pp. 6:1-6:8, 2013.
- **Early Experiences Scaling VMD Molecular Visualization and Analysis Jobs on Blue Waters.** J. E. Stone, B. Isralewitz, and K. Schulten. In proceedings, Extreme Scaling Workshop, 2013.
- **Lattice Microbes: High-performance stochastic simulation method for the reaction-diffusion master equation.** E. Roberts, J. E. Stone, and Z. Luthey-Schulten. J. Computational Chemistry 34 (3), 245-255, 2013.





# GPU Computing Publications

<http://www.ks.uiuc.edu/Research/gpu/>

- **Fast Visualization of Gaussian Density Surfaces for Molecular Dynamics and Particle System Trajectories.** M. Krone, J. E. Stone, T. Ertl, and K. Schulten. *EuroVis Short Papers*, pp. 67-71, 2012.
- **Fast Analysis of Molecular Dynamics Trajectories with Graphics Processing Units – Radial Distribution Functions.** B. Levine, J. Stone, and A. Kohlmeyer. *J. Comp. Physics*, 230(9):3556-3569, 2011.
- **Immersive Out-of-Core Visualization of Large-Size and Long-Timescale Molecular Dynamics Trajectories.** J. Stone, K. Vandivort, and K. Schulten. G. Bebis et al. (Eds.): *7th International Symposium on Visual Computing (ISVC 2011)*, LNCS 6939, pp. 1-12, 2011.
- **Quantifying the Impact of GPUs on Performance and Energy Efficiency in HPC Clusters.** J. Enos, C. Steffen, J. Fullop, M. Showerman, G. Shi, K. Esler, V. Kindratenko, J. Stone, J Phillips. *International Conference on Green Computing*, pp. 317-324, 2010.
- **GPU-accelerated molecular modeling coming of age.** J. Stone, D. Hardy, I. Ufimtsev, K. Schulten. *J. Molecular Graphics and Modeling*, 29:116-125, 2010.
- **OpenCL: A Parallel Programming Standard for Heterogeneous Computing.** J. Stone, D. Gohara, G. Shi. *Computing in Science and Engineering*, 12(3):66-73, 2010.



# GPU Computing Publications

<http://www.ks.uiuc.edu/Research/gpu/>

- **An Asymmetric Distributed Shared Memory Model for Heterogeneous Computing Systems.** I. Gelado, J. Stone, J. Cabezas, S. Patel, N. Navarro, W. Hwu. *ASPLOS '10: Proceedings of the 15<sup>th</sup> International Conference on Architectural Support for Programming Languages and Operating Systems*, pp. 347-358, 2010.
- **GPU Clusters for High Performance Computing.** V. Kindratenko, J. Enos, G. Shi, M. Showerman, G. Arnold, J. Stone, J. Phillips, W. Hwu. *Workshop on Parallel Programming on Accelerator Clusters (PPAC)*, In Proceedings IEEE Cluster 2009, pp. 1-8, Aug. 2009.
- **Long time-scale simulations of in vivo diffusion using GPU hardware.** E. Roberts, J. Stone, L. Sepulveda, W. Hwu, Z. Luthey-Schulten. In *IPDPS'09: Proceedings of the 2009 IEEE International Symposium on Parallel & Distributed Computing*, pp. 1-8, 2009.
- **High Performance Computation and Interactive Display of Molecular Orbitals on GPUs and Multi-core CPUs.** J. Stone, J. Saam, D. Hardy, K. Vandivort, W. Hwu, K. Schulten, *2nd Workshop on General-Purpose Computation on Graphics Processing Units (GPGPU-2)*, ACM International Conference Proceeding Series, volume 383, pp. 9-18, 2009.
- **Probing Biomolecular Machines with Graphics Processors.** J. Phillips, J. Stone. *Communications of the ACM*, 52(10):34-41, 2009.
- **Multilevel summation of electrostatic potentials using graphics processing units.** D. Hardy, J. Stone, K. Schulten. *J. Parallel Computing*, 35:164-177, 2009.



# GPU Computing Publications

<http://www.ks.uiuc.edu/Research/gpu/>

- **Adapting a message-driven parallel application to GPU-accelerated clusters.** J. Phillips, J. Stone, K. Schulten. *Proceedings of the 2008 ACM/IEEE Conference on Supercomputing*, IEEE Press, 2008.
- **GPU acceleration of cutoff pair potentials for molecular modeling applications.** C. Rodrigues, D. Hardy, J. Stone, K. Schulten, and W. Hwu. *Proceedings of the 2008 Conference On Computing Frontiers*, pp. 273-282, 2008.
- **GPU computing.** J. Owens, M. Houston, D. Luebke, S. Green, J. Stone, J. Phillips. *Proceedings of the IEEE*, 96:879-899, 2008.
- **Accelerating molecular modeling applications with graphics processors.** J. Stone, J. Phillips, P. Freddolino, D. Hardy, L. Trabuco, K. Schulten. *J. Comp. Chem.*, 28:2618-2640, 2007.
- **Continuous fluorescence microphotolysis and correlation spectroscopy.** A. Arkhipov, J. Hüve, M. Kahms, R. Peters, K. Schulten. *Biophysical Journal*, 93:4006-4017, 2007.

