Visualizing the Atomic Detail Dynamics of Biomolecular Complexes in Our Compute-Rich but I/O-Constrained Future

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/

Session: Visualization and HPC
17:20-17:40, Panorama 3, Forum Messe,
Frankfurt, Germany, Monday June 25th, 2018

VMD – “Visual Molecular Dynamics”

- 100,000 active users worldwide
- Visualization and analysis of:
  - Molecular dynamics simulations
  - Lattice cell simulations
  - Quantum chemistry calculations
  - Cryo-EM densities, volumetric data
- User extensible scripting and plugins
- http://www.ks.uiuc.edu/Research/vmd/

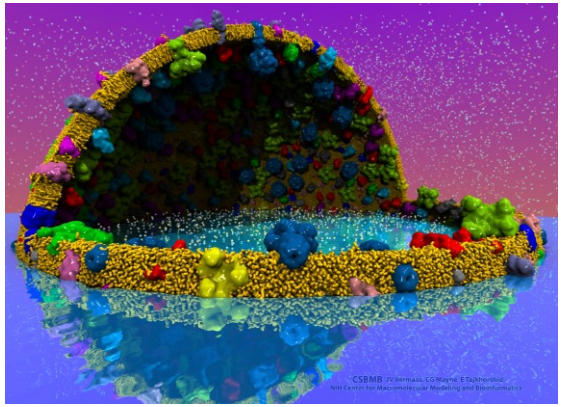

Cell-Scale Modeling

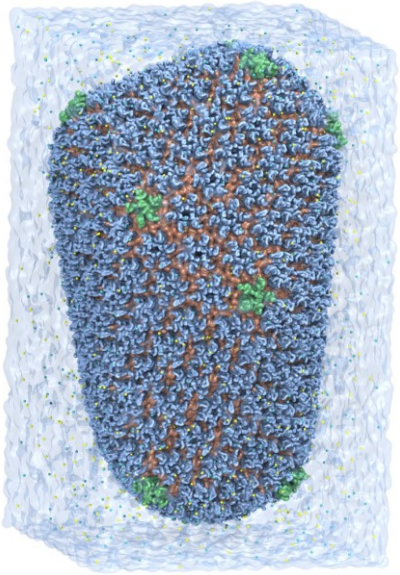

MD Simulation

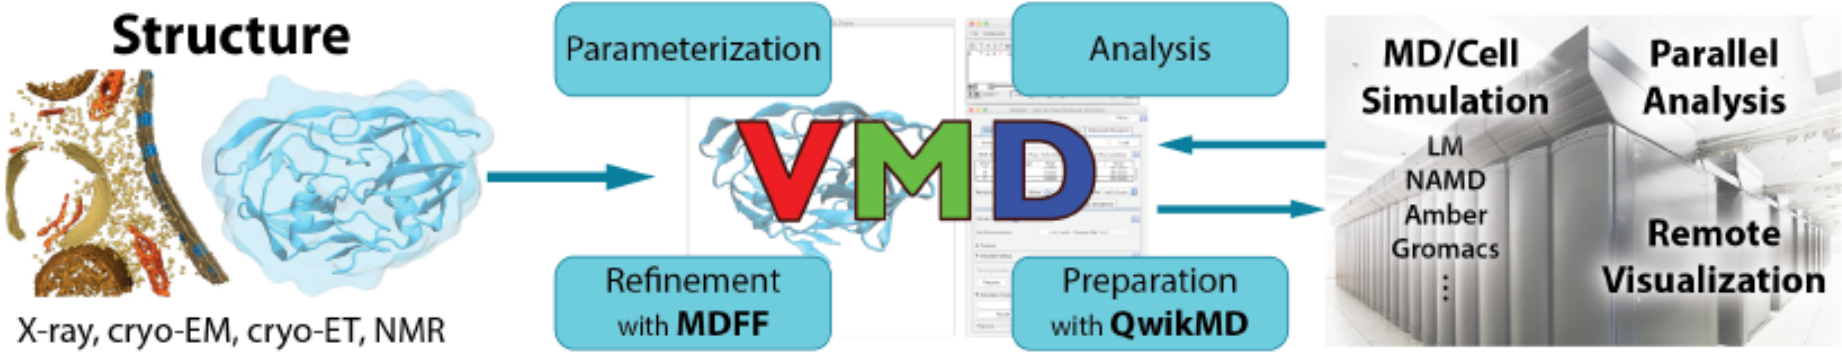

Goal: A Computational Microscope

Study the molecular machines in living cells

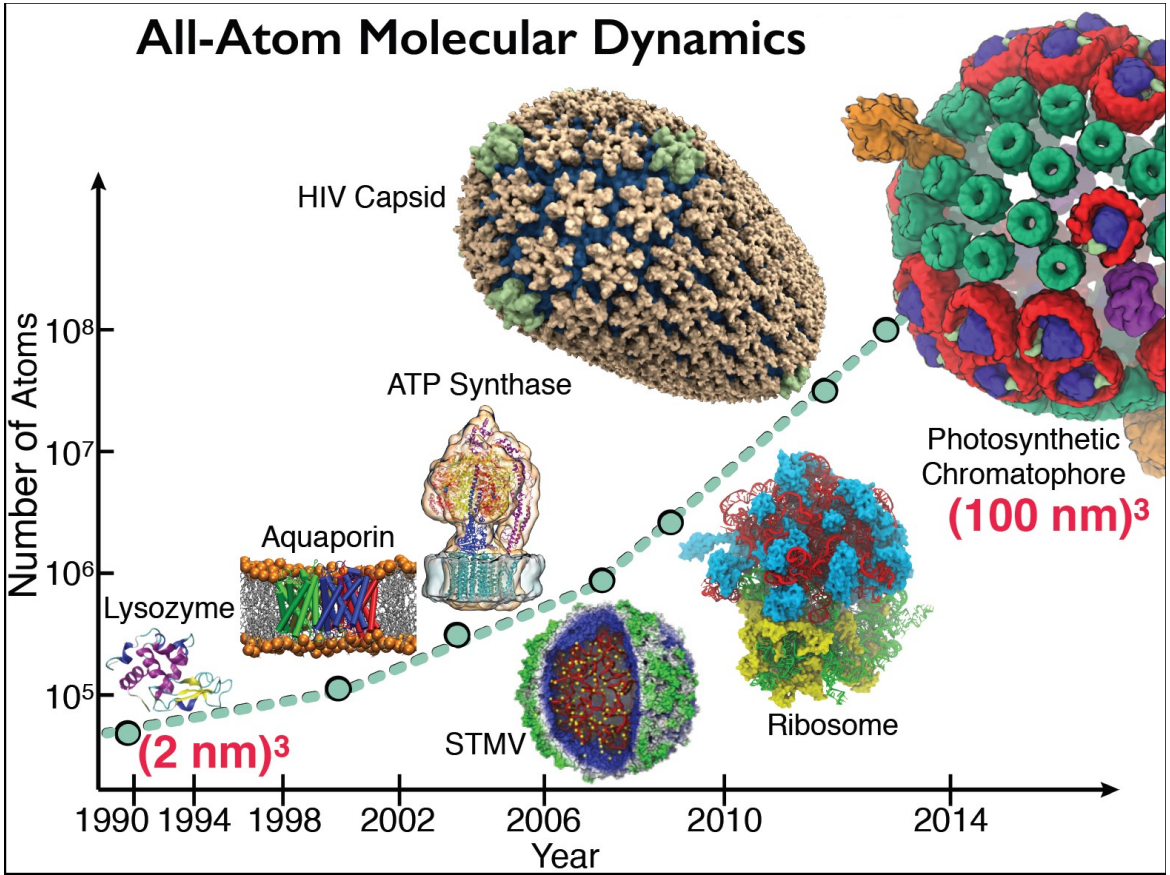

What Drives Increasing Molecular Dynamics System Size and Timescale?

- Working to gain insight into structure and dynamics of molecular basis of disease
- Many **health-relevant biomolecular complexes are large**, and key **processes often occur at long timescales**, presenting many computational challenges…
- New **hybrid modeling approaches** that combine the best structure information from **multiple modalities of experimental imaging**, physics, e.g. from MD force fields:
  - **"Computational Microscopy"**
- **Parallel computing** provides the resources required to keep pace with advances in structure determination and modeling

IBM AC922 Summit Node

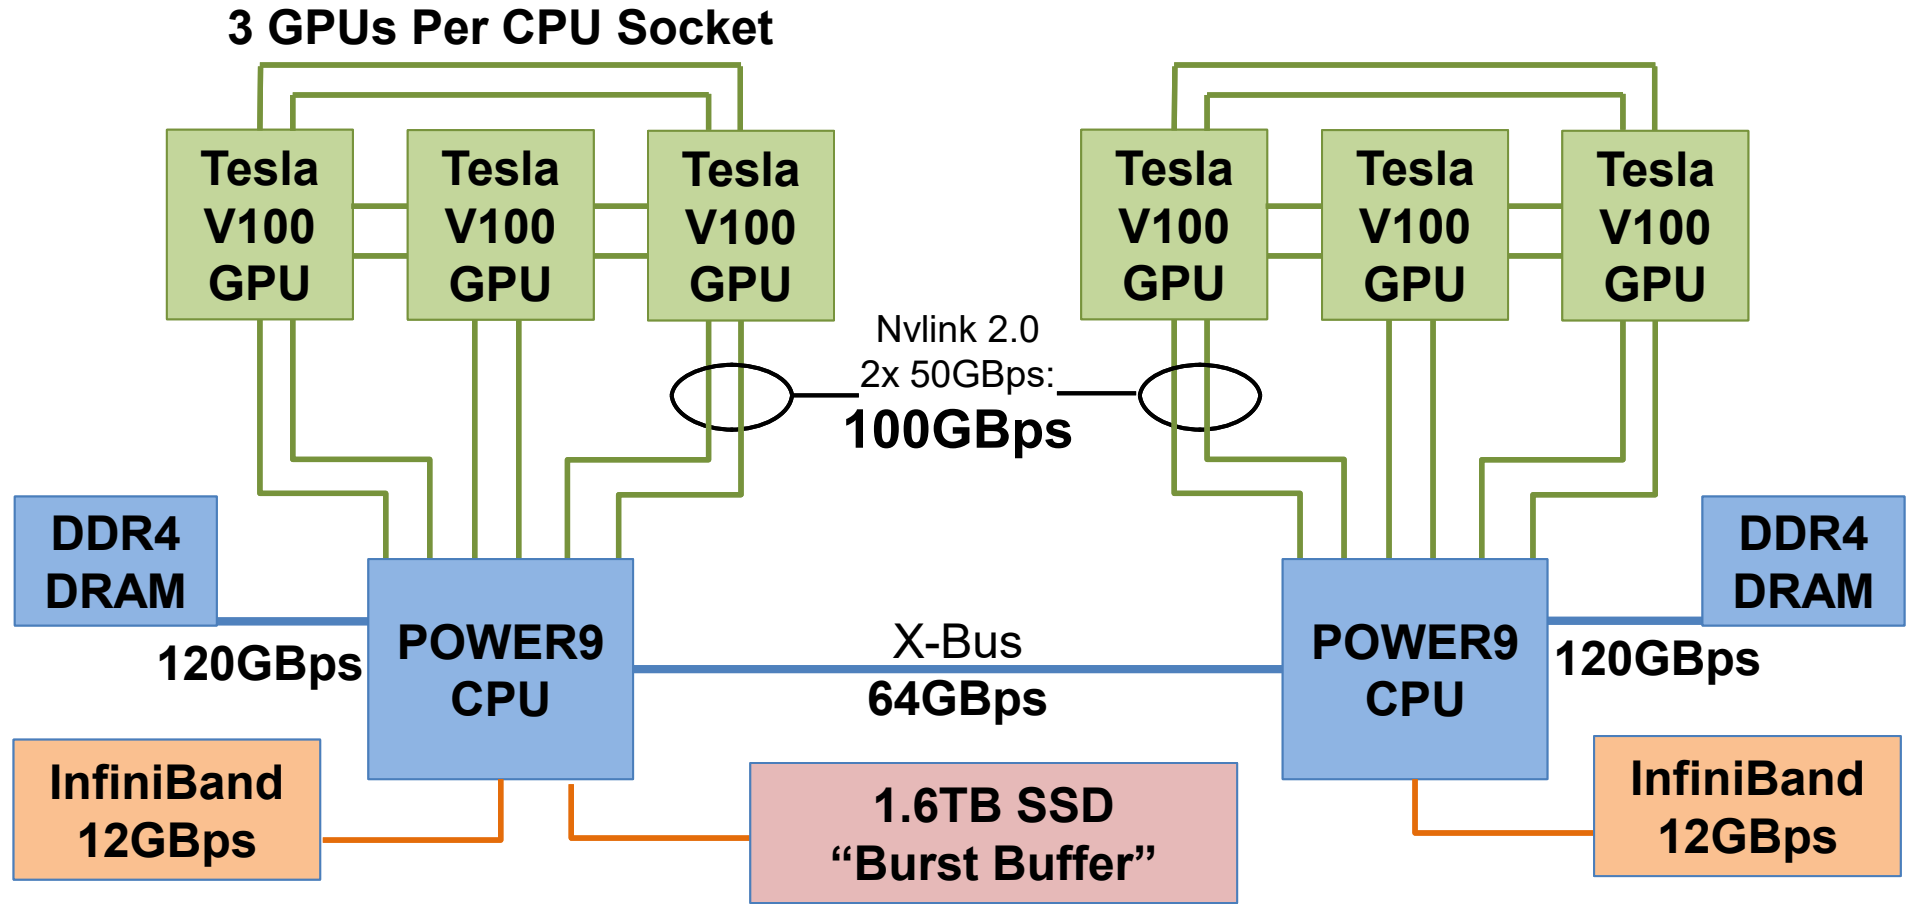

NAMD on Summit, May 2018

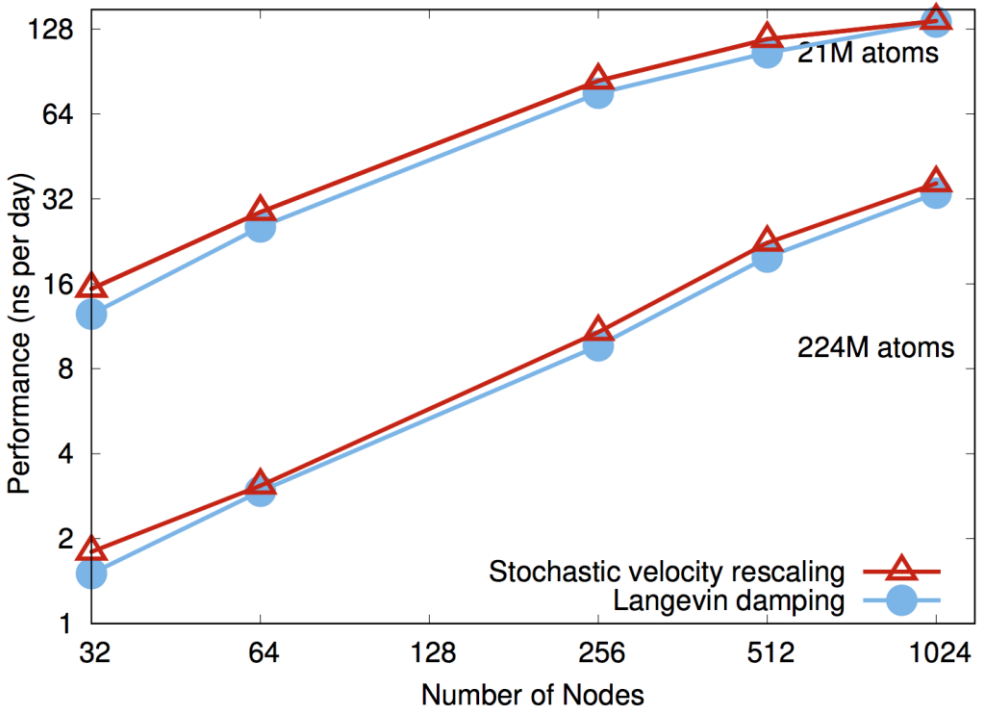

NAMD simulations can generate up to 10TB of output per day on 20% of Summit

Molecular Dyamics Trajectory Analysis

- MD simulations sample femtosecond timescales
- Millions of timesteps stored per trajectory
- Dynamics of biomolecular complexes are main interest, but **solvent often accounts for half or more of the simulation content**

  Skip I/O for regions of bulk solvent where possible [1]
- Modern MD tools, e.g., VMD, NAMD, LAMMPS, HOOMD, employ extensive **embedded scripting (Python, Tcl, etc)** to permit simulation preparation, custom simulation protocols, **analysis, and visualization**
- Unified collective variables module allows identical analytical computations to be performed within LAMMPS, NAMD, and VMD, during **pre-simulation modeling, in-situ, and post-hoc [2]**

[1] Immersive Out-of-Core Visualization of Large-Size and Long-Timescale Molecular Dynamics Trajectories. J. Stone, K. L. Vandivort, and K. Schulten. G. Bebis et al. (Eds.): *7th International Symposium on Visual Computing (ISVC 2011)*, LNCS 6939, pp. 1-12, 2011.

[2] Using collective variables to drive molecular dynamics simulations. G. Fiorin, M. L. Klein, and J. Hénin. *Molecular Physics*, 111:22-23, 3345-3362, 2013.

Petascale Molecular Dynamics I/O and Storage Challenges

- NAMD simulations can produce up to **10TB/day @ 1024 nodes (~20%) of ORNL Summit**, more as ongoing performance optimizations raise NAMD performance further
- Petascale science campaigns require **months of simulation runs**
- **Long-term storage of large-fractional petabytes impractical**
- **Historical "download output files for analysis and visualization" approach is a non-starter at this scale**
- Demands visualization and analysis operate on the data **in-place** on the HPC system, **whether *post-hoc, in-transit, or in-situ***
- Analyses must **identify salient features of structure, dynamics**, cull data that don't contribute to biomolecular processes of interest

VMD Petascale Visualization and Analysis

- Combination of growing system sizes and timescales of simulation trajectories poses a major data size challenge for molecular visualization and analysis
- Parallel I/O rates up to **275 GB/sec** on 8192 Cray XE6 nodes – can read in **231 TB in 15 minutes!**
- Analyze/visualize large trajectories **too large to transfer off-site:**
  - User-defined parallel analysis operations, data types
  - Parallel rendering, movie making
- Supports GPU-accelerated compute nodes for both visualization and analysis tasks:
  - GPU accelerated trajectory analysis w/ CUDA
  - OpenGL and GPU ray tracing for visualization and movie rendering

NCSA Blue Waters Hybrid Cray XE6 / XK7
22,640 XE6 dual-Opteron CPU nodes
4,224 XK7 nodes w/ Telsa K20X GPUs

Parallel VMD currently available on: ORNL Summit and Titan, NCSA Blue Waters, IU Big Red II, CSCS Piz Daint, many similar systems

VMD EGL Rendering: Supports full VMD GLSL shading features
Vulkan support coming soon...

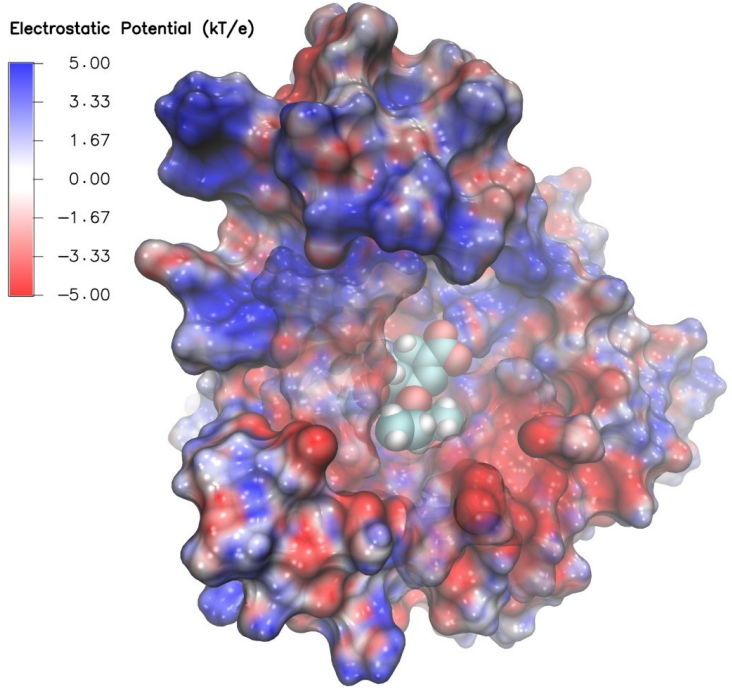

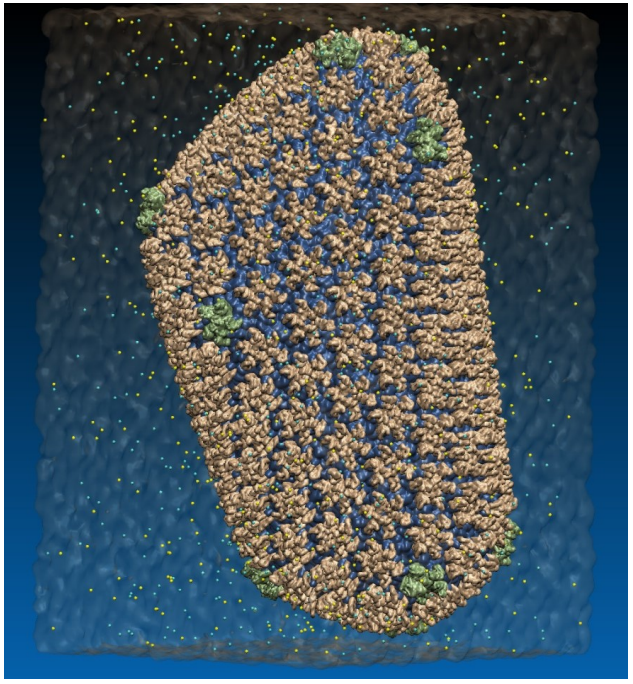

Swine Flu A/H1N1 neuraminidase bound to Tamiflu

64M atom HIV-1 capsid simulation

High Performance Molecular Visualization: In-Situ and Parallel Rendering with EGL.
J. E. Stone, P. Messmer, R. Sisneros, and K. Schulten. High Performance Data Analysis and Visualization Workshop, IEEE IPDPSW, pp. 1014-1023, 2016.

NIH

Next Generation: Simulating a Proto-Cell

- Emulate aspects of the *Mycoplasma mycoides* bacterium
- 200nm diameter
- ~1 billion atoms w/ solvent
- ~1400 proteins in membrane

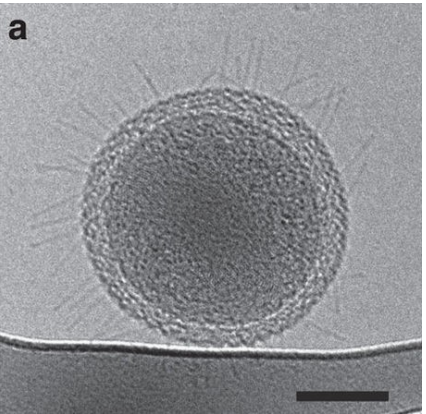

Cryo-ET image of ultra-small bacteria (scale bar 100nm) Luef et al. Nature Comm., 6:6372, 2015.

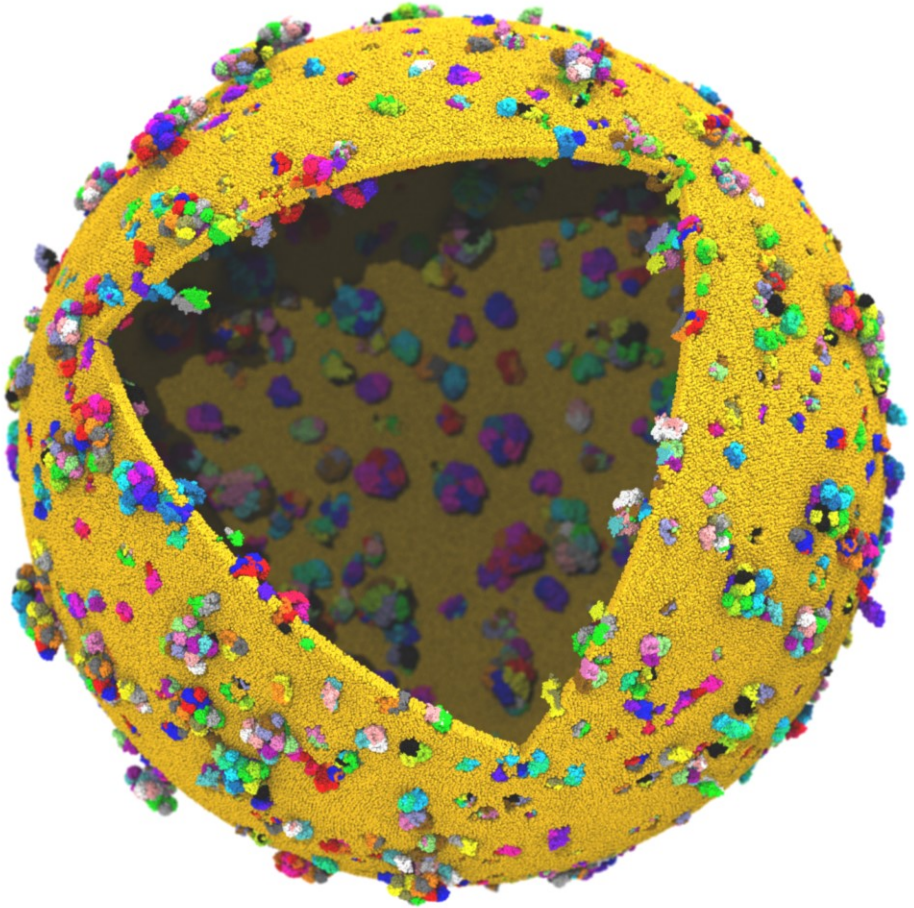

Proto-Cell Data Challenges

- **1B-atom proto-cell requires nodes with more than TB RAM to build complete model…**
- **1B-atom proto-cell binary structure file: 63GB**
- **Trajectory frame atomic coordinates: 12GB, 1.2TB/ns of simulation (1 frame per 10ps)**
- **Routine modeling and visualization tasks are a big challenge at this scale**
  - **Models contain thousands of atomic-detail components** that must work together in harmony
  - **Exploit persistent memory technologies** to enable “instant on” operation on massive cell-scale models – eliminate several minutes of startup during analysis/visualization of known structure
  - **Sparse output of results at multiple timescales** will help ameliorate visualization and analysis I/O
  - **Data quantization, compression, APIs like ZFP**

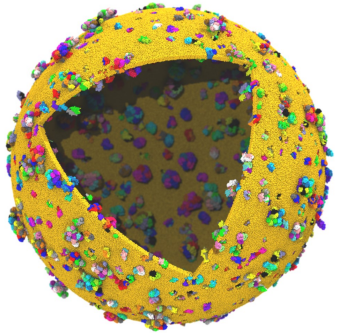

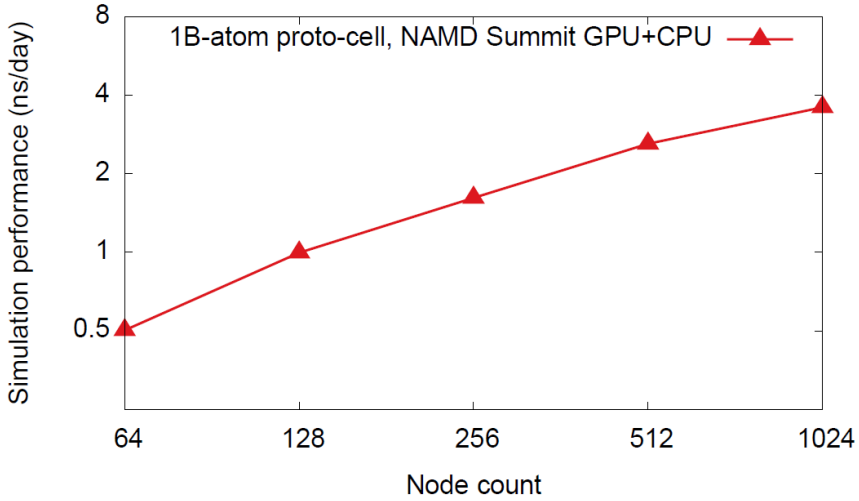

Clustering Analysis of Molecular Dynamics Trajectories:
Requires I/O+Memory for All-Pairs of Trajectory Frames

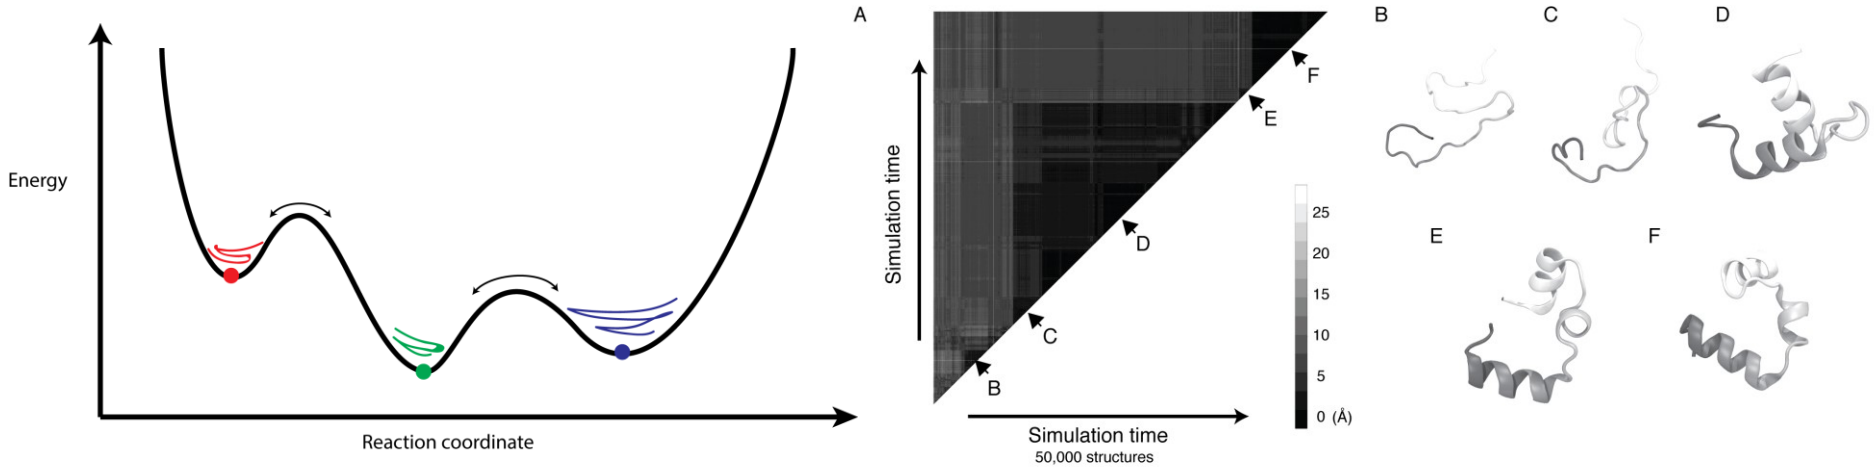

GPU-Accelerated Molecular Dynamics Clustering Analysis with OpenACC. J.E. Stone, J.R. Perilla, C. K. Cassidy, and K. Schulten. In, Robert Farber, ed., Parallel Programming with OpenACC, Morgan Kaufmann, Chapter 11, pp. 215-240, 2016.

Use of Node-Local Burst Buffers and Non-Volatile Memory DIMMs

- **Perform viz+analysis in-transit in node-local SSDs, persistent memory NVDIMMs**
- ORNL Summit I/O:
  - Parallel FS: 2.5 TB/s
  - **Node-local PCIe "burst buffer" SSDs: 10+ TB/sec, 7PB capacity**
- Plenty of capacity for full-detail MD trajectories, **could enable ~100x increase in temporal resolution** in cases where it would be valuable to the science
- **Enable all-pairs trajectory clustering analyses and resulting visualizations**
- Future systems with NVDIMMs (3D Xpoint, phase change memory) could eventually provide bandwidths approaching DRAM
- Use NVDIMMs w/ **mmap(), APIs like PMDK** to perform formerly-out-of-core calculations using persistent memory:

  https://github.com/pmem/pmdk
- **Imagine future Summit-like machines w/ NVLink-connected GPUs w/ access to high-bandwidth persistent memory on each node**

Trade FLOPS for Reduced I/O

ORNL Summit compute node:

- 6x Tesla V100 GPUs, 2x POWER9 CPUs
- GPUs Peak: ~46 DP TFLOPS, ~96 SP TFLOPS
- Peak IB rate per node: ~23GB/sec
- **Ratio of FLOPS vs. I/O:**

  ~2,000 DP FLOPS/byte, ~4000 SP FLOPS/byte

  ~16K FLOPS per FP word

Unconventional approach: Recompute to avoid I/O

Computing+Visualizing Molecular Orbitals

- Movies of simulation trajectories provides insight into results
- QM, and hybrid (QM/MM) MO visualizations historically done from huge “cube” files, impractical
- Store QM wavefunctions + Gaussian basis set, only 10s of KB per stored timestep compared to 100s of MB
- **Recompute MO grid on-the-fly from QM basis set, huge decrease in RAM+I/O in exchange for heavy FP arithmetic**

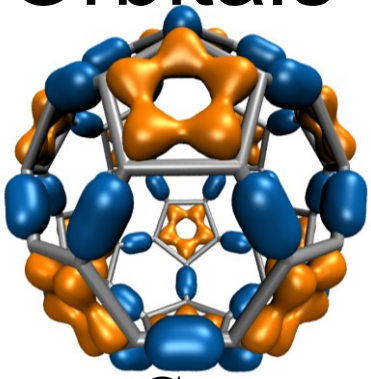

C_{60}

NAMD goes quantum: An integrative suite for hybrid simulations. Melo, M. C. R.; Bernardi, R. C.; Rudack T.; Scheurer, M.; Riplinger, C.; Phillips, J. C.; Maia, J. D. C.; Rocha, G. D.; Ribeiro, J. V.; Stone, J. E.; Neese, F.; Schulten, K.; Luthey-Schulten, Z.; **Nature Methods, 2018.**

http://dx.doi.org/10.1038/nmeth.4638

High Performance Computation and Interactive Display of Molecular Orbitals on GPUs and Multi-core CPUs. J. E. 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.

VMD C_{60} MO Viz. Perf, 516x519x507 Grid:
@ .13s/frame, avoids 3.8GB/s I/O per-node

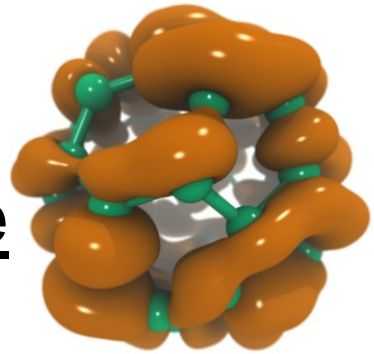

Hardware platform		Runtime,	Speedup
IBM Power8 (ORNL 'crest') + 1x Tesla K40 [1]		3.49s,	**1.0x**
Intel Xeon E5-2697Av4	**+ 1x Tesla V100**	**0.610s,**	**5.7x**
Intel Xeon E5-2697Av4	**+ 2x Tesla V100**	**0.294s,**	**11.8x**
Intel Xeon E5-2697Av4	**+ 3x Tesla V100**	**0.220s,**	**15.9x**
IBM Power9 "Newell"	**+ 1x Tesla V100**	**0.394s,**	**8.8x**
IBM Power9 "Newell"	**+ 2x Tesla V100**	**0.207s,**	**16.8x**
IBM Power9 "Newell"	**+ 3x Tesla V100**	**0.151s,**	**23.1x**
IBM Power9 "Newell"	**+ 4x Tesla V100**	**0.130s,**	**26.8x**

NVLink perf. boost w/ no code tuning (YET)

[1] Early Experiences Porting the NAMD and VMD Molecular Simulation and Analysis Software to GPU-Accelerated OpenPOWER Platforms. J. E. Stone, A.-P. Hynninen, J. C. Phillips, K. Schulten. International Workshop on OpenPOWER for HPC (IWOPH'16), LNCS 9945, pp. 188-206, 2016.

NIH

Omnidirectional Stereoscopic Ray Tracing

- **Ray trace 360° images and movies for Desk and VR HMDs: Oculus, Vive, Cardboard**
- Stereo spheremaps or cubemaps allow very high-frame-rate interactive OpenGL display
- **AO lighting, depth of field**, shadows, transparency, curved geometry, …
- **Summit 6x Tesla V100 GPU nodes:**
  - **Render many omni-stereo viewpoints, no acceleration structure rebuilds, tens of frames/sec per-node!**
  - **OptiX multi-GPU rendering, NVLink compositing and data distribution, etc…**
  - **Future: AI for warping between views**

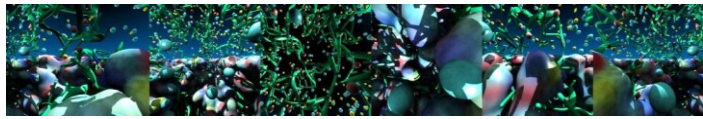

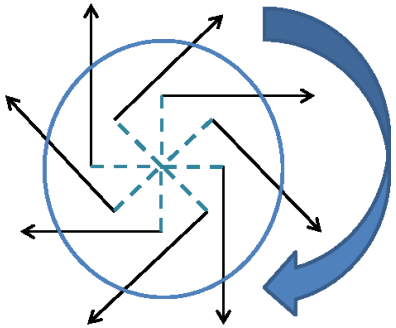

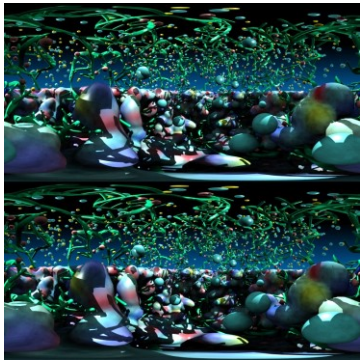

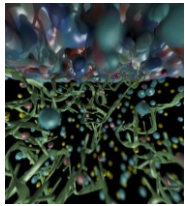

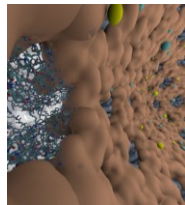

Atomic Detail Visualization of Photosynthetic Membranes with GPU-Accelerated Ray Tracing. J. E. Stone, et al. J. Parallel Computing, 55:17-27, 2016.

Immersive Molecular Visualization with Omnidirectional Stereoscopic Ray Tracing and Remote Rendering. J. E. Stone, W. R. Sherman, and K. Schulten. High Performance Data Analysis and Visualization Workshop, IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW), pp. 1048-1057, 2016.

Acknowledgements

- Theoretical and Computational Biophysics Group, University of Illinois at Urbana-Champaign
- IBM POWER team
- NVIDIA CUDA and OptiX teams
- Funding:
  - NIH support: P41GM104601
  - DOE INCITE, ORNL Titan: DE-AC05-00OR22725
  - NSF Blue Waters: Special thanks to James C. Phillips
    NSF OCI 07-25070, PRAC “The Computational Microscope”,
    ACI-1238993, ACI-1440026

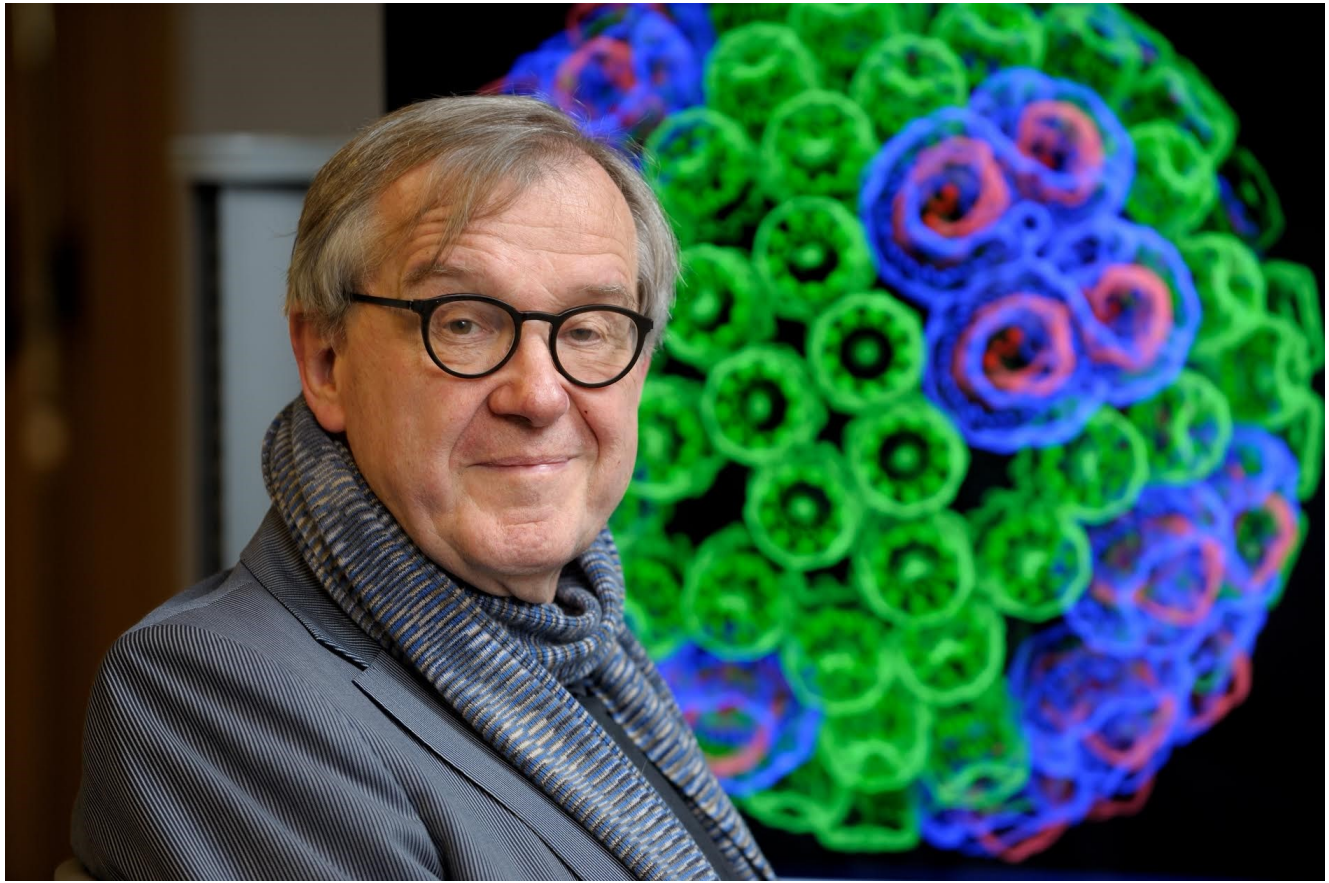

“When I was a young man, my goal was to look with mathematical and computational means at the inside of cells, one atom at a time, to decipher how living systems work. That is what I strived for and I never deflected from this goal.” – Klaus Schulten