Visualization Challenges and Opportunities Posed by 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/

Workshop on Molecular Graphics and Visual Analysis of Molecular Data

Co-located with EuroVis 2018

14:25-15:20, Hotel International,

Brno, Czech Republic, Thursday June 4th, 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/

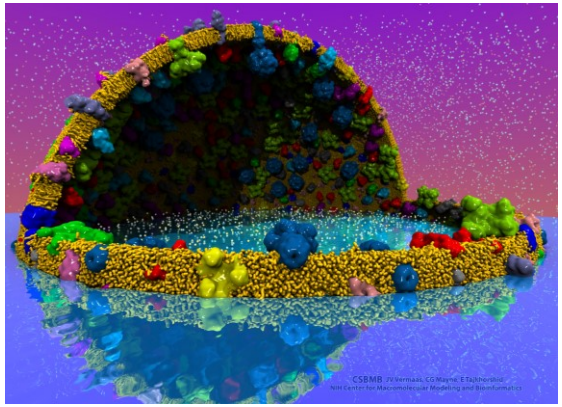

Cell-Scale Modeling

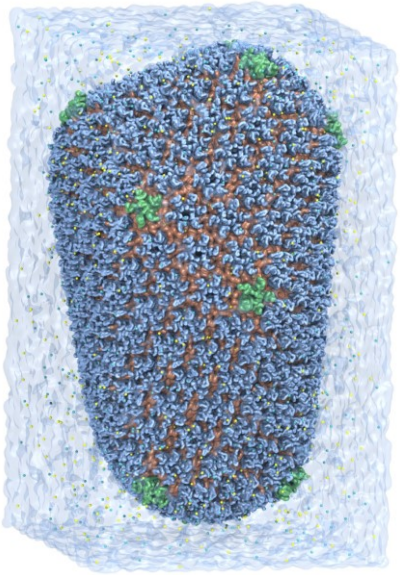

MD Simulation

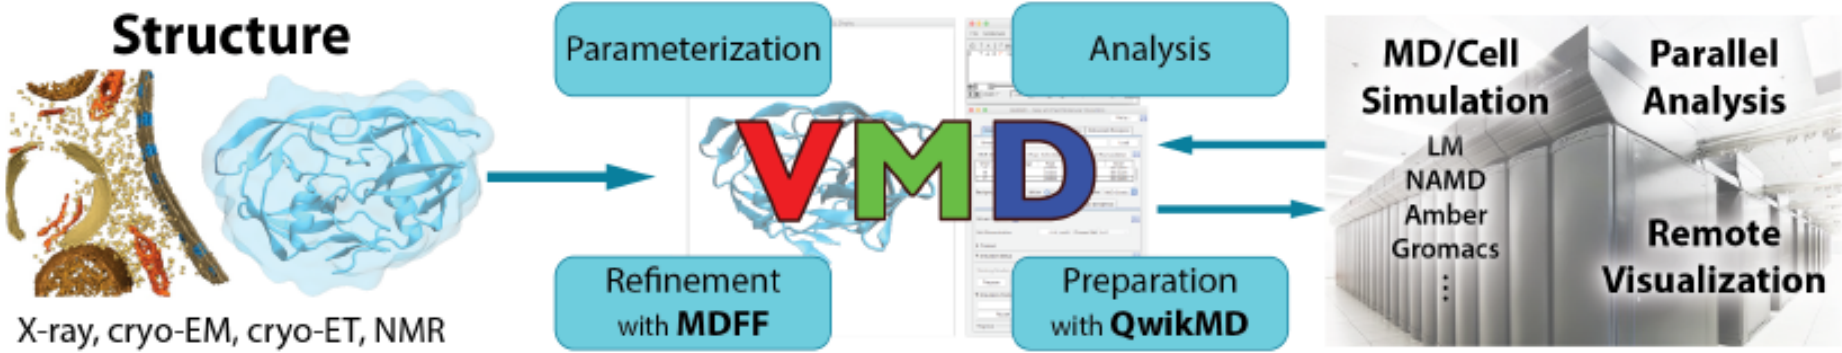

Goal: A Computational Microscope

Study the molecular machines in living cells

Ribosome: target for antibiotics

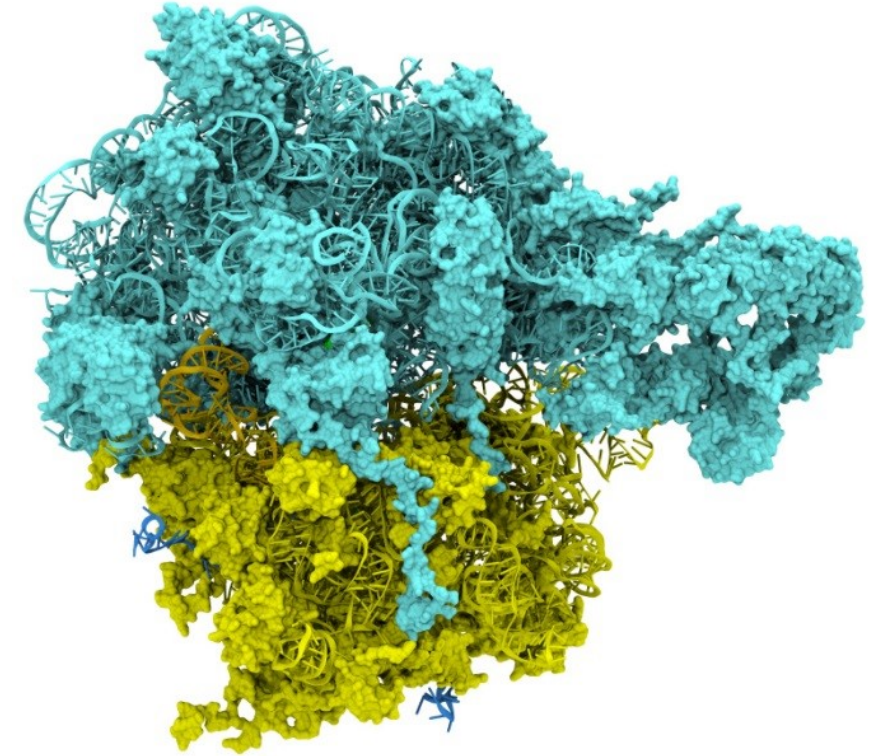

Poliovirus

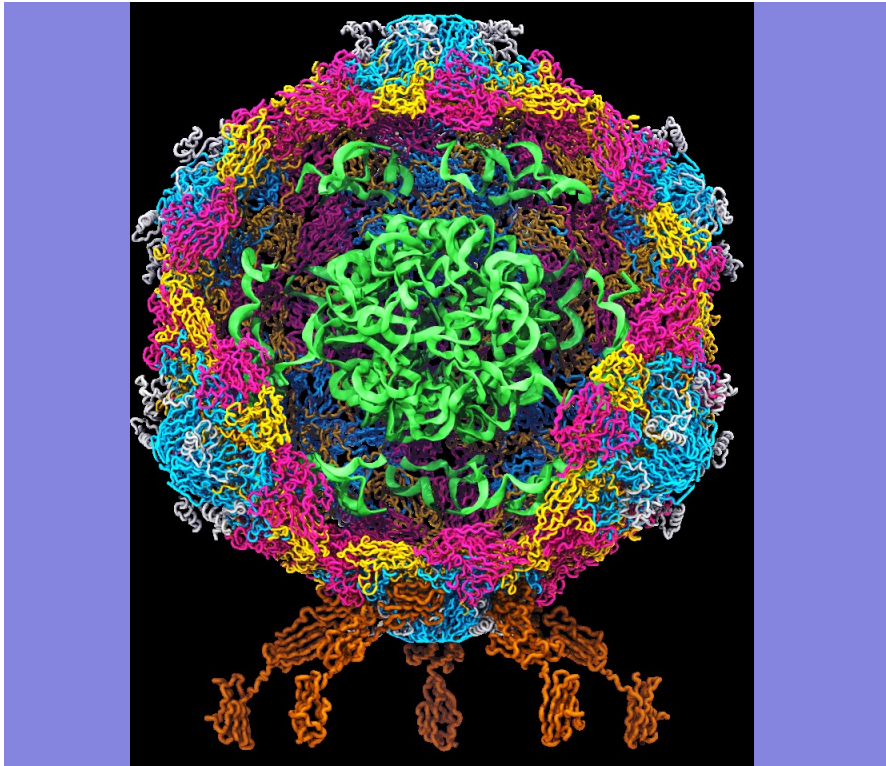

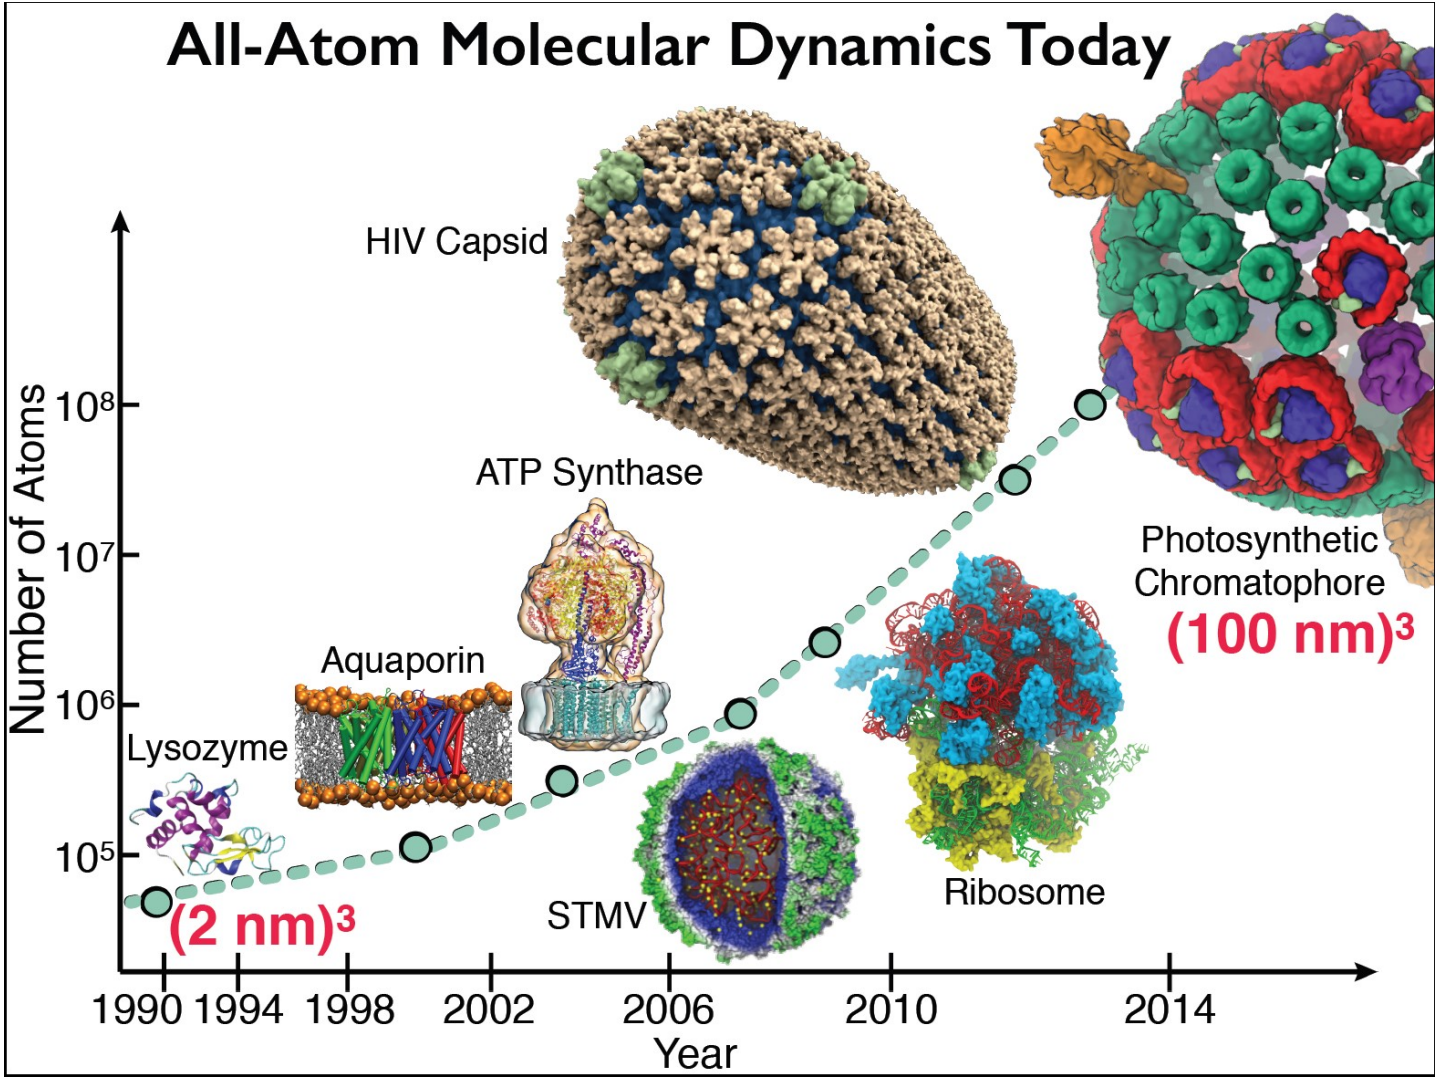
All-Atom Molecular Dynamics Today
HIV Capsid
ATP Synthase
Photosynthetic
Chromatophore
(100 nm)3
Aquaporin
Lysozyme
STMV
Ribosome
(2 nm)3
Number of Atoms
10^8
10^7
10^6
10^5
1990
1994
1998
2002
2006
2010
2014
Year

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

Petascale Simulations Driving NAMD/VMD Development

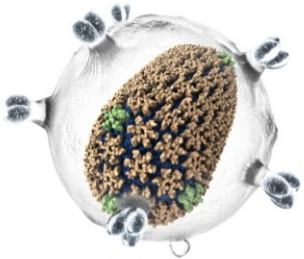

Viral Infection

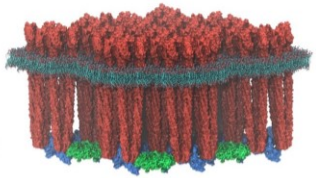

Symbiont Bacteria

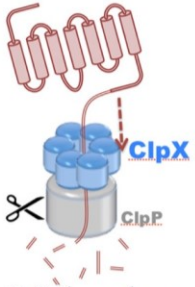

Molecular Motors

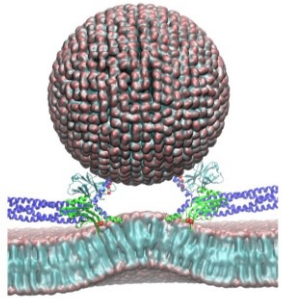

Neurons and Synapses

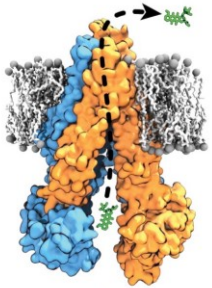

Membrane Transporters

NCSA ORNL	Blue Waters (4,228 XK7 nodes) Titan (18,688 XK7 nodes)	AMD Opteron + K20X Kepler GPU	16 CPU cores / GPU
TACC	Stampede 2 (4200 KNL nodes, 1736 Skylake nodes)	Intel Knights Landing Intel Xeon Skylake	68 CPU cores 48 CPU cores
ORNL	**Summit (~4600 nodes, 27,000 GPUs)**	**2 IBM Power9 CPUs + 6 Tesla V100 GPUs**	**7 CPU cores / GPU**

IBM AC922 Summit Node

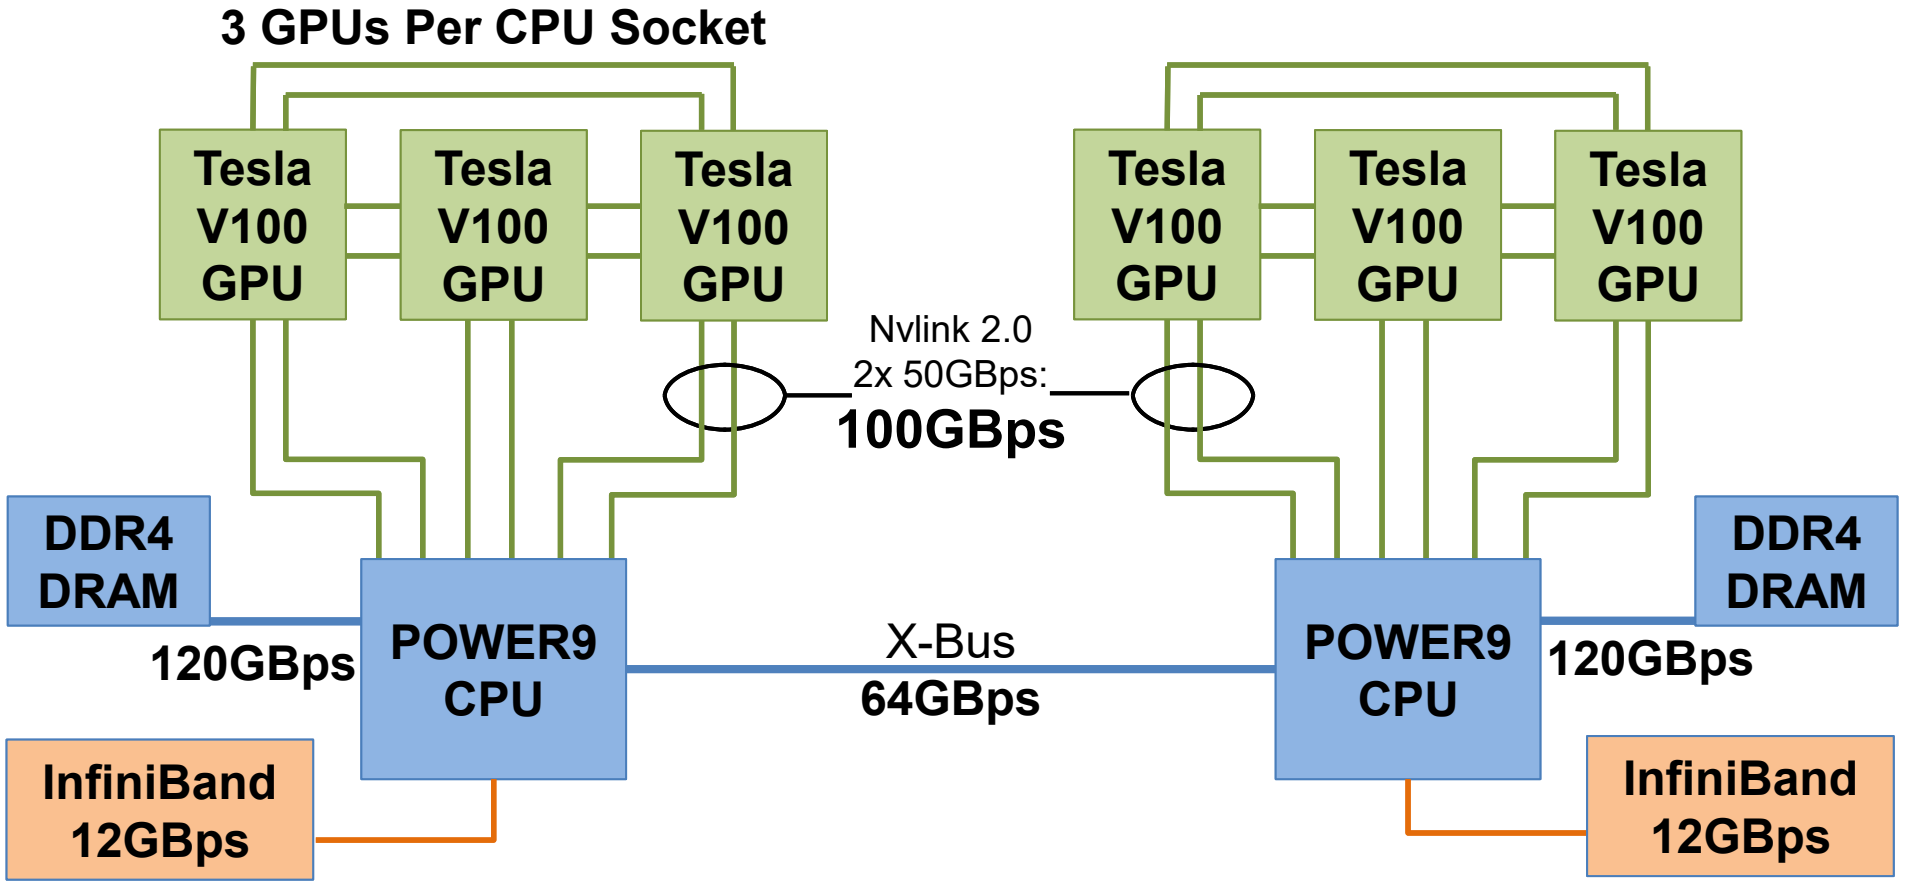

NAMD Performance on Summit, May 2018

- **NAMD Summit GPU performance somewhat CPU-bound (at present)**
- **NAMD scaling efficiency is currently affected by system noise** in the early developmental OS/software
- Stochastic velocity rescaling yields up to **20% performance benefit** vs. Langevin damping
- **Reduction in RNG/entropy required per timestep** yields immediate performance benefits
- **Reduces CPU-bound components of NAMD timestep integration**

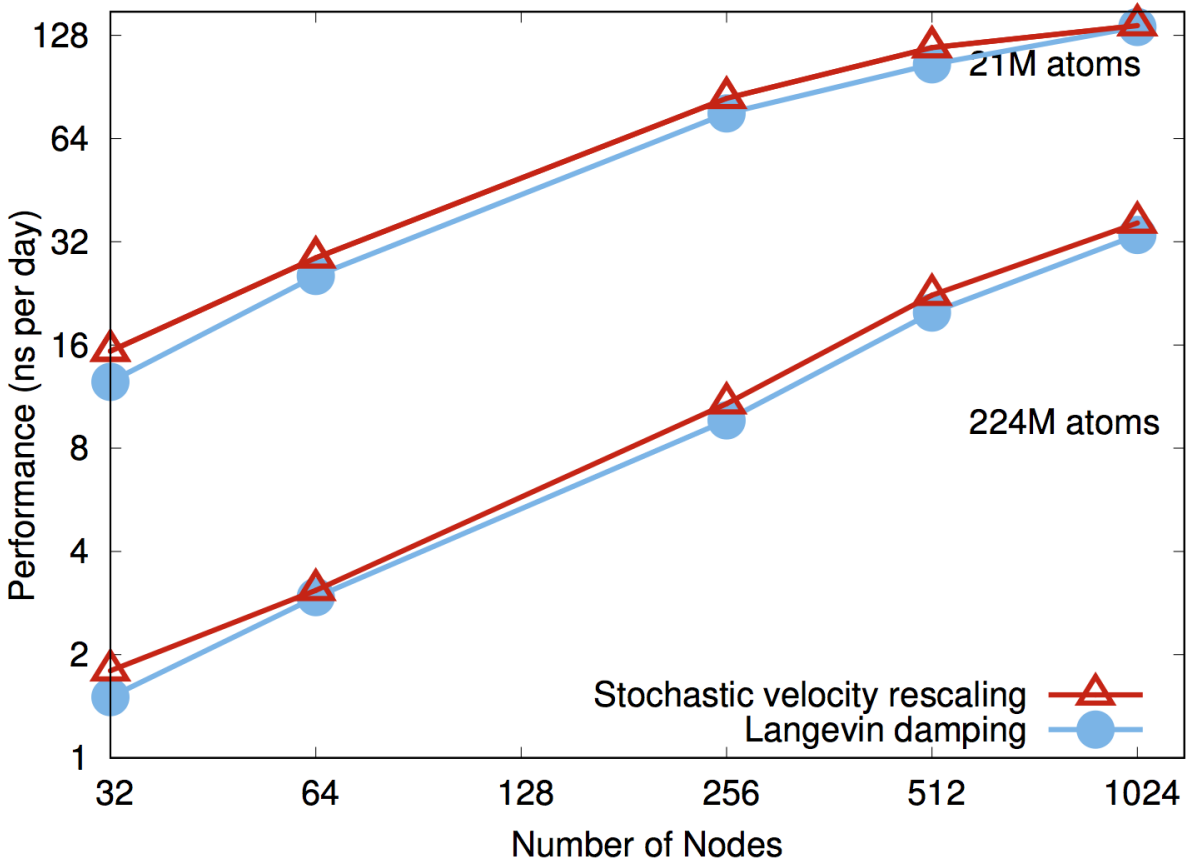

Petascale Molecular Dynamics I/O and Storage Challenges

- **Petascale science campaigns require months of simulation runs**
- **Historical "download output files for interactive visualization" approach is a non-starter at this scale**
- 12 bytes/atom/timestep for common trajectory file formats
- NAMD performance levels on Summit can generate **more than 10TB/day @ 1024 nodes (20% of final machine)**, w/ trajectory frames written once per 10ps (once per 5,000 steps w/ 2fs steps)
- **Long-term storage of large-fractional petabytes is impractical**
- 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 biomolecular processes of interest

Next Generation: Simulating a Protocell

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

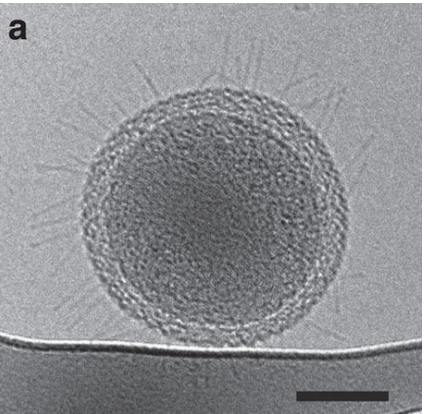

Cryo-ET image of ultra-small bacteria (scale bar 100nm) Luef et al. Nature Comm., 6:6372, 2015.

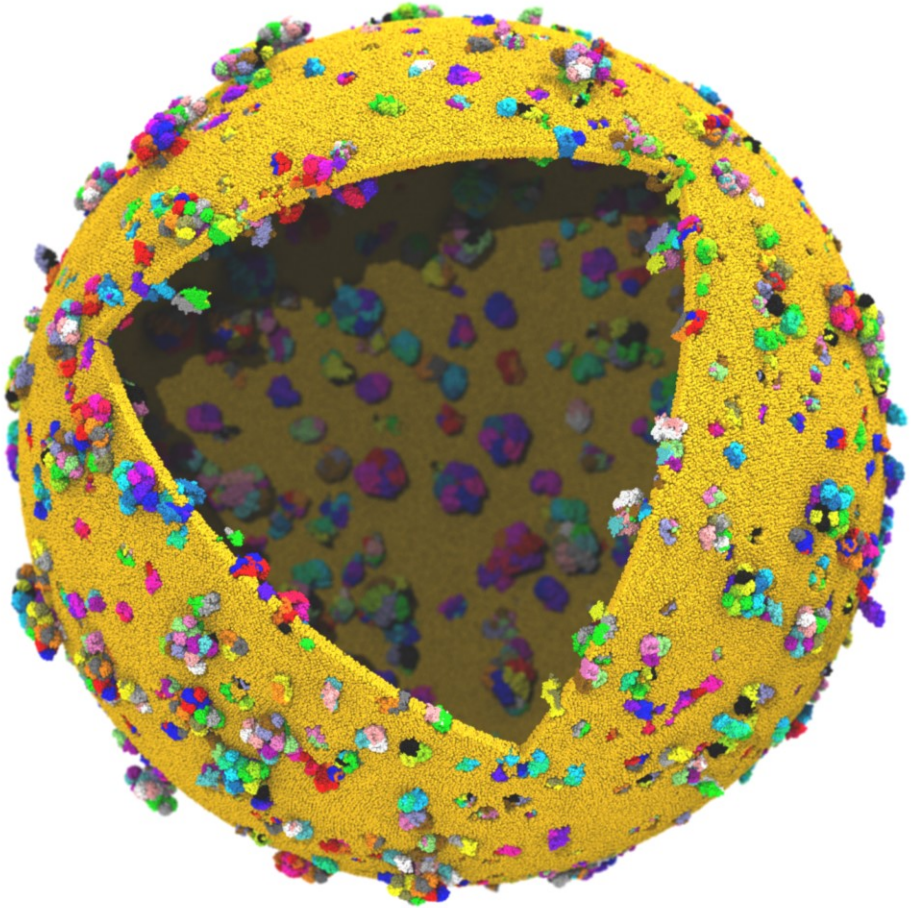

Protocell Data Challenges

- **Requires machine w/ ~1TB RAM to build complete protocell model from scratch…**
- **Complete protocell binary structure file: ~63GB**
- **Trajectory frame atomic coordinates: 12GB, 1.2TB/ns of simulation (1 frame per 10ps)**
- At start, VMD (currently) reads+analyzes the structure entirely from scratch:
  - **Existing file formats don't yet encompass cell-scale organization**, requires many new levels of additional structure information
  - **8 minutes to do structure checking**, identify and validate everything from atoms, to residues, to segments, classify protein vs. nucleic vs. carbohydrate, identify+record key structure components, etc.
  - With **better file formats**, VMD would trust inputs and avoid reanalyzing the structure at load time
- **Interactive** modeling and visualization tasks are a big challenge at this scale
  - **Models contain thousands of atomic-detail components** that must all work together in harmony
  - **New interactive visual analysis schemes are needed** to help scientists work on models that reach the sub-cellular and minimal-cell scale…
  - **Exploit persistent memory technologies** to enable "instant on" operation on massive cell-scale models

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

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

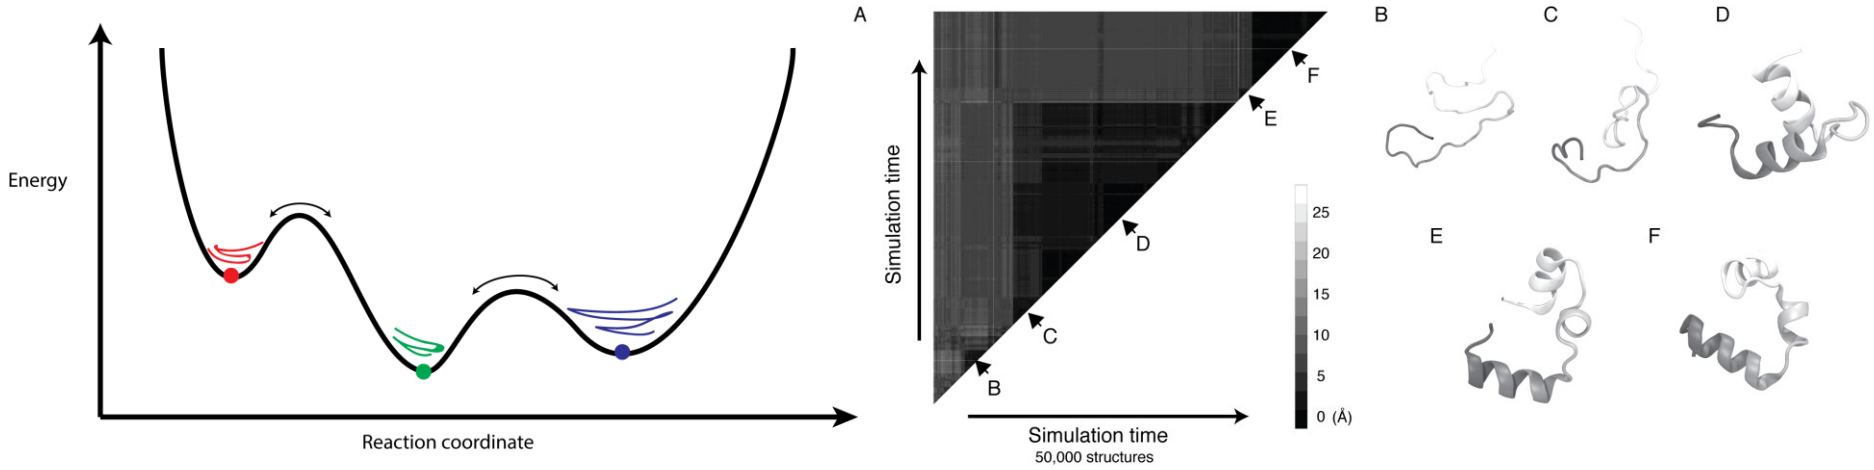

GPU-Accelerated Molecular Dynamics Clustering Analysis with OpenACC. J.E. Stone, J.R. Perilla, C. K. Cassidy, and K. Schulten. In, Robert Farber, ed., Parallel Programming with OpenACC, Morgan Kaufmann, Chapter 11, pp. 215-240, 2016.

Molecular Dynamics Flexible Fitting (MDFF)

X-ray crystallography → MDFF ← Electron microscopy

APS at Argonne

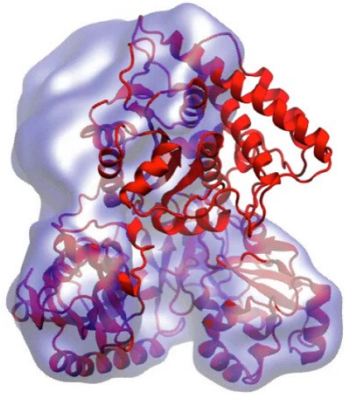

FEI microscope

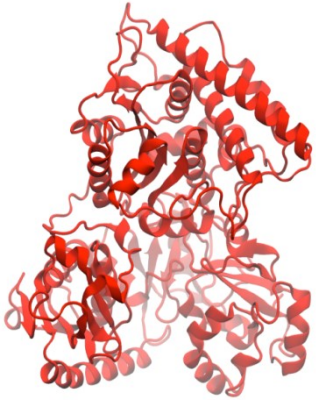

ORNL Titan

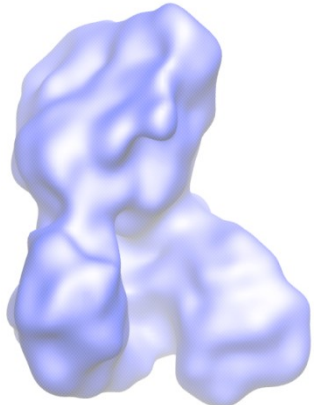

Molecular dynamics-based refinement and validation for sub-5Å cryo-electron microscopy maps. A. Singharoy, I. Teo, R. McGreevy, J. E. Stone, J. Zhao, and K. Schulten. eLife 2016;10.7554/eLife.16105

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

Compute Pearson correlation to evaluate **quality-of-fit** between a reference cryo-EM density map and a **simulated density map** from an **all-atom structure**.

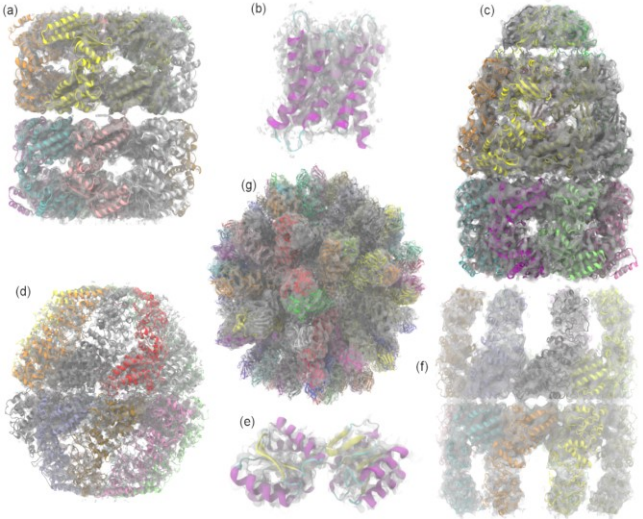

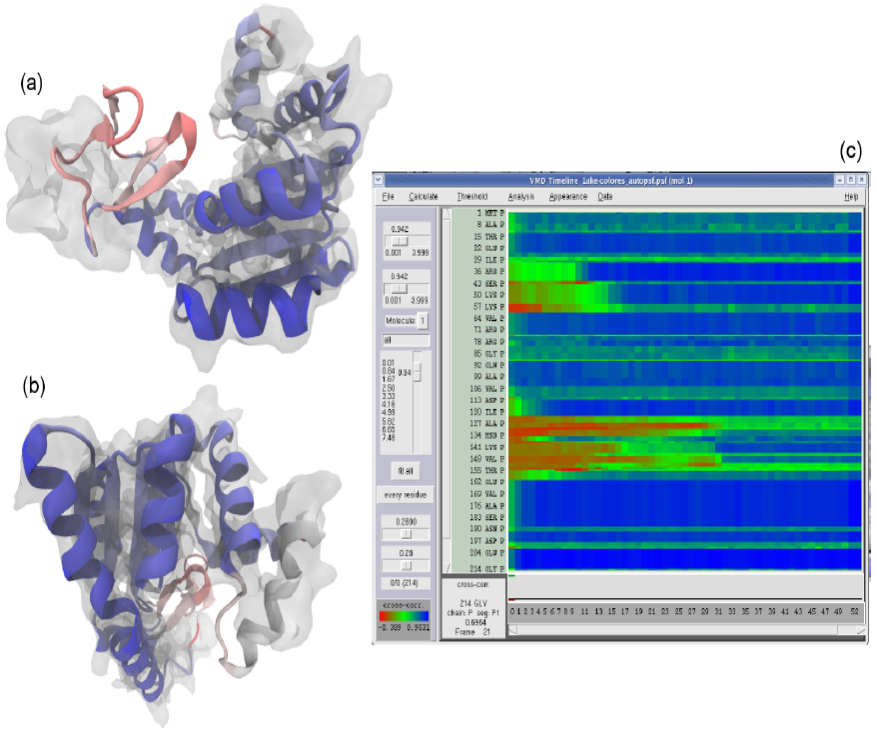

MDFF Cross Correlation Timeline

Regions with poor fit **Regions with good fit**

Parallel MDFF Cross Correlation Analysis on Cray XK7

Rabbit Hemorrhagic Disease Virus (RHDV)	
Traj. frames	10,000
Structure component selections	720
Single-node XK7 (projected)	336 hours (14 days)
128-node XK7	3.2 hours 105x speedup
2048-node XK7	19.5 minutes 1035x speedup

Calculation of 7M CCs would take **5 years** using serial CPU algorithm!

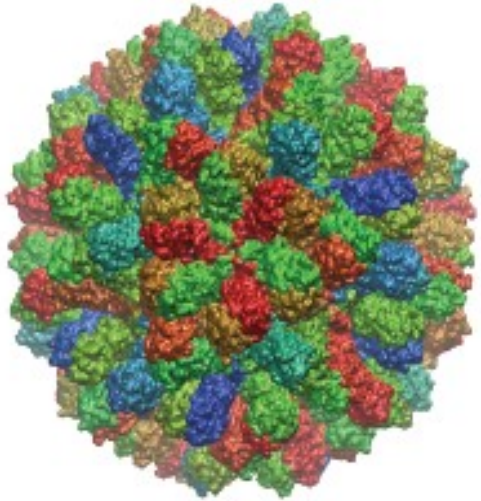

RHDV colored by relative CC

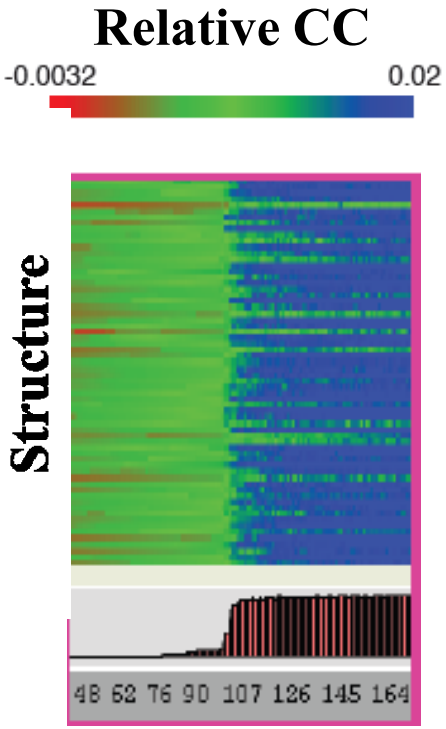

Stone et al., Faraday Discuss., 169:265-283, 2014.

VMD Tesla V100 Cross Correlation Performance

Rabbit Hemorrhagic Disease Virus: 702K atoms, 6.5Å resolution
Volta GPU architecture almost 2x faster than previous gen Pascal:

Application and Hardware platform	Runtime, Speedup vs. Chimera,		VMD+GPU
Chimera Xeon E5-2687W (2 socket) [1]	15.860s,	1x	
VMD-CUDA IBM Power8 + 1x Tesla K40 [2]	0.488s,	32x	**0.9x**
VMD-CUDA Intel Xeon E5-2687W + 1x Quadro K6000 [1,2]	0.458s,	35x	**1.0x**
VMD-CUDA Intel Xeon E5-2698v3 + 1x Tesla P100	**0.090s,**	176x	**5.1x**
VMD-CUDA IBM Power8 "Minsky" + 1x Tesla P100	**0.080s,**	198x	**5.7x**
VMD-CUDA Intel Xeon E5-2697Av4 + 1x Tesla V100	**0.050s,**	317x	**9.2x**
VMD-CUDA IBM Power9 "Newell" + 1x Tesla V100	**0.049s,**	**323x**	**9.3x**

[1] 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 Discussions 169:265-283, 2014.
[2] 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.

Challenges Posed by Next-Gen GPU-Dense Workstation/Node Designs

Application and Hardware platform	Runtime,	VMD+GPU
VMD-CUDA Intel Xeon E5-2687W + 1x Quadro K6000 [1,2]	0.458s,	**1.0x**
VMD-CUDA IBM Power9 "Newell" + 1x Tesla V100	**0.049s,**	**9.3x**

- **~9x performance gain from Kepler to Volta GPUs in 4 years**
- CPUs and PCIe have not matched this rate of performance gain
- **Challenge: new GPU-dense nodes** have far less CPU available to manage GPUs and execute non-GPU code
  - **More GPUs per CPU socket, fewer CPU threads per GPU**
  - **ORNL Summit 3 GPUs/socket, 7 CPU cores per GPU**
  - **Remaining CPU code needs to be FAST, otherwise it must move to GPU**

Cryo-EM / Cryo-ET Image Segmentation

Evaluate 3-D volumetric electron density maps and segment them, to identify key structural components

Index/label components so they can be referred to, colored, analyzed, and simulated…

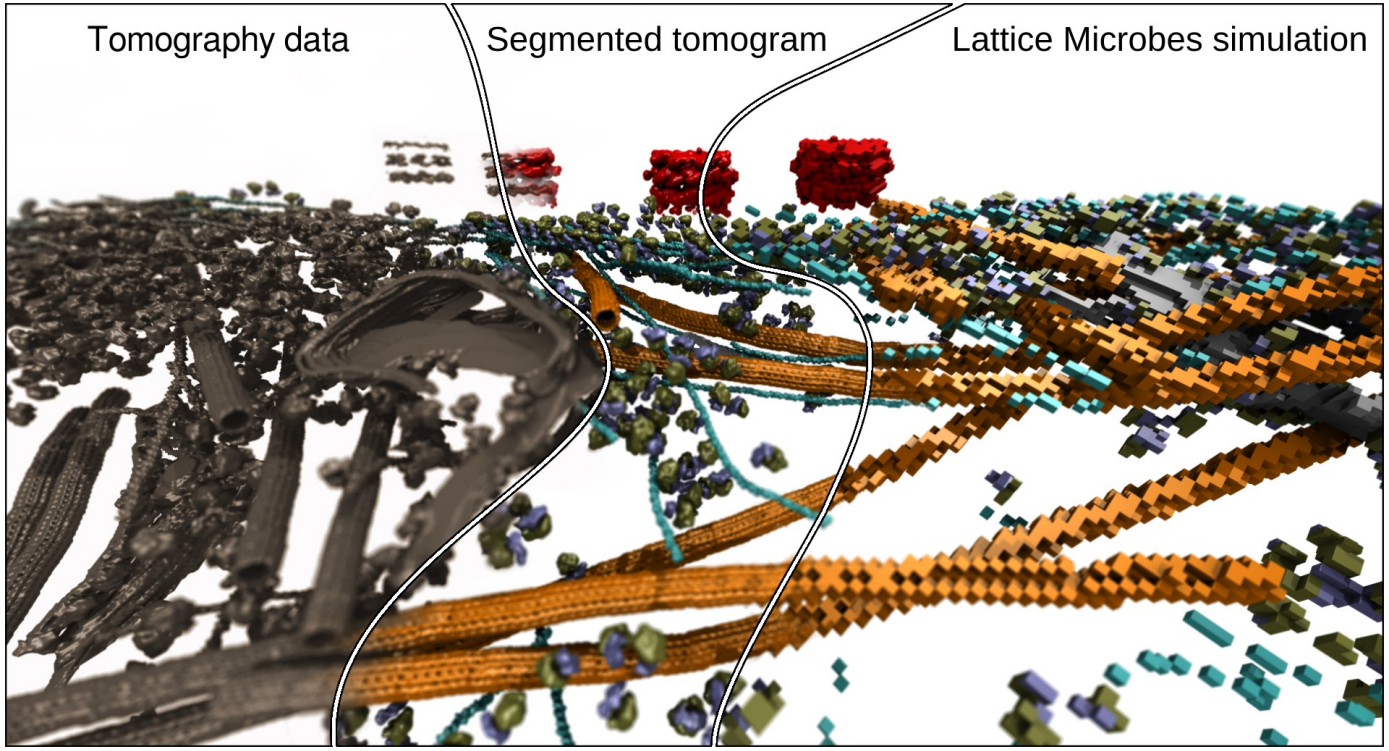

Cryo-EM Density Map Segmentation Approach, goals

Watershed segmentation:

- Smooth/denoise image (e.g. blur)
- Find local minima of image/gradients
- Connect minimum voxels with neighbors of similar intensity, marking them with the same "group" number
- "Grow" each group (merging groups where rules allow) until no more updates occur

Scale-space segmentation variant does further blurring and group merging to reach a user-specified target segment count

Goals:

- **Reach interactive performance rates (under 1 second)** for common density map sizes between 128^3 to 256^3 voxels
- Handle **next-generation problem sizes (768^3 to 2048^3)** smoothly with only a brief wait

Key methods:

- **Watershed: Tile-based early-exit schemes** pervasively used for all iterative segmentation update/merge kernels
- **Privatization, shared memory atomic counters** for segmentation group index kernels
- **Significantly faster (12x-25x) than other algorithm designs we are aware of**

Density Map Segmentation Optimization Opportunities

- Optimize numerical precision for 3-D density maps:
  - **Improved memory bandwidth, lower arithmetic cost**
  - FP16: half-precision EM density map representation
  - INT8: byte density map representation for EM tomograms
- **NVIDIA Tensor Core (~120 TFLOPS FP16 MMA)** for iterative scale-space segmentation merge/blur convolutions, initial noise filtering steps:
  - **Difficult to prevent TC from becoming mem bandwidth-bound**
  - Some dimensionalities and matrix-based problem decompositions perform much better than others – this is an area of ongoing exploration

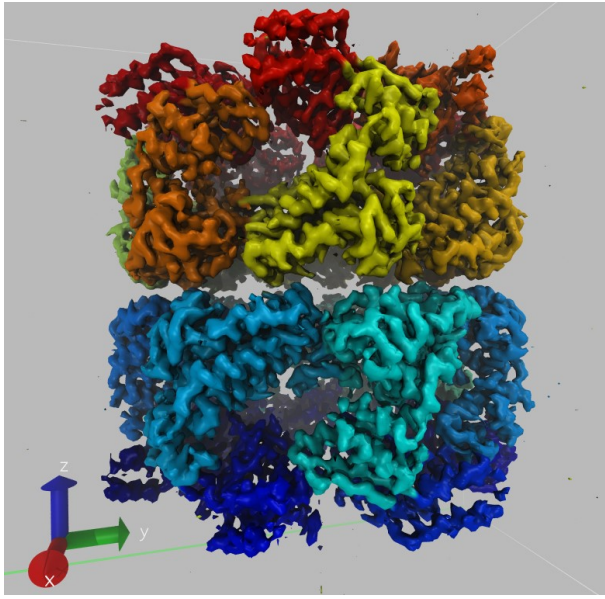

VMD GPU density map segmentation of GroEL

Computing+Visualizing Molecular Orbitals

- Animation of (classical mechanics) molecular dynamics trajectories provides insight into simulation results
- To do the same for QM or hybrid QM/MM simulations one must compute MOs at ~**5-10 FPS** or more
- Re-compute MO grid on-the-fly from QM basis set, **huge decrease in RAM+I/O in exchange for heavy FP arithmetic**

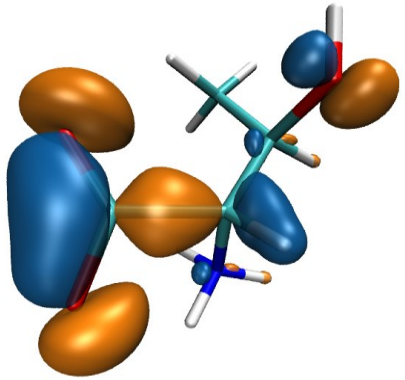

NAMD goes quantum: An integrative suite for hybrid simulations. Melo, M. C. R.; Bernardi, R. C.; Rudack T.; Scheurer, M.; Riplinger, C.; Phillips, J. C.; Maia, J. D. C.; Rocha, G. D.; Ribeiro, J. V.; Stone, J. E.; Neese, F.; Schulten, K.; Luthey-Schulten, Z.; **Nature Methods, 2018.**

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.

MO Kernel for One Grid Point (Naive C)

```
…
for (at=0; at<numatoms; at++) {
   int prim_counter = atom_basis[at];
   calc_distances_to_atom(&atompos[at], &xdist, &ydist, &zdist, &dist2, &xdiv);
   for (contracted_gto=0.0f, shell=0; shell < num_shells_per_atom[at]; shell++) {
      int shell_type = shell_symmetry[shell_counter];
      for (prim=0; prim < num_prim_per_shell[shell_counter];  prim++) {
         float exponent         = basis_array[prim_counter      ];
         float contract_coeff = basis_array[prim_counter + 1];
         contracted_gto += contract_coeff * expf(-exponent*dist2);
         prim_counter += 2;
      }
      for (tmpshell=0.0f, j=0, zdp=1.0f; j<=shell_type; j++, zdp*=zdist) {
         int imax = shell_type - j;
         for (i=0, ydp=1.0f, xdp=pow(xdist, imax); i<=imax; i++, ydp*=ydist, xdp*=xdiv)
            tmpshell += wave_f[ifunc++] * xdp * ydp * zdp;
      }
      value += tmpshell * contracted_gto;
      shell_counter++;
   }
} …..
```

Loop over atoms

Loop over shells

Loop over primitives: largest component of runtime, due to **expf()**

Loop over angular momenta

(unrolled in real code)

VMD Tesla V100 Performance for C_{60} Molecular Orbitals, 516x519x507 grid

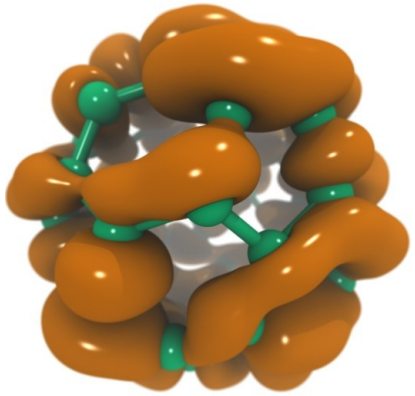

Hardware platform		Runtime,	Speedup
IBM Power8 (ORNL 'crest') + 1x Tesla K40 [1]		3.49s,	**1.0x**
Intel Xeon E5-2697Av4	**+ 1x Tesla V100**	**0.610s,**	**5.7x**
Intel Xeon E5-2697Av4	**+ 2x Tesla V100**	**0.294s,**	**11.8x**
Intel Xeon E5-2697Av4	**+ 3x Tesla V100**	**0.220s,**	**15.9x**
IBM Power9 "Newell"	**+ 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**
IBM Power9 "Newell"	**+ 6x Tesla V100**	**0.156s,**	**22.3x**

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

Need tune

[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

Molecular Orbital Alg. Behavior on Summit

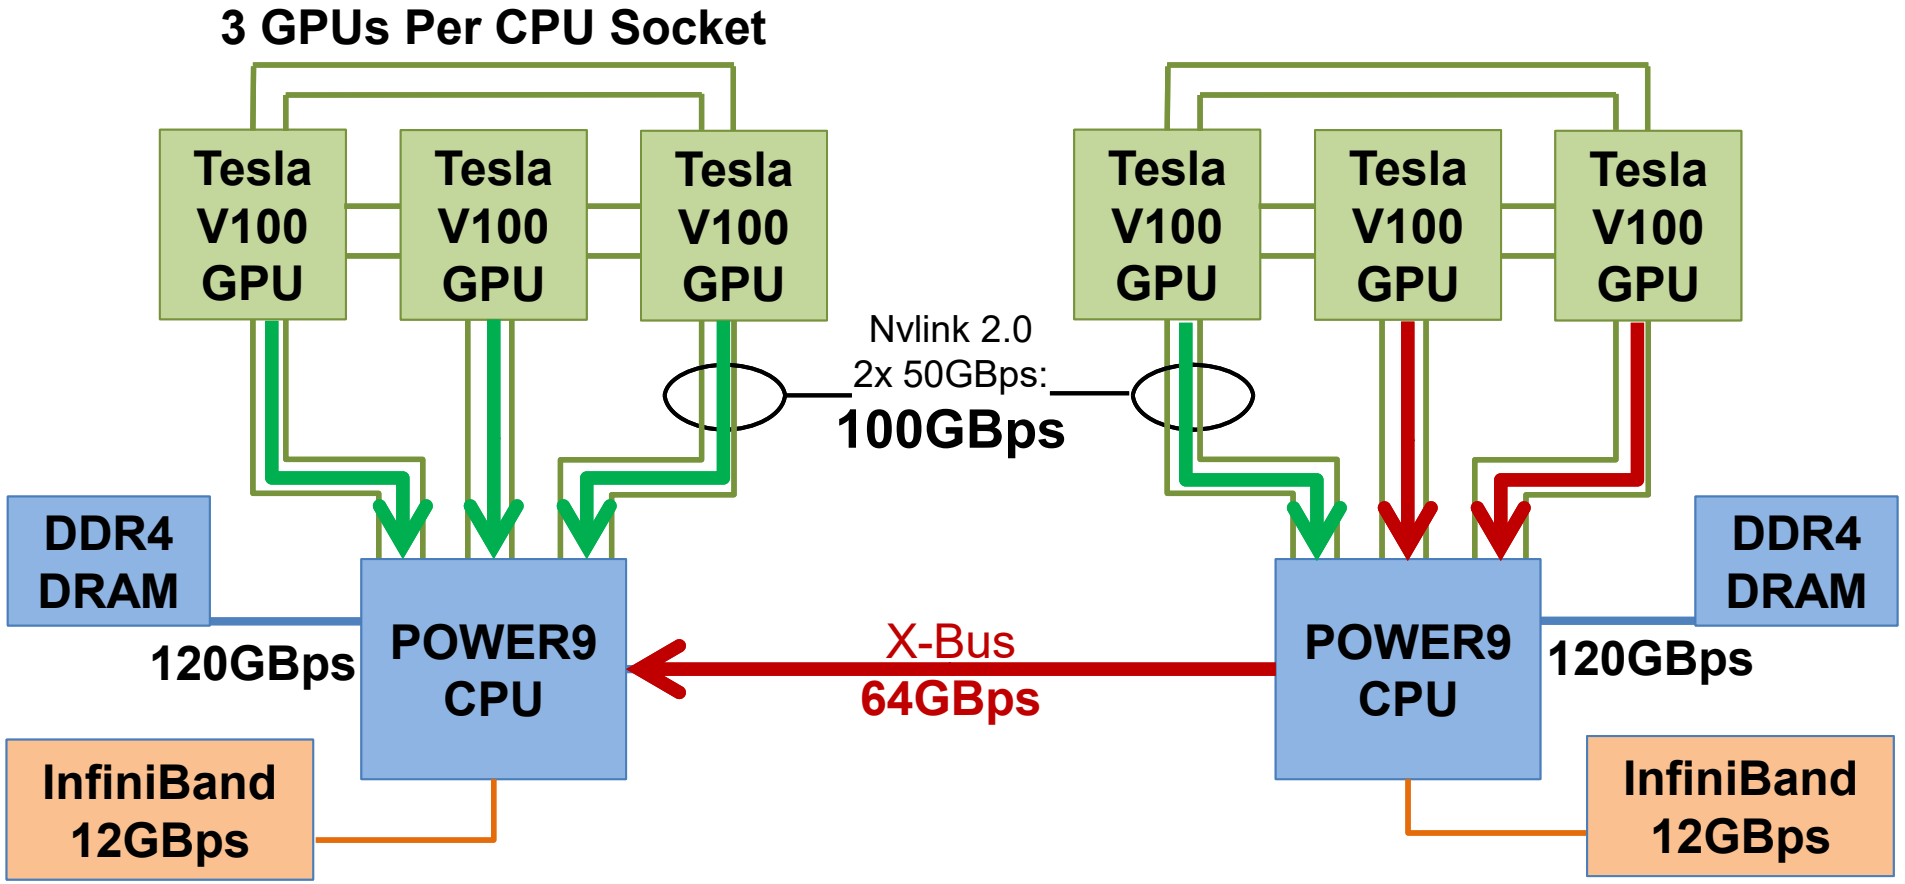

Molecular Orbital Computation and Display Process

Runtime Kernel Generation, NVRTC Just-In-Time **(JIT) Compilation: 1.8x Faster**

One-time initialization

Initialize Pool of GPU Worker Threads

Read QM simulation log file, trajectory

Preprocess MO coefficient data
eliminate duplicates, sort by type, etc…

Generate/compile basis set-specific CUDA kernel

For each trj frame, for each MO shown

For current frame and MO index,
retrieve MO wavefunction coefficients

Compute 3-D grid of MO wavefunction amplitudes using basis set-specific CUDA kernel

Extract isosurface mesh from 3-D MO grid

Render the resulting surface

VMD Off-Screen Raserization w/ EGL

- **Eliminate requirement for windowing system**, support HPC visualization workloads, in-situ, and remote visualization
- VMD on HPC systems w/ EGL:
  - Cray XC50, CSCS Piz Daint
  - **ORNL Summit in progress now**
  - IBM OpenPOWER, NVIDIA drivers 375.66 and later support both GLX and EGL
- Containers+Cloud+Workstations with recent NVIDIA drivers

VMD EGL Rendering: Supports full VMD GLSL shading features, multisample antialiasing, ray cast spheres, 3-D tex mapping, ...

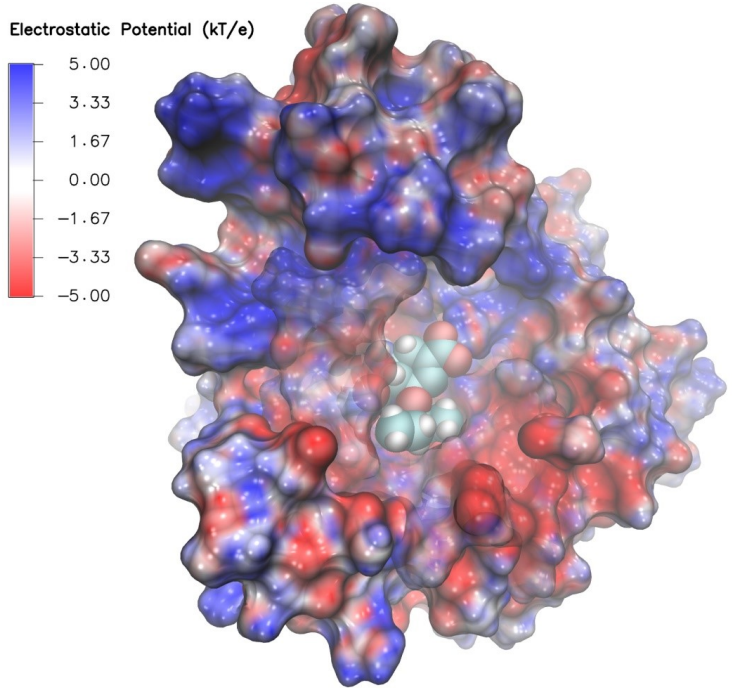

Swine Flu A/H1N1 neuraminidase bound to Tamiflu

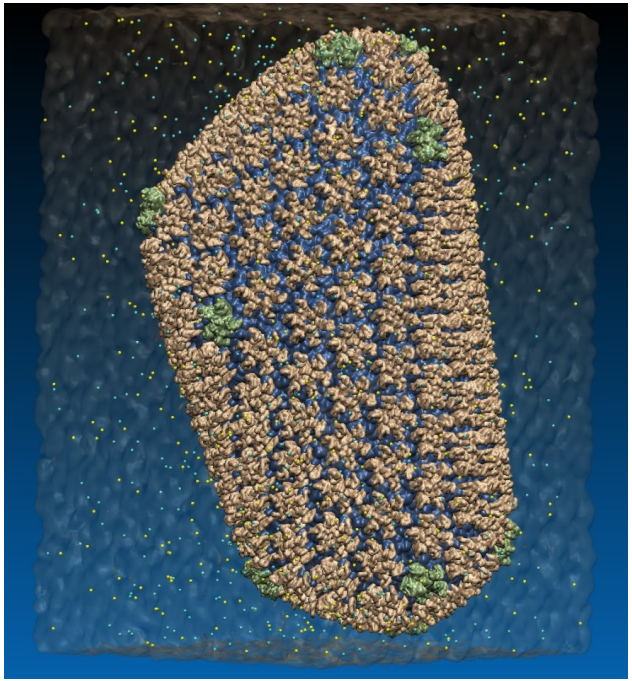

64M atom HIV-1 capsid simulation

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.

Beckman Institute,
U. Illinois at Urbana-Champaign

Molecular Structure Data and Global VMD State

Scene Graph

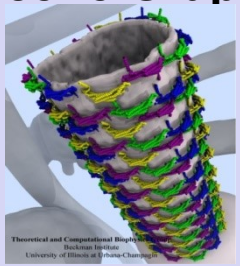

Graphical Representations

- DrawMolecule
- Non-Molecular Geometry

User Interface Subsystem

- Tcl/Python Scripting
- Mouse + Windows
- 6DoF Input “Tools”

Display Subsystem

- VMDDisplayList
- DisplayDevice
- OpenGLRenderer

GLX+X11+Drv

- OpenGL Pbuffer/FBO
- Windowed OpenGL

EGL+Drv

- OpenGL Pbuffer/FBO

Benefits of EGL Platform Interfaces for HPC Systems

- **EGL interfaces make it EASY to bind a GPU** to a thread with optimal CPU affinity **with respect to NUMA topology, NVLink GPU topology**
  - High-perf. multi-GPU image compositing, video streaming
  - EGL plays nicely with MPI, CUDA/OpenCL, OptiX, NVENC, etc
  - NVIDIA EGL supports multiple GPU indexing schemes, e.g. **PCIe ordering**
  - **Exploit NVLink interconnect topology on IBM OpenPOWER platforms, DOE/ORNL "Summit" system**

VMD on Vulkan

- **High performance, low-overhead, low-abstraction rasterization API**
- ***In-progress:*** **Vulkan-based rasterization path for VMD:**
  - **Vulkan ideally suited** as the API for **high-end rasterization approaches**
  - Maintain existing OpenGL renderer to support integrated/legacy GPUs
  - **Parallel Vulkan command buffer generation** will allow deep multithreading of time-varying VMD graphical representation updates
  - VMD Vulkan rendering path will be able to go all-in on techniques that are only viable on high-end GPUs
  - **Headless operation supported**, akin to EGL and GLX Pbuffer APIs
    - **Headless Vulkan works now on CSCS Piz Daint Cray XC50 test and development system with a few Cray-specific tweaks**
    - **Testing on ORNL Summit and POWER9 CPUs ongoing**

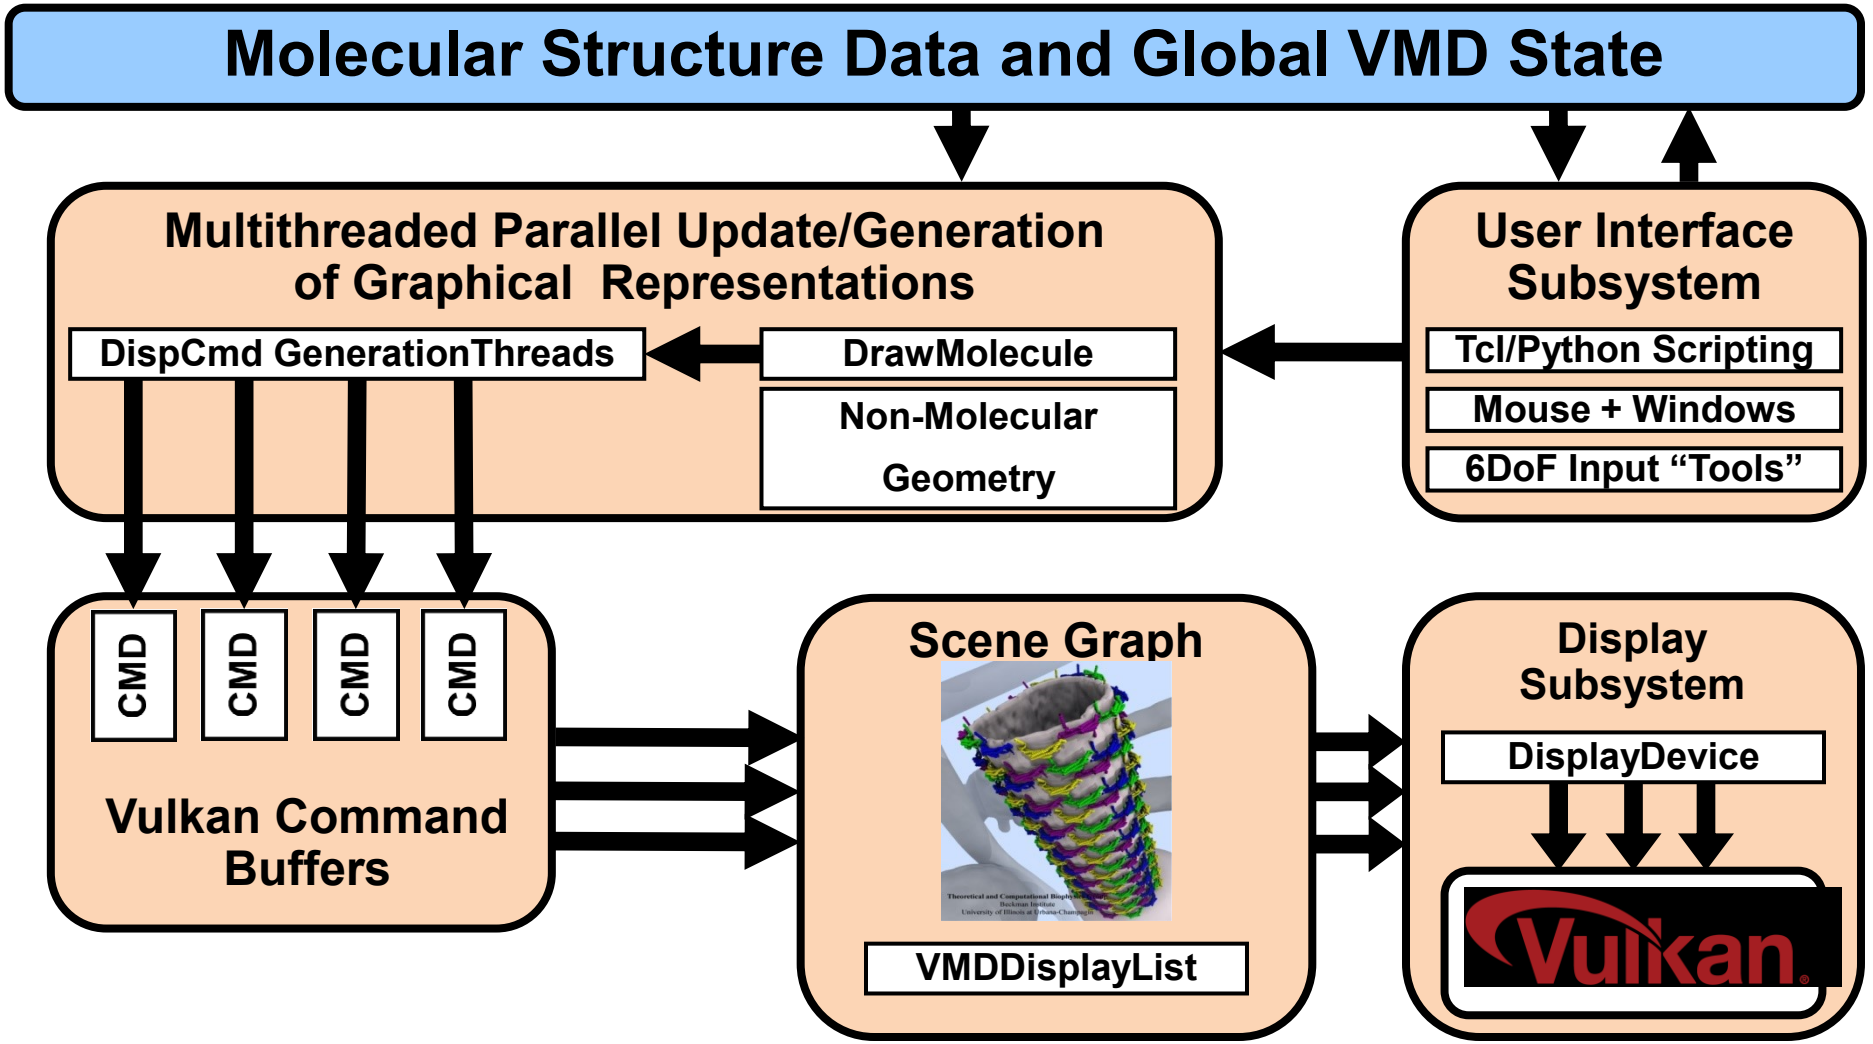
Molecular Structure Data and Global VMD State
Multithreaded Parallel Update/Generation of Graphical Representations
DispCmd GenerationThreads
DrawMolecule
Non-Molecular Geometry
User Interface Subsystem
Tcl/Python Scripting
Mouse + Windows
6DoF Input "Tools"
CMD
CMD
CMD
CMD
Vulkan Command Buffers
Scene Graph
VMDDisplayList
Display Subsystem
DisplayDevice
Vulkan

VMD Interactive Ray Tracing

- **Uses hardware-optimized RT frameworks for both GPUs and CPUs**
- **Interactive RT** on laptops, desktops, cloud, and HPC clusters or supercomputers
- Large-scale parallel rendering: in situ or post hoc visualization
- Remote RT on NVIDIA GPU clusters
- Stereoscopic panoramic and full-dome projections
- **Omnidirectional VR for YouTube, VR HMDs**
- **GPU memory sharing via NVLink, e.g., Tesla V100, Quadro GV100**
- **In-progress:**
  - **OptiX denoising support: fast turnaround w/ AO, DoF, etc**
  - **Denoising to enable practical use of path tracing in VMD**

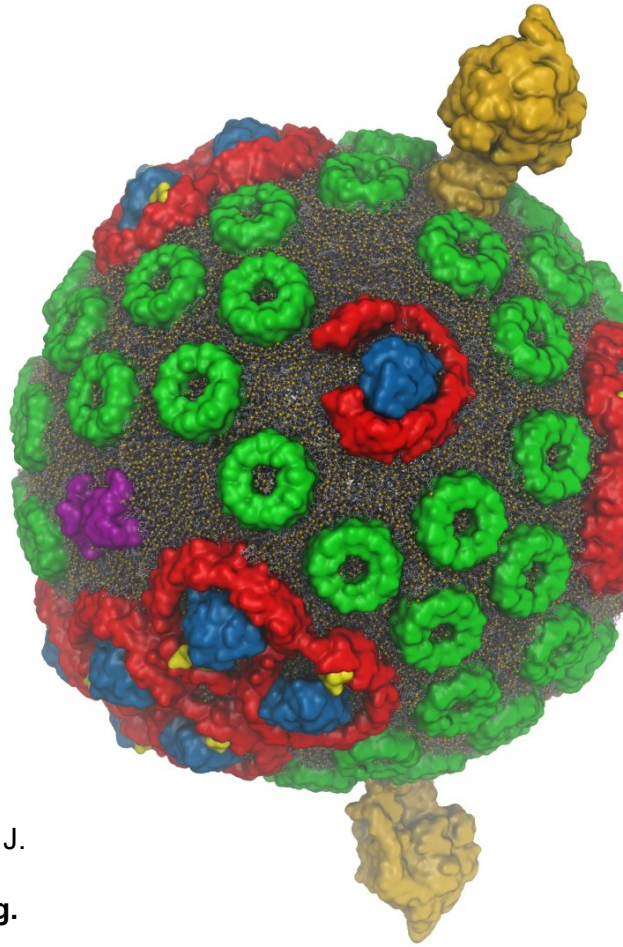

VMD/OptiX GPU Ray Tracing of all-atom Chromatophore w/ lipids.

GPU-Accelerated Molecular Visualization on Petascale Supercomputing Platforms. J. E. Stone, K. L. Vandivort, and K. Schulten. UltraVis'13, pp. 6:1-6:8, 2013.

Visualization of Energy Conversion Processes in a Light Harvesting Organelle at Atomic Detail. M. Sener, et al. SC'14 Visualization and Data Analytics Showcase, 2014.

Chemical Visualization of Human Pathogens: the Retroviral Capsids. J. R. Perilla, B.-C. Goh, J. E. Stone, and K. Schulten. SC'15 Visualization and Data Analytics Showcase, 2015.

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. HPDAV, IPDPSW, pp. 1048-1057, 2016.

Lighting Comparison, STMV Capsid

Two lights, no shadows

Two lights, hard shadows, 1 shadow ray per light

Ambient occlusion + two lights, 144 AO rays/hit

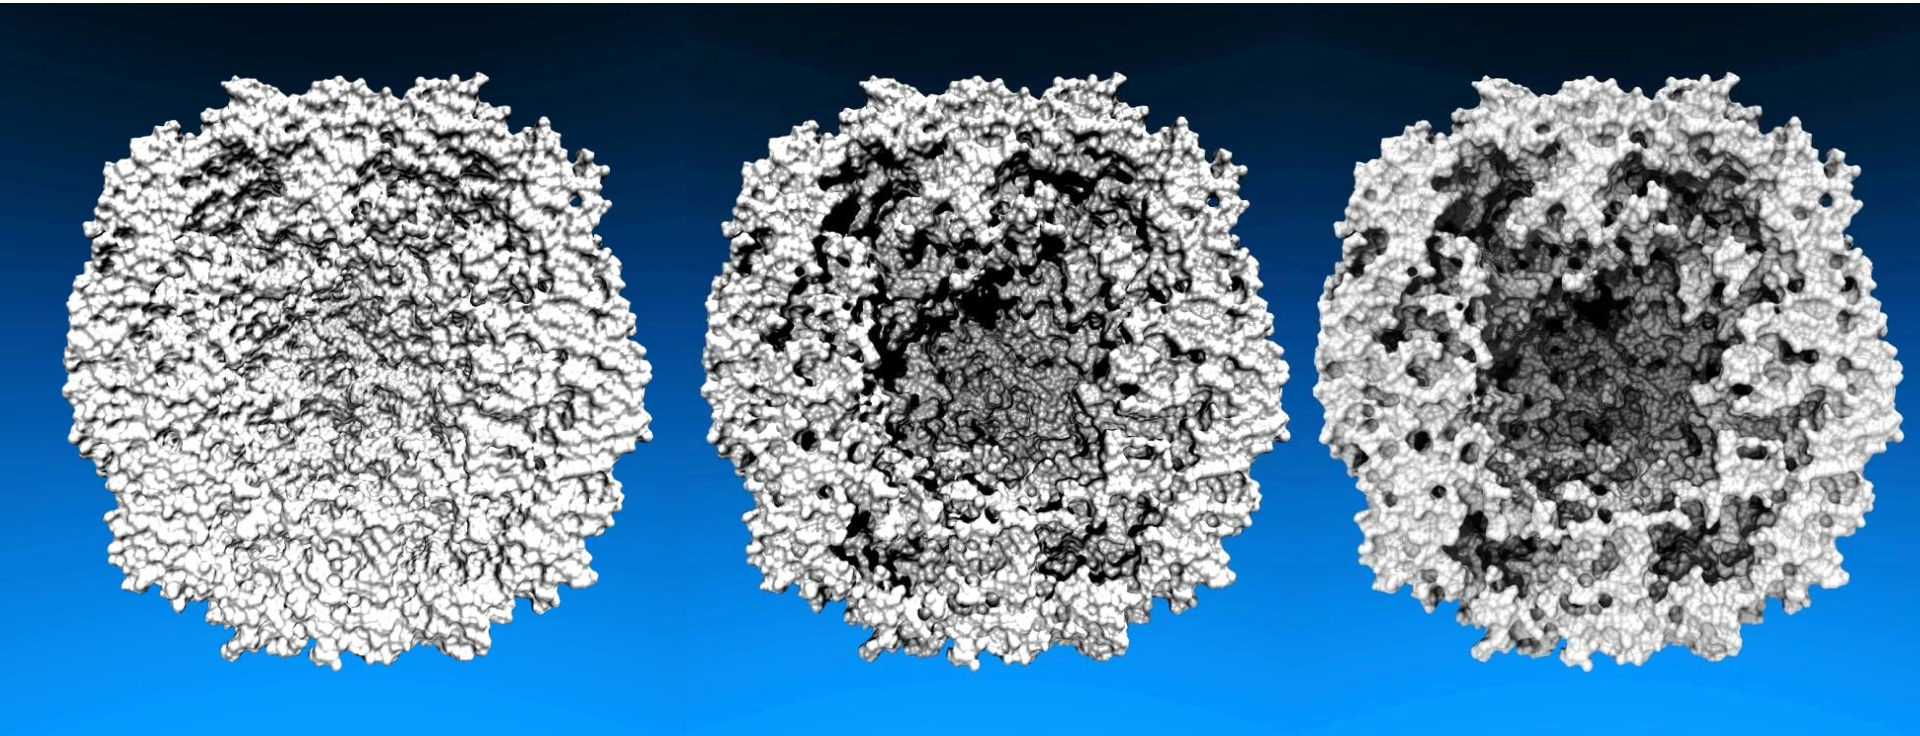

VMD “QuickSurf” Representation, Ray Tracing

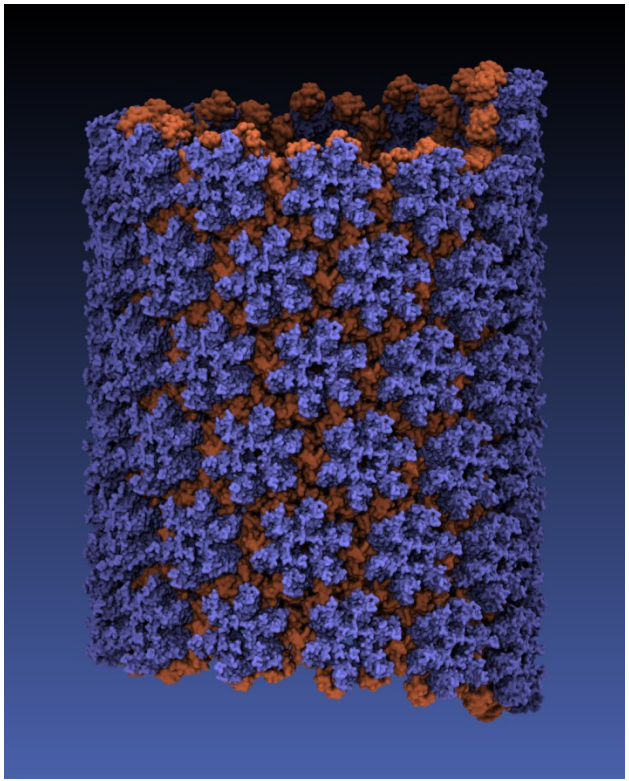

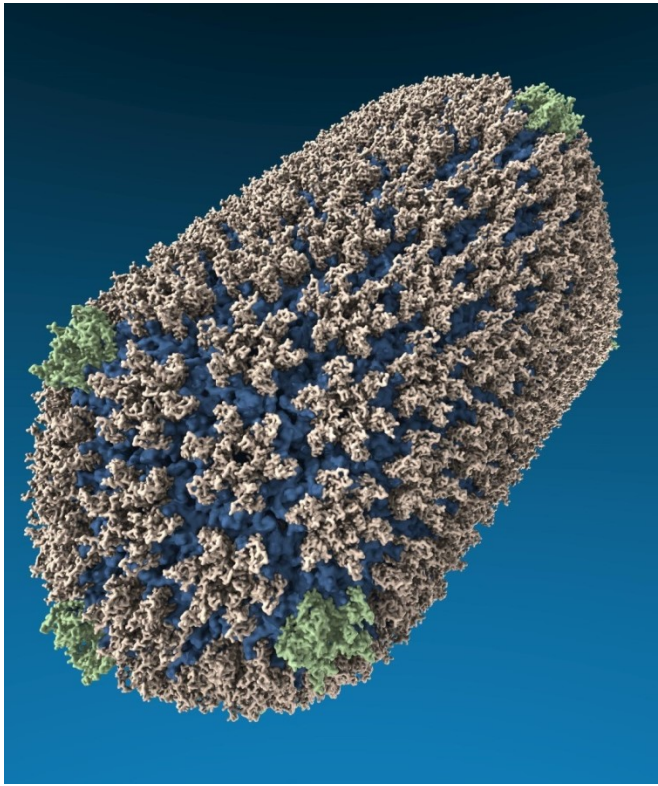

All-atom HIV capsid simulations w/ up to 64M atoms on Blue Waters

Interactive Ray Tracing of Cells

- High resolution cellular tomograms, **billions of voxels**
- Even isosurface or lattice site graphical representations involve ~100M geometric primitives
- **Quadro GP100 / GV100 GPUs benefit from OptiX, NVLink distribution of molecular scene over multiple GPUs**

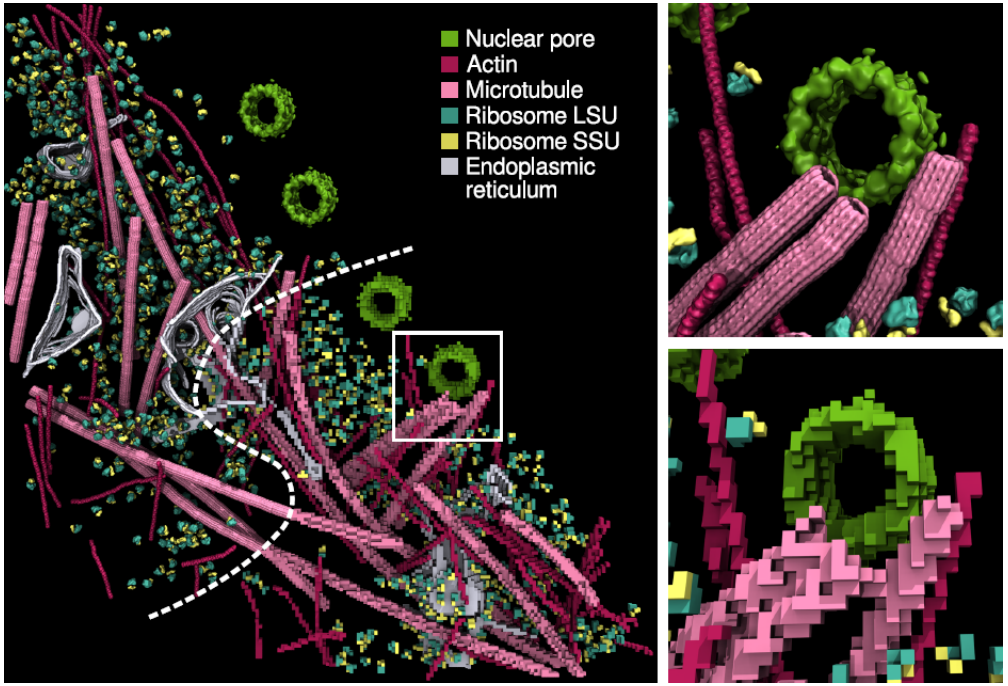

Earnest, et al. J. Physical Chemistry B, 121(15): 3871-3881, 2017.

VMD Atomic Detail Visualization of Cell Architecture w/ Instancing

- VMD 1.9.4 instancing of graphical representations
- Exploit **VBO caching** in OpenGL to eliminate host-GPU geometry transfers
- **RT Instancing** of geometry, minimize memory footprint for cell-scale scenes w/ atomic detail structures

Tomogram showing dense packing of photosynthetic chromatophores (circles) in bacterial cell

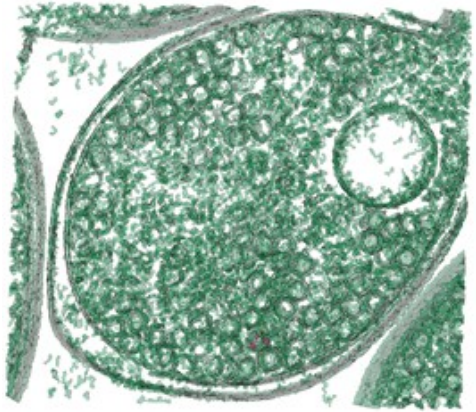

Immersive Viz. w/ VMD

- VMD began as a CAVE app (1993)
- Use of immersive viz by molecular scientists limited due to cost, complexity, lack of local availability, convenience
- Commoditization of HMDs excellent opportunity to overcome cost/availability
- This leaves many challenges still to solve:
  - Incorporate support for **remote visualization**
  - UIs, **multi-user collaboration**/interaction
  - **Rendering perf for large molecular systems**
  - Accomodate limitations idiosyncracies of commercial HMDs

VMD running in a CAVE w/ VR Juggler

Goal: Intuitive interactive viz. in crowded molecular complexes

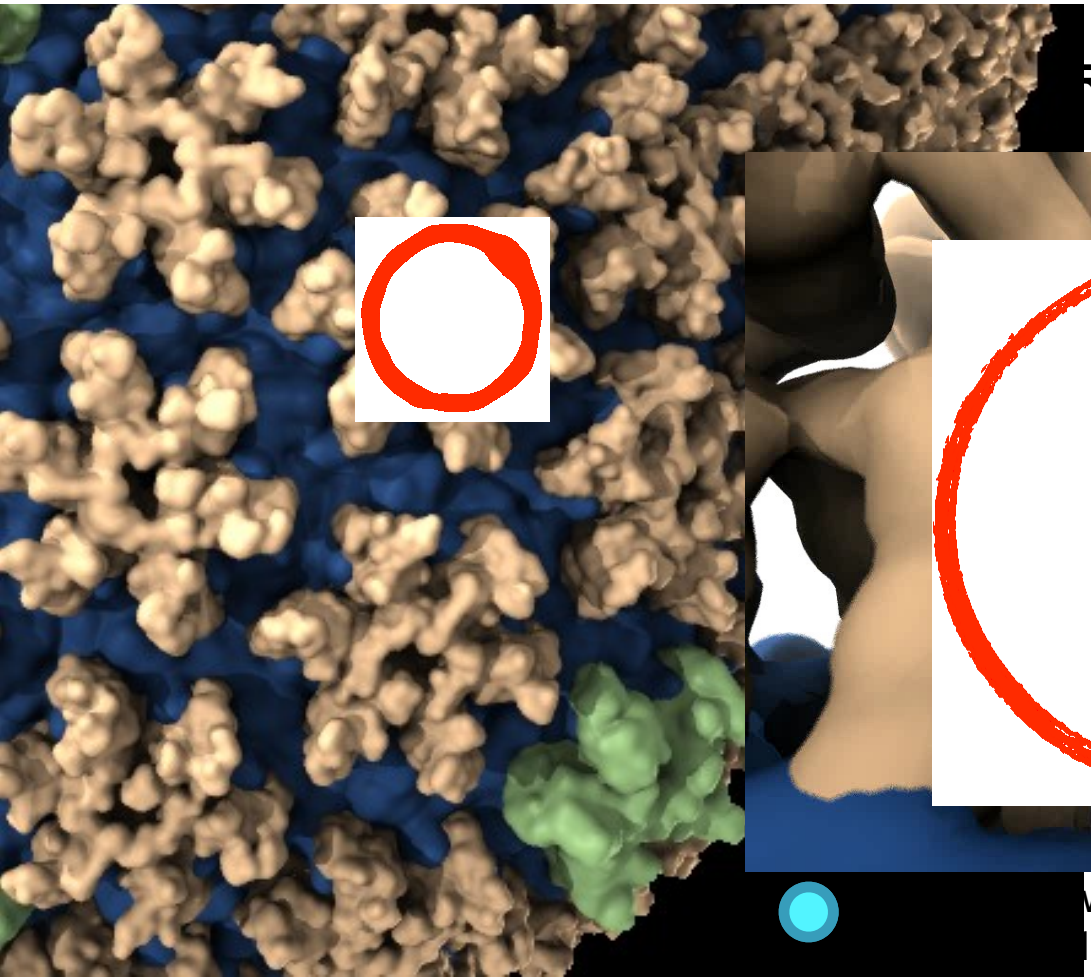

Results from 64 M atom, 1 μs sim!

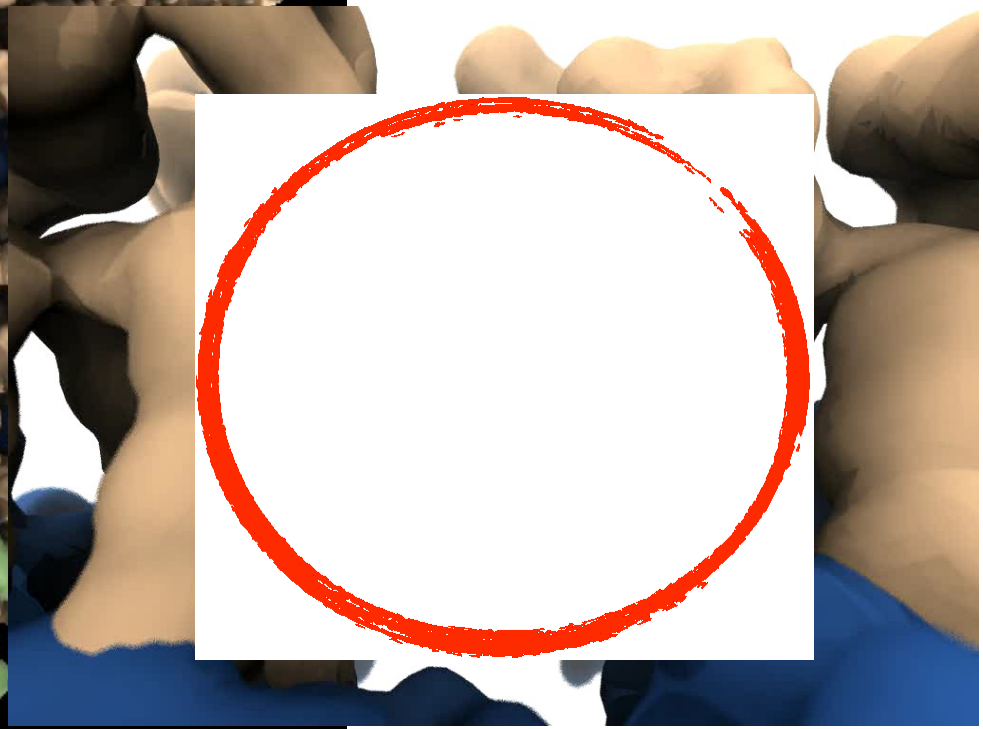

v of chloride ions permeating
capsid hexameric centers

In-Progress VMD VR Development

VMD VR ray tracing:
- Google Cardboard [1]
- Demo w/ Indiana U., SC'15 [2]

Prototype of VR user interaction with VMD models in **room-scale VR** with NVIDIA @ SC'16

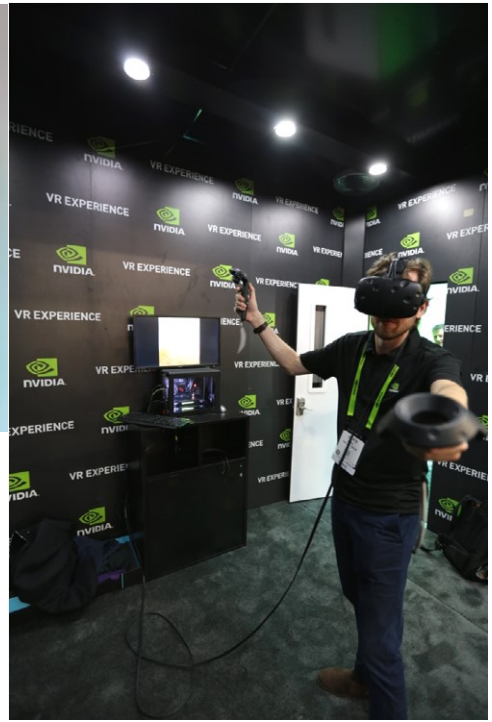

VMD Chromatophore Demo, NVIDIA VR Room at SC'16

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

[2] Immersive Molecular Visualization with Omnidirectional Stereoscopic Ray Tracing and Remote Rendering. J.E. Stone, W.R. Sherman, K. Schulten. IEEE HPDAV (IPDPSW), pp. 1048-1057, 2016.

HMD Ray Tracing Challenges

- HMDs require high frame rates **(90Hz or more)** and minimum latency between IMU sensor reads and presentation on the display
- Multi-GPU workstations fast enough to direct-drive HMDs at required frame rates for simple scenes with direct lighting, hard shadows
- Advanced RT effects such as AO lighting, depth of field require much **larger sample counts**, impractical for direct-driving HMDs
- **Remote viz. required for** many HPC problems due to **large data**
- **Remote viz. latencies too high for direct-drive of HMD**
- **Our two-phase approach: moderate-FPS remote RT combined with local high-FPS view-dependent HMD reprojection w/ OpenGL**

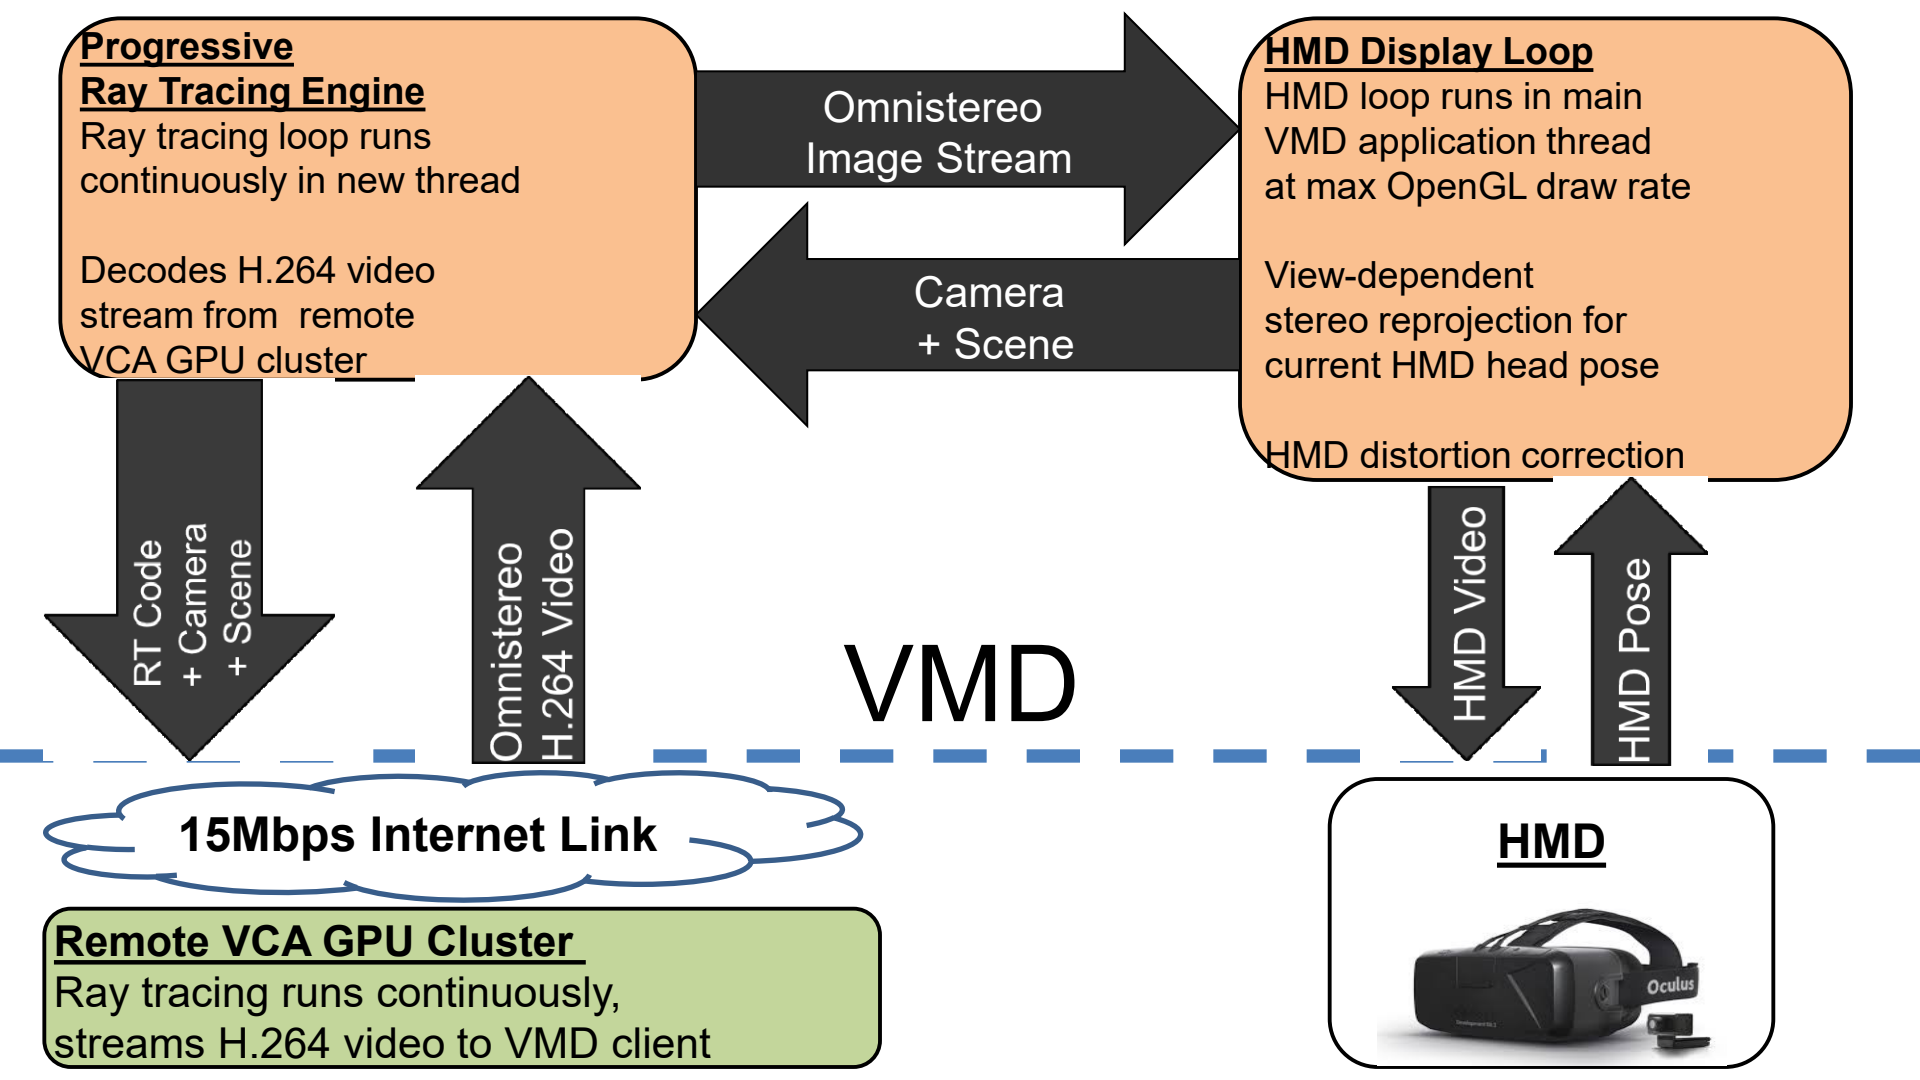
Progressive
Ray Tracing Engine
Ray tracing loop runs continuously in new thread
Decodes H.264 video stream from remote VCA GPU cluster
Omnistereo Image Stream
Camera + Scene
HMD Display Loop
HMD loop runs in main VMD application thread at max OpenGL draw rate
View-dependent stereo reprojection for current HMD head pose
HMD distortion correction
RT Code + Camera + Scene
Omnistereo H.264 Video
HMD Video
HMD Pose
VMD
15Mbps Internet Link
Remote VCA GPU Cluster
Ray tracing runs continuously, streams H.264 video to VMD client
HMD
Oculus

Stereoscopic Panorama Ray Tracing w/ OptiX

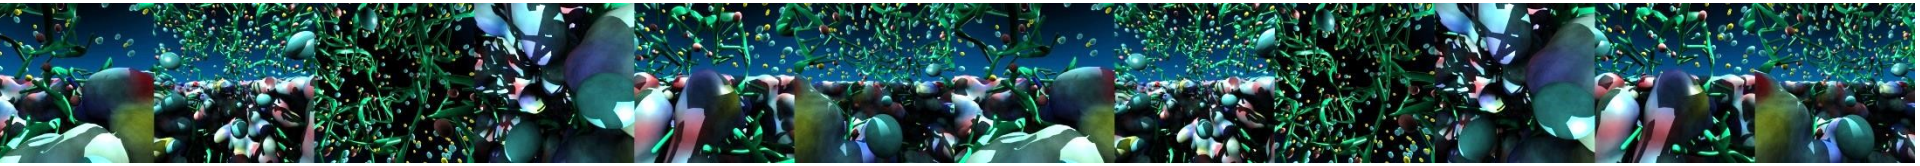

- **Render 360° images and movies for VR headsets such as Oculus Rift, Google Cardboard**
- Ray trace panoramic stereo spheremaps or cubemaps for very high-frame-rate display via OpenGL texturing onto simple geometry
- Stereo requires spherical camera projections **poorly suited to rasterization**
- Benefits from OptiX multi-GPU rendering and load balancing, **remote visualization**

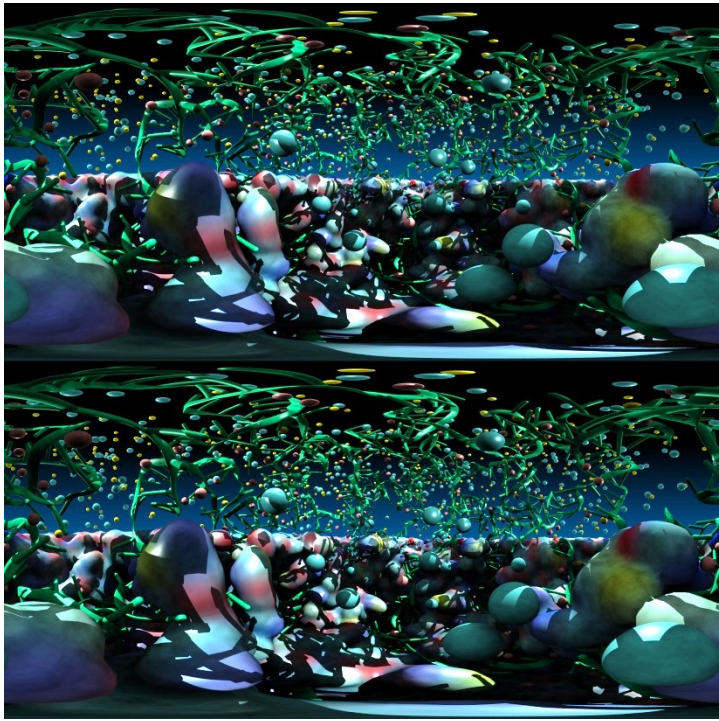

Immersive Molecular Visualization with Omnidirectional Stereoscopic Ray Tracing and Remote Rendering. J. E. Stone, W. R. Sherman, and K. Schulten. High Performance Data Analysis and Visualization Workshop, IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW), pp. 1048-1057,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

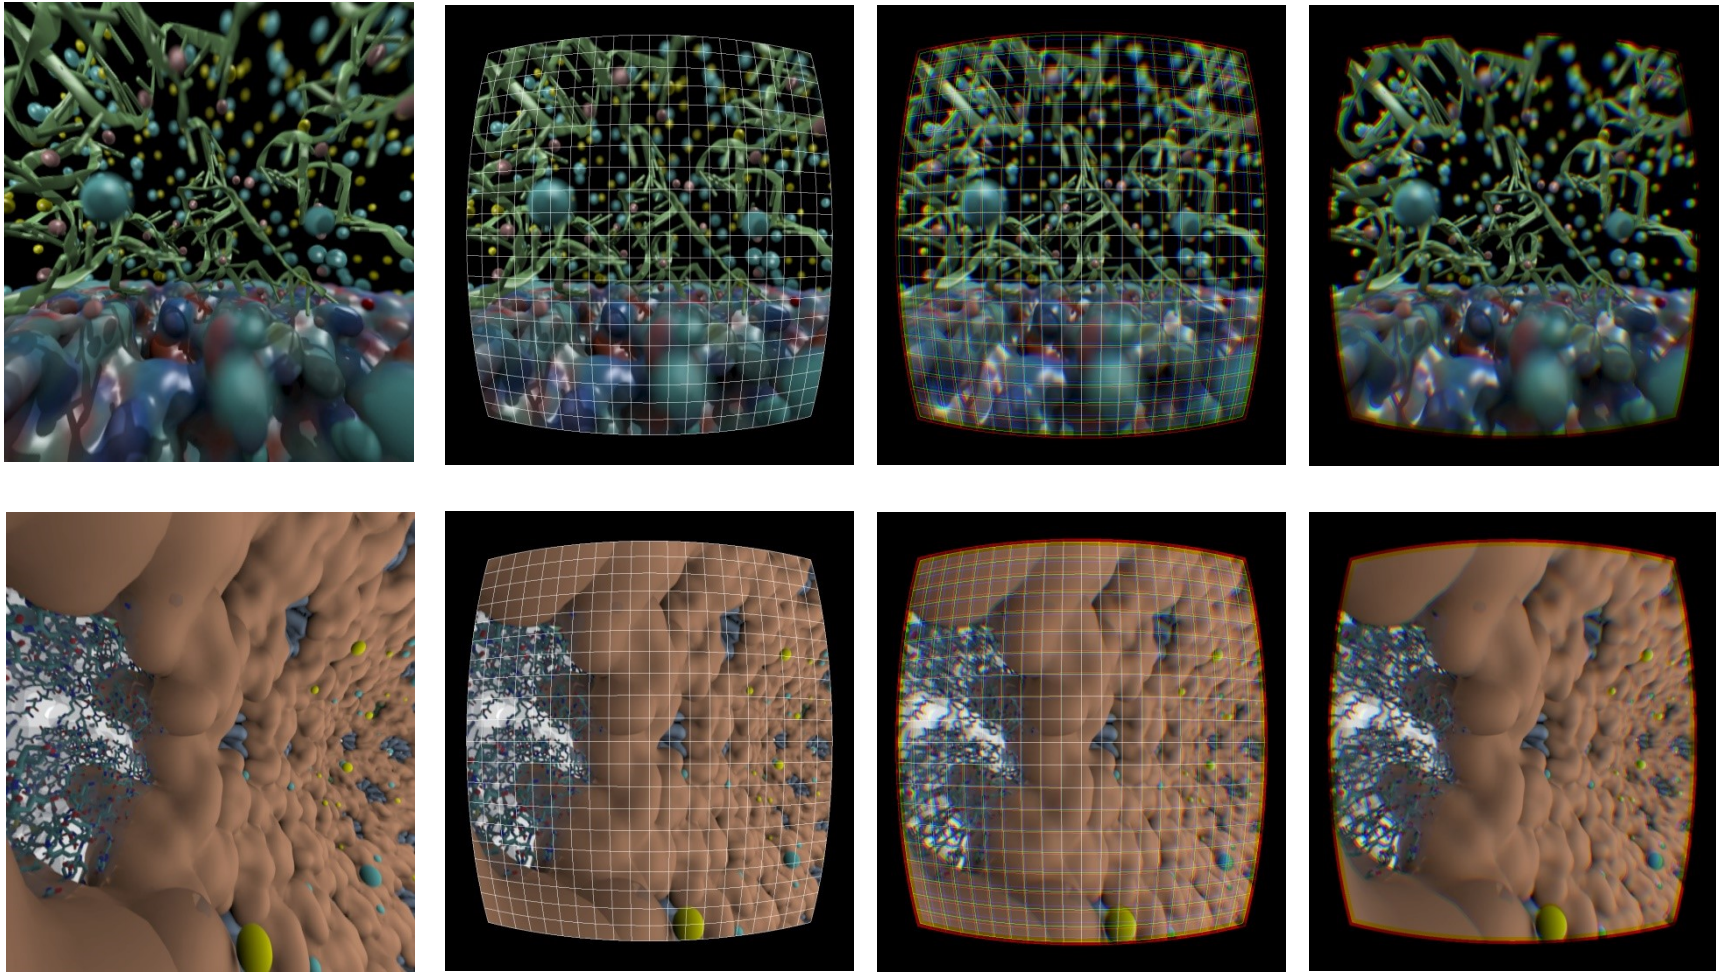

Ongoing VR Work

- OpenXR – cross platform muti-vendor HMD support
- Ray tracing engine and optimizations:
  - **AI denoising for better average quality**
  - Interactive RT stochastic sampling strategies to improve interactivity
  - Improved omnidirectional cubemap/spheremap sampling approaches
  - **AI multi-view warping to allow rapid in-between view generation amid multiple HMD head locations**
  - **H.265 for high-res omnidirectional video streaming**
  - **Multi-node parallel RT and remote viz. on general clusters and supercomputers, e.g. NCSA Blue Waters, ORNL Titan**
- Tons of work to do on VR user interfaces, multi-user collaborative visualization, …

Acknowledgements

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

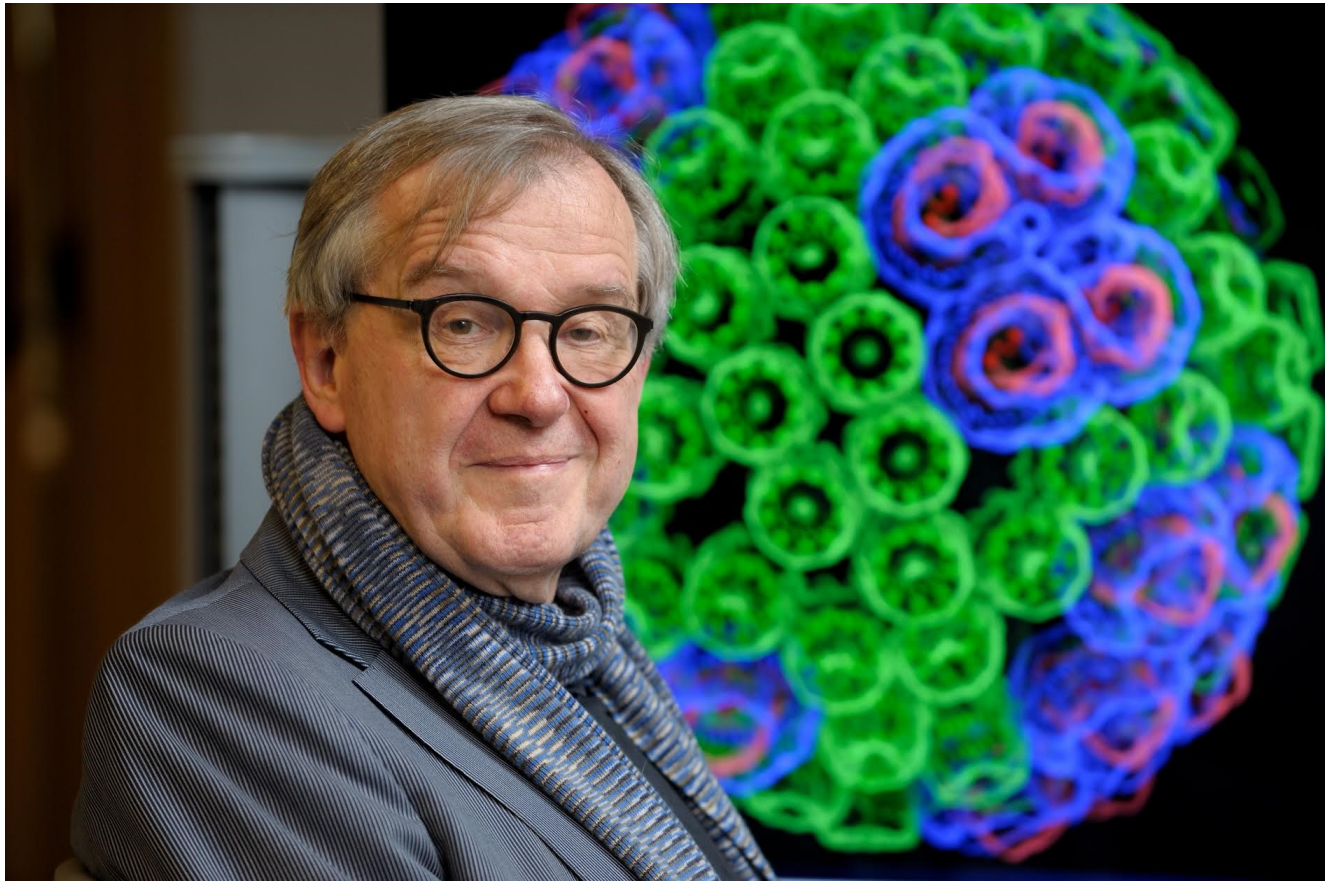

“When I was a young man, my goal was to look with mathematical and computational means at the inside of cells, one atom at a time, to decipher how living systems work. That is what I strived for and I never deflected from this goal.” – Klaus Schulten

Related Publications

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

- **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. (**In press)**
- **Challenges of Integrating Stochastic Dynamics and Cryo-electron Tomograms in Whole-cell Simulations.** T. M. Earnest, R. Watanabe, J. E. Stone, J. Mahamid, W. Baumeister, E. Villa, and Z. Luthey-Schulten. J. Physical Chemistry B, 121(15): 3871-3881, 2017.
- **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, and K. Schulten. International Workshop on OpenPOWER for HPC (IWOPH'16), LNCS 9945, pp. 188-206, 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 Workshop (IPDPSW), pp. 1048-1057, 2016.
- **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 International Parallel and Distributed Processing Symposium Workshop (IPDPSW), pp. 1014-1023, 2016.
- **Evaluation of Emerging Energy-Efficient Heterogeneous Computing Platforms for Biomolecular and Cellular Simulation Workloads.** J. E. Stone, M. J. Hallock, J. C. Phillips, J. R. Peterson, Z. Luthey-Schulten, and K. Schulten.25th International Heterogeneity in Computing Workshop, IEEE International Parallel and Distributed Processing Symposium Workshop (IPDPSW), pp. 89-100, 2016.

Related Publications

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

- **Atomic Detail Visualization of Photosynthetic Membranes with GPU-Accelerated Ray Tracing.** J. E. Stone, M. Sener, K. L. Vandivort, A. Barragan, A. Singharoy, I. Teo, J. V. Ribeiro, B. Isralewitz, B. Liu, B.-C. Goh, J. C. Phillips, C. MacGregor-Chatwin, M. P. Johnson, L. F. Kourkoutis, C. Neil Hunter, and K. Schulten. J. Parallel Computing, 55:17-27, 2016.
- **Chemical Visualization of Human Pathogens: the Retroviral Capsids.** Juan R. Perilla, Boon Chong Goh, John E. Stone, and Klaus Schulten. SC'15 Visualization and Data Analytics Showcase, 2015.
- **Visualization of Energy Conversion Processes in a Light Harvesting Organelle at Atomic Detail.** M. Sener, J. E. Stone, A. Barragan, A. Singharoy, I. Teo, K. L. Vandivort, B. Isralewitz, B. Liu, B. Goh, J. C. Phillips, L. F. Kourkoutis, C. N. Hunter, and K. Schulten. SC'14 Visualization and Data Analytics Showcase, 2014.
  ***Winner of the SC'14 Visualization and Data Analytics Showcase**
- **Runtime and Architecture Support for Efficient Data Exchange in Multi-Accelerator Applications.** J. Cabezas, I. Gelado, J. E. Stone, N. Navarro, D. B. Kirk, and W. Hwu. IEEE Transactions on Parallel and Distributed Systems, 26(5):1405-1418, 2015.
- **Unlocking the Full Potential of the Cray XK7 Accelerator.** M. D. Klein and J. E. Stone. Cray Users Group, Lugano Switzerland, May 2014.
- **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 Discussions, 169:265-283, 2014.
- **Simulation of reaction diffusion processes over biologically relevant size and time scales using multi-GPU workstations**. M. J. Hallock, J. E. Stone, E. Roberts, C. Fry, and Z. Luthey-Schulten. Journal of Parallel Computing, 40:86-99, 2014.

Related Publications

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

- **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. 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. Stone, and Z. Luthey-Schulten. J. Computational Chemistry 34 (3), 245-255, 2013**.**
- **Fast Visualization of Gaussian Density Surfaces for Molecular Dynamics and Particle System Trajectories.** M. Krone, J. Stone, T. Ertl, and K. Schulten. *EuroVis Short Papers,* pp. 67-71, 2012.
- **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.
- **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.

Related Publications

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

- **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.
- **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 15th International Conference on Architectural Support for Programming Languages and Operating Systems,* pp. 347-358, 2010.

Related Publications

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

- **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. E. Stone, J. Saam, D. Hardy, K. Vandivort, W. Hwu, K. Schulten, *2nd Workshop on General-Purpose Computation on Graphics Pricessing 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.

Related 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.

Related Publications

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

- **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. E. Stone, J. Saam, D. Hardy, K. Vandivort, W. Hwu, K. Schulten, *2nd Workshop on General-Purpose Computation on Graphics Pricessing 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.

Related 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.