Preparing and Analyzing Large Molecular Simulations with VMD

John Stone, Senior Research Programmer

NIH Center for Macromolecular Modeling and Bioinformatics,

University of Illinois at Urbana-Champaign

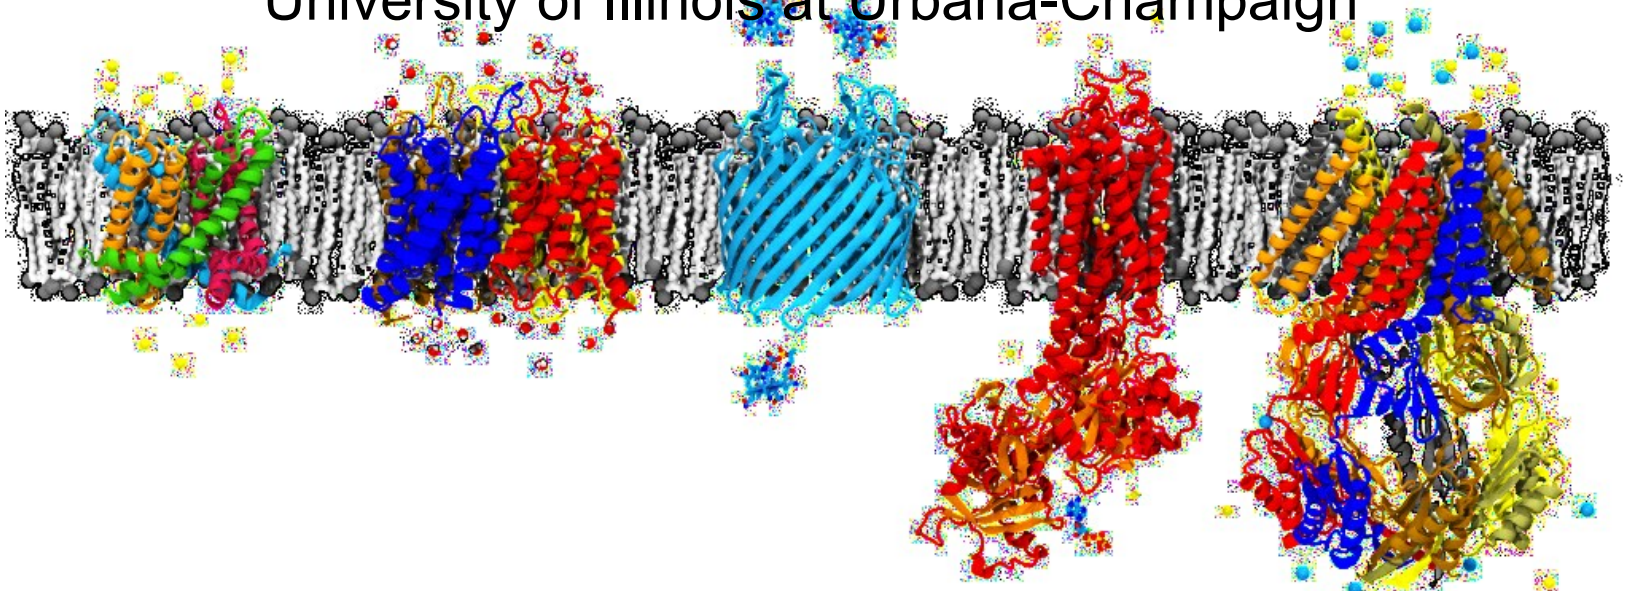

VMD Tutorials Home Page

- http://www.ks.uiuc.edu/Training/Tutorials/
  - Main VMD tutorial
  - VMD images and movies tutorial
  - QwikMD simulation preparation and analysis plugin
  - Structure check
  - VMD quantum chemistry visualization tutorial
  - Visualization and analysis of CPMD data with VMD
  - Parameterizing small molecules using ffTK

Overview

- Introduction
- Data model
- Visualization
- Scripting and analytical features
- Trajectory analysis and visualization
- High fidelity ray tracing
- Plugins and user-extensibility
- Large system analysis and visualization

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/

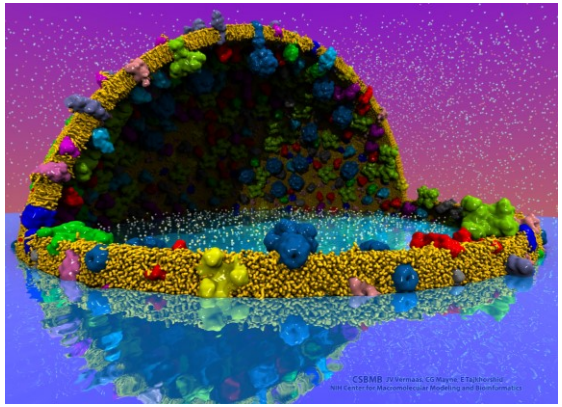

Cell-Scale Modeling

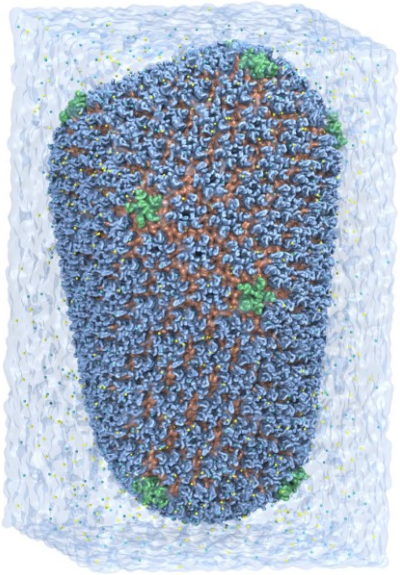

MD Simulation

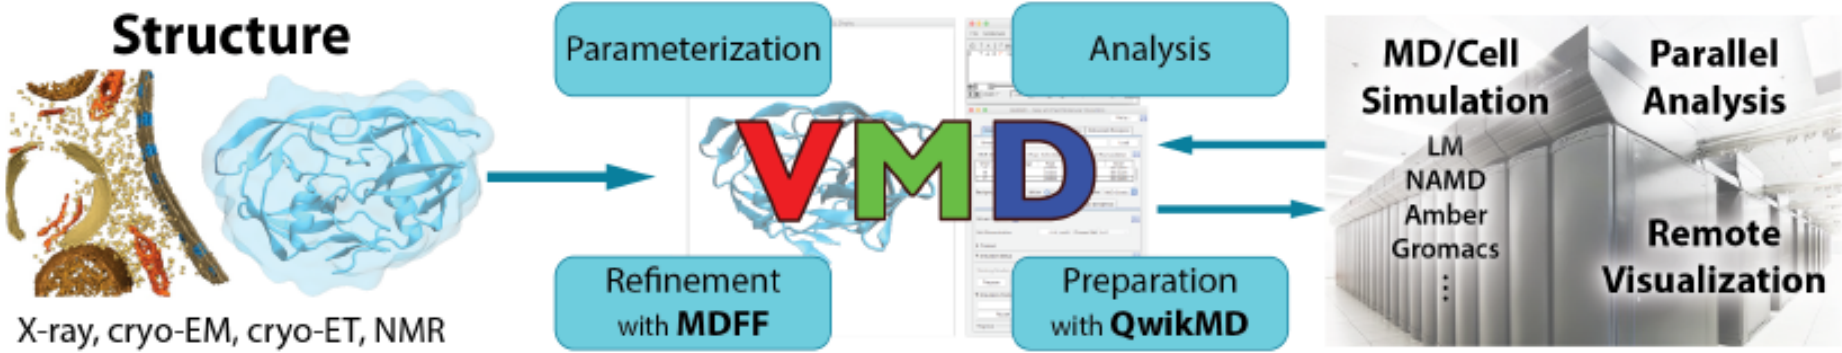

VMD – “Visual Molecular Dynamics”

- Unique capabilities:
  - Trajectories are fundamental to VMD
  - Support for very large systems, now reaching billions of particles
  - Extensive GPU acceleration
  - Parallel analysis/visualization with MPI
- Visualization and analysis of:
  - Molecular dynamics simulations
  - “Particle” systems and whole cells
  - Cryo-EM densities, volumetric data
  - Quantum chemistry calculations
  - Sequence information

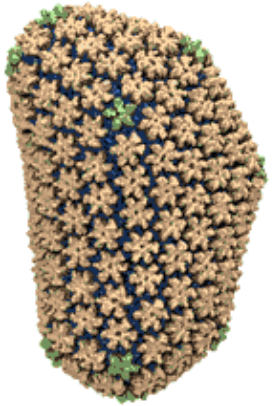

MD Simulations

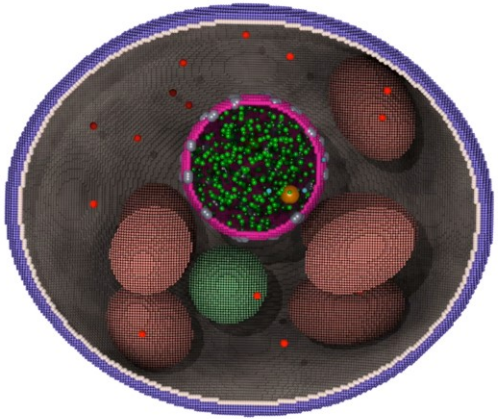

Cell-Scale Simulation

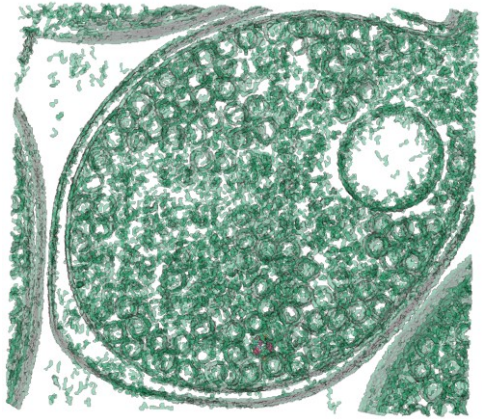

Integrate w/ Cryo-EM/ET

VMD Serves Many Communities

- VMD user statistics:
  - 103,000 unique registered users; 17,000 (16%) are NIH funded researchers
  - 25,000 citations, over 3,000 citations per year
- Supports key data types, file formats, and databases
- **User extensible** to support new tools, data types, **custom analyses**

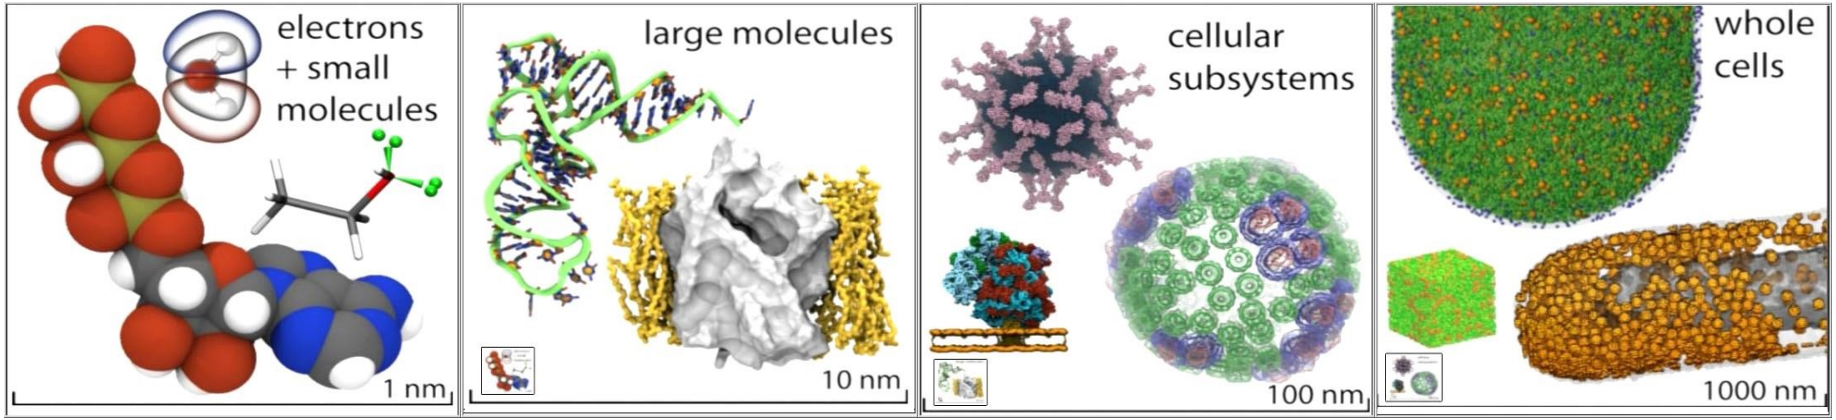

Goal: A Computational Microscope

Study the molecular machines in living cells

Ribosome: target for antibiotics

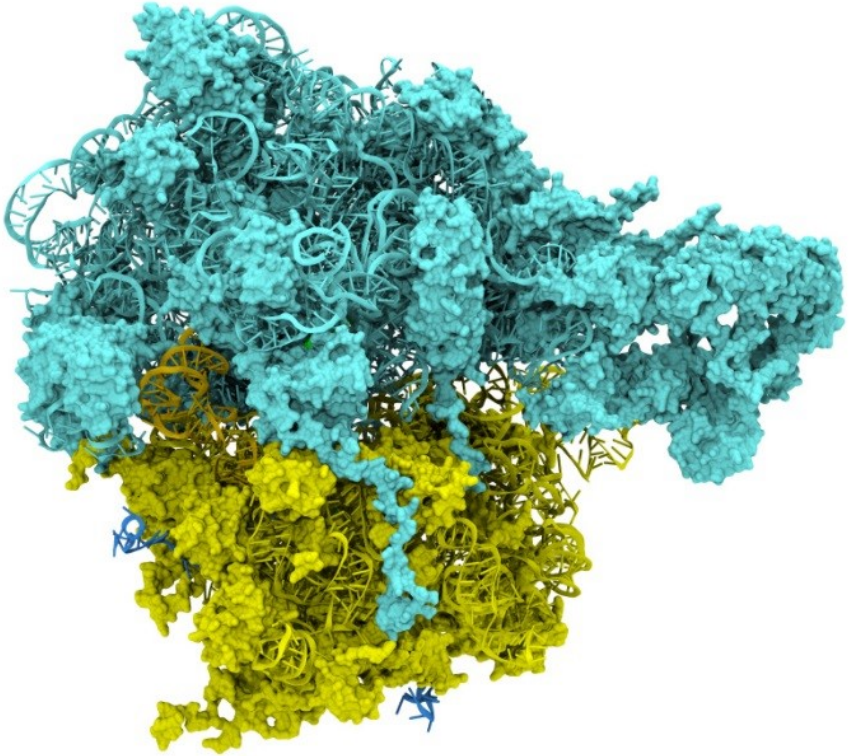

Poliovirus

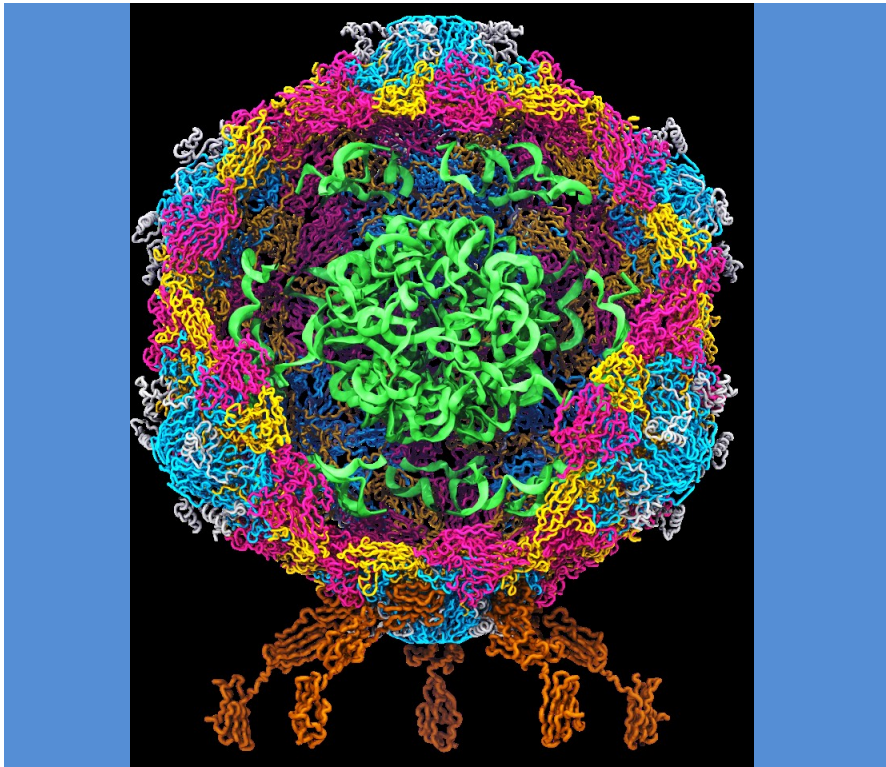

VMD is a Tool for Accessing Information about the Cell

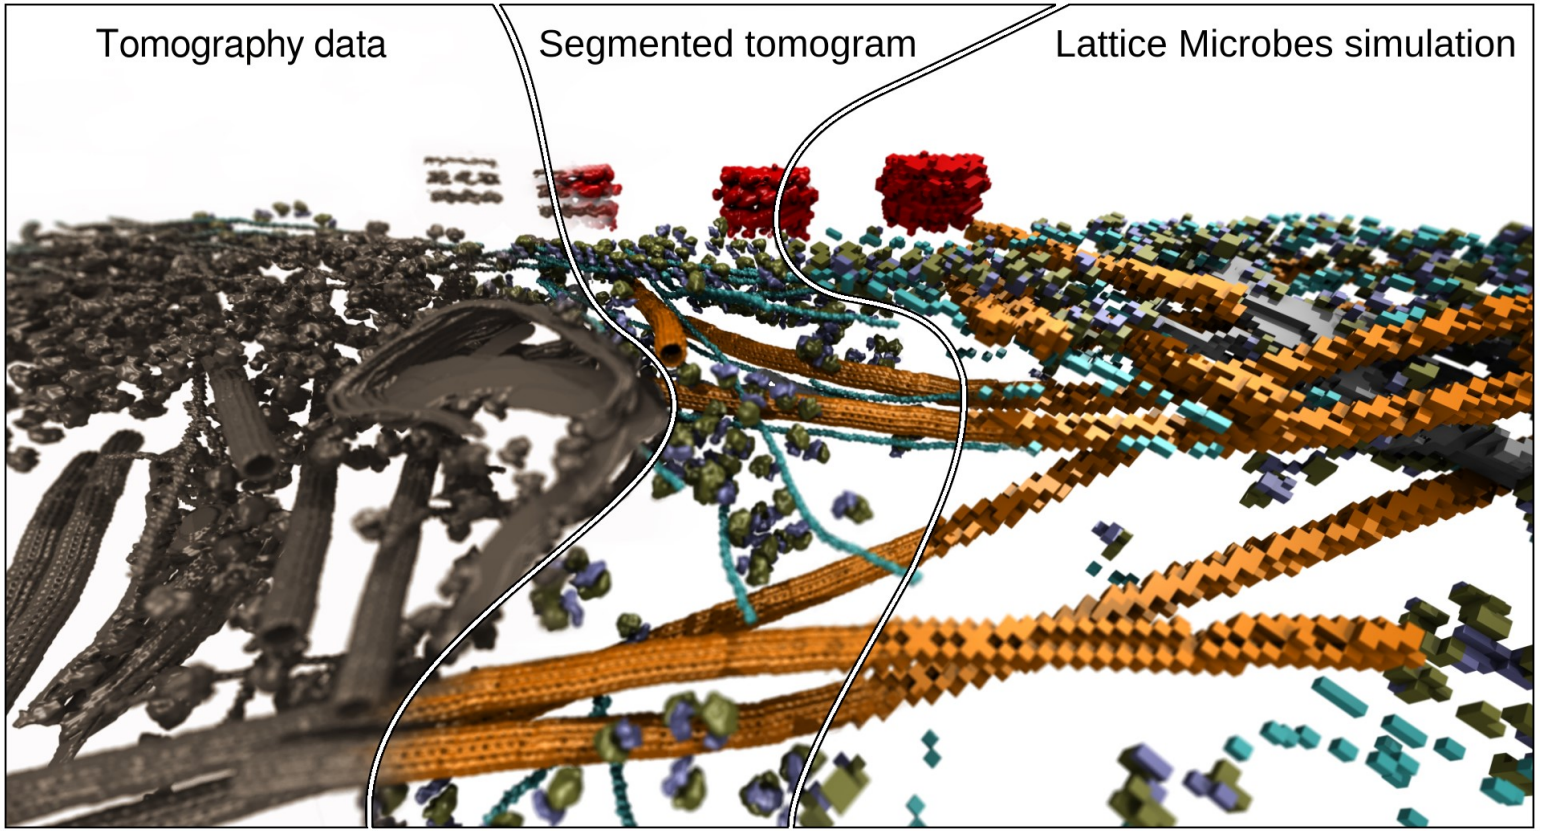

VMD Interoperates with Mainstream Research Tools

- Provides tools for simulation preparation, visualization, and analysis
- Interpret and process multi-modal structural information
- Connects with key software tools to enable state-of-the-art simulations
- Openness, extensibility, and interoperability are VMD hallmarks
- Uses advanced algorithms and hardware technologies to address data size challenges posed by cutting-edge experimental imaging and simulation

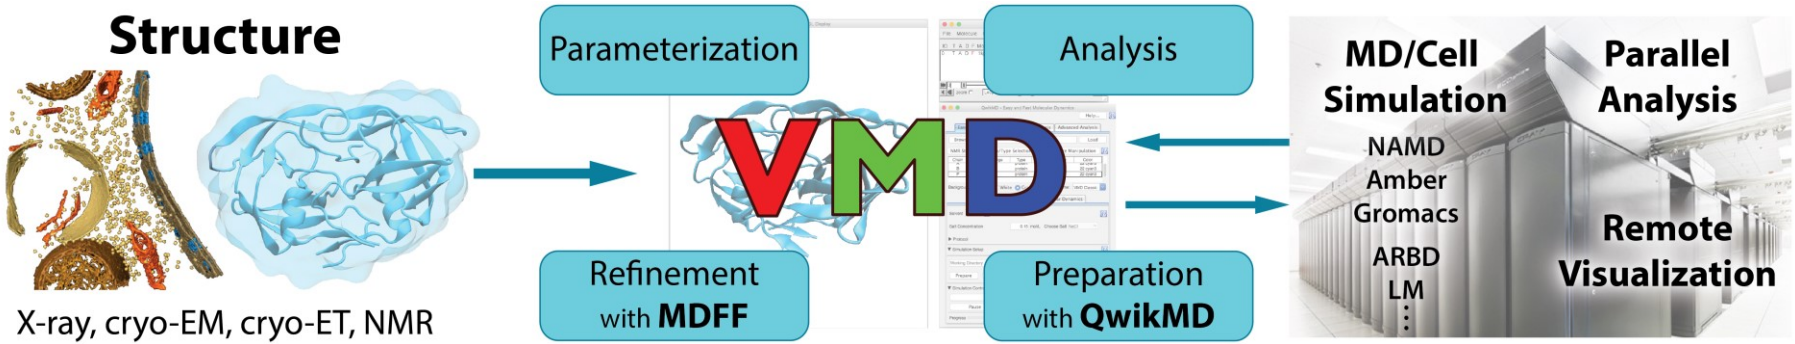

VMD is a Platform for Developing Research Tools
Over 110 VMD Plugins, Half Developed by Users

- VMD user-extensible scripting w/ Tcl/Tk, Python
- User-developed plugins:
  - Alanine Scanning
  - Collective Variable Analyzer
  - Clustering Tool
  - Carbon Nanostructure Builder
  - TorsionPlot
  - RMSD Trajectory Tool
  - Many others…

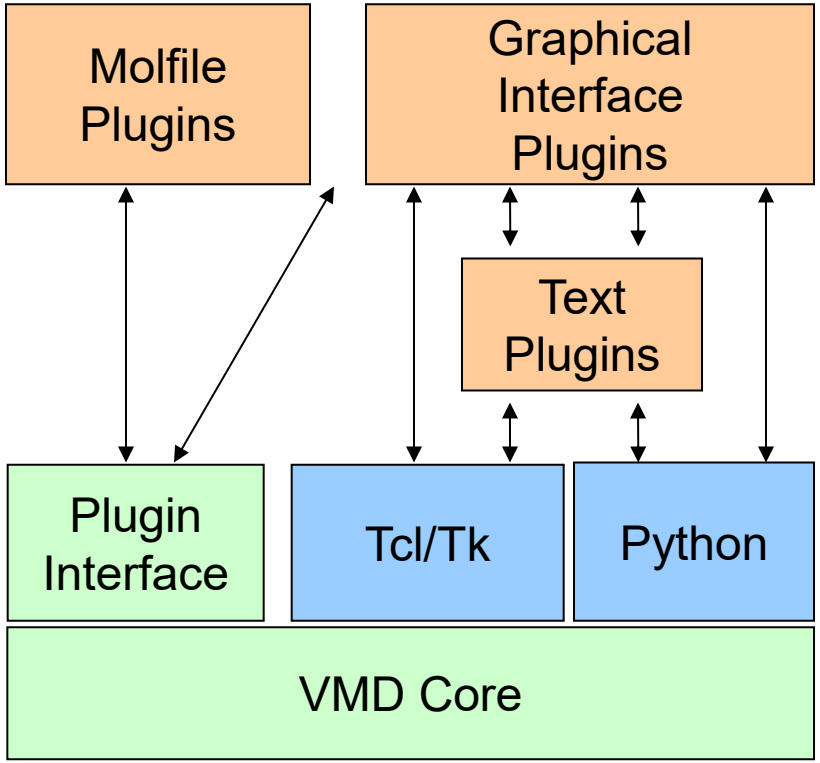

Selected VMD Plugins: Center Developed, and User Developed

Analysis

- APBSRun
- CatDCD
- Contact Map
- GofRGUI
- HeatMapper
- ILSTools
- IRSpecGUI
- MultiSeq
- NAMD Energy
- NAMD Plot
- NetworkView
- NMWiz
- ParseFEP
- PBCTools
- PMEpot
- PropKa GUI
- RamaPlot
- RMSD Tool
- RMSD Trajectory Tool
- RMSD Visualizer Tool
- Salt Bridges
- Sequence Viewer
- Symmetry Tool
- Timeline
- TorsionPlot
- VolMap

Modeling

- AutoIonize
- AutoPSF
- Chirality
- Cionize
- Cispeptide
- CGTools
- Dowser
- ffTK
- Inorganic Builder
- MDFF
- Membrane
- Merge Structs
- Molefacture
- Mutator
- Nanotube
- Psfgen
- RESPTool
- RNAView
- Solvate
- SSRestraints
- Topotools

Visualization

- Clipping Plane Tool
- Clone Rep
- DemoMaster
- Dipole Watcher
- Intersurf
- Navigate
- NavFly
- MultiMolAnim
- Color Scale Bar
- Remote
- Palette Tool
- ViewChangeRender
- ViewMaster
- Virtual DNA Viewer
- VMD Movie Maker

Simulation

- AlaScan
- AutoIMD
- IMDMenu
- NAMD GUI
- NAMD Server
- QMTool

Collaboration

- Remote Control

Data Import and Plotting

- Data Import
- Multiplot
- PDBTool
- MultiText

Externally Hosted Plugins and Extensions

- Check sidechains
- MultiMSMS
- Interactive Essential Dynamics
- Mead Ionize
- Clustering Tool
- iTrajComp
- Swap RMSD
- Intervor
- SurfVol
- vmdICE

75 MolFile I/O Plugins:
structure, trajectory, sequence, and density map

http://www.ks.uiuc.edu/Research/vmd/plugins/

Goal: A Computational Microscope

Study the molecular machines in living cells

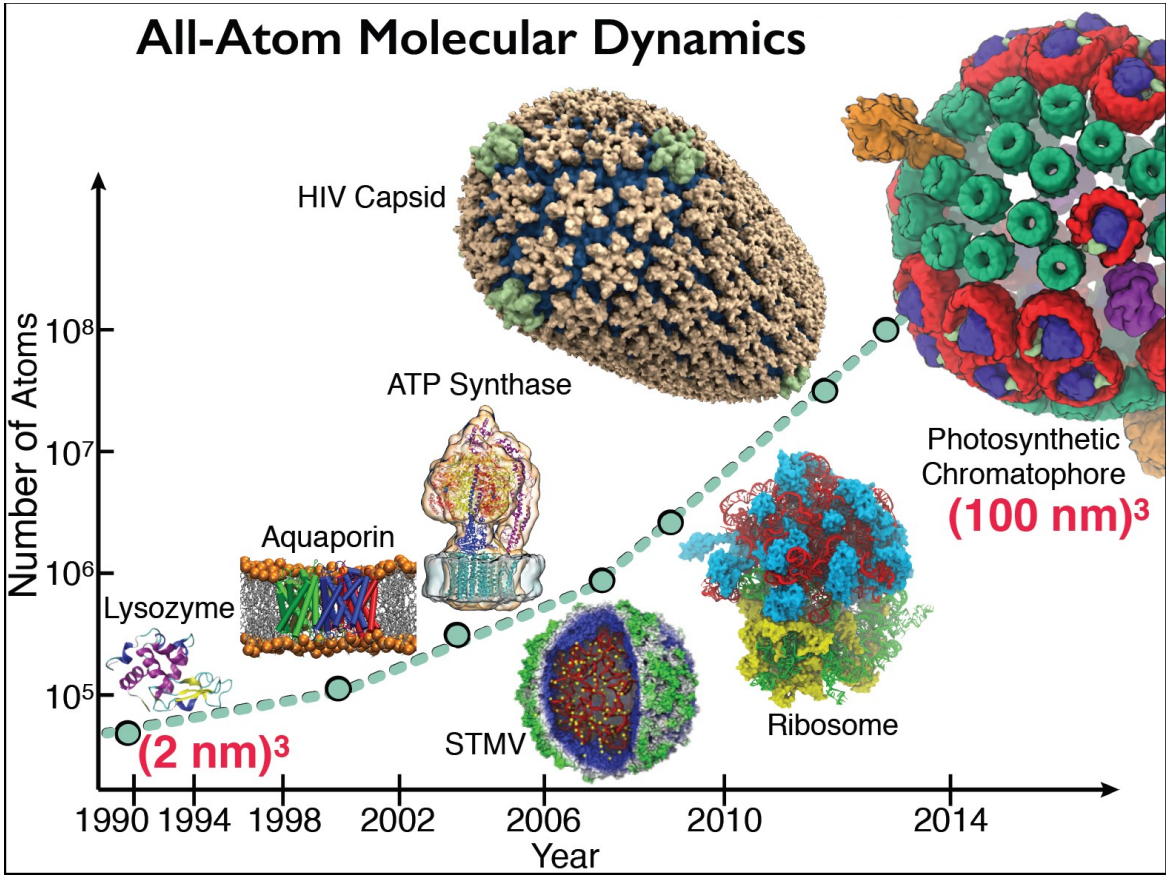

Technology Opportunities and Collaborations

- Supercomputer Centers, Cray, IBM
  - Remote visualization
  - Performance, power profiling and optimization
- NVIDIA
  - **GPU computing**
  - **Ray tracing**
  - **Remote visualization**
  - ARM, Tablets, power profiling and optimization
- Intel
  - **x86, Xeon Phi optimization, ray tracing**
- Amazon
  - Cloud deployment of VMD/NAMD, related tools
  - Remote visualization
- Universities:
  - **G. Fiorin, J. Henin, Toni Giorgino, collective variables**
  - **T. Ertl, U. Stuttgart: visualization algorithms**
  - M. Kuttel, U. Cape Town: visualization and analysis
  - W. Sherman, Indiana U.: VR HMDs, visualization

GPU computing,
Ray tracing,
Remote viz.

VR HMDs, 6DoF input devices

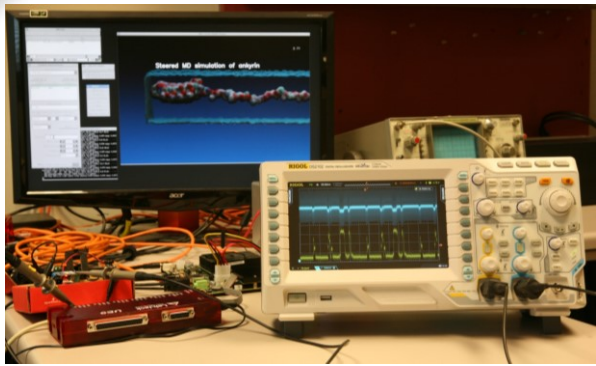

Energy efficiency: ARM+GPU

Making Our Research Tools Easily Accessible

- Cloud based deployment
  - Full virtual machines (known as “AMI” in Amazon terminology)
  - Amazon AWS EC2 GPU-accelerated instances:
    http://www.ks.uiuc.edu/Research/cloud/

- Docker “container” images available in NVIDIA NGC registry
  - Users obtain Docker images via registry, download and run on the laptop, workstation, cloud, or supercomputer of their choosing
  - **https://ngc.nvidia.com/registry/**
  - **https://ngc.nvidia.com/registry/hpc-vmd**

Clusters, Supercomputers

Workstations,
Servers,
Cloud

Molecular dynamics-based refinement and validation for sub-5 Å cryo-electron microscopy maps. Abhishek Singharoy, Ivan Teo, Ryan McGreevy, John E. Stone, Jianhua Zhao, and Klaus Schulten. *eLife*, 10.7554/eLife.16105, 2016. (66 pages).
QwikMD-integrative molecular dynamics toolkit for novices and experts. Joao V. Ribeiro, Rafael C. Bernardi, Till Rudack, John E. Stone, James C. Phillips, Peter L. Freddolino, and Klaus Schulten. *Scientific Reports*, 6:26536, 2016.
High performance molecular visualization: In-situ and parallel rendering with EGL. John E. Stone, Peter Messmer, Robert Sisneros, and Klaus Schulten. *2016 IEEE International Parallel and Distributed Processing Symposium Workshop (IPDPSW)*, pp. 1014-1023, 2016.

VMD Data Model

VMD “Molecule”

- Collection of self-consistent molecular information
  - Atomic structure, time-varying atomic coordinates, velocities, …
  - Volumetric data
  - Graphics objects
- Key missing information required for visualization is auto-generated heuristically – “guessed”, if not provided in input files
  - Bond topology
- Limitations:
  - Fixed atom count per-timestep, requires use of “dummy atoms” for so-called “open system” simulations
  - VMD doesn’t store or use MD force field parameters, various other data that are typically specific to particular simulation tools, force fields, etc.

Overcoming the Biomolecular Orientation of VMD

- VMD molecular data structures are optimized for biomolecules
  - Expects atom/residue names that (roughly) follow RCSB Protein Databank conventions
  - Significant up-front structure analysis enables increased interactivity later
  - Cost of this is decreased generality in various cases:
    - Fixed atom count per timestep
    - Unrecognized atom naming conventions can inhibit use of some of the high-level atom selections
- Nanodevice and materials science data can be displayed in VMD
  - Assign per-atom name/type fields after loading
  - Enable VMD to better recognize structural elements
  - Example: xyz or LAMMPS dump files lack some expected fields, so it helps to use special features of LAMMPS plugin to assign them

Atomic Structure

- Per-Atom attributes:
  - Name
  - Type
  - Residue ID, Unique residue ID, Residue Name, Residue Type
  - Bond list
  - Atomic number
  - Alternate location identifier
  - Insertion code
  - Chain
  - Segment name
  - Atom/Residue classification (protein, nucleic, hydrogen, water, “other”)

Other, Optional Per-Atom Fields

- Optional per-atom attributes:
  - Mass
  - Charge
  - Radius
  - Occupancy
  - B-factor
  - Fields primarily used by simulation preparation tools:
    - Bond Orders
    - Angles, Dihedrals, Impropers, Cross-Term Maps

Time-Varying Trajectory Data

- Atomic coordinates
- Atomic velocities
- PBC unit cell
- QM, QM/MM molecular orbital data
- User-defined time-varying scalar quantities:
  - user, user2, user3, user4

VMD LAMMPS Molfile Plugin

https://sites.google.com/site/akohlmey/software/lammps-plugin

- Reads LAMMPS text mode trajectory “dump” files
- Map LAMMPS per-atom fields to VMD molfile API:
  - LAMMPSREMAPFIELDS environment variable allows user to control mapping of per-atom fields to VMD molfile API
  - List of mappings of LAMMPS fields into VMD fields, e.g. “VMD=LAMMPS”
  - Example to map LAMMPS forces into VMD velocity fields:

    set env(LAMMPSREMAPFIELDS) “vx=fx,vy=fy,vz=fz”
- Adapt LAMMPS trajectories with varying atom counts to VMD
  - LAMMPSMAXATOMS environment variable sets peak atom count for VMD, adding “dummy” atoms as placeholders:

    Dummy atoms initialized with properties: name = @, type = X, resname = @@@, segid = @@@, chain = @, mass =0, charge=0, radius=0, element=X.
  - LAMMPSDUMMYPOS sets default atomic coords for dummy atoms that don’t exist for a given frame:

    set env(LAMMPSDUMMYPOS) {0.0 0.0 -10.0}

Exporting Molecular Data

- VMD can be used to prepare MD simulations in conjunction with plugins and user scripting
- Need to watch out for assumptions made by heuristics, structure building tools
- Example: bond topology may need to be modified or deleted for some force fields
- Use the most appropriate structure manipulation tools for the MD package in use, e.g. for LAMMPS, the best choice in VMD is the “topotools” plugin

VMD Visualization Concepts

VMD Approach to Visualization

- Molecular scene is composed of “graphical representations”
- Each representation encapsulates a group of selected atoms, a drawing style, coloring style, and other parameters
- Representations are independent of each other, can be toggled on/off easily, allowing molecular scenes to be built-up incrementally
- VMD atom selection language is shared with its analytical and scripting interfaces

Selection, Filtering

- Most viz tools allow interactive visual picking, menu-driven selections of structure components to display or operate on
- VMD also extensively uses a text-based selection language (think google):
  - “water within 10 of protein and z > 0”
  - Allows selection on user-defined fields
  - ***Promotes synergy between interactive and scripting interfaces***
  - Works very well when dealing with huge time-varying structures

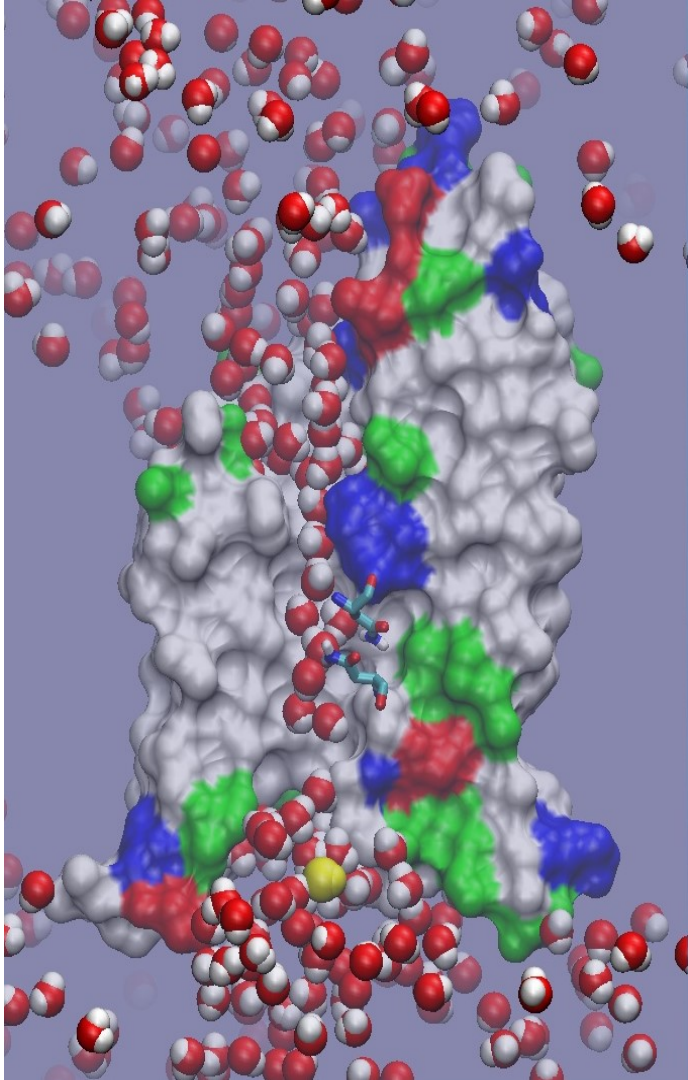

Structure Visualization

Molecular representations provide different levels of abstraction, atomic detail vs. higher level organizational information

- Atoms, VdW spheres, bonds, ball-stick, …
- Coarse-grained “beads”
- Ribbons, secondary structure, “cartoon” reps, RNA/DNA
- Molecular surfaces
- Molecular orbitals (quantum chemistry)

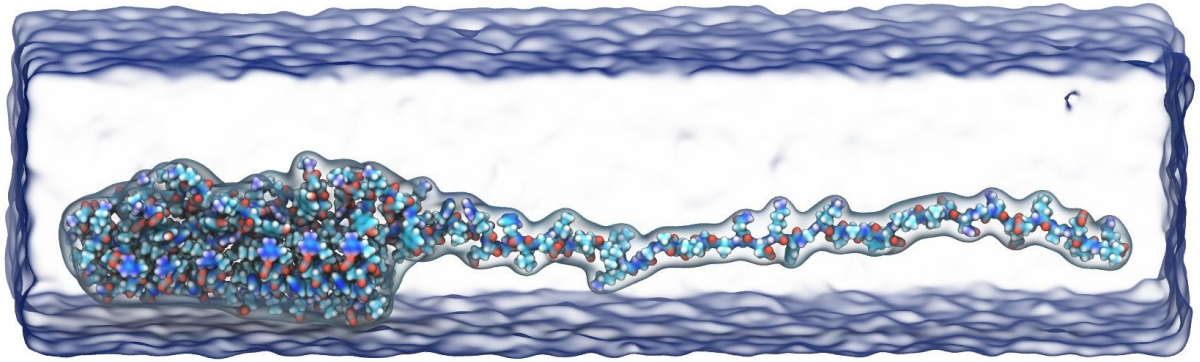

Computed Properties

- Smoothing of thermal noise
- Secondary structure
- Hydrogen bonds, salt bridges
- Forces, energies, stress, strain
- Time averaging of electrostatic fields, occupancy maps
- Quality-of-fit cross correlation with cryo-EM density maps
- Normal modes, principal component analysis, essential dynamics
- Cluster simulation trajectory timesteps by structural similarity

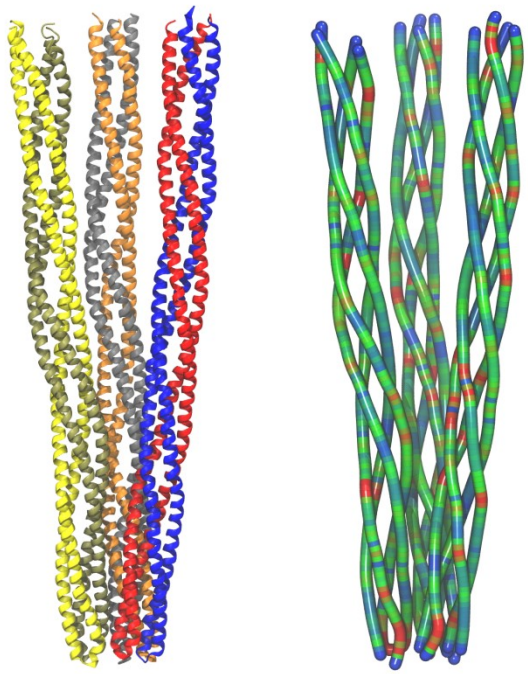

Chemoreceptor trimer-of-dimers analysis with Bendix plugin in VMD

Display of Computed Properties on Structures

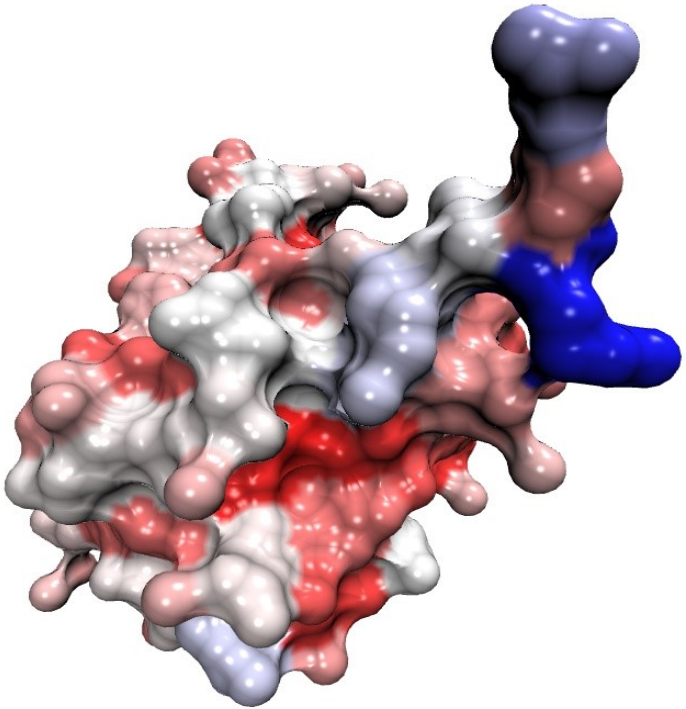

Per-residue solvent-accessible surface area of Ubiquitin

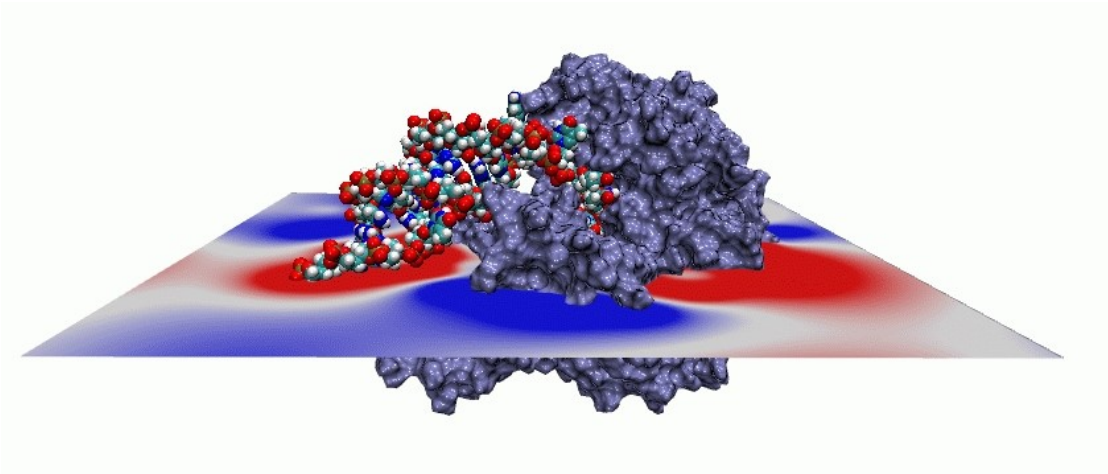

PME electrostatic potential contour for a helicase on a volumetric slice plane

CheA kinase PCA: first principal component porcupine plot

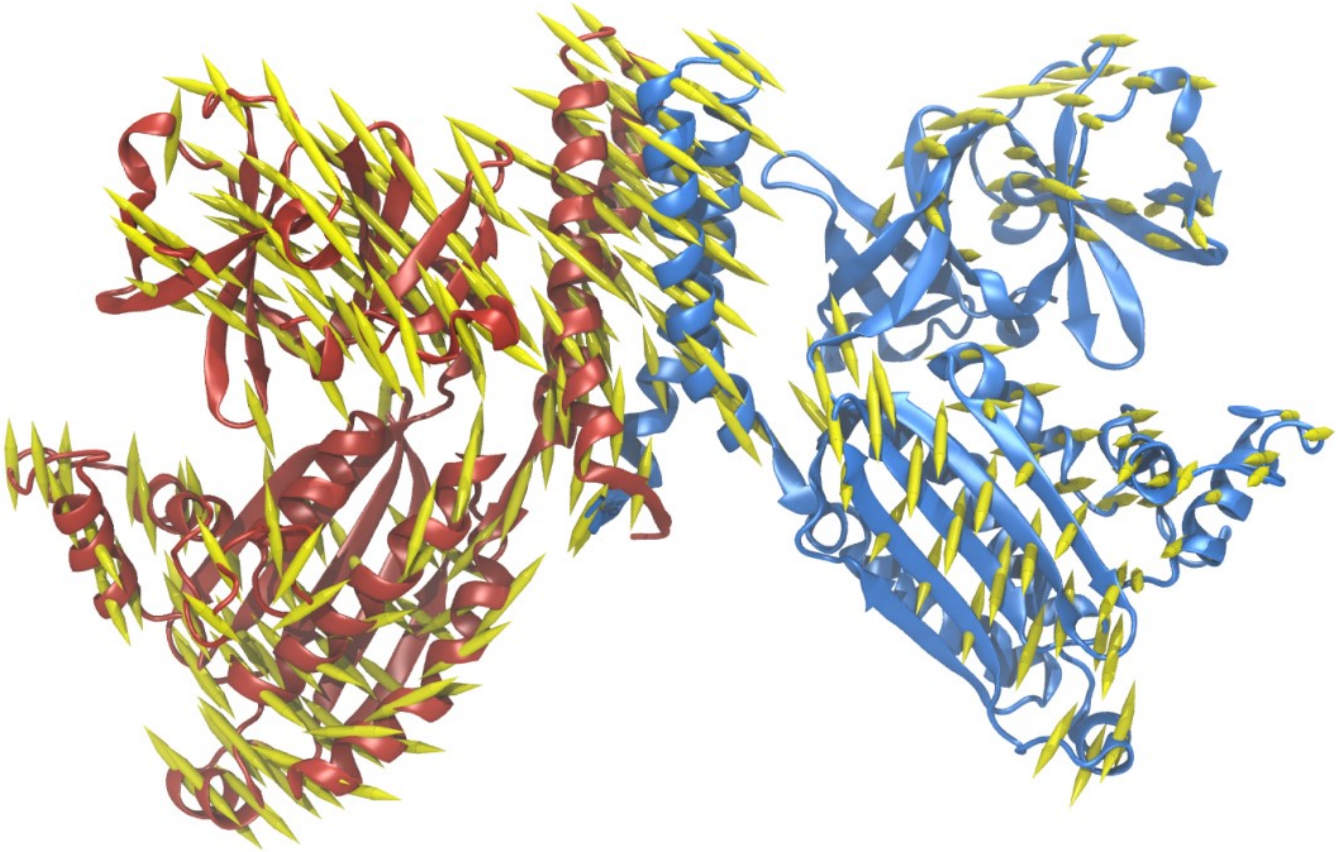

Molecular Visualization Challenges

- Geometrically complex scenes
- Often show a mix of structural and spatial properties
- Time varying!
- Spatial relationships important to see clearly: depth cueing, shadows, AO helpful

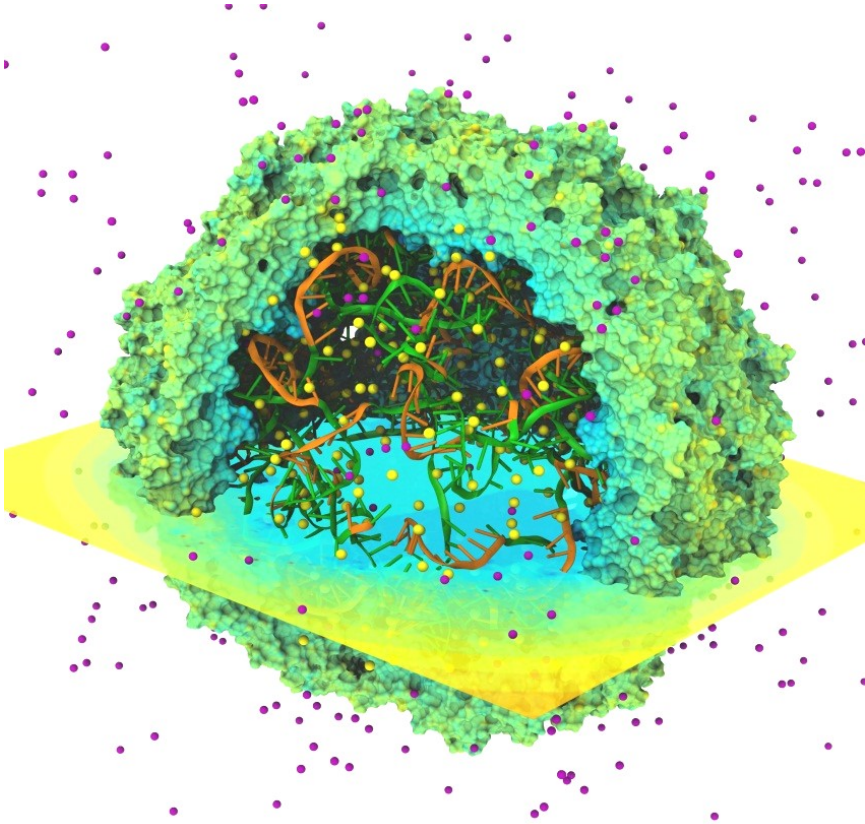

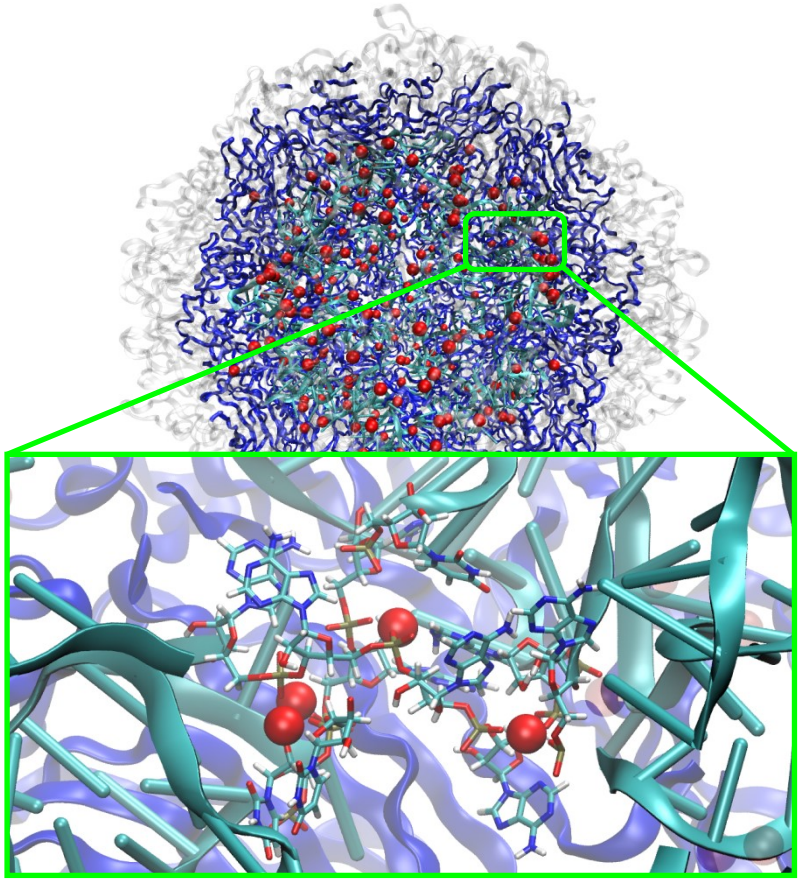

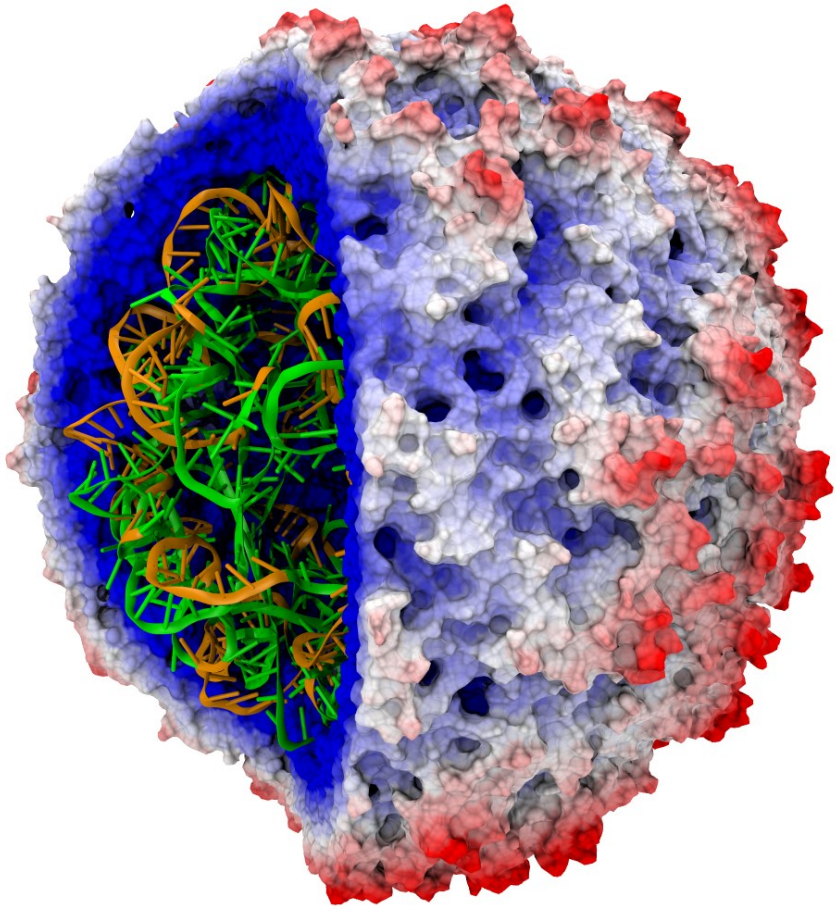

Satellite Tobacco Mosaic Virus

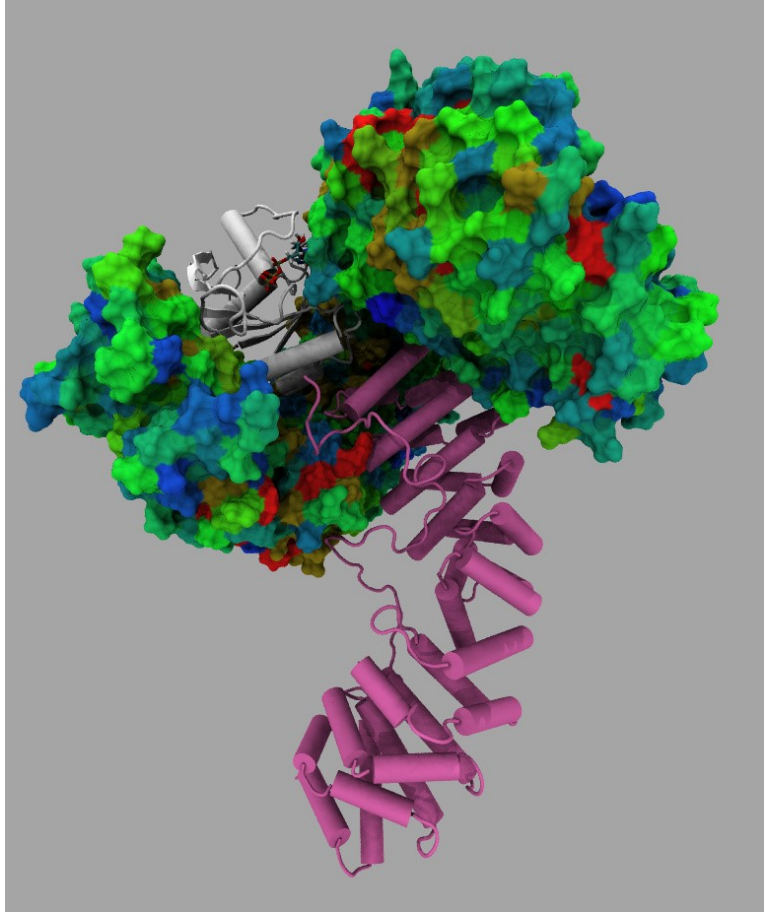

Exportin Cse1p

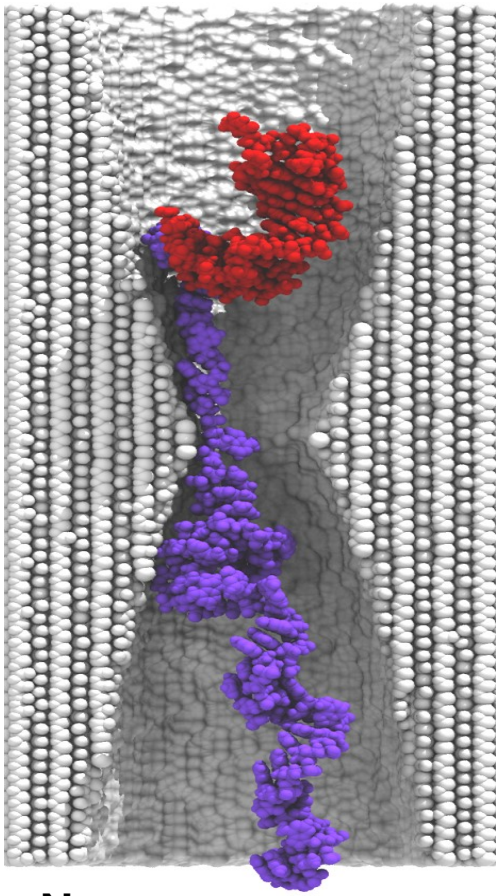

Nanopore

VMD Shading Comparison: EF-Tu

Outline Shader

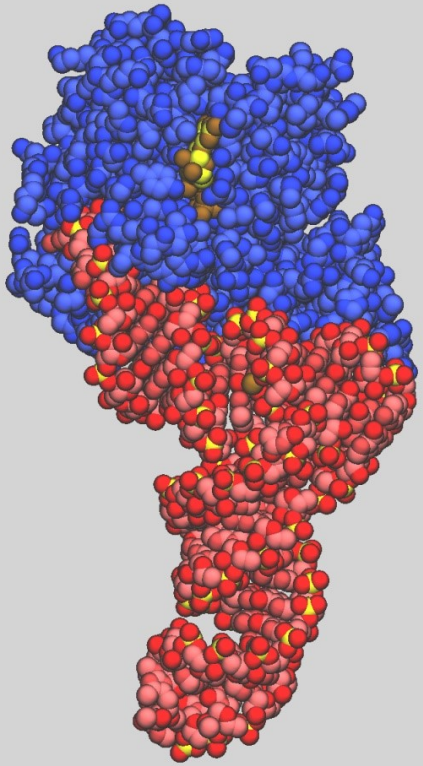

“Goodsell” Shader

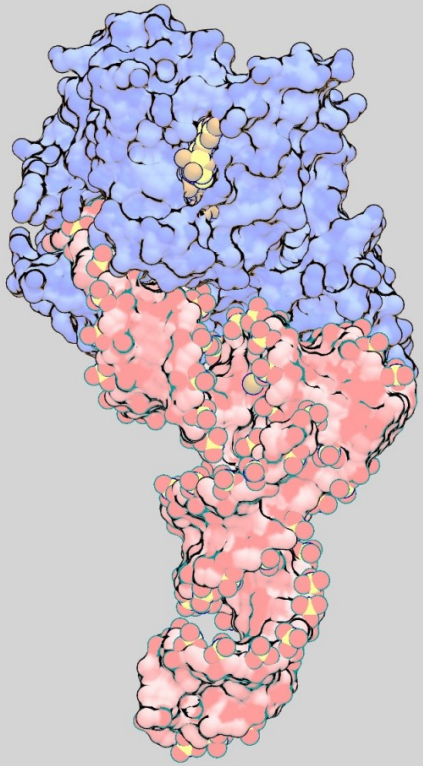

Glossy Shader

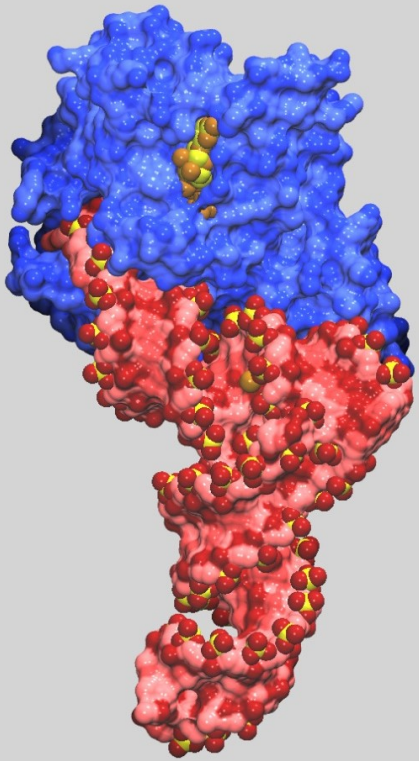

Ambient Occlusion, Shadowing

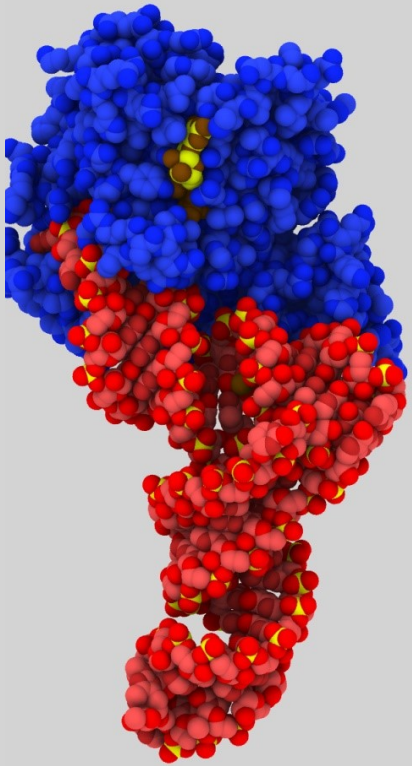

VMD Scripting

VMD Scripting

- Built-in Tcl interpreter is standard in all builds
- Custom builds can contain Python interpreter too
- Collection of simple VMD scripts:

  http://www.ks.uiuc.edu/Research/vmd/script_library/

Atom Selections

```
set ionsel [atomselect top “ions”]
set watsel [atomselect top “water”]
set solute [atomselect top “not water and not ions”]

set numions [$ionsel num]
set coords [$solute get {x y z}]
```

Built-in “Measure” Analysis Routines

- Provide high-performance routines for common analysis functions
- Operate on atom selections, sometimes over a frame range

```
set com [measure center $sel weight mass]
set avpos [measure avpos $sel first 0 last 999 step 1]
set c [measure contacts 1.0 $sel1 $sel2]

set txmat [measure fit $sel1 $sel2]
$sel2 move $txmat
```

```
# Print RMSD between two selections between the first timestep and each
# later timestep for the given molecule id (default: top)
proc print_rmsd_trajectory {{mol top}} {
    set reference  [atomselect $mol "protein" frame 0]        # use frame 0 for the reference
    set compare    [atomselect $mol "protein"]                # the frame being compared
    set numsteps [molinfo $mol get numframes]                 # query number of trajectory frames
    for {set frame 0} {$frame < $numsteps} {incr frame} {
        $compare frame $frame                                 # set the correct comparison frame
        set transmat [measure fit $compare $reference]        # compute the transformation
        $compare move $transmat                               # align comparison to reference
        set rmsd [measure rmsd $compare $reference]           # compute the RMSD
        puts "RMSD of $frame is $rmsd"                        # print the RMSD
     }
}
```

Trajectory Analysis and Visualization

Visualization of MD Trajectories

- Allow researchers study trajectories by analyzing force profiles, energies, structural changes, etc.
- **Visualization selections, graphics, structure properties, colors can all be recomputed for each trajectory timestep!**

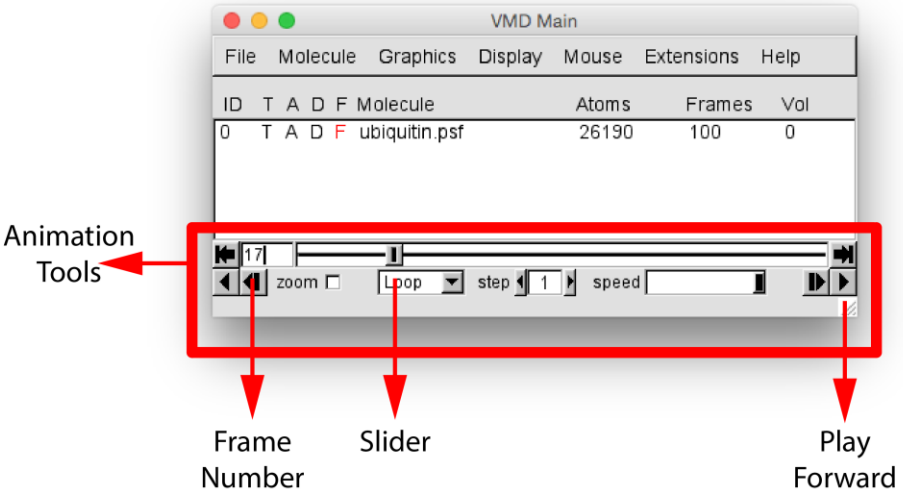

Folding Dynamics of Villin Headpiece Unveiled

6.9μs folding simulation of 30K atoms: 380GB trajectory

Schulten et al. *Biophys J* **94**:L75, 2008, **97**: 2009

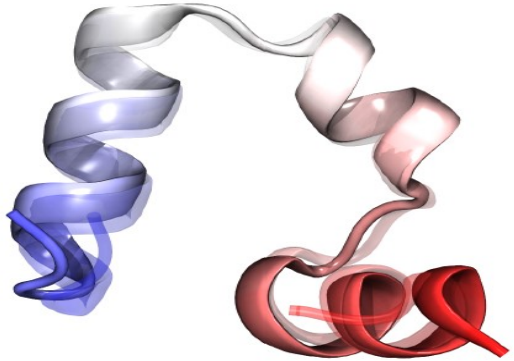

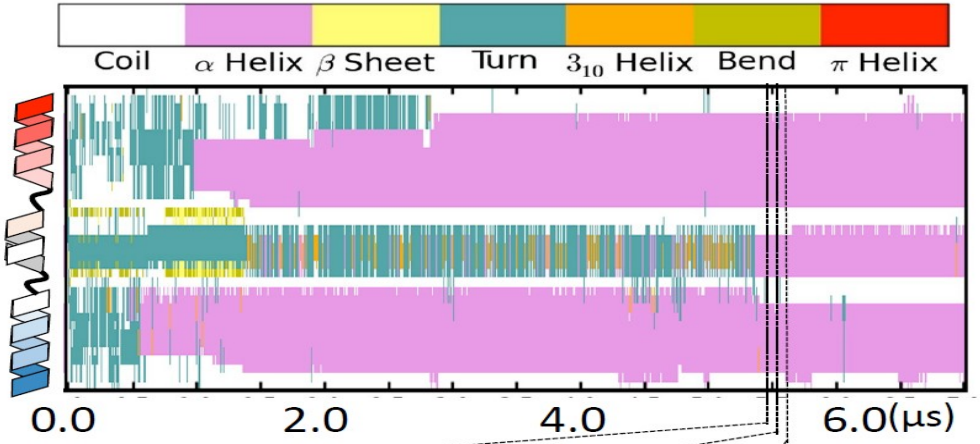

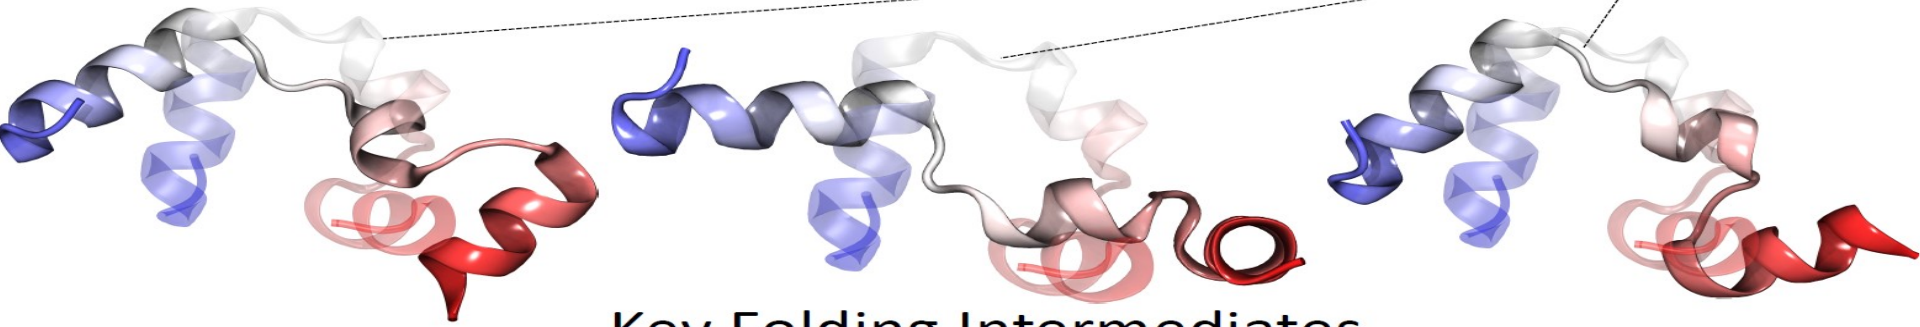

Key Folding Intermediates

Time-Averaged Volumetric Properties

- Compute density, distance, occupancy, potential maps for a frame or averaged over a trajectory
- Example: display binding sites for diffusively bound ions as probability density isosurfaces

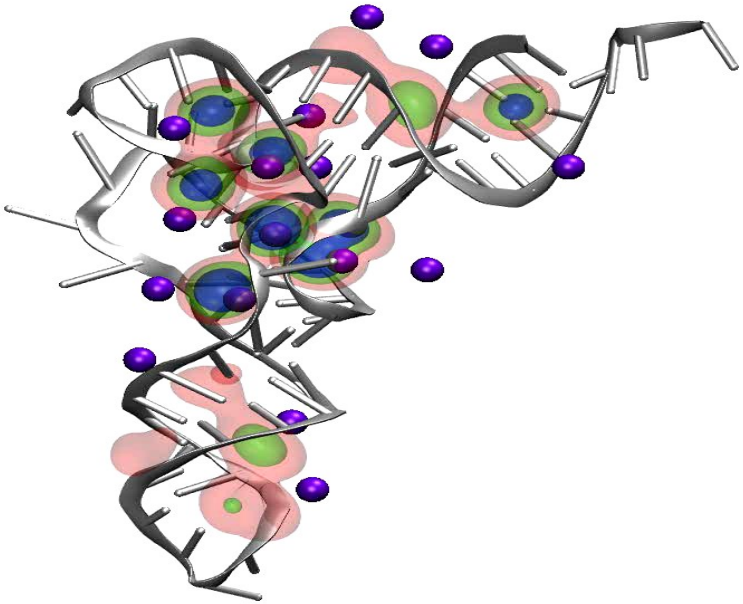

tRNA magnesium ion occupancy: VMD volmap plugin

Petascale Computing - A Key Instrument for Life Science

MDFF Solves Structures from X-ray Crystallography and Cryo-EM

Electron microscopy

FEI microscope

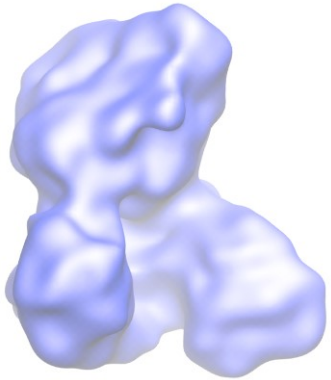

Electron density of protein in action at low resolution

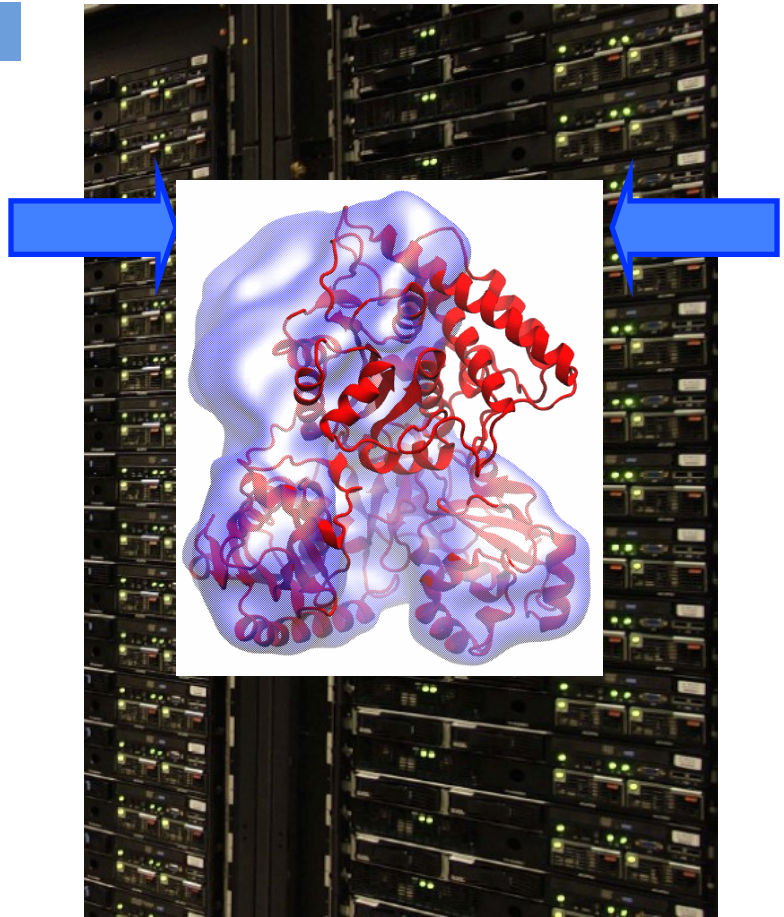

X-ray crystallography

APS at Argonne

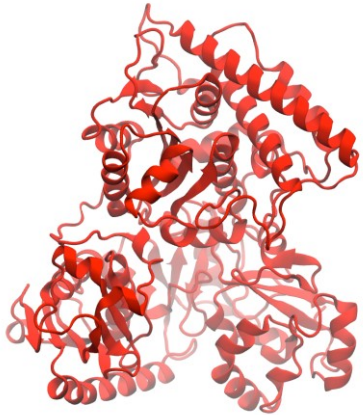

Ideal protein structure at high resolution
Acetyl – CoA Synthase

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**.

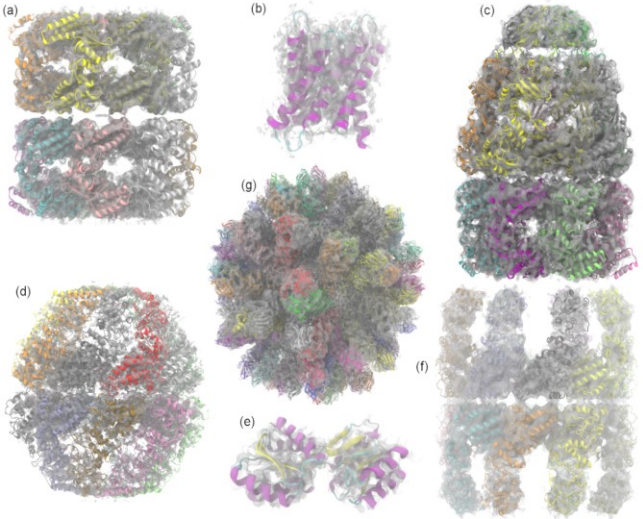

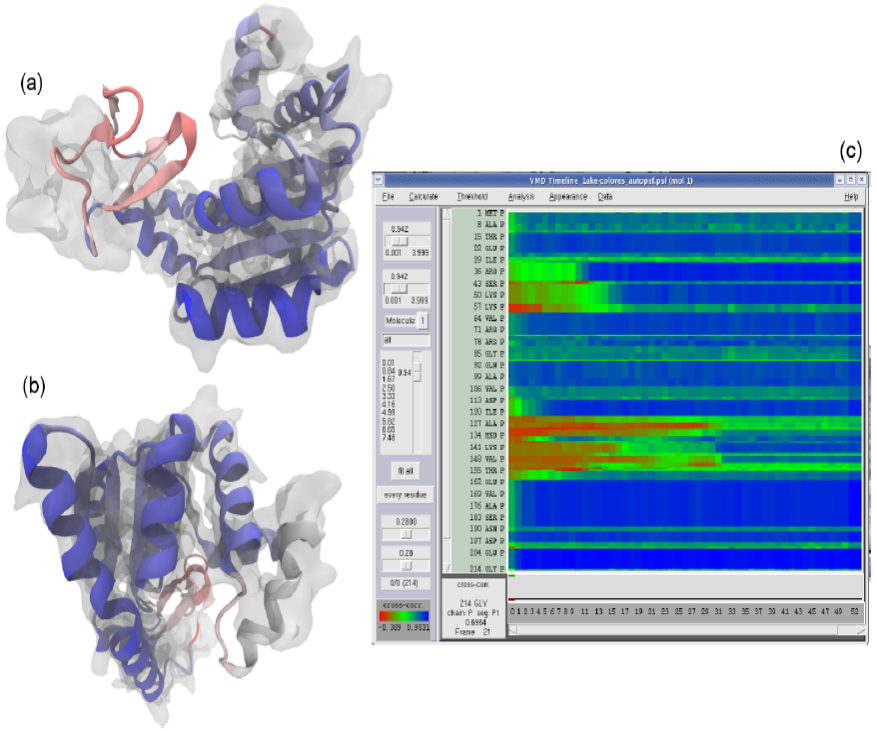

MDFF Cross Correlation Timeline

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!

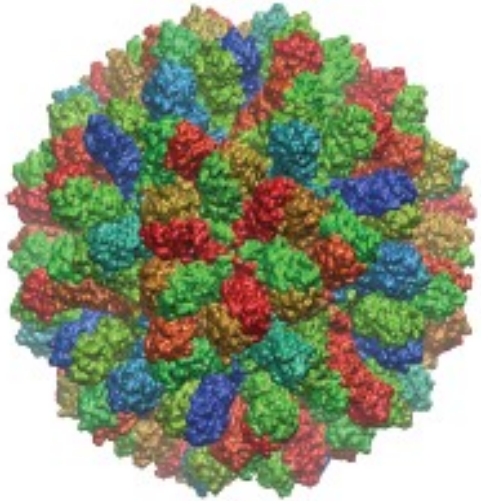

RHDV colored by relative CC

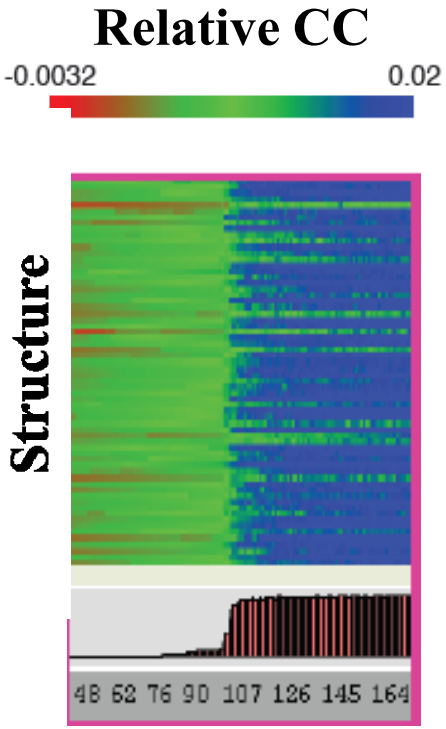

Stone et al., Faraday Discuss., 169:265-283, 2014.

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.

Clustering Analysis of Molecular Dynamics Trajectories

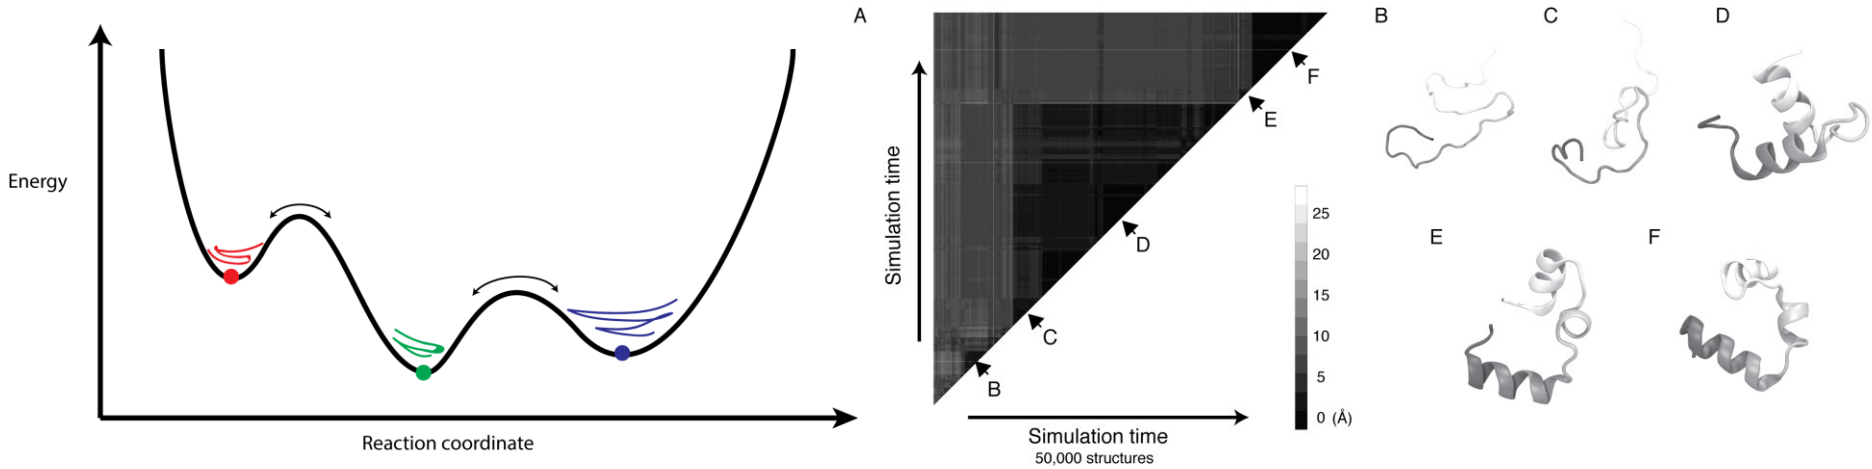

GPU-Accelerated Molecular Dynamics Clustering Analysis with OpenACC. J.E. Stone, J.R. Perilla, C. K. Cassidy, and K. Schulten. In, Robert Farber, ed., Parallel Programming with OpenACC, Morgan Kaufmann, Chapter 11, pp. 215-240, 2016.

High Fidelity Ray Tracing of Molecular Scenes

Ray Tracing in VMD

- Support for ray tracing of VMD molecular scenes began in 1995
- Tachyon parallel RT engine interfaced with VMD (1999)
- Tachyon embedded as an internal VMD rendering engine (2002)
- Built-in support for large scale parallel rendering (2012)
- Refactoring of VMD to allow fully interactive ray tracing as an alternative to OpenGL (2014)

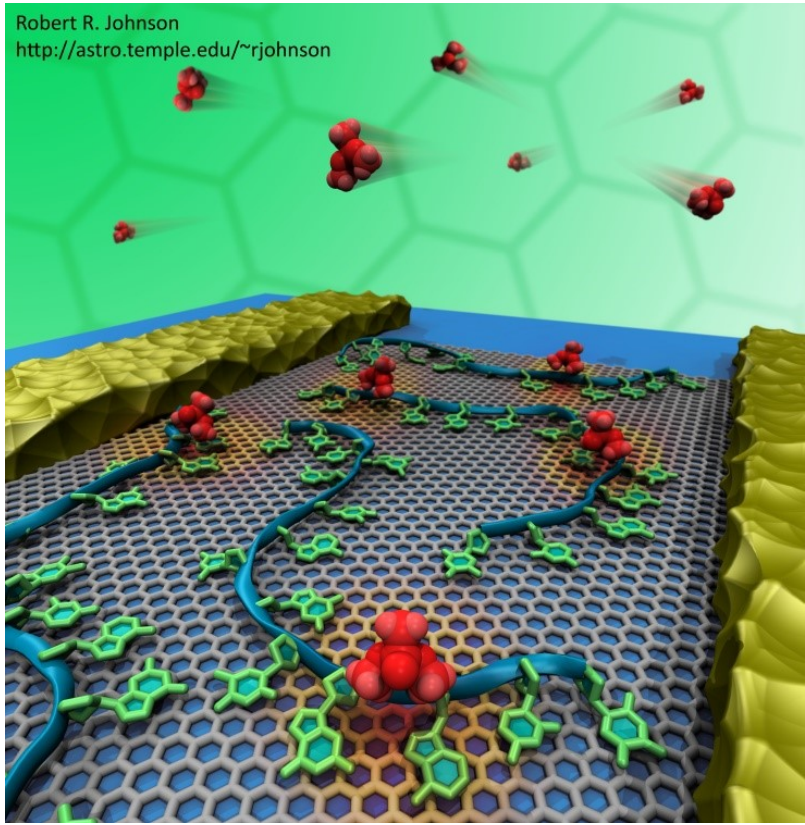

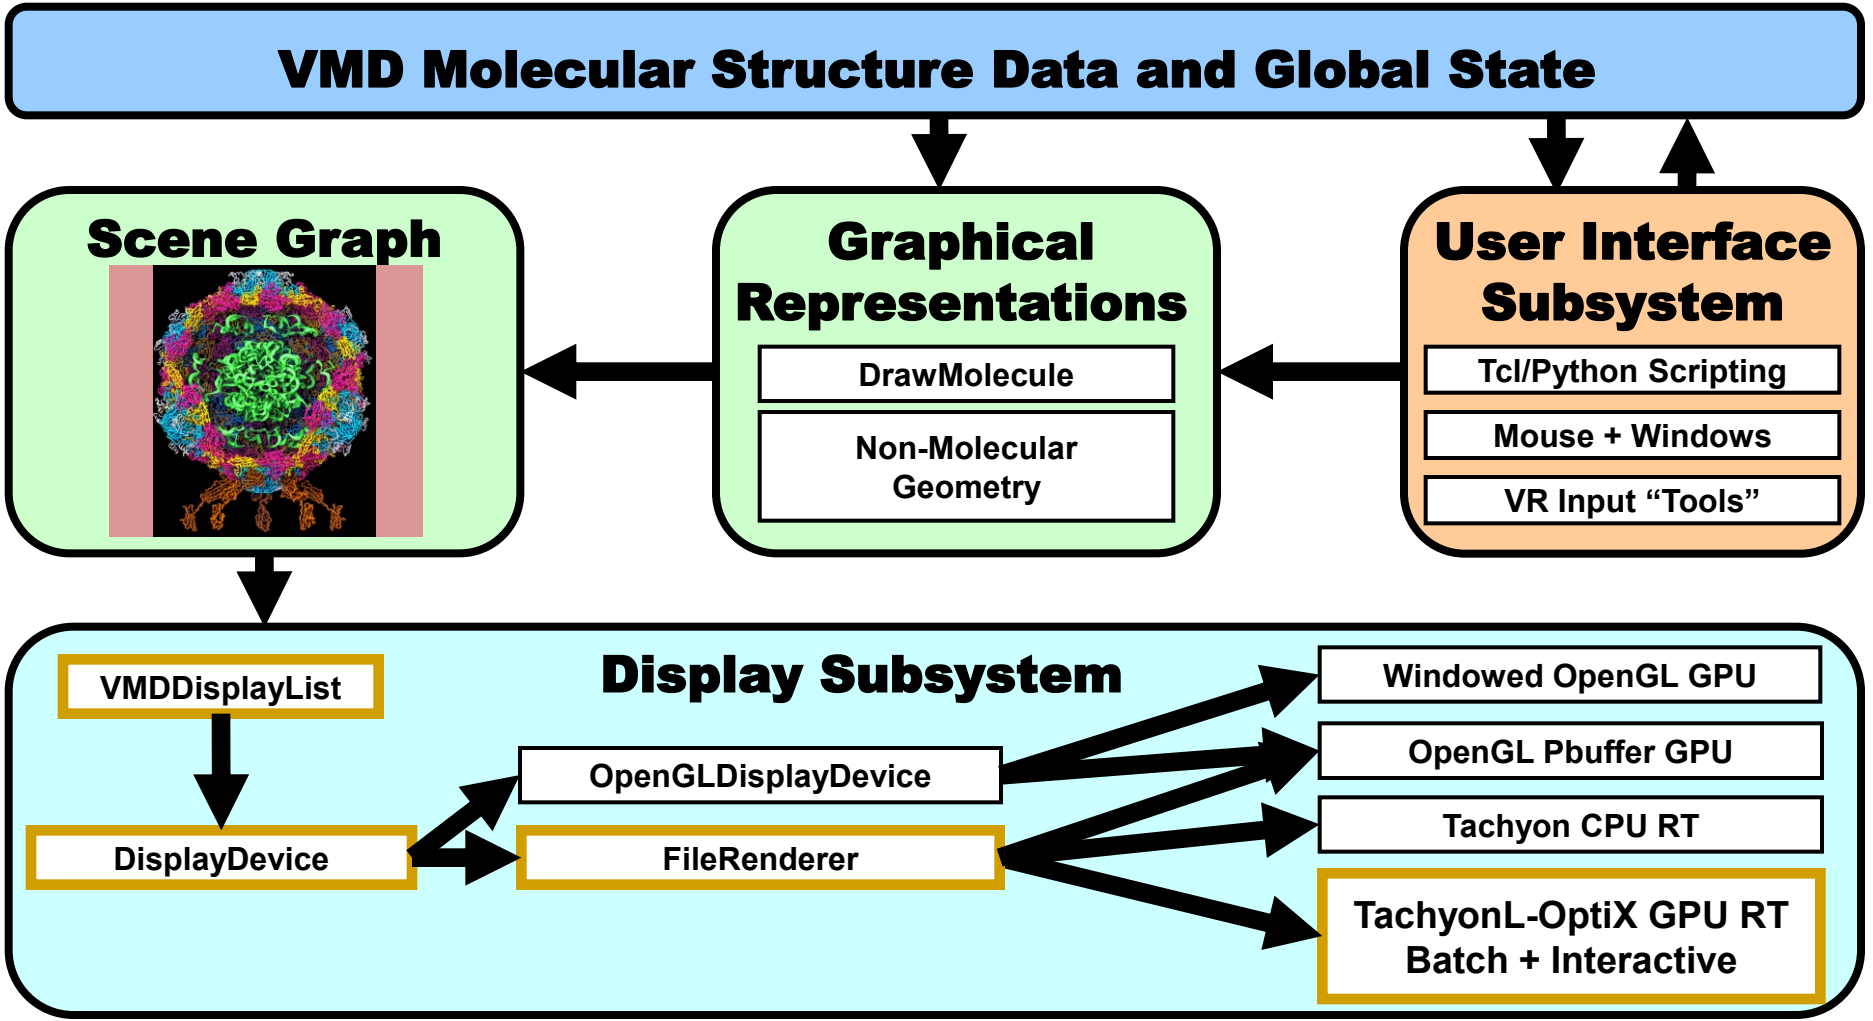
VMD Molecular Structure Data and Global State
Scene Graph
Graphical Representations
DrawMolecule
Non-Molecular Geometry
User Interface Subsystem
Tcl/Python Scripting
Mouse + Windows
VR Input "Tools"
Display Subsystem
VMDDisplayList
DisplayDevice
OpenGLDisplayDevice
FileRenderer
Windowed OpenGL GPU
OpenGL Pbuffer GPU
Tachyon CPU RT
TachyonL-OptiX GPU RT
Batch + Interactive

Interactive Ray Tracing, Lighting Comparison: STMV Capsid

Two lights, no shadows (e.g. as used by OpenGL)

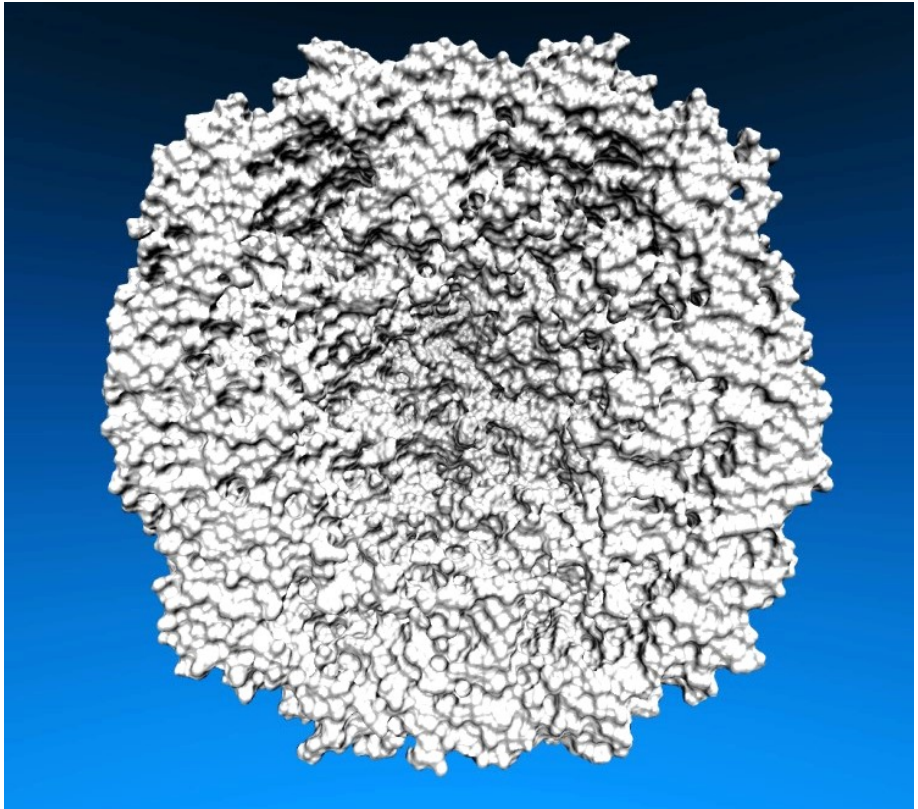

Ambient occlusion lighting and shadows w/ RT

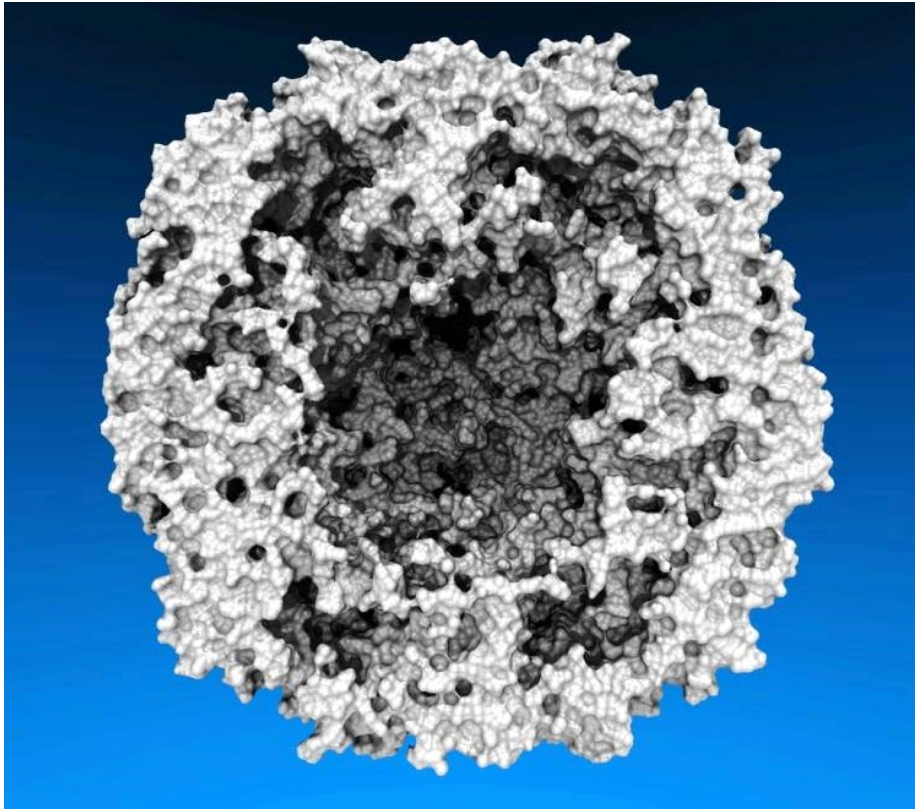

Diverse Shading and Lighting Approaches

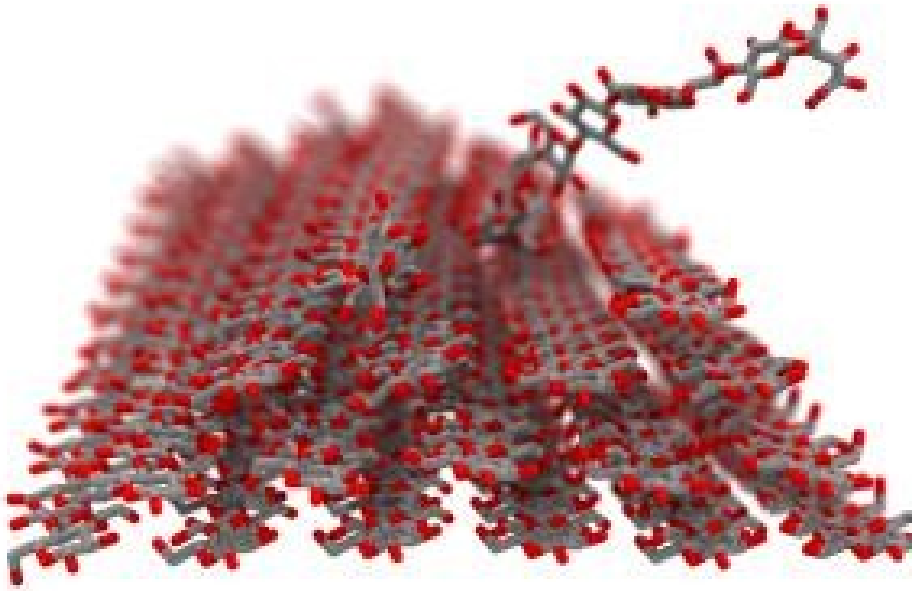

Decrystallization: Interactive Ray Tracing w/ Ambient Occlusion Lighting, Depth of Field Focal Blur

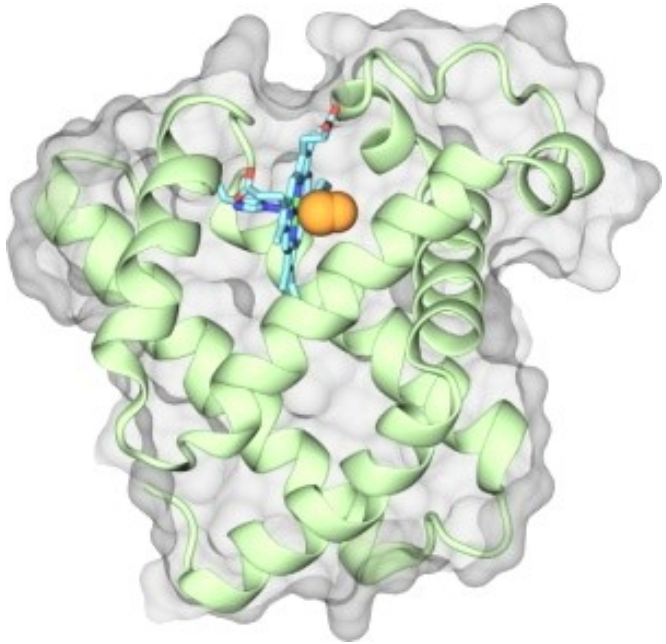

Myoglobin

Geometrically Complex Scenes

Ray tracing techniques well matched to molecular viz. needs:

- Curved geometry, e.g. spheres, cylinders, toroidal patches, easily supported
- Greatly reduced memory footprint vs. polygonalization
- Runtime scales only moderately with increasing geometric complexity
- Occlusion culling is “free”, RT acceleration algorithms do this and much more

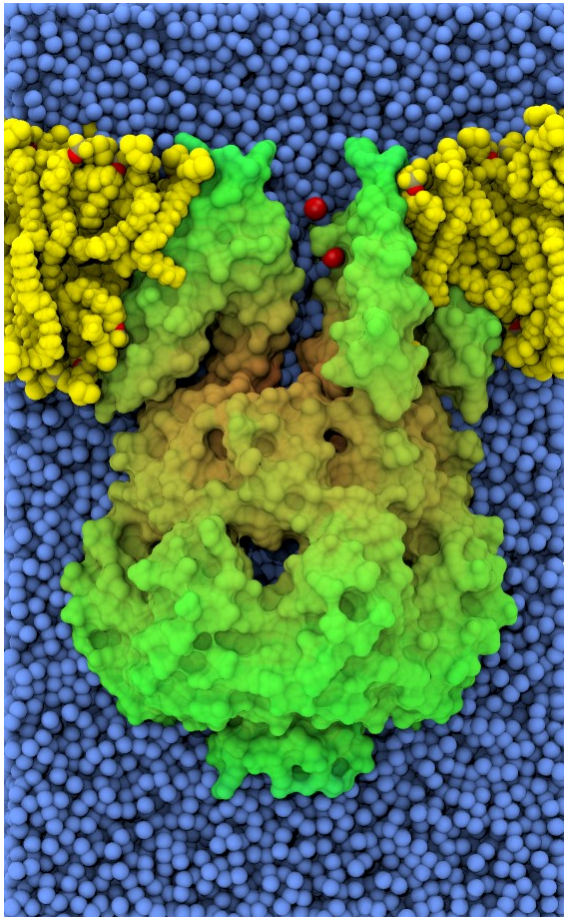

Benefits of Advanced Lighting and Shading Techniques

- Exploit visual intuition
- Spend computer time in exchange for scientists' time, make images that are more easily interpreted

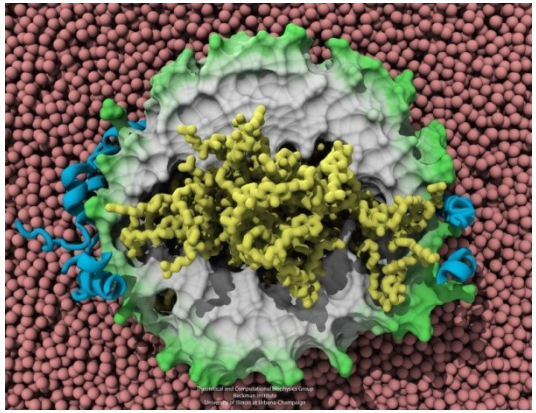

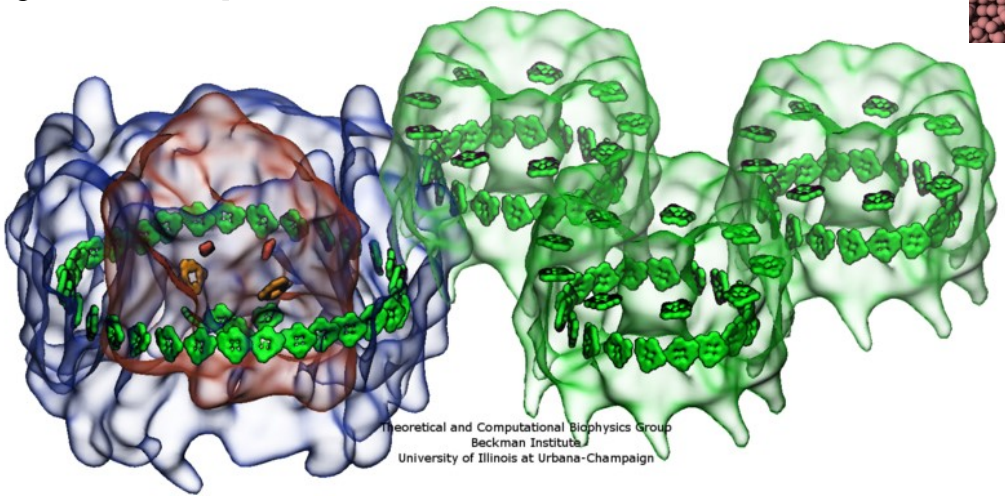

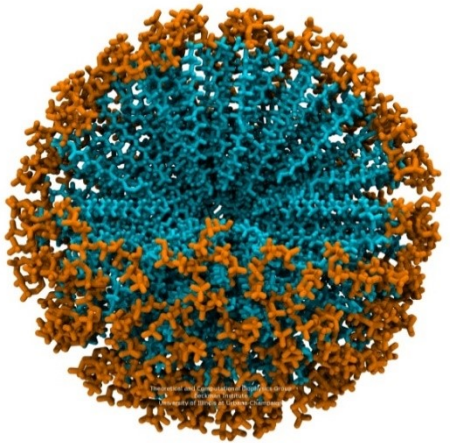

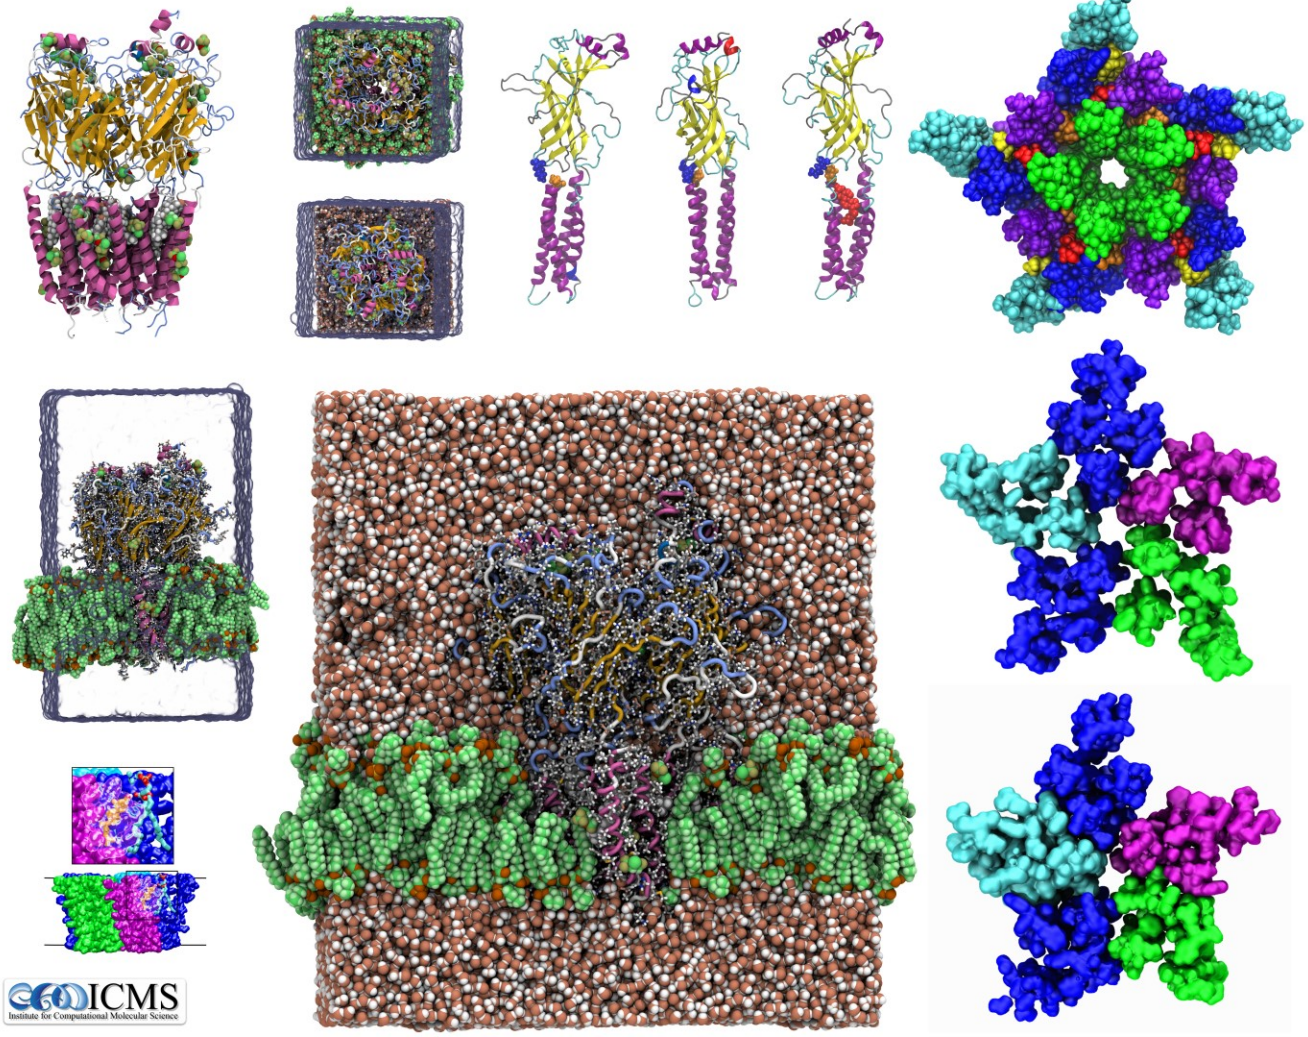
ICMS
Institute for Computational Molecular Science

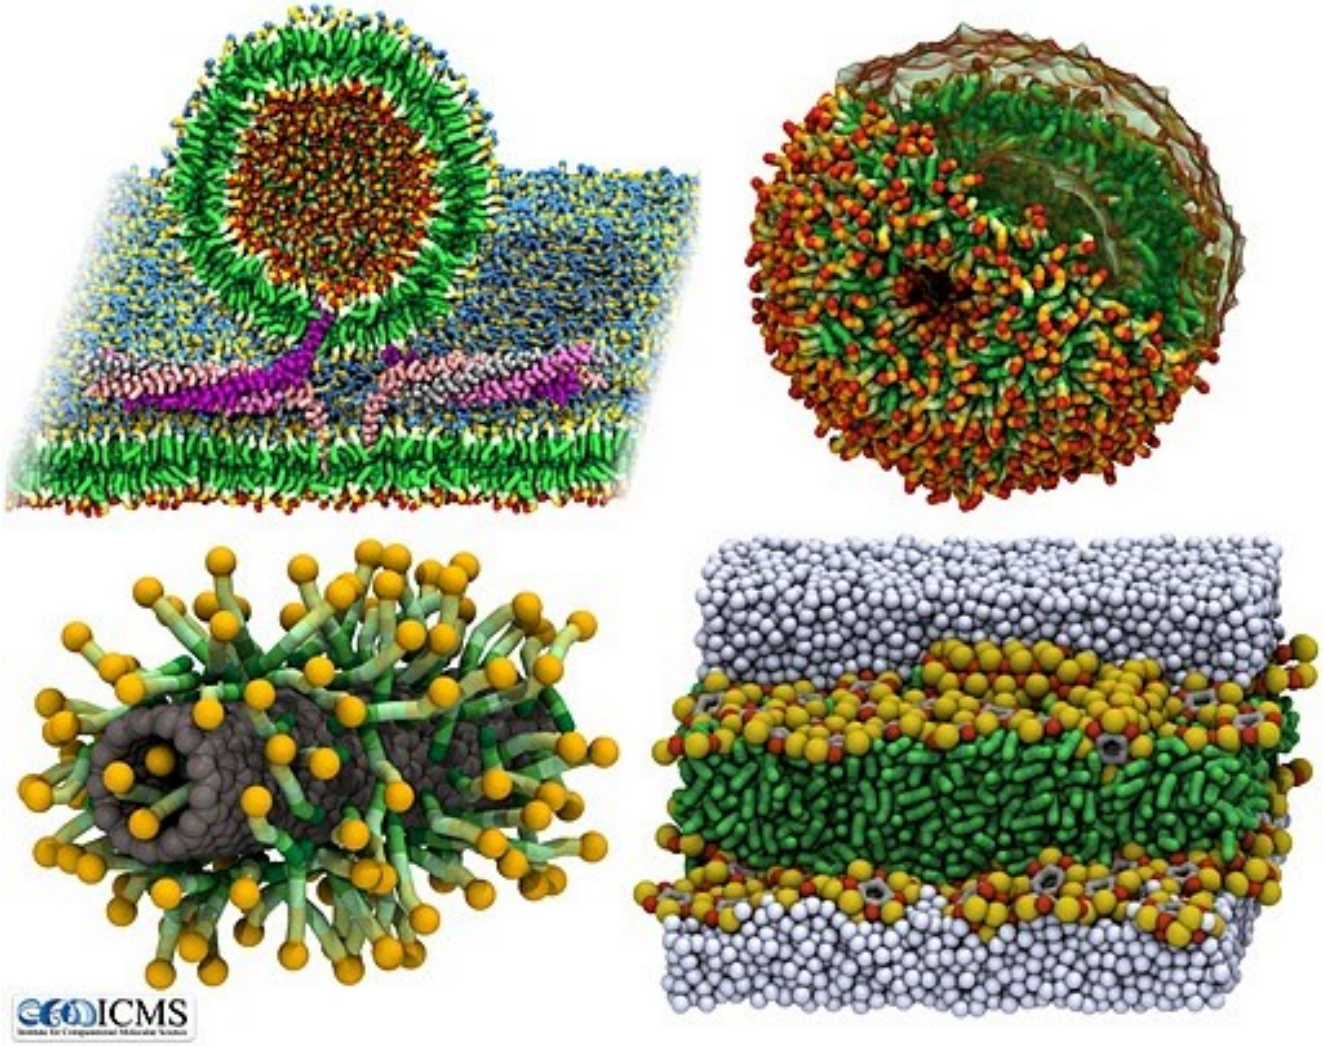
ICMS

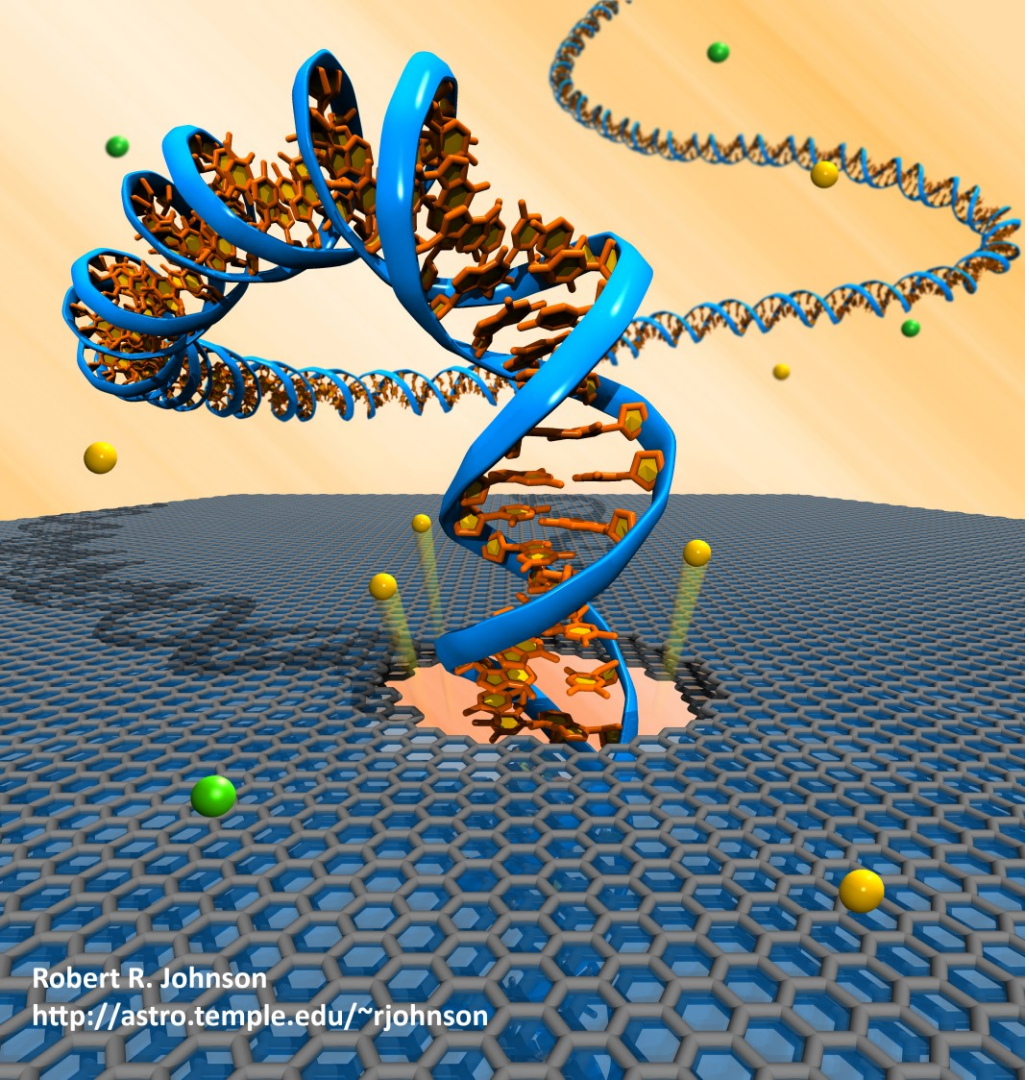
Robert R. Johnson
http://astro.temple.edu/~rjohnson

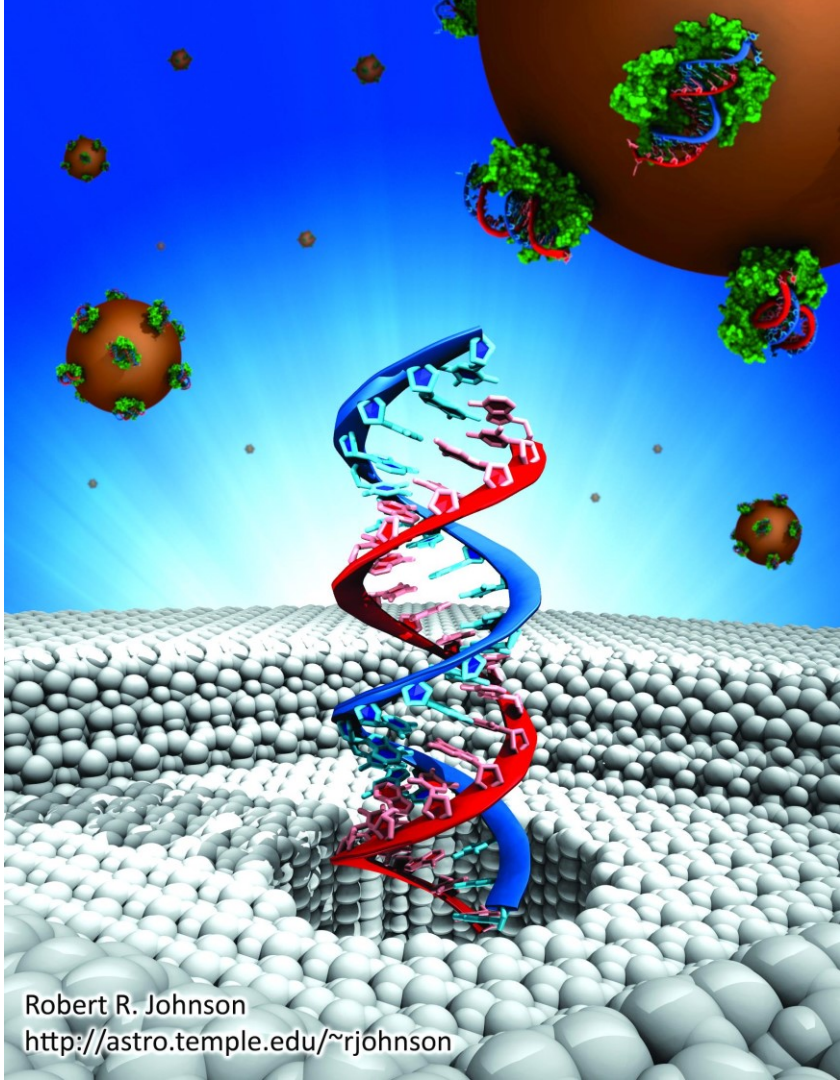
Robert R. Johnson
http://astro.temple.edu/~rjohnson

High Fidelity Interactive Visualization

- VMD interactive ray tracing
  - Interactive ray tracing on GPUs with progressive refinement of image and lighting quality
  - Fully interactive rendering of large structures with advanced lighting features, and WYSIWYG “What you see is what you get” final image output

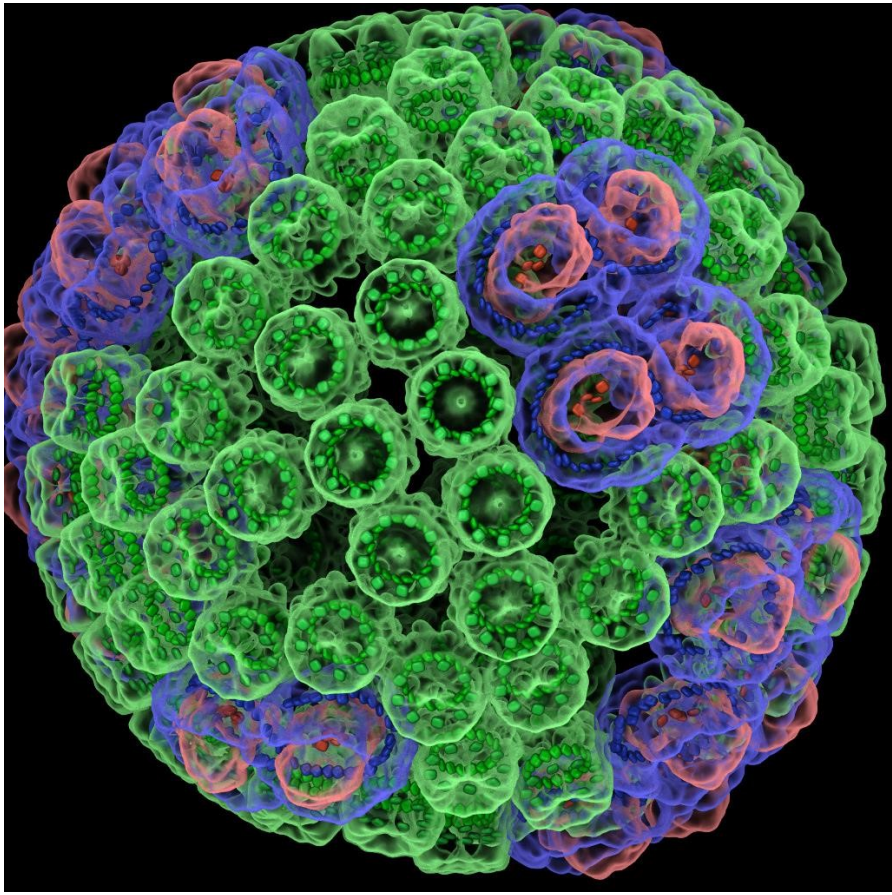

VMD Atomic Detail Visualization of Cellular Architecture with Instancing

- VMD 1.9.4 supports instancing of graphical representations associated with molecules
- Exploit **VBO caching** in OpenGL to eliminate host-GPU geometry transfers
- **OptiX instancing** of geometry buffers to minimize GPU memory footprint for cell-scale scenes w/ atomic structures

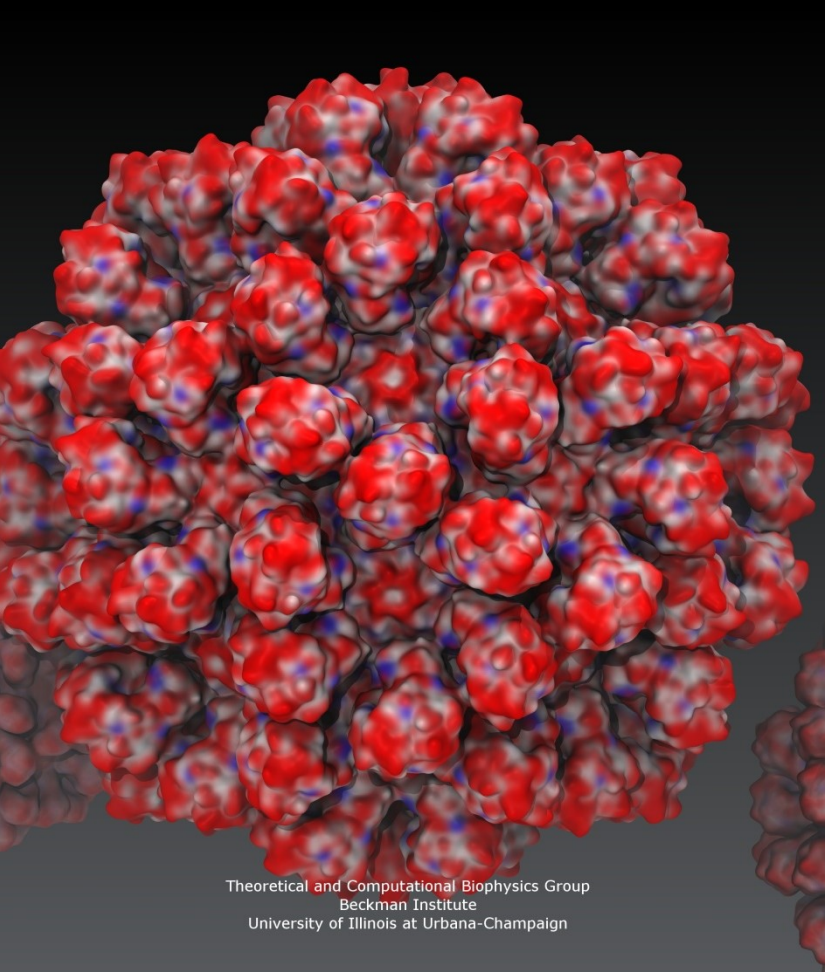
Theoretical and Computational Biophysics Group
Beckman Institute
University of Illinois at Urbana-Champaign

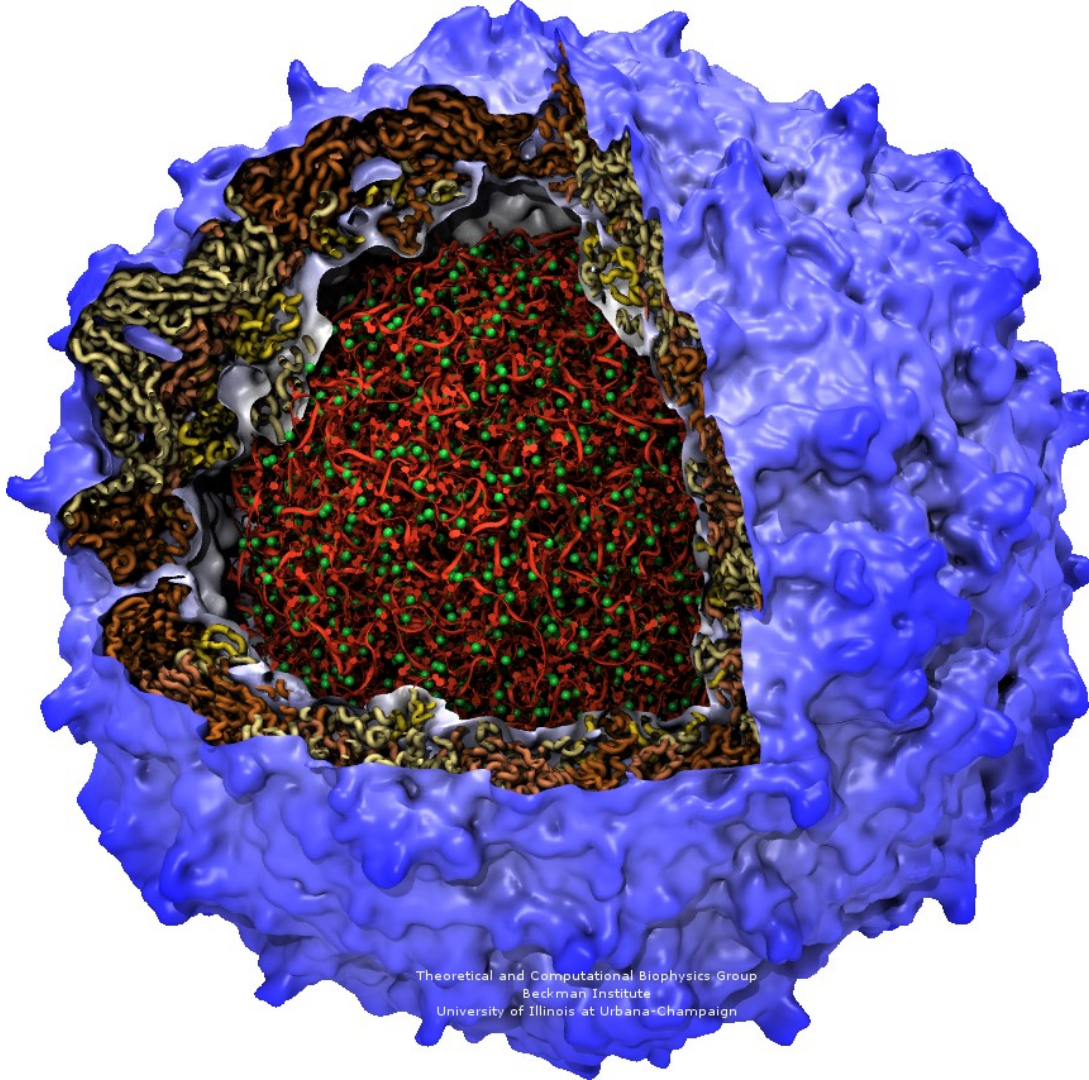
Theoretical and Computational Biophysics Group
Beckman Institute
University of Illinois at Urbana-Champaign

Capsid acts as an osmotic regulator

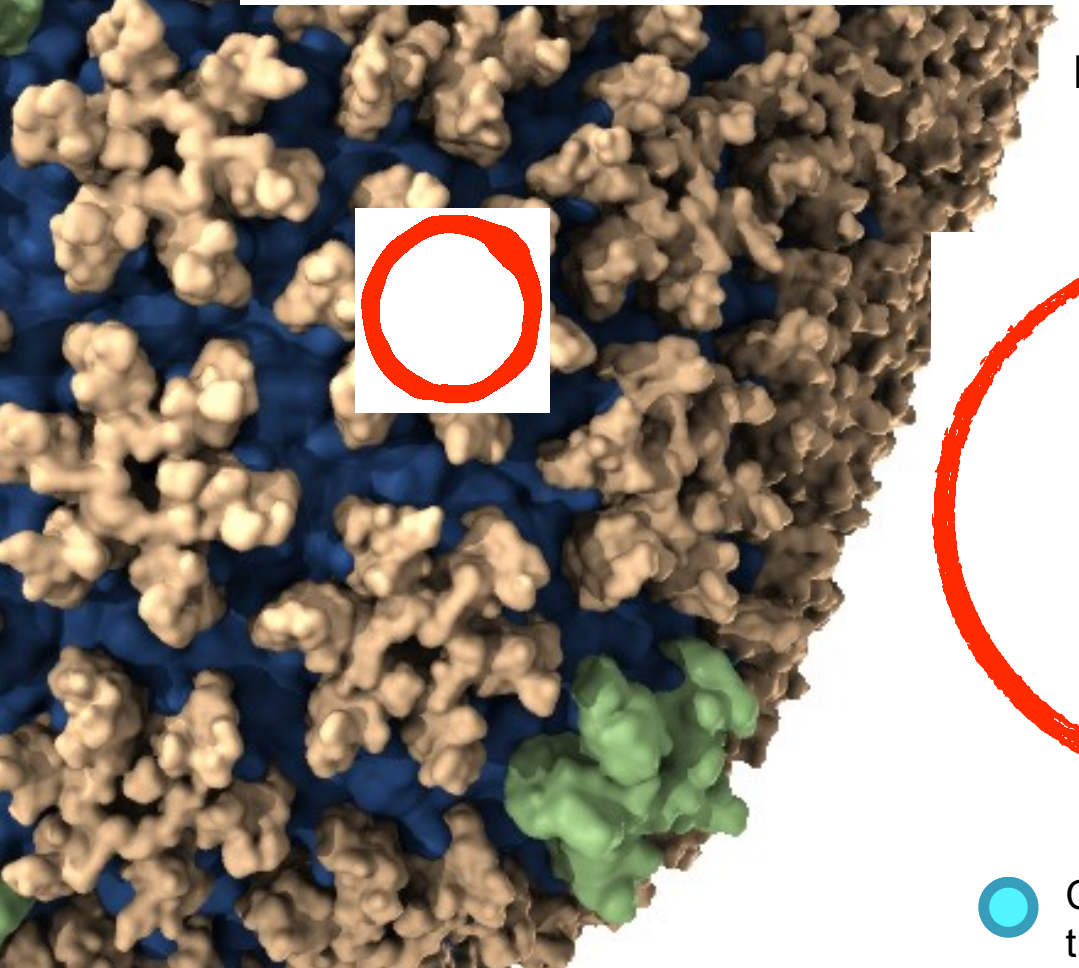

Results from 64 M atom, 1 μs sim

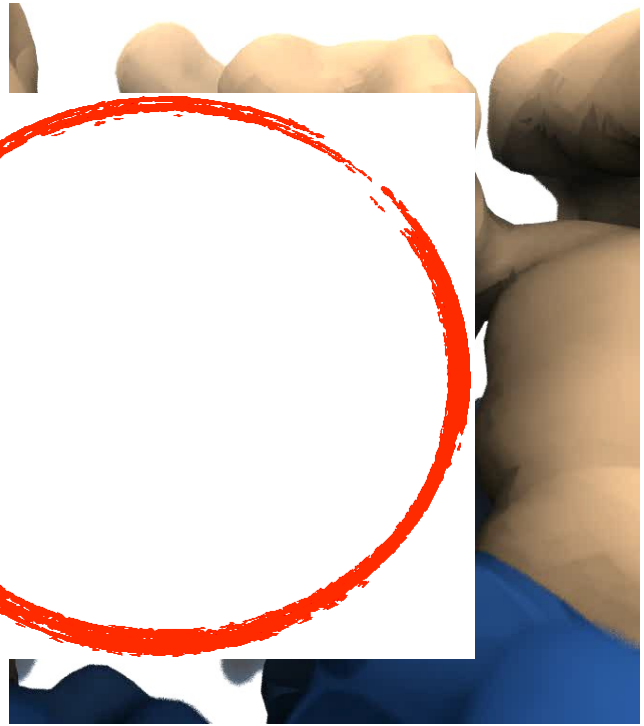

Chloride ions permeate through the hexameric center

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**

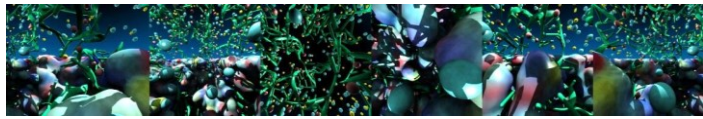

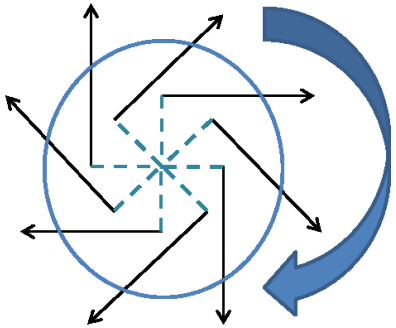

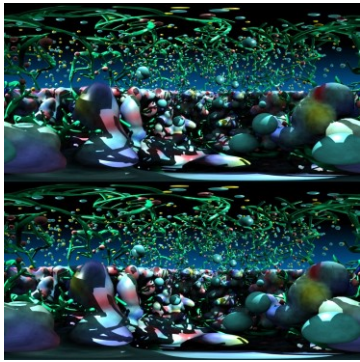

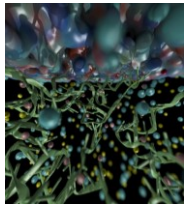

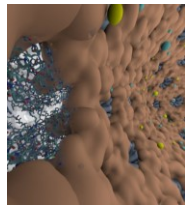

Atomic Detail Visualization of Photosynthetic Membranes with GPU-Accelerated Ray Tracing. J. E. Stone, et al. J. Parallel Computing, 55:17-27, 2016.

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.

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, …

VMD Planetarium Dome Master Camera

- RT-based dome projection -- rasterization poorly suited to non-planar projections
- Fully interactive RT with ambient occlusion, shadows, depth of field, reflections, and so on
- Both mono and stereoscopic
- No further post-processing required

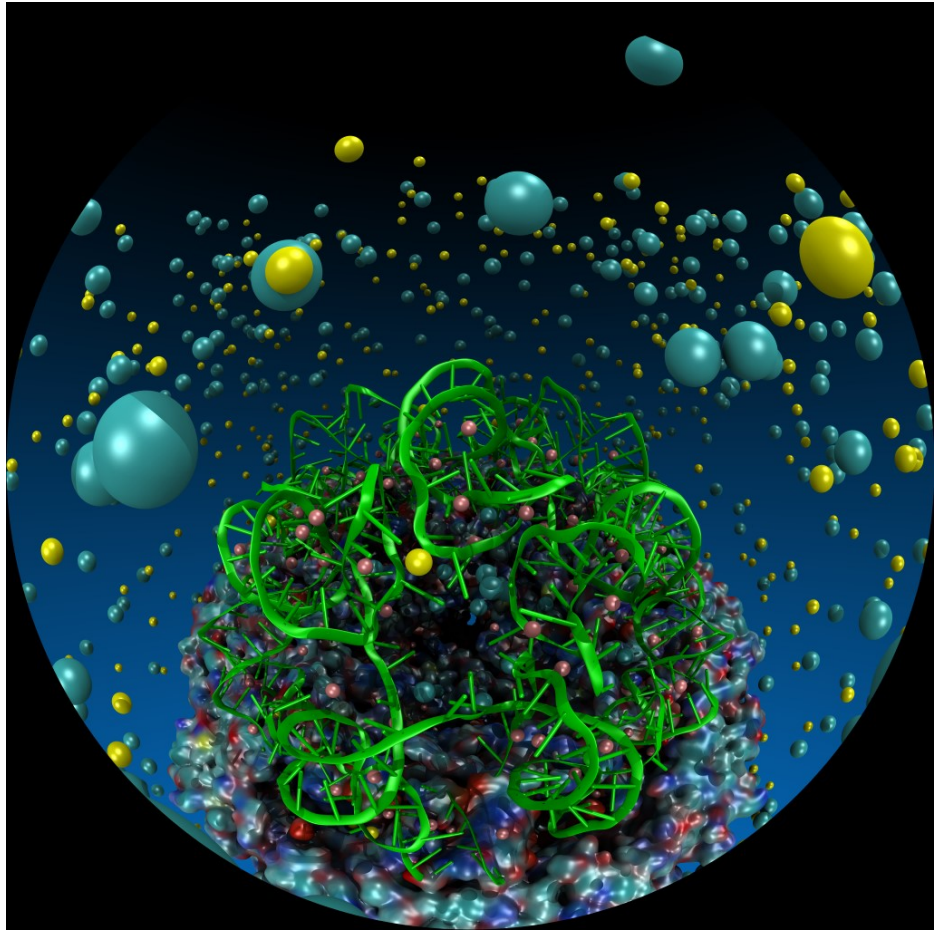

VMD Plugins and Support for User-Extensibility

QwikMD: Guided MD Simulation and Training

Smooths initial learning curve (non-expert users)

Speed up tedious simulation preparation tasks (expert users)

Reproducibility: detailed log of all steps

Interactive preparation, simulation, and analysis

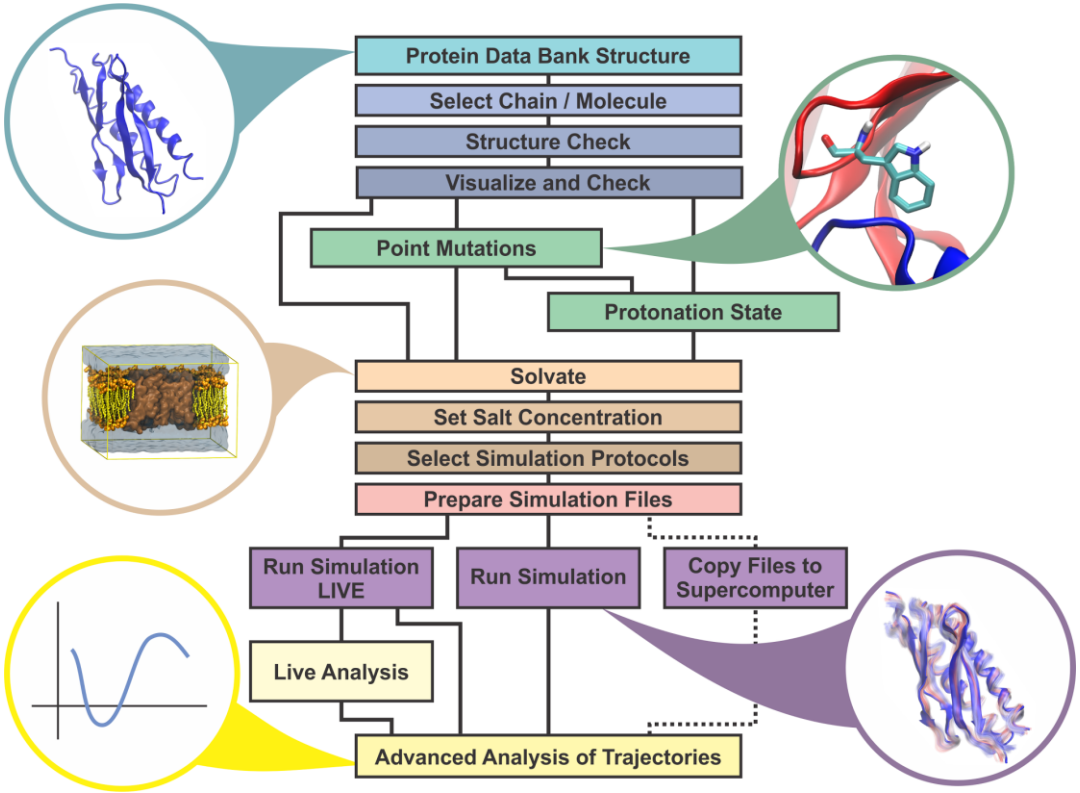

Ribeiro, et al. QwikMD - Integrative Molecular Dynamics Toolkit for Novices and Experts. ***Sci. Rep***. 6, 26536, 2016

Carbon Nanostructure Builder

http://www.ks.uiuc.edu/Research/vmd/plugins/nanotube/

- Build single-wall carbon / boron nitride nanotubes, graphene-like sheets or stacks

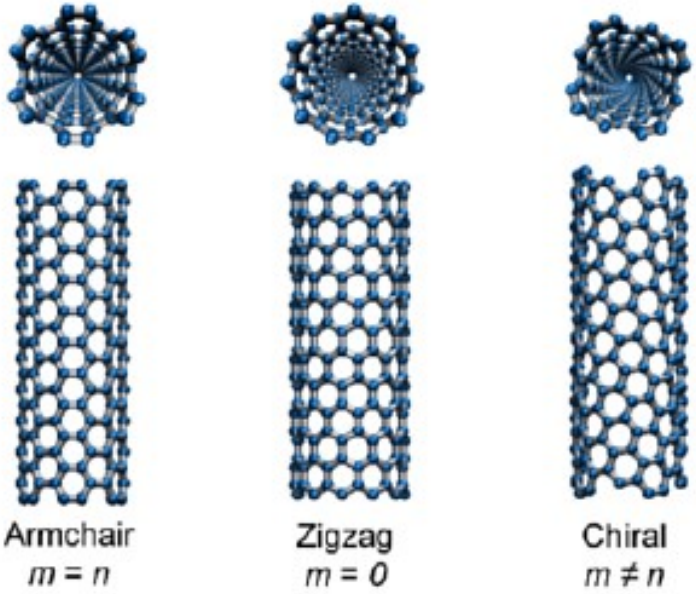

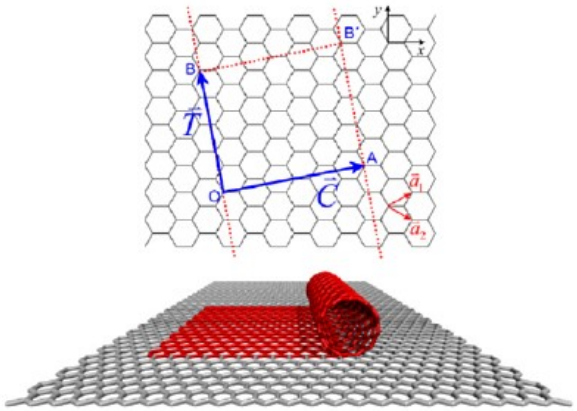

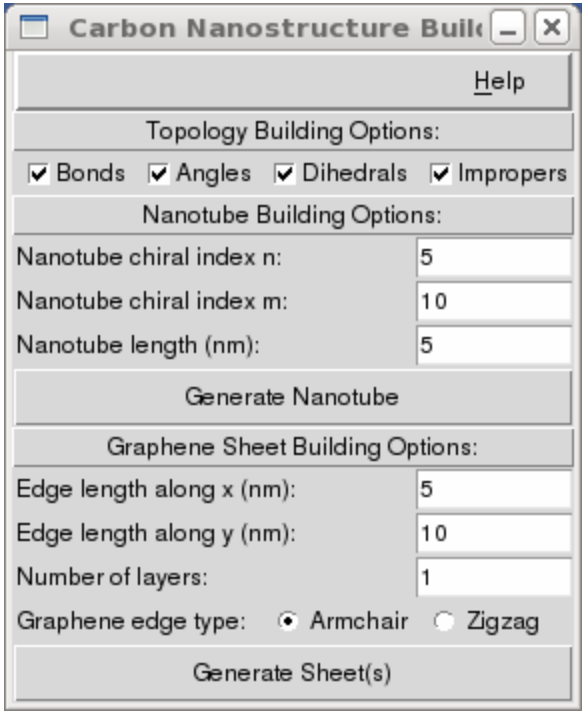

Parametrization with the Force Field Toolkit (FFTK)

A plugin enabling the simulation of drug-like small molecules by lowering force field development barriers

A graphical interface organized to follow the standardized CGenFF workflow

**** NEW in VMD 1.9.3 ****

Import data from the CGenFF Program webserver to assign initial parameters by analogy

Leverages VMD's powerful graphics capabilities to visualize parameter-related data within the context of molecular structure

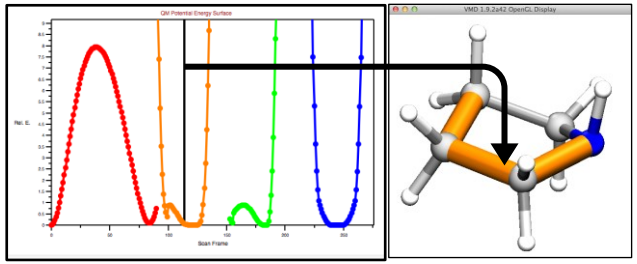

Interactively explore energy profiles

Future Plans:

- Support free-for-academics QM software, ORCA (Neese, Max Planck)
- Parametrization tools for additional MM force fields
  - AMBER force field (Gumbart, Georgia Tech.)
  - Drude polarizable force fields (Roux, U. Chicago; MacKerell, U. Maryland)

Extensive training material available from the Center's website

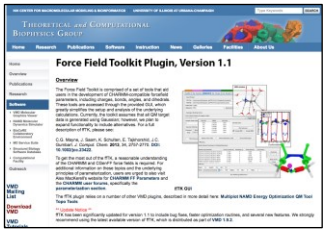

Documentation Website

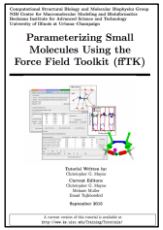

Workshop Tutorial

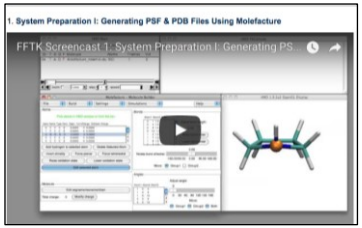

Screencast Demos

Example VMD Visualization and Analysis Plugins

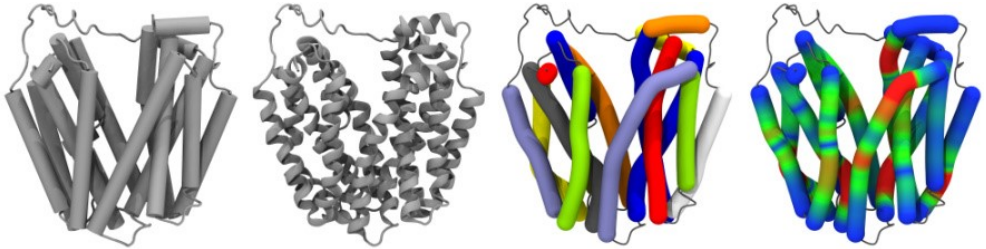

Bendix

Dahl ACE, Chavent M and Sansom MSP Bendix: intuitive helix geometry analysis and abstraction. *Bioinformatics* 2012 28(16): 2193-2194.

Normal Mode Wizard

Bakan A, Meireles LM, Bahar I ProDy: Protein Dynamics Inferred from Theory and Experiments. *Bioinformatics* 2011 27(11):1575-1577.

Large System Analysis and Visualization

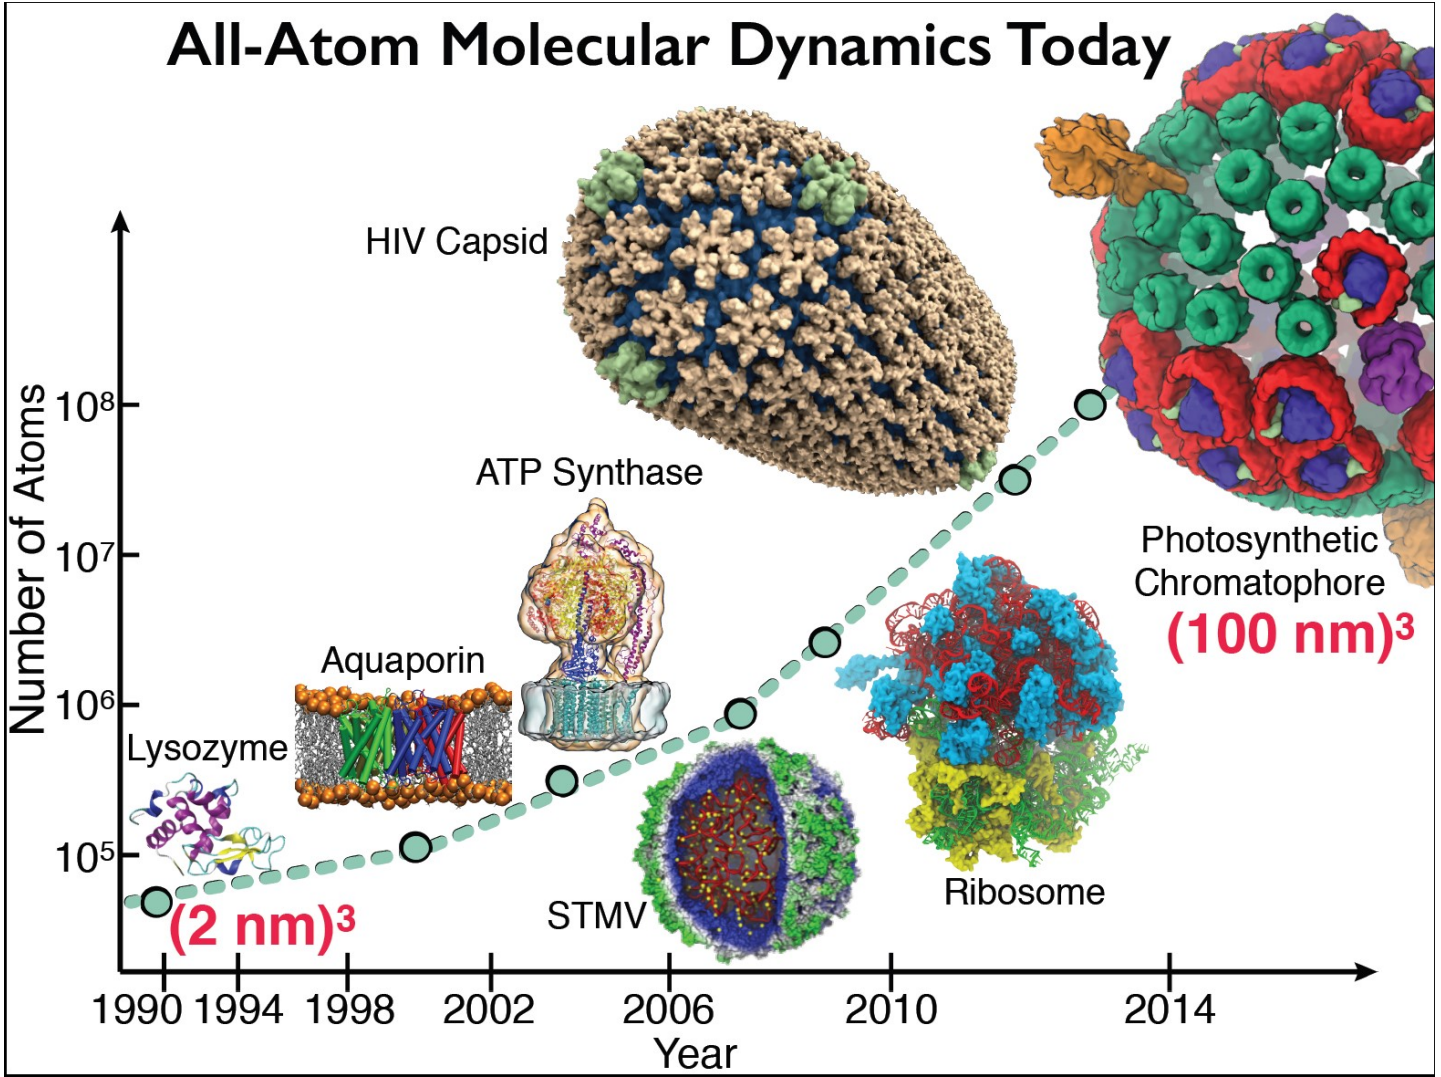
All-Atom Molecular Dynamics Today
HIV Capsid
ATP Synthase
Aquaporin
Lysozyme
STMV
Ribosome
Photosynthetic Chromatophore
(100 nm)3
(2 nm)3
Number of Atoms
10^8
10^7
10^6
10^5
1990
1994
1998
2002
2006
2010
2014
Year

NAMD Titan XK7 Performance August 2013

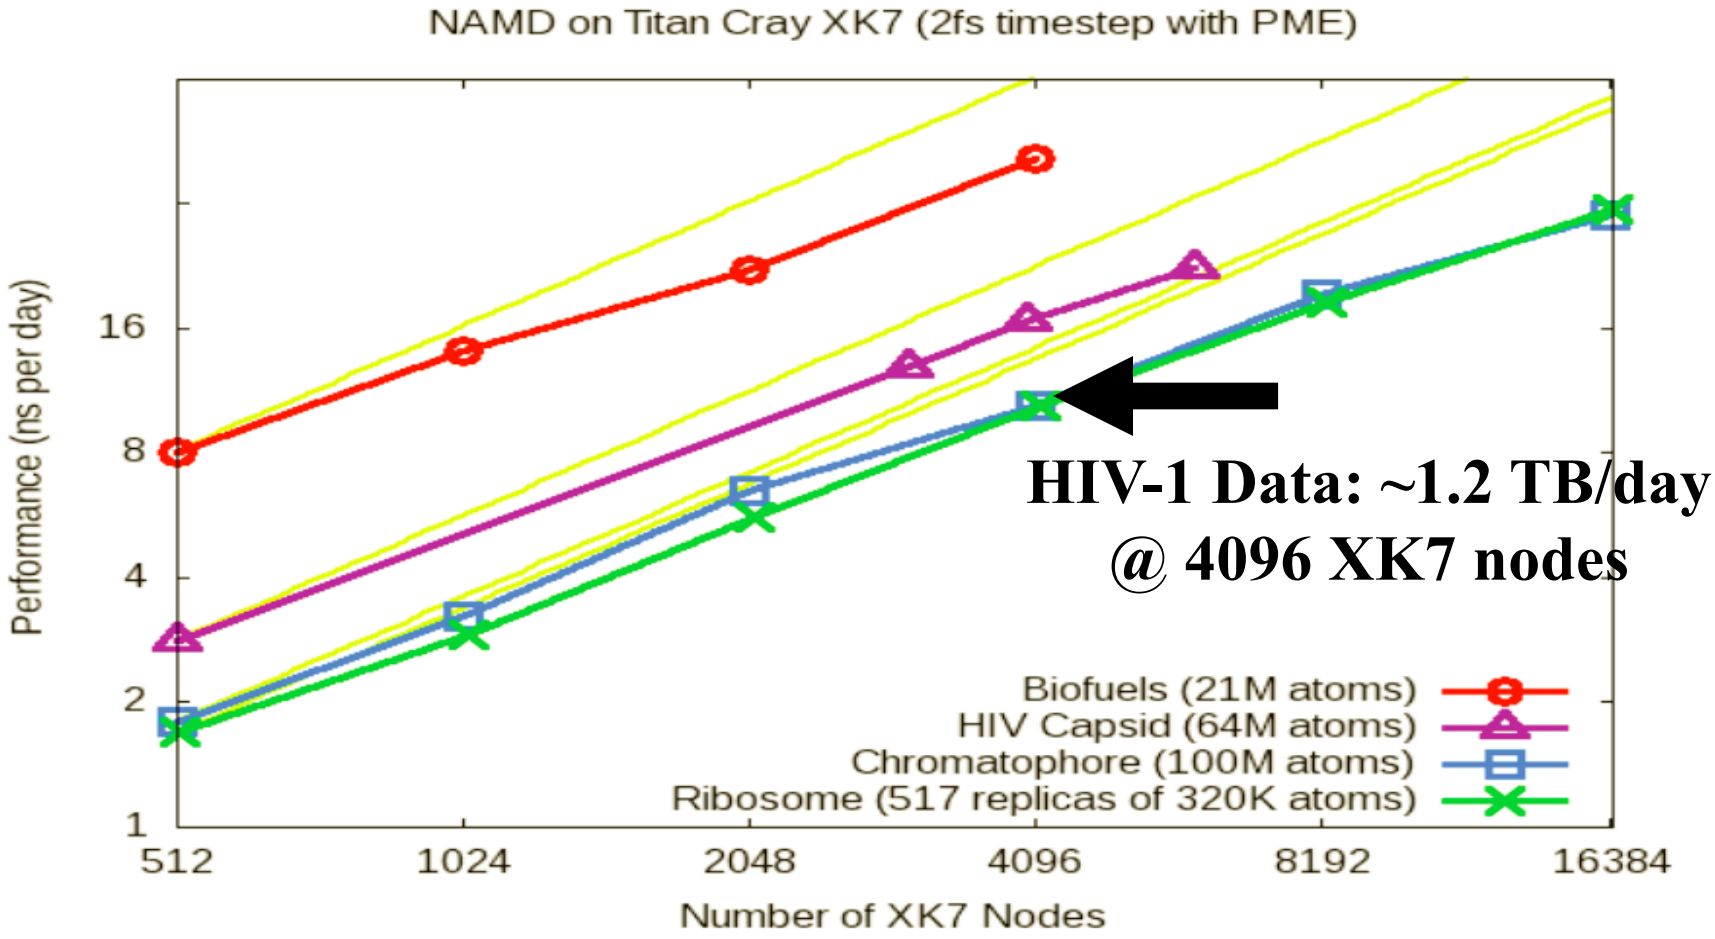

NAMD on Summit, May 2018

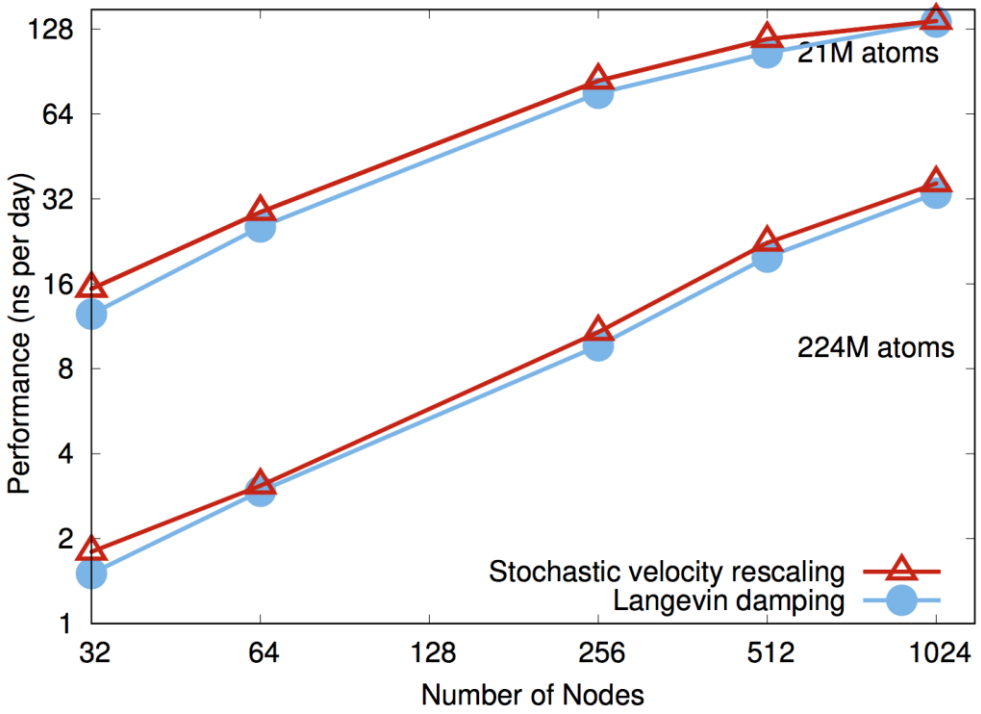

NAMD simulations can generate up to 10TB of output per day on 20% of Summit

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

Remote Visualization and Analysis

In-development: VMD w/ built-in remote visualization

- Access **large data** located anywhere in the world
- Enable capabilities not possible with conventional laptops or workstations
- VMD session available to any device, browser

Clusters, Supercomputers

Workstations,
Servers,
Cloud

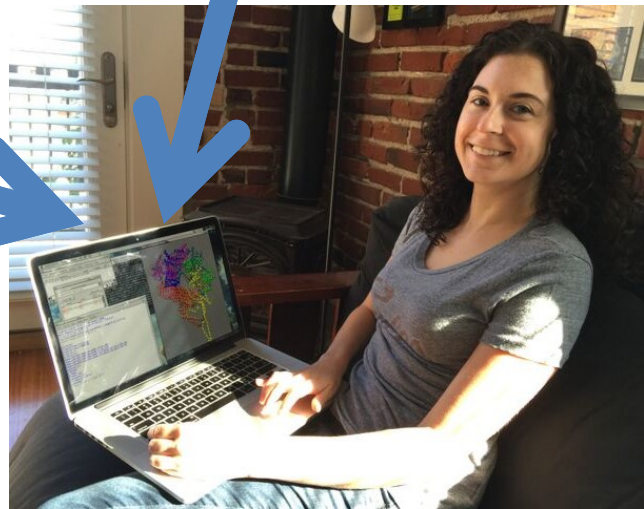

Chemical Visualization of Human Pathogens: the Retroviral Capsids.
J.R. Perilla, B.-C. Goh, J.E. Stone, K. Schulten. SC'15 Visualization and Data Analytics Showcase, 2015.
Atomic Detail Visualization of Photosynthetic Membranes with GPU-Accelerated Ray Tracing.
J.E. Stone, …, K. Schulten, J. Parallel Computing, 55:17-27, 2016.
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.

VMD “QuickSurf” Representation, Ray Tracing

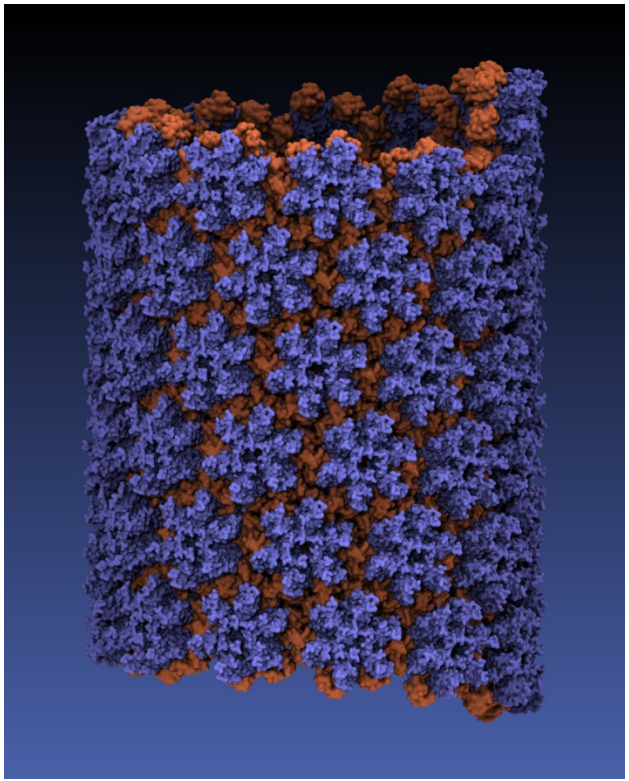

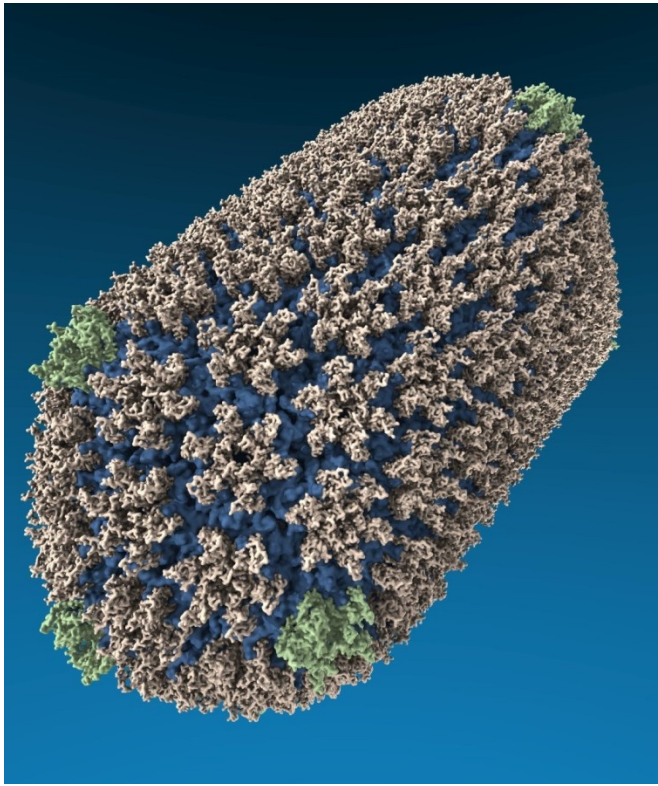

All-atom HIV capsid simulations w/ up to 64M atoms on Blue Waters

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

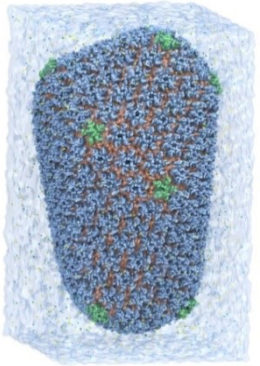

New VMD TachyonL-OptiX on XK7 vs. Tachyon on XE6:
K20X GPUs yield **up to twelve times** geom+ray tracing speedup

Ray Tracer Version	Node Type and Count	Script Load	State Load	Geometry + Ray Tracing	Total Time
New TachyonL-OptiX	64 XK7 Tesla K20X GPUs	2 s	39 s	435 s	476 s
New TachyonL-OptiX	**128 XK7 Tesla K20X GPUs**	**3 s**	**62 s**	**230 s**	**295 s**
TachyonL-OptiX [1]	64 XK7 Tesla K20X GPUs	2 s	38 s	655 s	695 s
TachyonL-OptiX [1]	128 XK7 Tesla K20X GPUs	4 s	74 s	331 s	410 s
TachyonL-OptiX [1]	**256 XK7 Tesla K20X GPUs**	**7 s**	**110 s**	**171 s**	**288 s**
Tachyon [1]	**256 XE6 CPUs**	7 s	160 s	**1,374 s**	**1,541 s**
Tachyon [1]	512 XE6 CPUs	13 s	211 s	808 s	1,032 s

[1] GPU-Accelerated Molecular Visualization on Petascale Supercomputing Platforms.
J. E. Stone, K. L. Vandivort, and K. Schulten. UltraVis'13: Proceedings of the 8th International Workshop on Ultrascale Visualization, pp. 6:1-6:8, 2013.

VMD Chromatophore Rendering on Blue Waters

- New representations, GPU-accelerated molecular surface calculations, memory-efficient algorithms for huge complexes
- VMD GPU-accelerated ray tracing engine w/ OptiX+CUDA+MPI+Pthreads
- ***Each revision:*** 7,500 frames render on ~96 Cray XK7 nodes in 290 node-hours, 45GB of images prior to editing

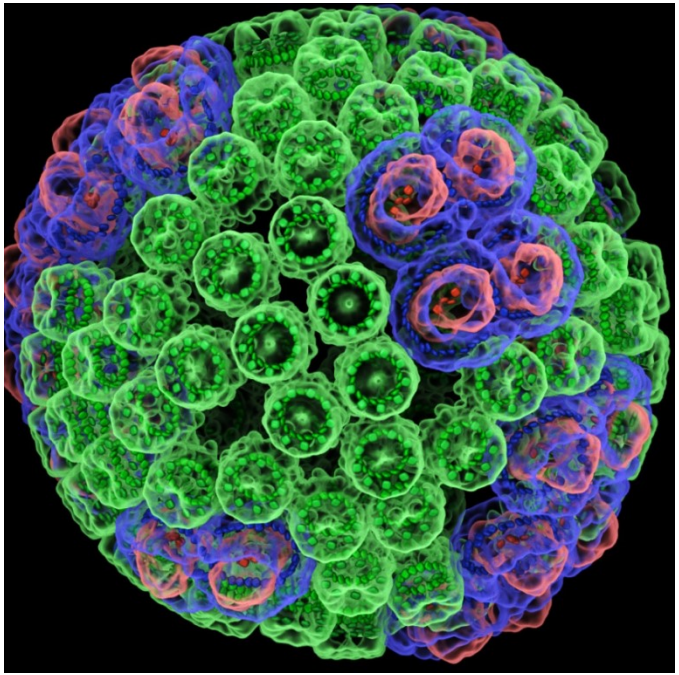

GPU-Accelerated Molecular Visualization on Petascale Supercomputing Platforms.
J. E. Stone, K. L. Vandivort, and K. Schulten. UltraVis'13, 2013.

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.
Winner of the SC'14 Visualization and Data Analytics Showcase

VMD w/ OptiX 5

- Interactive RT on laptops, desktops, and cloud
- 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 on Quadro GP100, Tesla P100**
- **VMD+OptiX 5, NVIDIA NGC container: https://ngc.nvidia.com/registry/**
- **In-progress:**
  - **OptiX denoising support: fast turnaround w/ AO, DoF, etc**
  - **Denoising to enable practical use of path tracing in VMD**

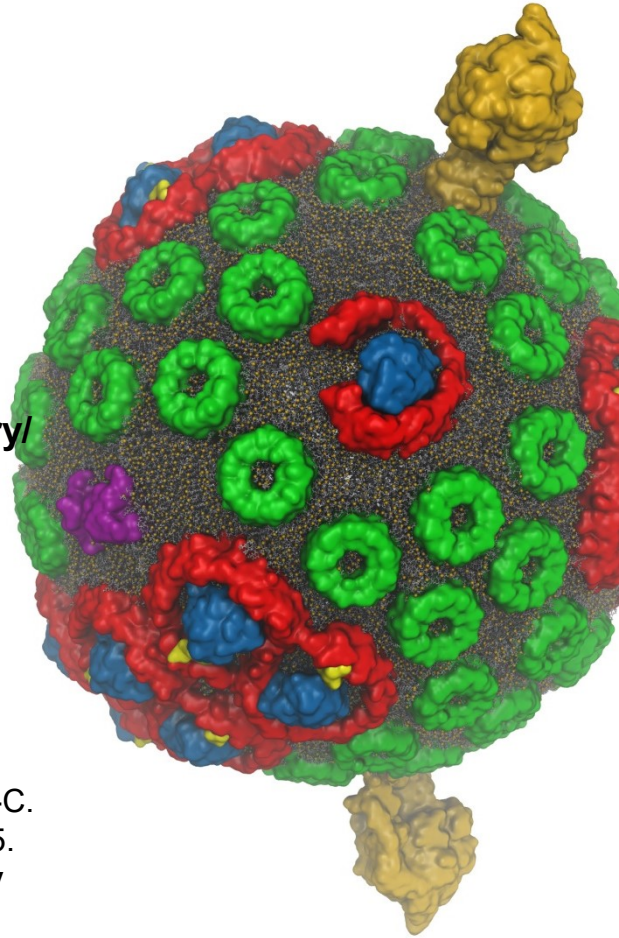

VMD/OptiX GPU Ray Tracing of all-atom Chromatophore w/ lipids.

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.

IBM AC922 w/ 6 GPUs

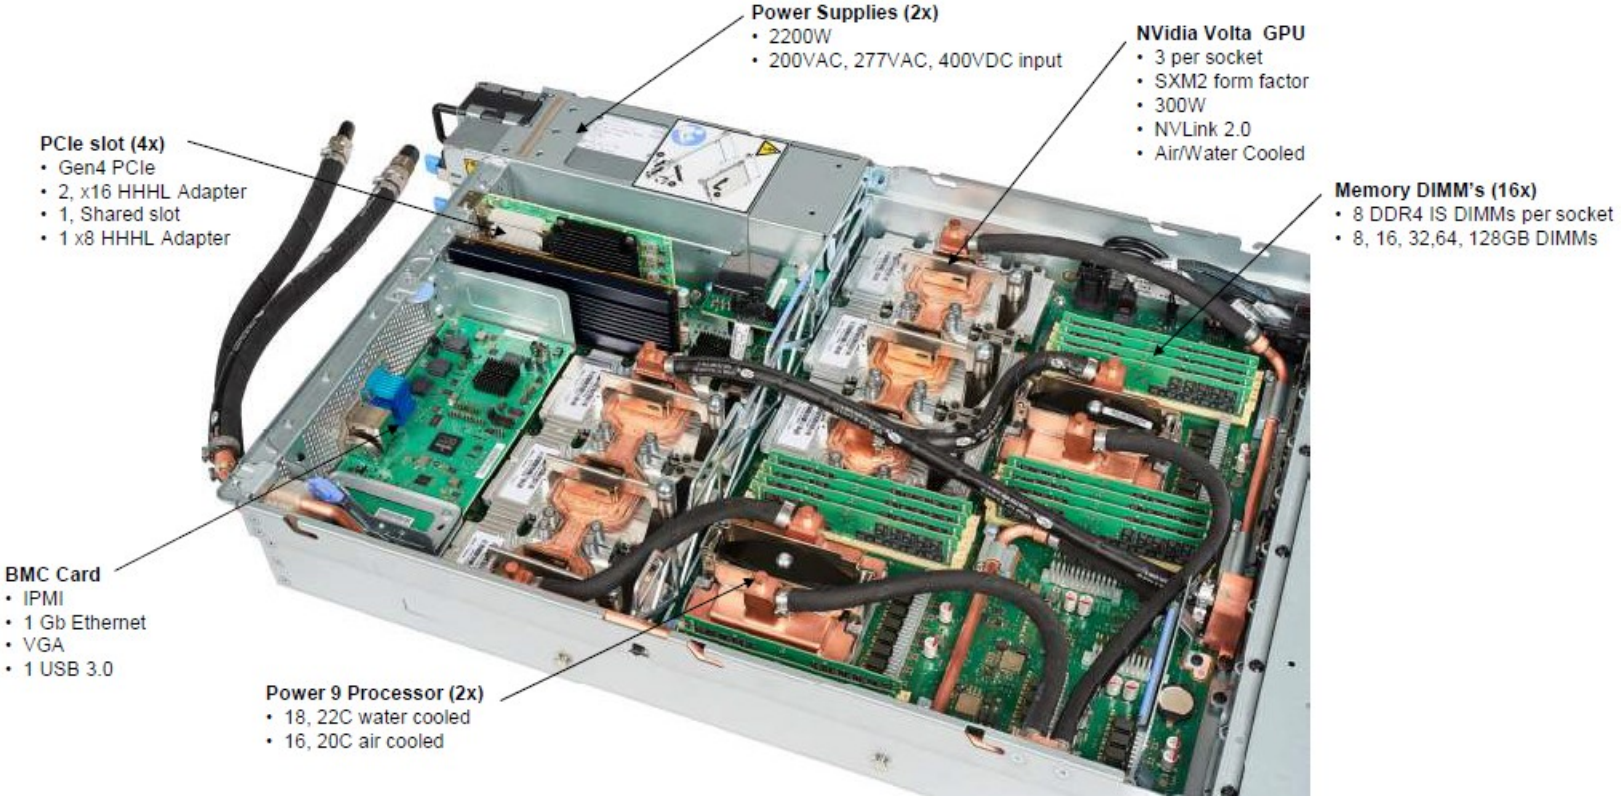

IBM AC922 Summit Node

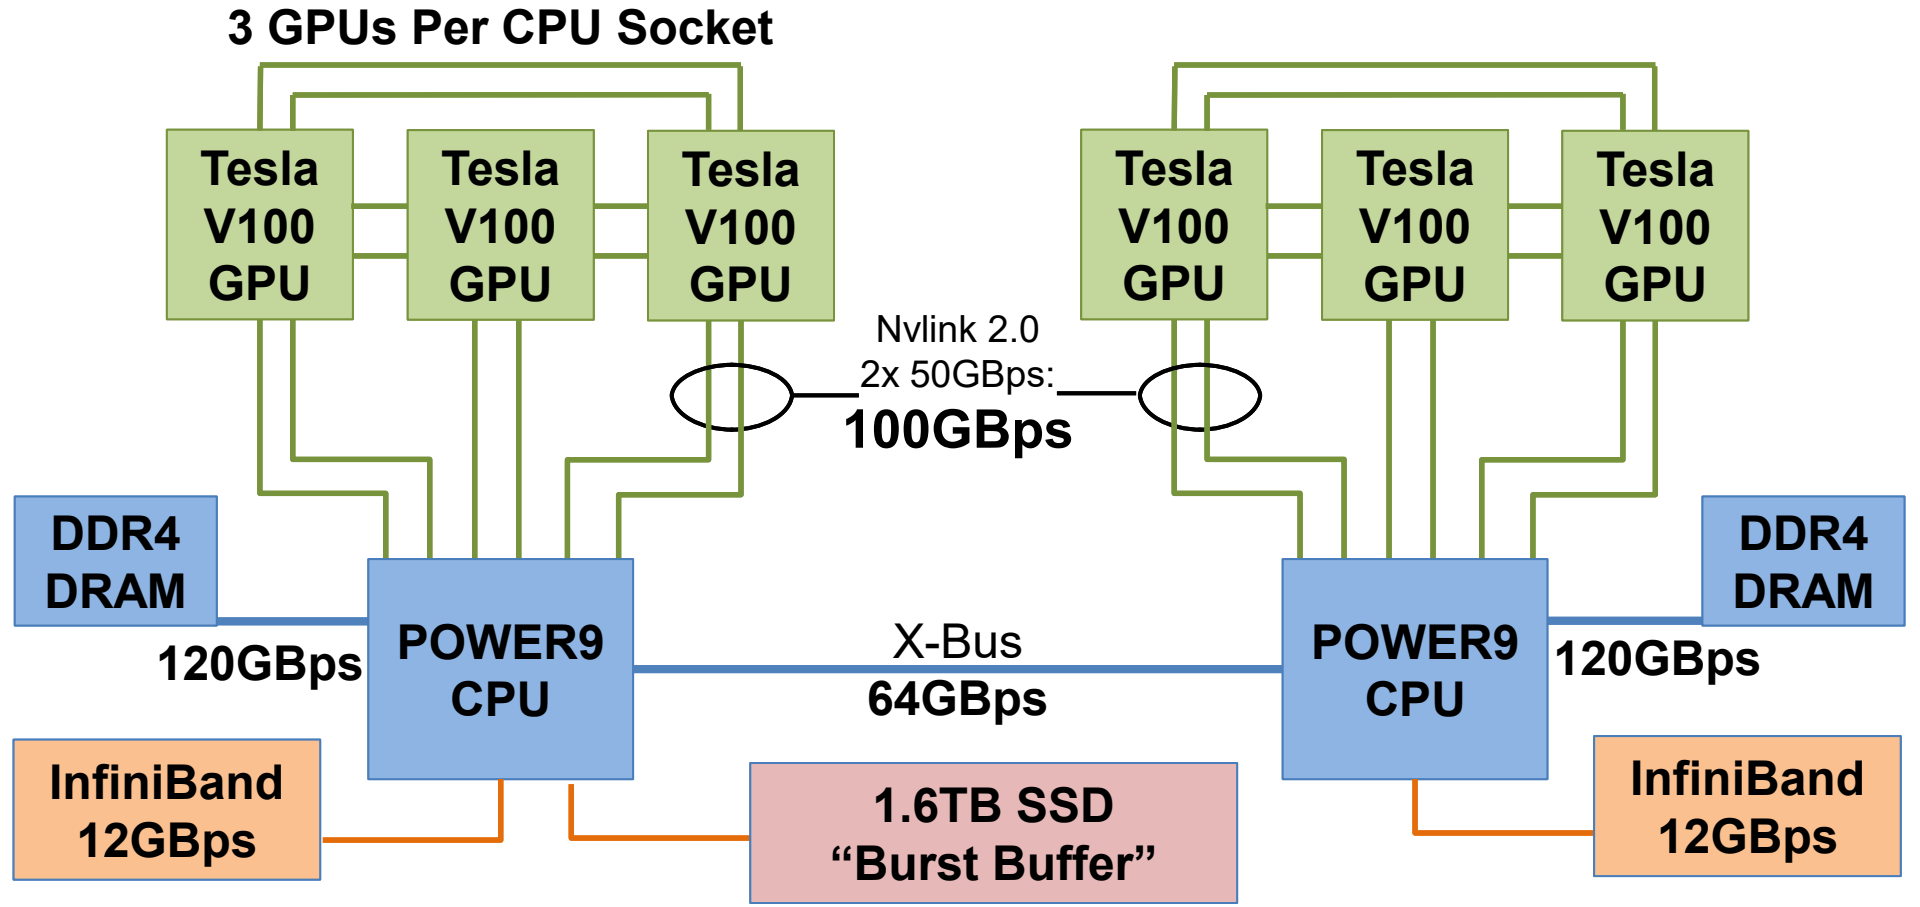

VMD Off-Screen Rendering w/ EGL

- Containers+Cloud+Workstations with recent NVIDIA drivers
- VMD on HPC systems w/ latest GPUs:
  - Cray XC50, CSCS Piz Daint
  - **ORNL Summit in progress now**
  - IBM OpenPOWER, drivers 375.66 and later support both GLX and EGL

OpenGL: GLX vs. EGL

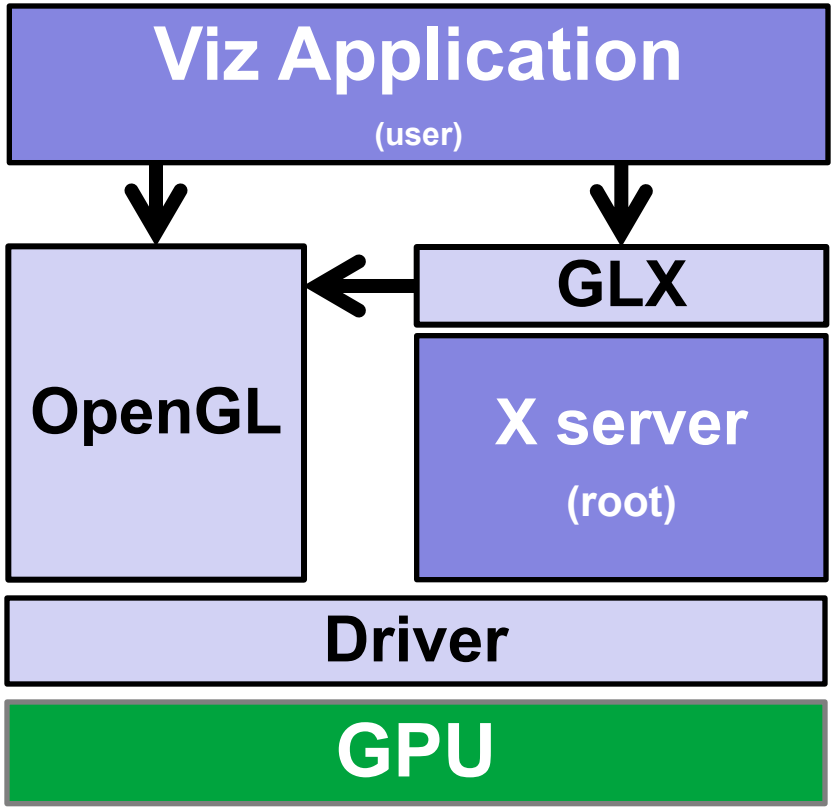

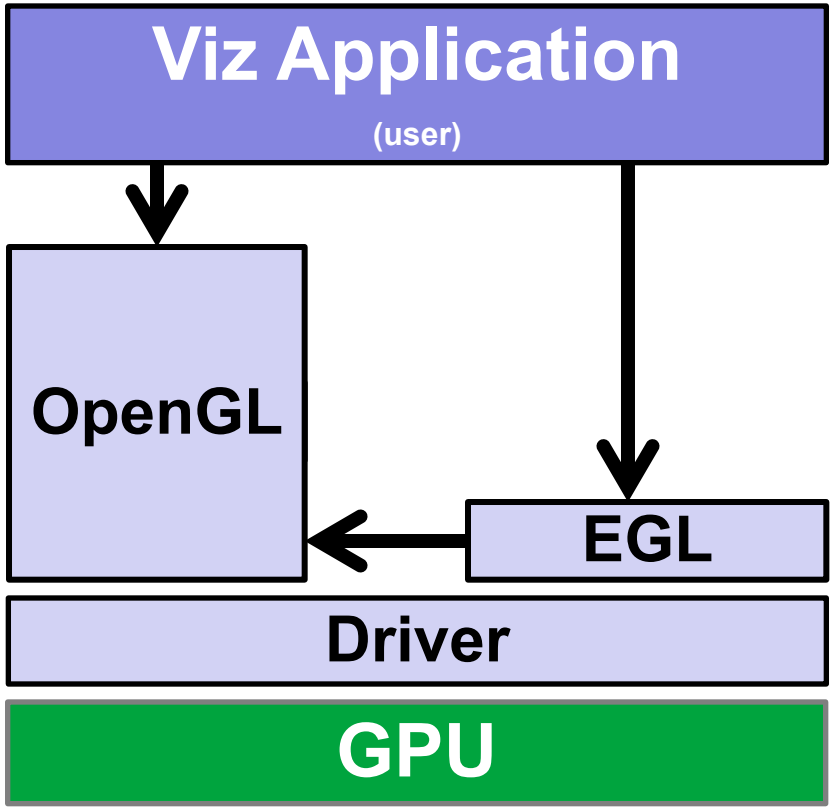

VMD EGL Rendering: Supports full VMD GLSL shading features
Vulkan support coming soon...

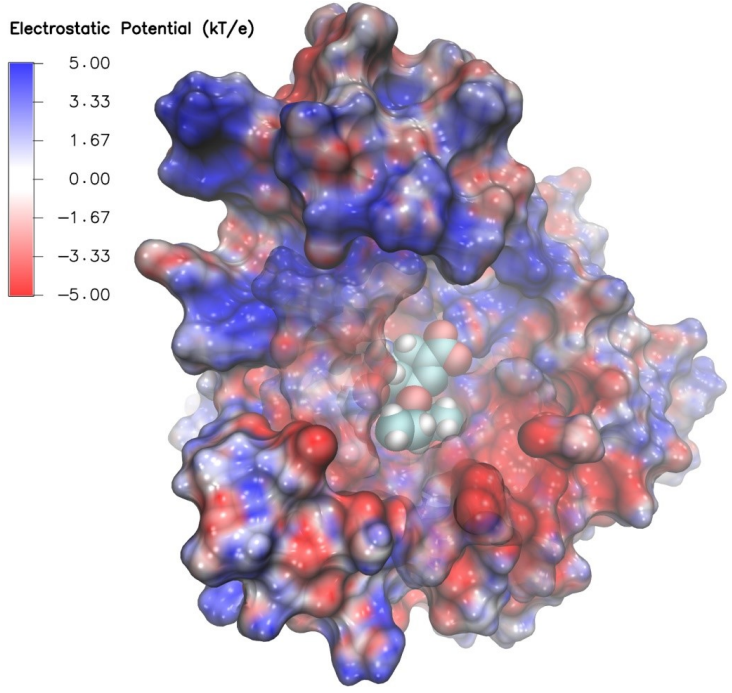

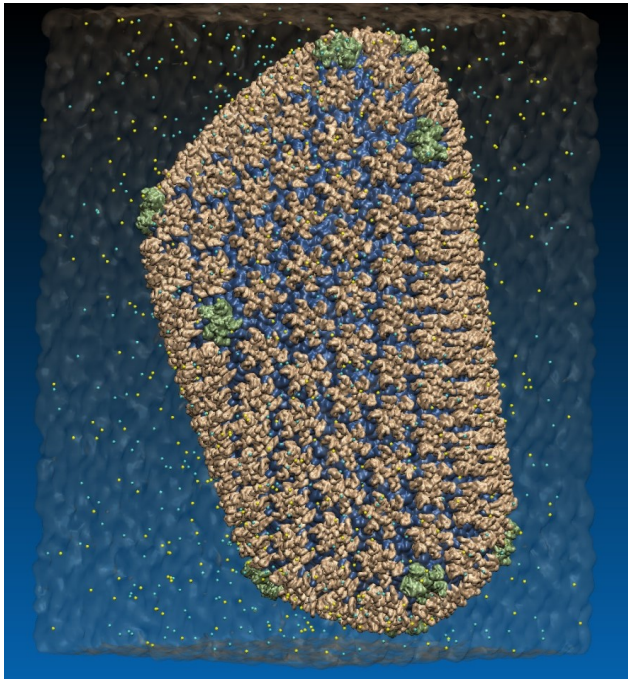

Swine Flu A/H1N1 neuraminidase bound to Tamiflu

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

VMD EGL Performance on Amazon EC2 Cloud

MPI Ranks	EC2 "G2.8xlarge" GPU Instances	HIV-1 movie rendering time (sec), (I/O %) 3840x2160 resolution
1	1	626s (10% I/O)
2	1	347s (19% I/O)
4	1	221s (31% I/O)
8	2	141s (46% I/O)
16	4	107s (64% I/O)
32	8	90s (76% I/O)

Performance at 32 nodes reaches ~48 FPS

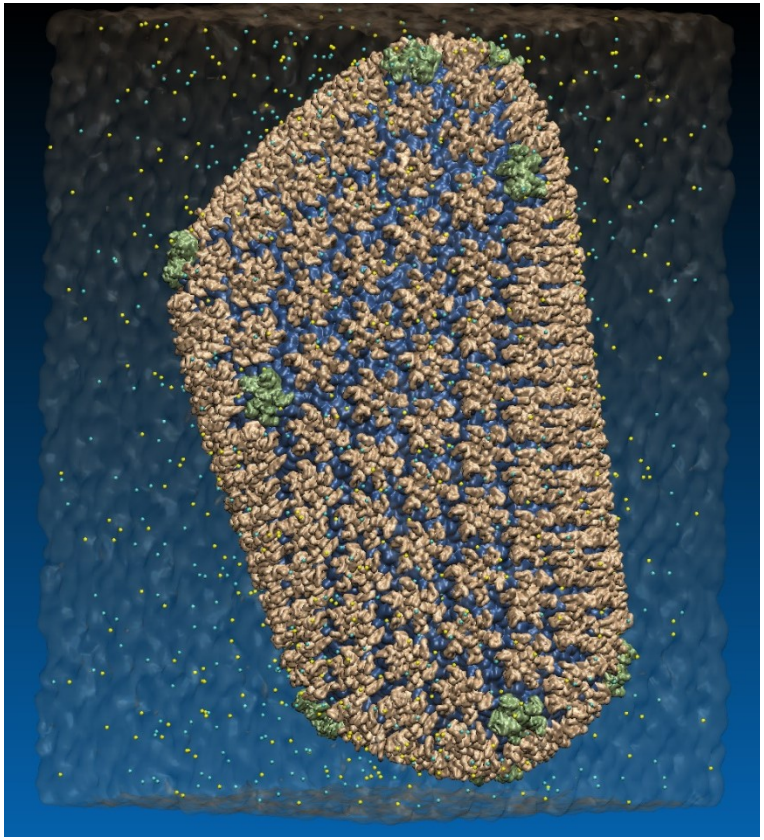

64M atom HIV-1 capsid simulation rendered via EGL

High performance molecular visualization: In-situ and parallel rendering with EGL. J. E. Stone, P. Messmer, R. Sisneros, and K. Schulten. *2016 IEEE International Parallel and Distributed Processing Symposium Workshop (IPDPSW)*, pp. 1014-1023, 2016.

VMD OptiX/EGL NGC Container

- **https://ngc.nvidia.com/registry/**
- **CUDA-accelerated** viz+analysis
- **EGL off-screen rendering** – no windowing system needed
- **OptiX high-fidelity GPU ray tracing** engine built in
- All dependencies included
- **Easy to deploy on a wide range of GPU accelerated platforms**

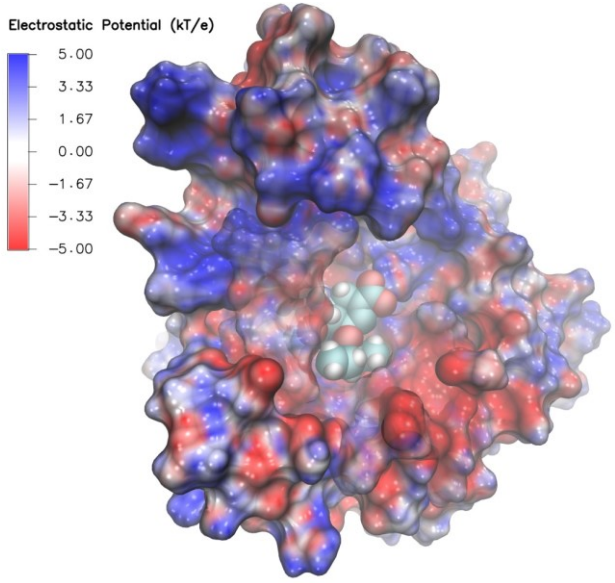

High performance molecular visualization: In-situ and parallel rendering with EGL. J. E. Stone, P. Messmer, R. Sisneros, and K. Schulten. *2016 IEEE International Parallel and Distributed Processing Symposium Workshop (IPDPSW)*, pp. 1014-1023, 2016.

Preparation, Visualization, Analysis of All-Atom Cell-Scale Simulations

- Parallel analysis, visualization w/ MPI
- Support for large memory (TB), up to **2 billion atoms per "molecule" now**
- Interactive rasterization w/ Vulkan, EGL
- **Interactive ray tracing on CPUs and GPUs**

- 200 nm spherical envelope
- Membrane with ~50% occupancy by proteins
- 63M atoms in envelope model

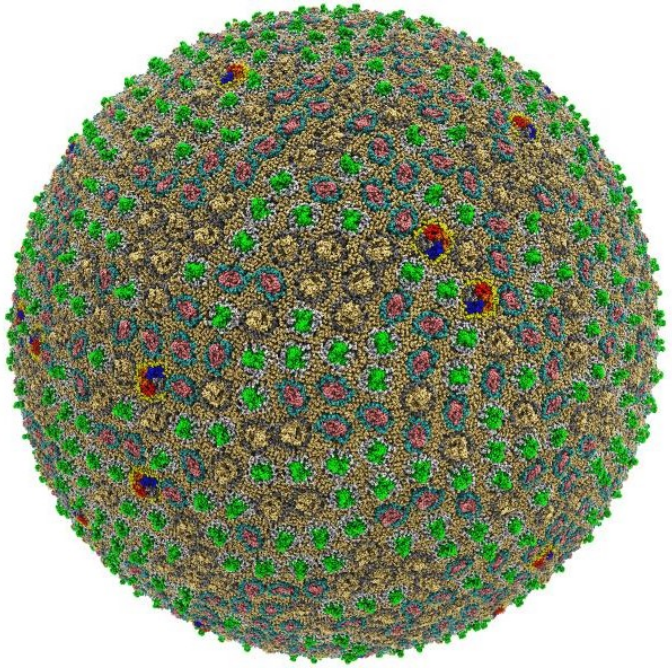

Atomic Detail Visualization of Photosynthetic Membranes with GPU-Accelerated Ray Tracing. J.E. Stone, …, K. Schulten, J. Parallel Computing, 55:17-27, 2016.

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

Next Generation: Simulating a Proto-Cell

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

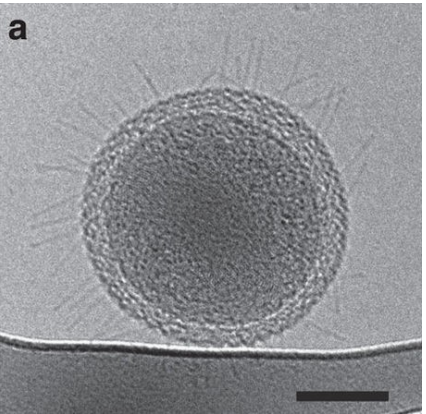

Cryo-ET image of ultra-small bacteria (scale bar 100nm) Luef et al. Nature Comm., 6:6372, 2015.

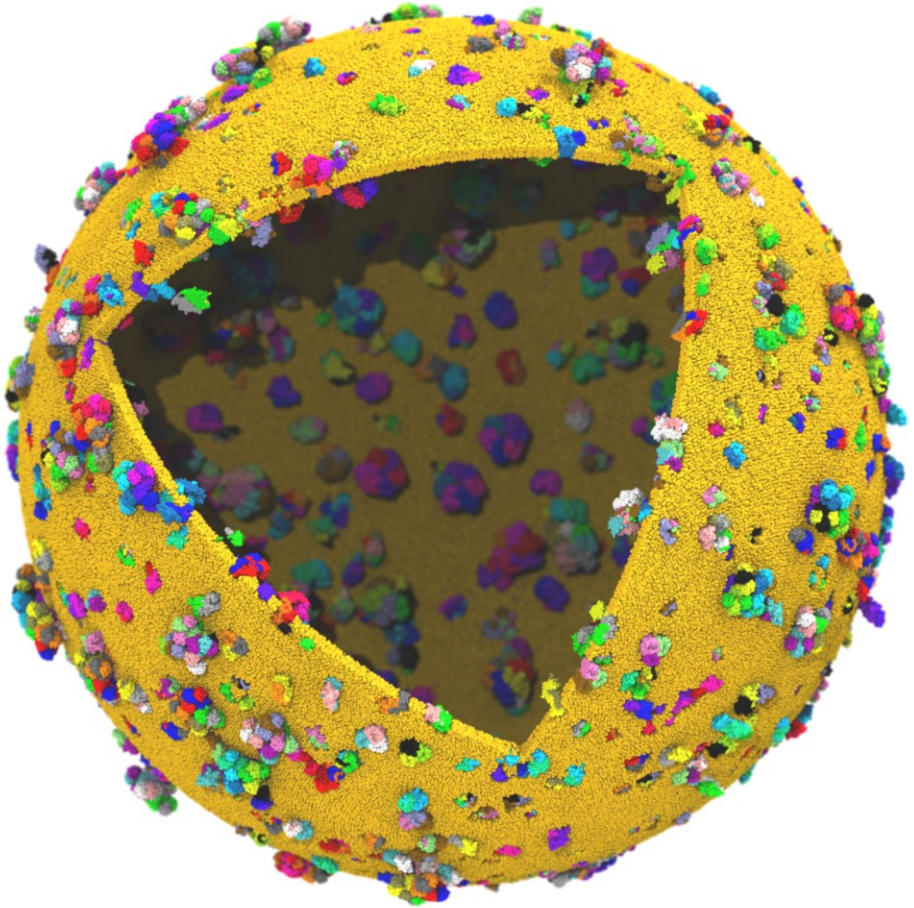

Proto-Cell Data Challenges

- **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**

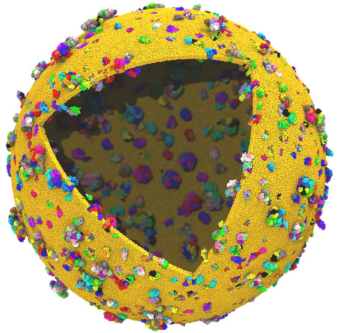

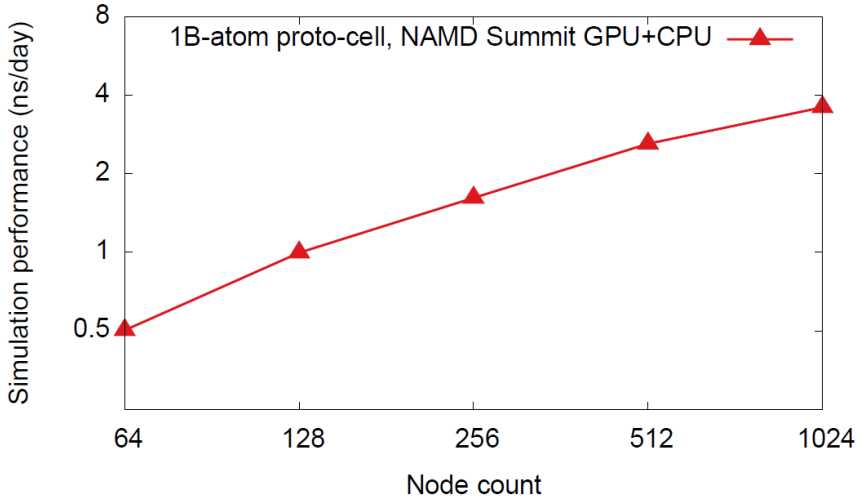

VMD Tesla P100 Performance for C_{60} Molecular Orbitals, 516x519x507 grid

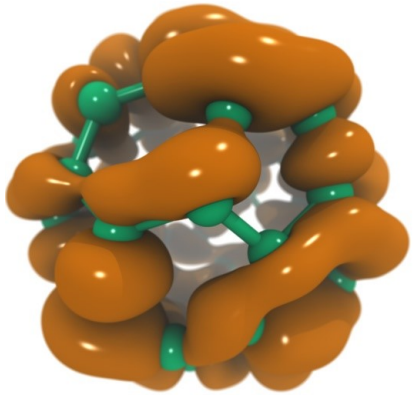

Hardware platform	Runtime,	Speedup
IBM Power8 (2 socket) (ORNL 'crest') [1]	8.03s,	0.4x
Intel Xeon E5-2660v3 (2 socket) [1]	7.14s,	0.5x
IBM Power8 (ORNL 'crest') + 1x Tesla K40 [1]	3.49s,	**1.0x**
Intel Xeon E5-2698v3 + 1x Tesla P100	**1.35s,**	**2.5x**
IBM Power8 "Minsky" + 1x Tesla P100	**1.09s,**	**3.3x**
IBM Power8 (ORNL 'crest') + 4x Tesla K40 [1]	0.91s,	3.8x
Intel Xeon E5-2698v3 + 4x Tesla P100	**0.37s,**	**9.4x**
IBM Power8 "Minsky" + 4x Tesla P100	**0.30s,**	**11.6x**

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.

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

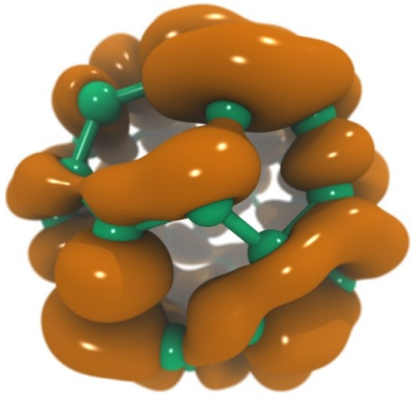

Hardware platform		Runtime,	Speedup
IBM Power8 (ORNL 'crest') + 1x Tesla K40 [1]		3.49s,	**1.0x**
Intel Xeon E5-2697Av4	**+ 1x Tesla V100**	**0.610s,**	**5.7x**
Intel Xeon E5-2697Av4	**+ 2x Tesla V100**	**0.294s,**	**11.8x**
Intel Xeon E5-2697Av4	**+ 3x Tesla V100**	**0.220s,**	**15.9x**
IBM Power9 "Newell"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

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

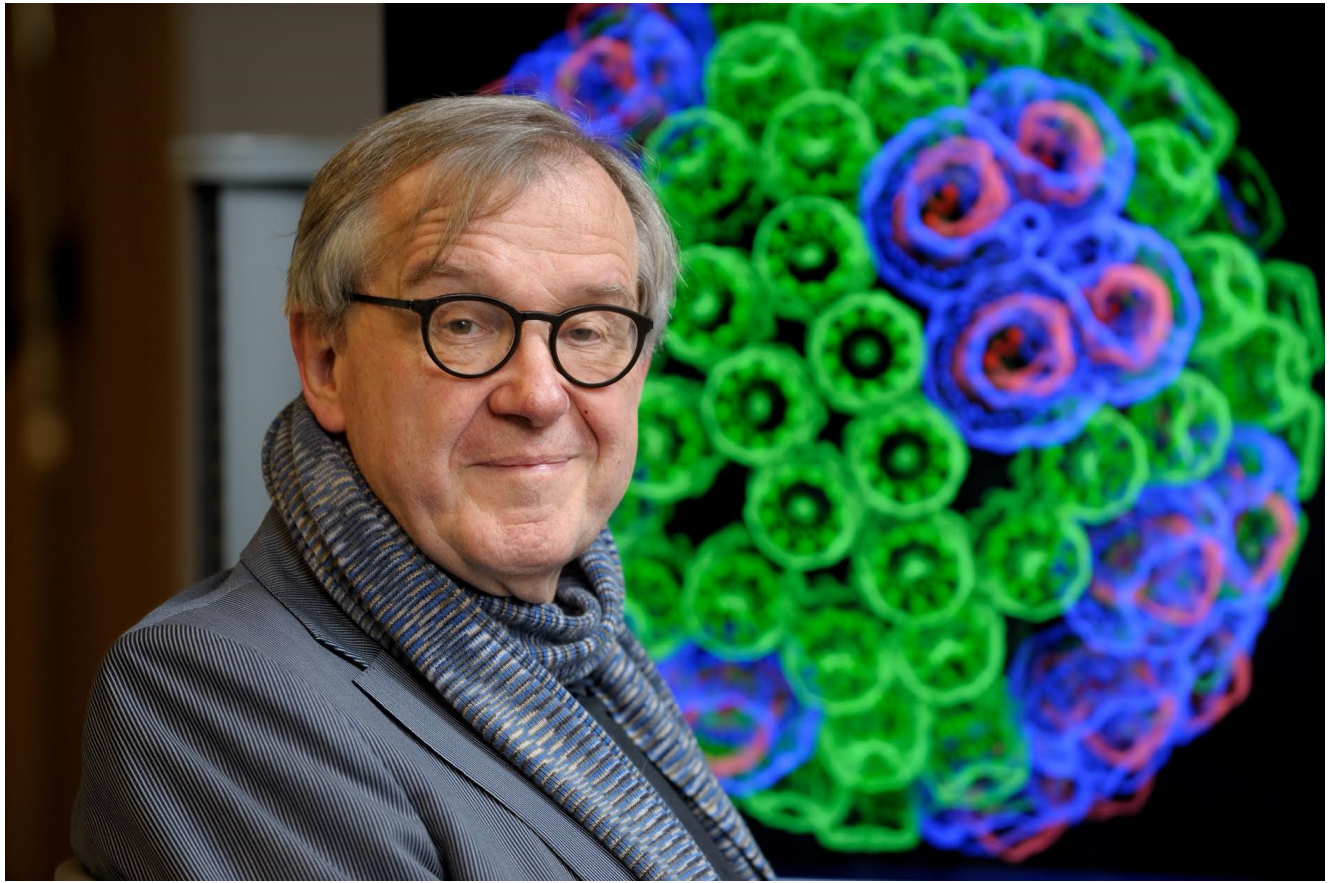

“When I was a young man, my goal was to look with mathematical and computational means at the inside of cells, one atom at a time, to decipher how living systems work. That is what I strived for and I never deflected from this goal.” – Klaus Schulten