Optimizing NAMD and VMD for the IBM Power9 Platform

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/gpu/

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

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

IBM Power User's Group, Supercomputing 2018,

Omni Hotel 3rd Floor Ballroom, Thursday November 15th, 2018

NAMD & VMD: Computational Microscope

Enable researchers to investigate systems described at the atomic scale

NAMD - molecular dynamics simulation

VMD - visualization, system preparation and analysis

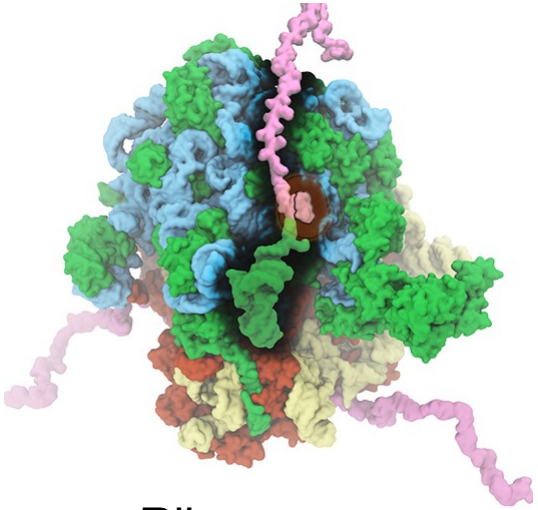

Ribosome

Neuron

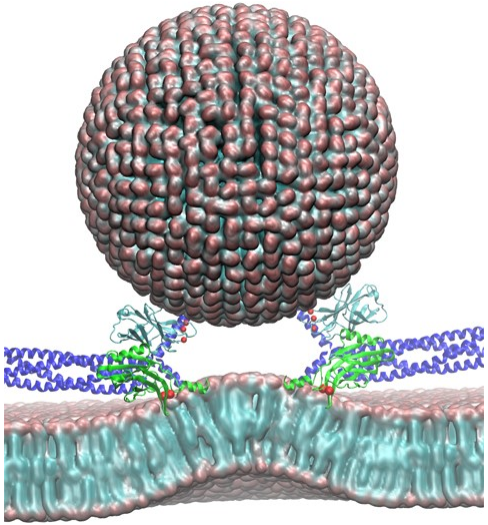

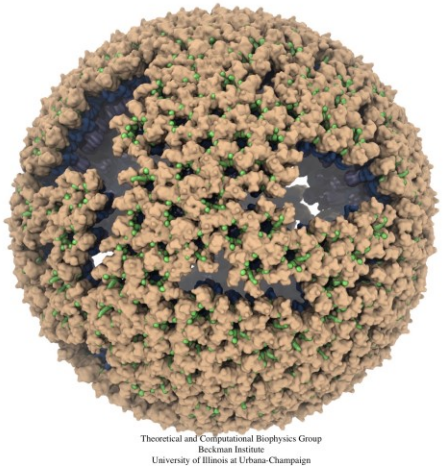

Virus Capsid

NAMD+VMD: Building A Next Generation Modeling Platform

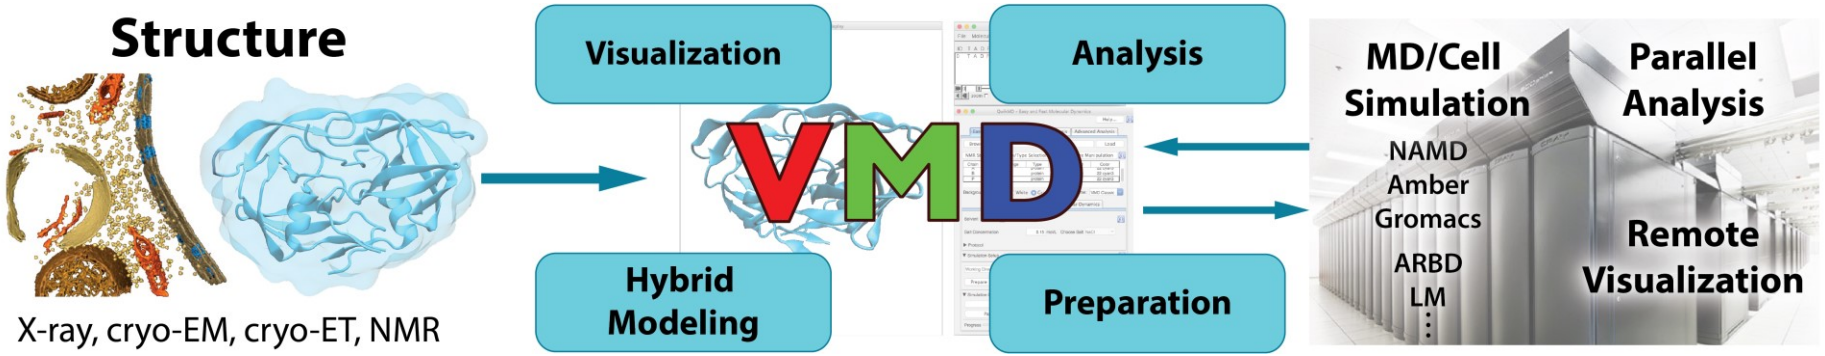

- Provide tools for preparation, simulation, visualization, and analysis
  - Reach cell-scale modeling w/ all-atom MD, coarse graining, Lattice Microbes
  - Improved performance, visual fidelity, exploit advanced technologies (GPUs, VR HMDs)
- Enable hybrid modeling and computational electron microscopy
  - Load, filter, process, interpret, visualize multi-modal structural information
- Connect key software tools to enable state-of-the-art simulations
  - Support new data types, file formats, software interfaces
  - Openness, extensibility, and interoperability are our hallmarks
  - Reusable algorithms made available in NAMD, for other tools

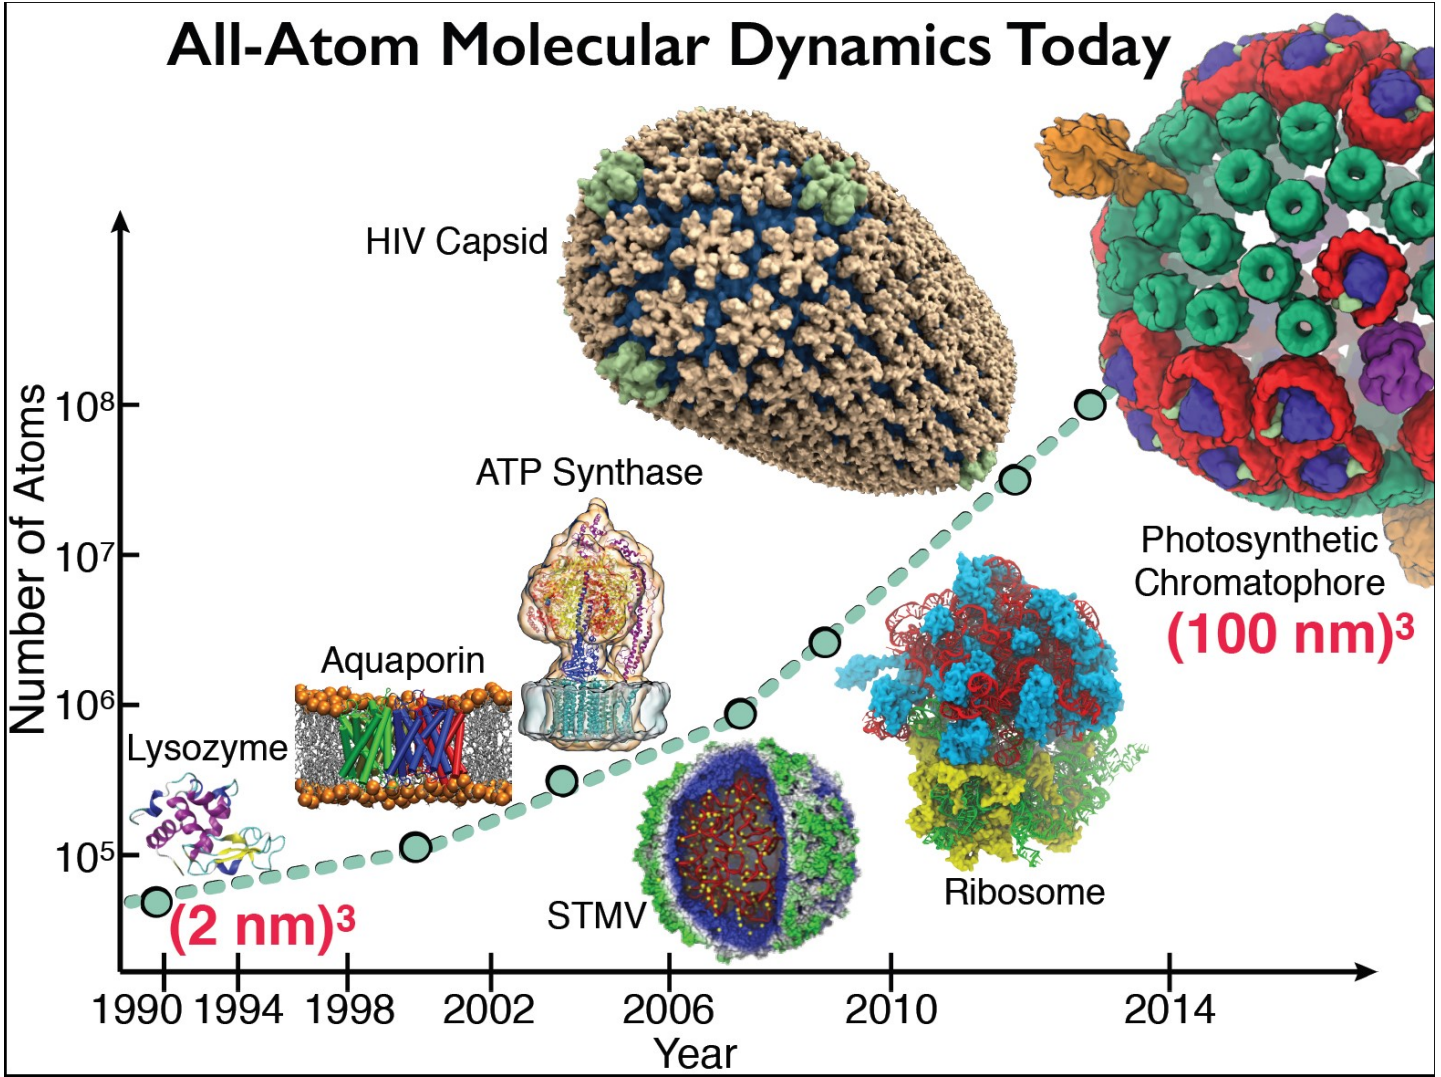
All-Atom Molecular Dynamics Today
HIV Capsid
ATP Synthase
Aquaporin
Lysozyme
STMV
Ribosome
Photosynthetic Chromatophore
(100 nm)3
(2 nm)3
Number of Atoms
10^8
10^7
10^6
10^5
1990
1994
1998
2002
2006
2010
2014
Year

History of NAMD and VMD on POWER

- NAMD + VMD have supported IBM POWER hardware since 1998!
  - Originally ported in big-endian mode under AIX 4.x
- 2016: Adapted to POWER8 w/ Linux in little-endian mode w/ P100 GPUs:
  - New NAMD GPU kernels improved overall P8+P100 performance [1]
  - Used VSX instructions for hand-coded and vectorized kernels [1]
  - Supported CUDA 7.x [1], and **CUDA 8.x w/ P100 and NVLink**
  - **First VMD support for OpenGL GLX+EGL on POWER ppc64le**

[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.

IBM AC922 Summit Node

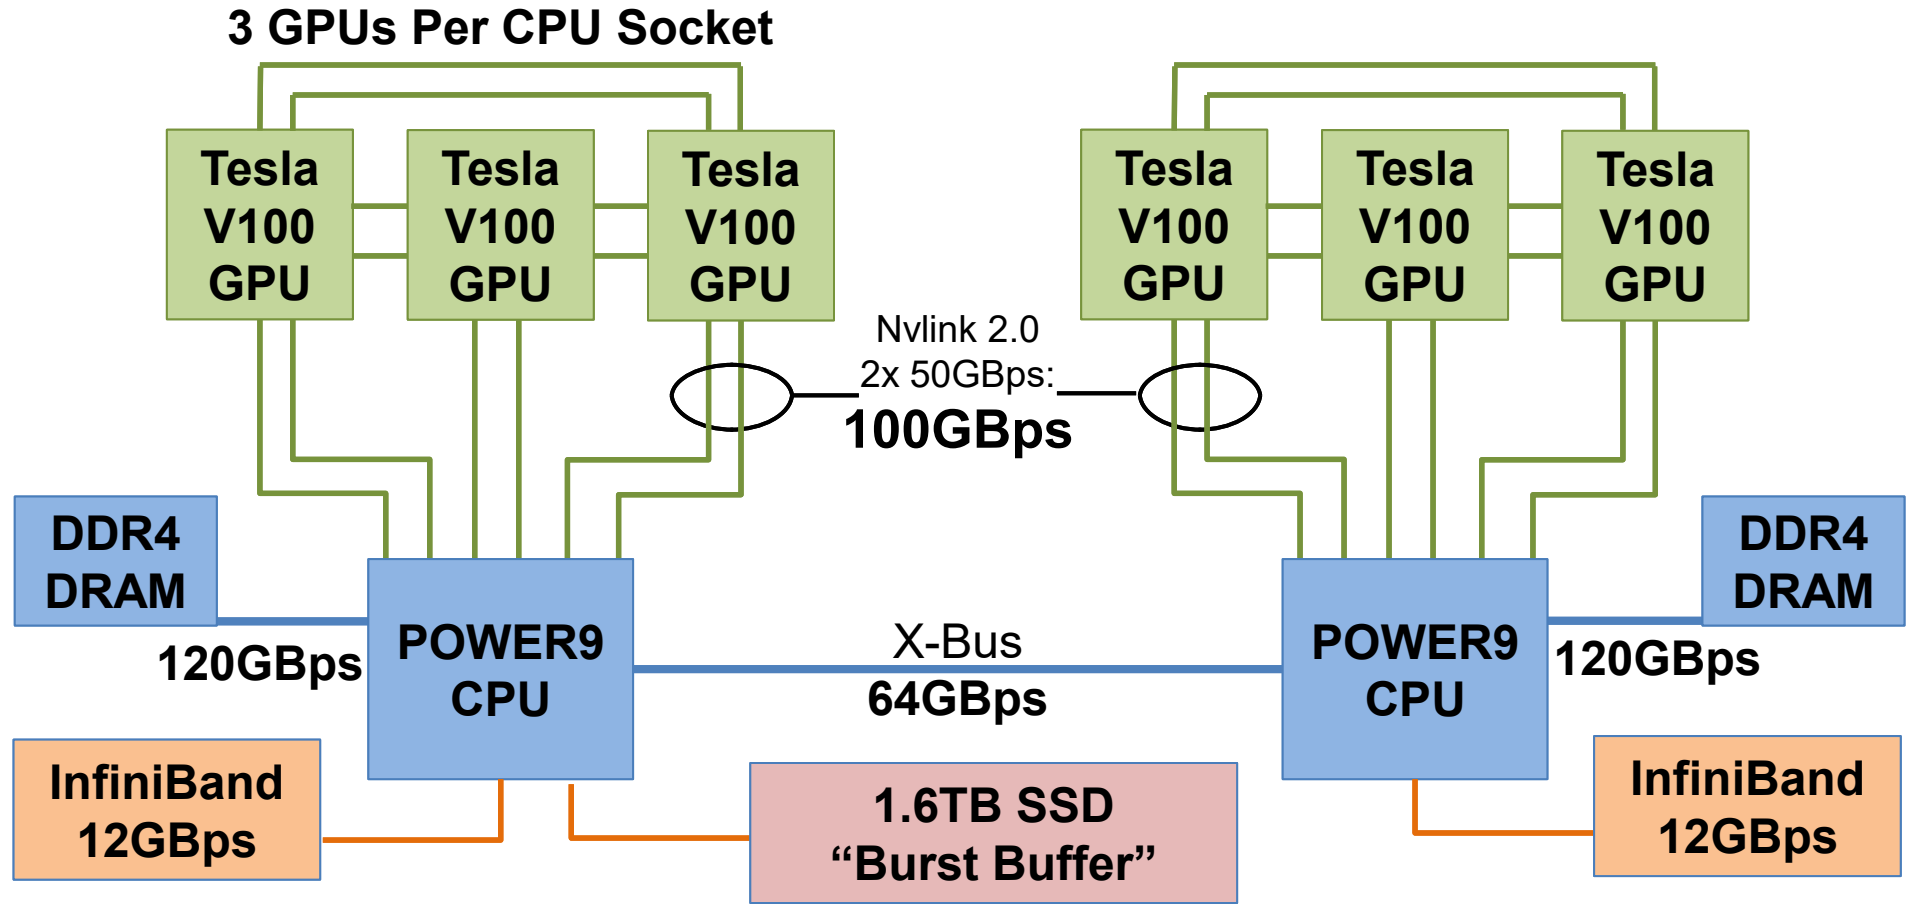

Earliest NAMD Runs on Summit

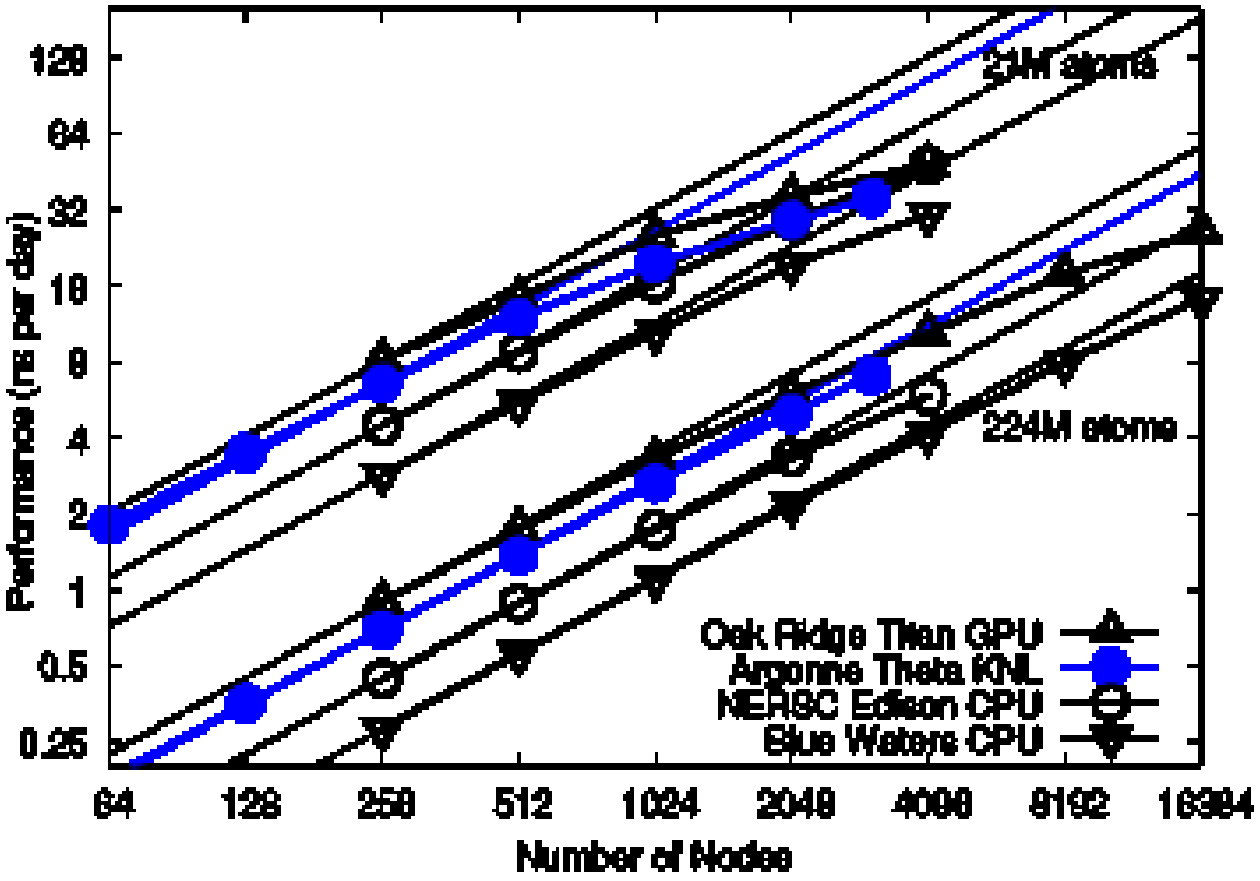

Earliest NAMD Runs on Summit

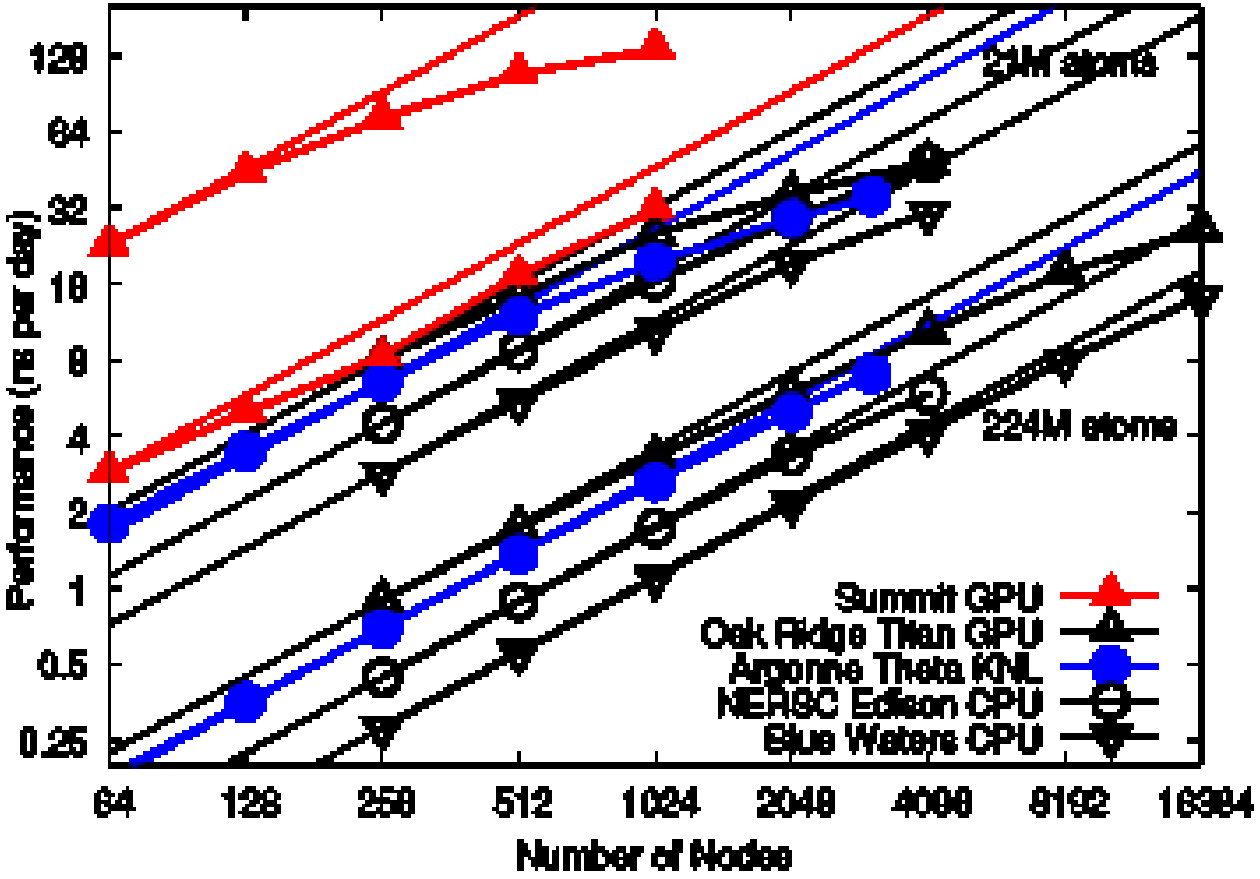

NAMD on Summit, May 2018:
~20% Performance Increase

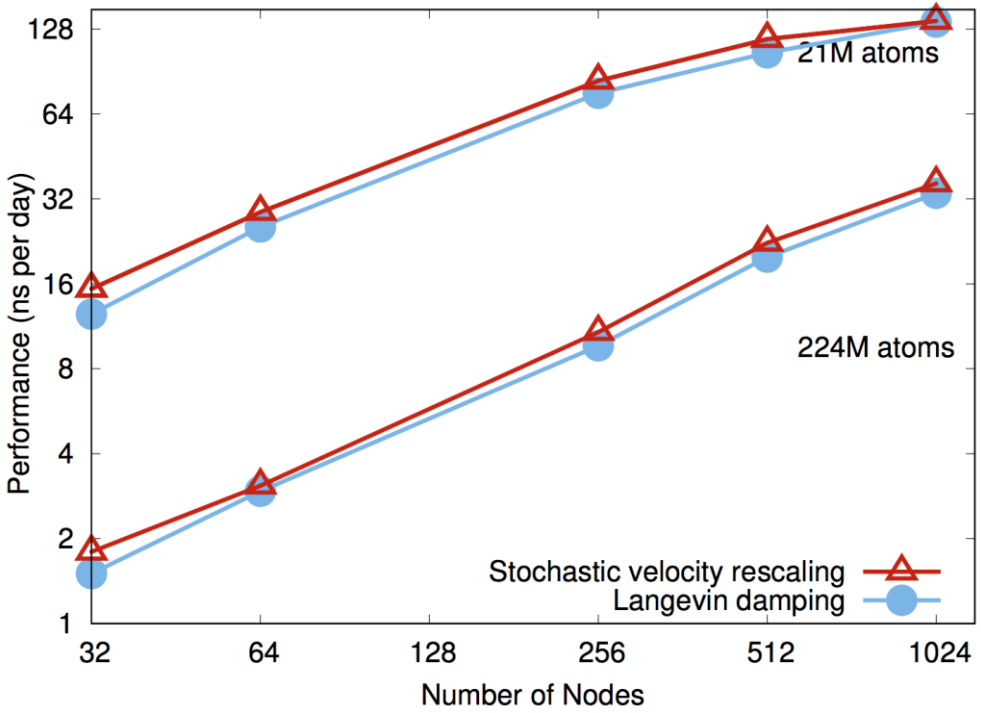

NAMD simulations can generate up to 10TB of output per day on 20% of Summit

NAMD 2 Billion Atom Benchmark on 20% of Summit

"Scalable Molecular Dynamics with NAMD on the Summit System"
IBM Journal of Research and Development, 2018. ***(In press)***

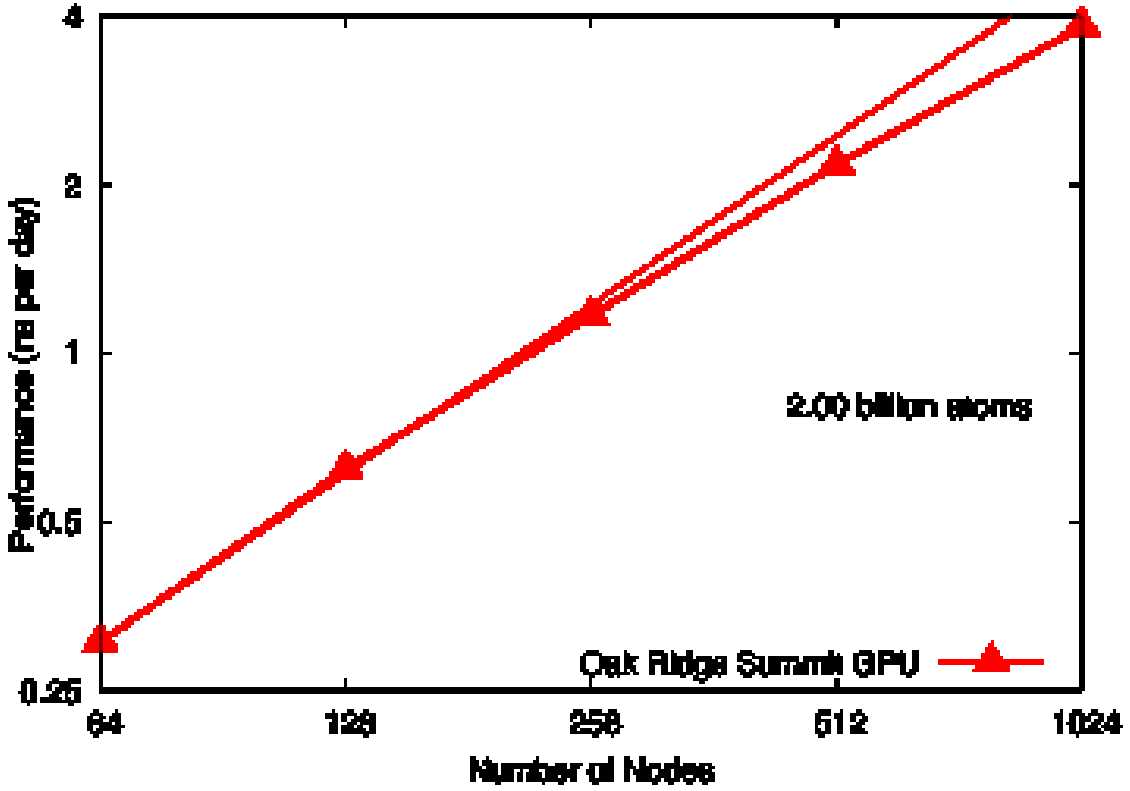

Density Map Segmentation

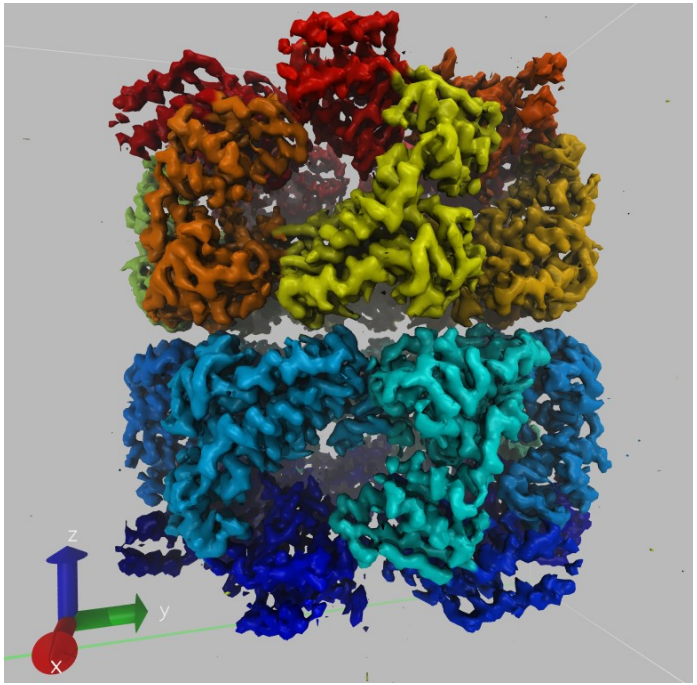

VMD GPU-accelerated density map segmentation of GroEL

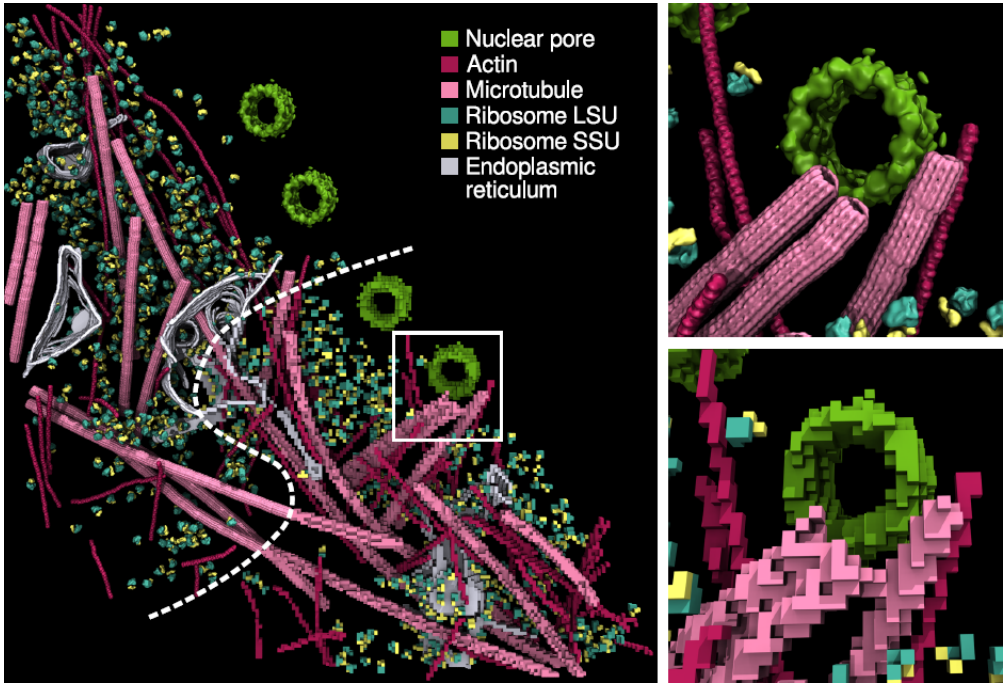

Earnest, et al. J. Physical Chemistry B, 121(15): 3871-3881, 2017.

VMD Tesla V100 Cross Correlation Performance

Rabbit Hemorrhagic Disease Virus: 702K atoms, 6.5Å resolution

VMD on Volta GPUs now **~9x faster** than Kepler GPUs

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.

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

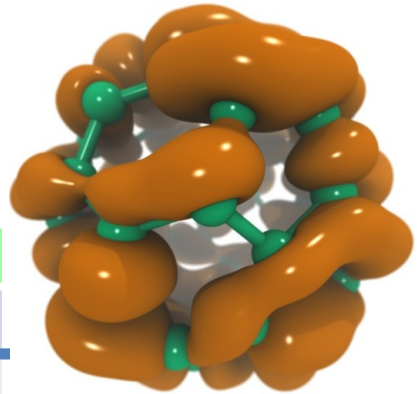

Hardware platform		Runtime,	Speedup
IBM Power8 (ORNL 'crest') + 1x Tesla K40 [1]		3.49s,	**1.0x**
Intel Xeon E5-2697Av4	**+ 1x Tesla V100 [2]**	**0.610s,**	**5.7x**
Intel Xeon E5-2697Av4	**+ 2x Tesla V100 [2]**	**0.294s,**	**11.8x**
Intel Xeon E5-2697Av4	**+ 3x Tesla V100 [2]**	**0.220s,**	**15.9x**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

NVLink perf. boost

[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.

[2] NAMD goes quantum: An integrative suite for hybrid simulations. Melo et al., Nature Methods, 2018.

Next Generation: Simulating a Proto-Cell

- **ORNL Summit: NVLink-connected Tesla V100 GPUs enable next-gen visualizations**
- 200nm diameter
- ~1 billion atoms w/ solvent
- ~1400 proteins in membrane

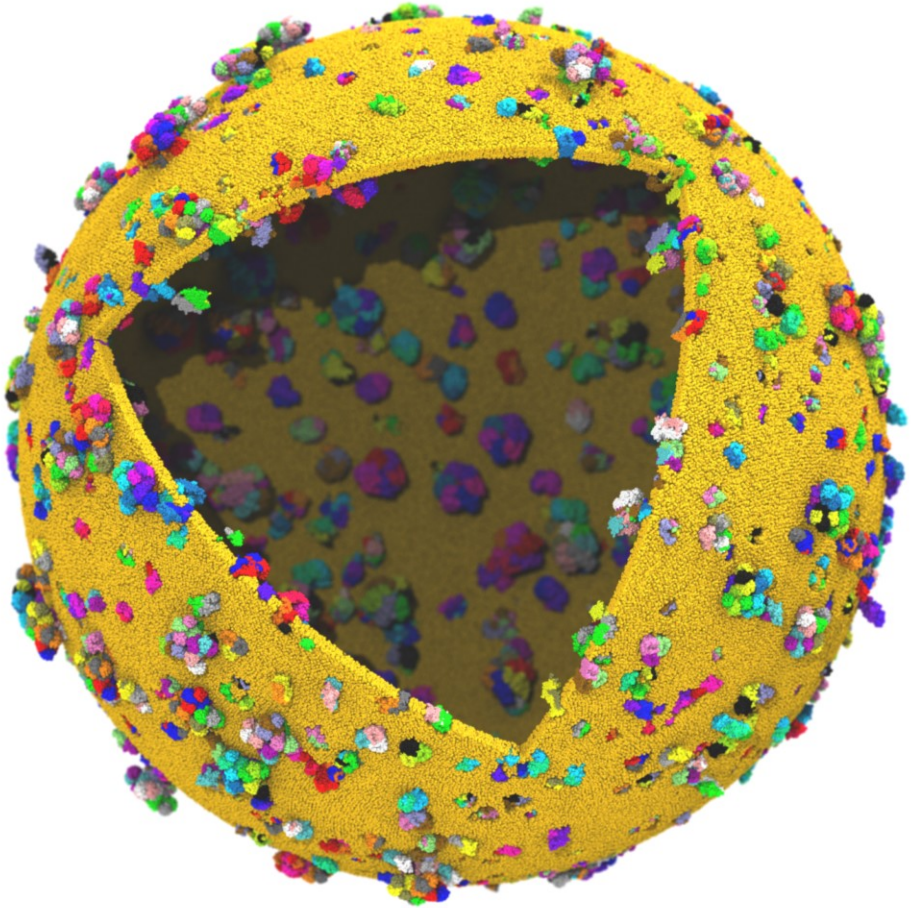

Proto-Cell Data Challenges

- **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**

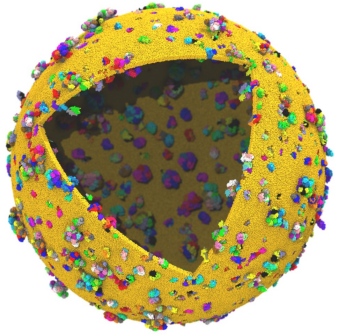

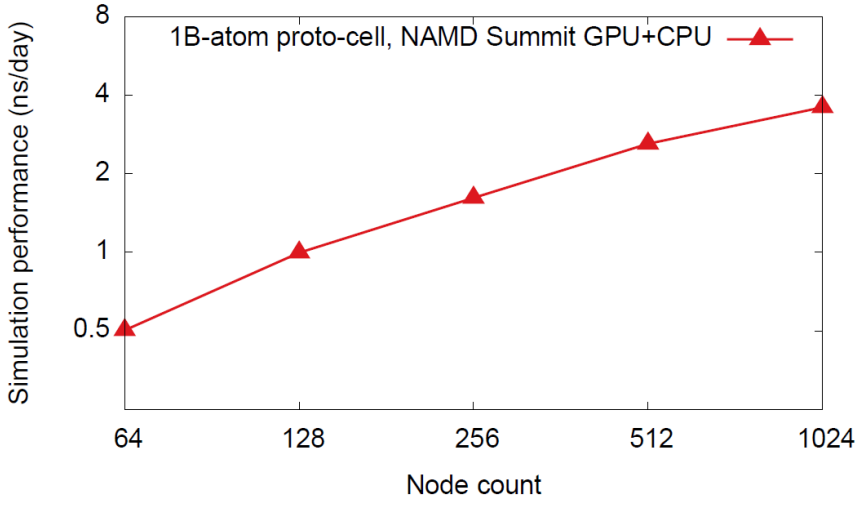

NEW: Power9+V100 Interactive Remote Visualization

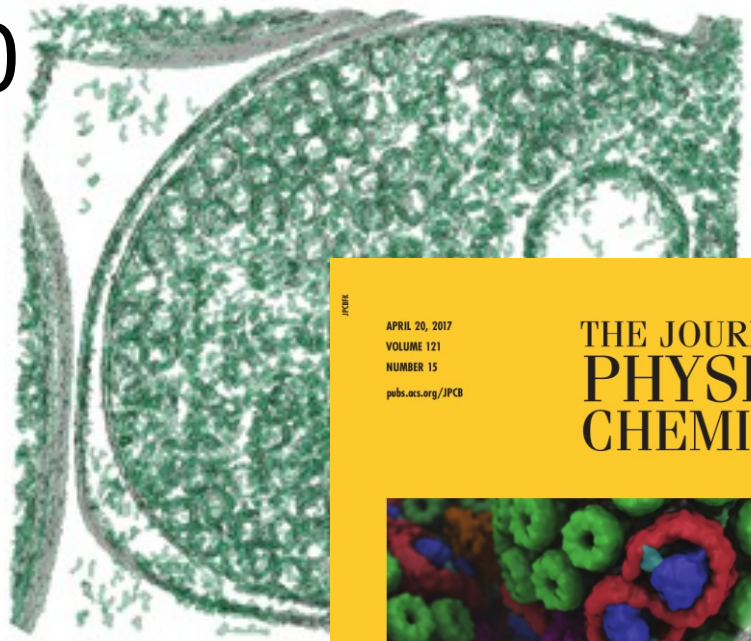

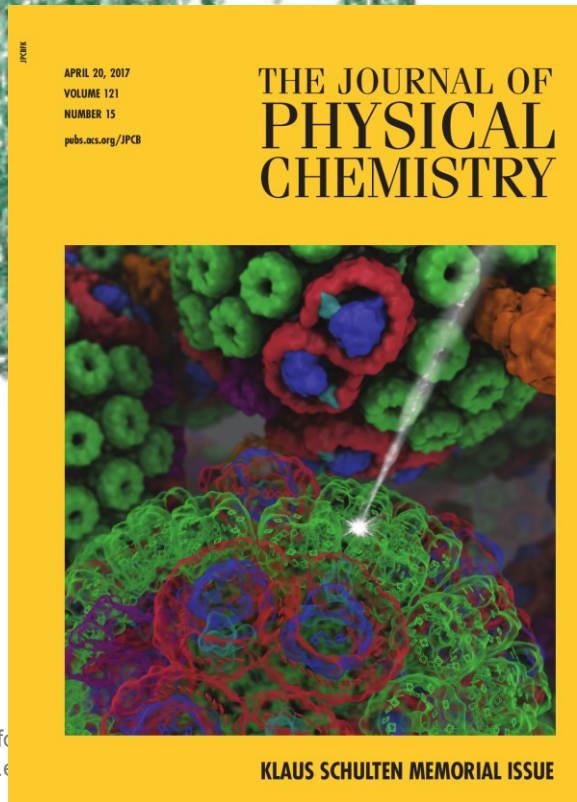

- Built-into VMD itself
- Enable access to massive data sets
- Uses GPU H.264 / HEVC hardware accelerated video encode/decode
- Supports interactive remote visualizations (both rasterization and ray tracing)
- Development ongoing, expected in next major VMD release, in 1H 2019…

Acknowledgements

- Theoretical and Computational Biophysics Group, University of Illinois at Urbana-Champaign
- NVIDIA CUDA and OptiX teams
- Funding:
  - NIH support: P41GM104601
  - ORNL Center for Advanced Application Readiness (CAAR)
  - IBM POWER team, IBM Poughkeepsie Customer Center
  - NVIDIA CUDA, OptiX, Devtech teams
  - UIUC/IBM C3SR
  - NCSA ISL

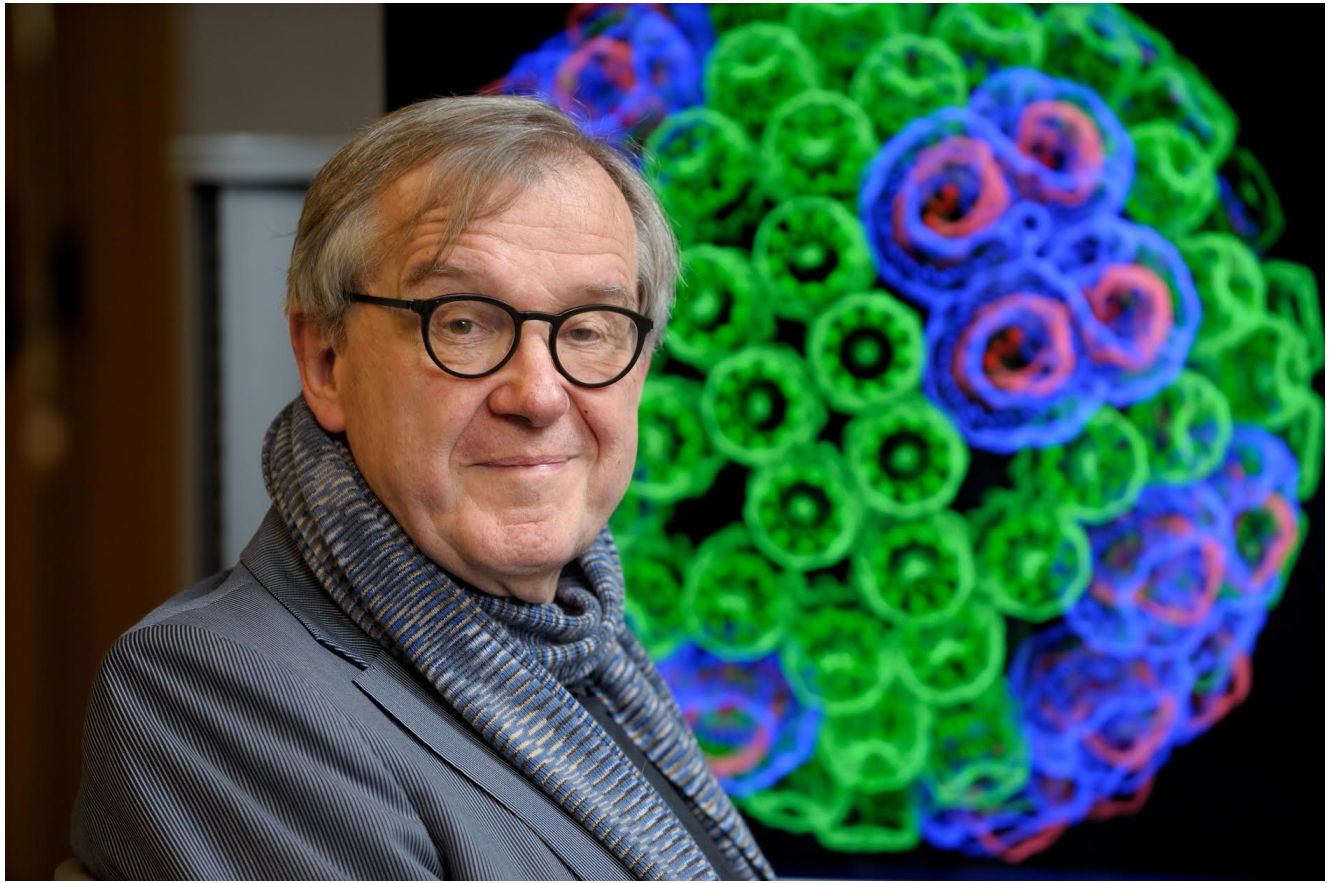

“When I was a young man, my goal was to look with mathematical and computational means at the inside of cells, one atom at a time, to decipher how living systems work. That is what I strived for and I never deflected from this goal.” – Klaus Schulten