NAMD
Scalable Molecular Dynamics

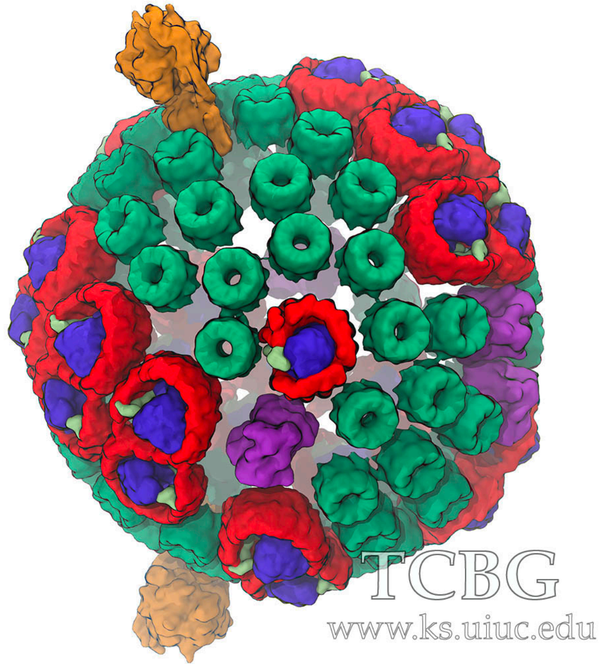

NAMD Molecular Dynamics on Frontera and Beyond

David J. Hardy
Senior Research Programmer
University Of Illinois at Urbana-Champaign

NIH National Institutes of Health

NAMD: Scalable Molecular Dynamics

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

- Code written in C++ with Charm++ parallel objects
  - CUDA for NVIDIA devices
  - HIP (via Hipify) for AMD devices
  - SYCL/oneAPI for Intel devices
- Simulate movements of biomolecules over time
- Enable parallel scaling
  - **Large systems (single-copy scaling)**
  - **Enhanced sampling (multi-copy scaling)**
- Over 44,000 registered users, over 20,000 citations

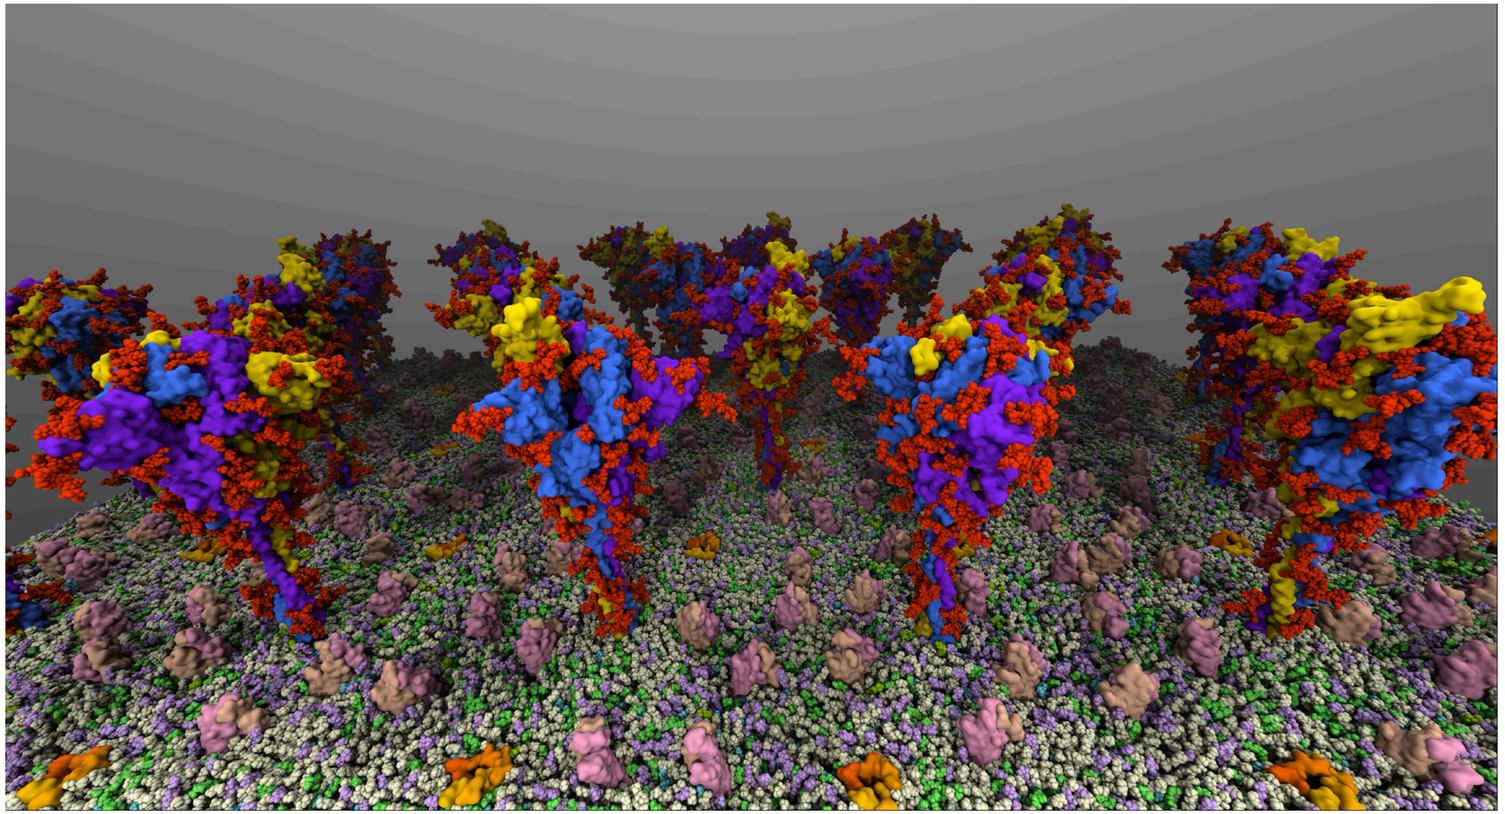

Investigations of coronavirus (SARS-CoV-2) spike dynamics.
Credit: Tianle Chen, Karanpal Kapoor, Emad Tajkhorshid (UIUC).
Simulations with NAMD, movie created with VMD.

Phillips, et al. *J. Comput. Chem.* 26, 1781-1802 (2005)
Phillips, et al. *J. Chem. Phys.* 153, 044130 (2020)

Molecular Dynamics Simulation

Integrate Newton's equations of motion:

$$m_i \frac{d^2 \vec{r}_i}{dt^2} = \vec{F}_i = -\vec{\nabla} U(\vec{R})$$

← Integrate for millions of time steps

Bonded

$$E_{total} = \sum_{bonds} K_r(r - r_{eq})^2 + \sum_{angles} K_\theta(\theta - \theta_{eq})^2 + \sum_{dihedrals} \frac{V_n}{2}[1 + \cos(n\phi - \gamma)]$$

Non-bonded

$$+ \sum_{i<j} \left[\frac{A_{ij}}{R_{ij}^{12}} - \frac{B_{ij}}{R_{ij}^{6}} + \frac{q_i q_j}{\epsilon R_{ij}} \right]$$

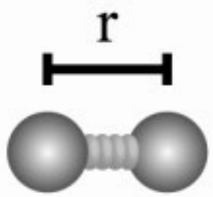

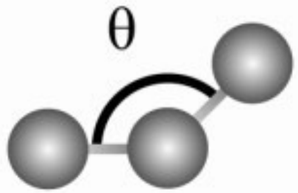

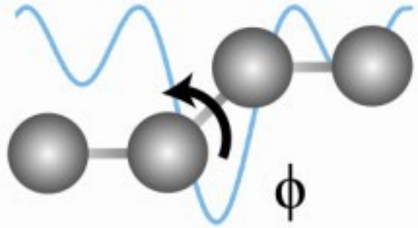

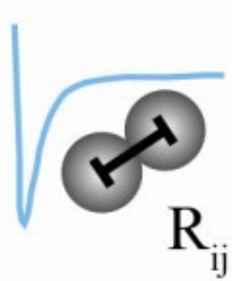

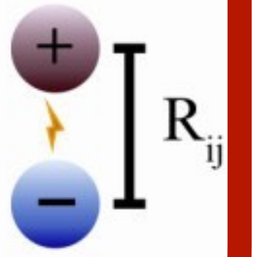

Most computationally intensive part →

NIH National Institutes of Health

NAMD Parallelizes Domain and Interaction Space

- Decompose atoms into equal volume ***patches***
- Calculate short-range pairwise forces between atoms, treated as **interactions between neighboring patches**
- Decompose patch-patch interactions into ***compute objects***
- Moving atoms: update spatial decomposition by ***migrating atoms*** between adjacent patches
- Load balancing: update work decomposition by ***migrating compute objects*** to keep processors consistently occupied
- Available parallelism: across each time step force calculation, but numerical integration steps must be calculated in sequence (making MD highly **sensitive to latency**)

Spatial decomposition of atoms into patches

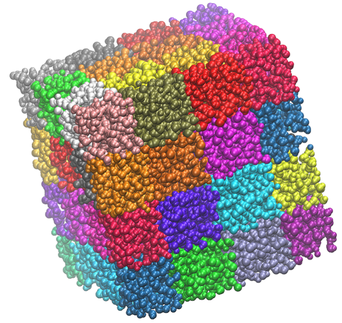

Work decomposition of patch-patch interactions into migratable compute objects

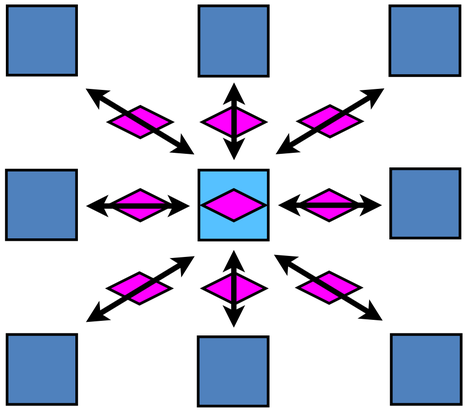

NAMD GPU-offload Approach for Multi-node Simulation

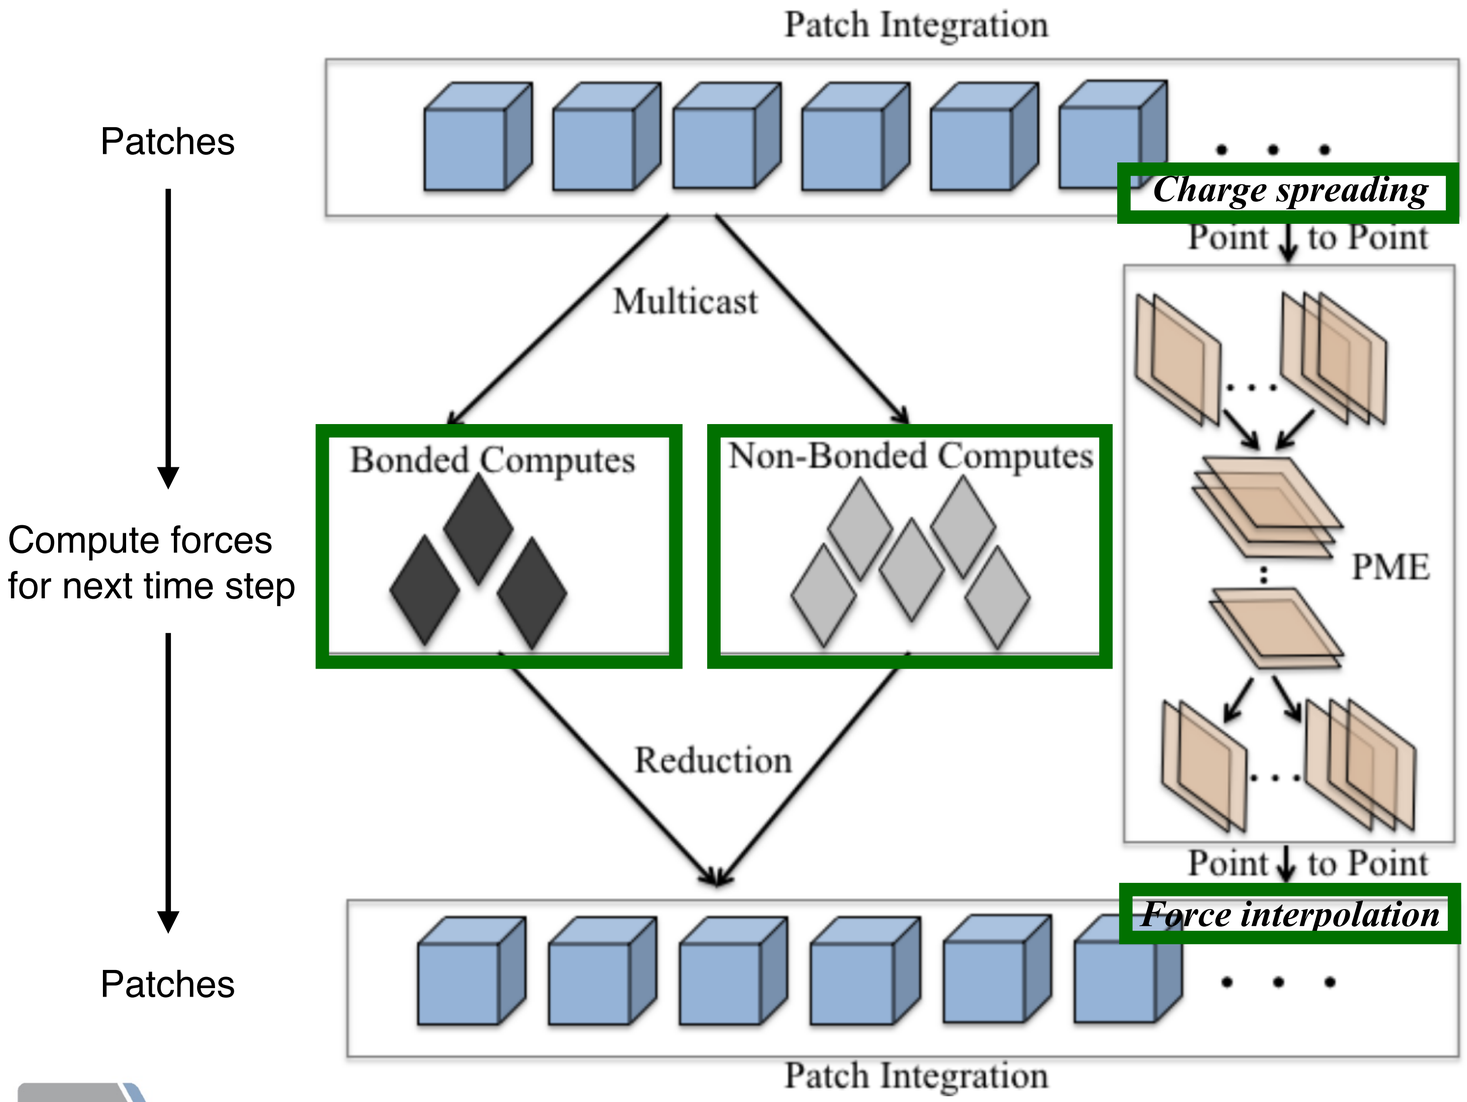

Offload force compute to GPU

Must aggregate positions

Outline

- NAMD 3.0 performance improvements
  - CPU optimizations via new AVX-512 kernels
  - New GPU-resident mode for single-GPU and single-node, multi-GPU simulation
- Science projects showcasing these performance improvements
  - Rommie Amaro's Lab (UCSD) studying SARS-CoV-2
  - Emad Tajkhorshid's Lab (UIUC) studying EAATs (Excitatory Amino Acid Transporters)
- Advantages running NAMD on Grace Hopper

AVX-512 Kernels Boost NAMD Performance

https://www.hpcwire.com/2020/08/12/intel-speeds-namd-by-1-8x-saves-xeon-processor-users-millions-of-compute-hours/

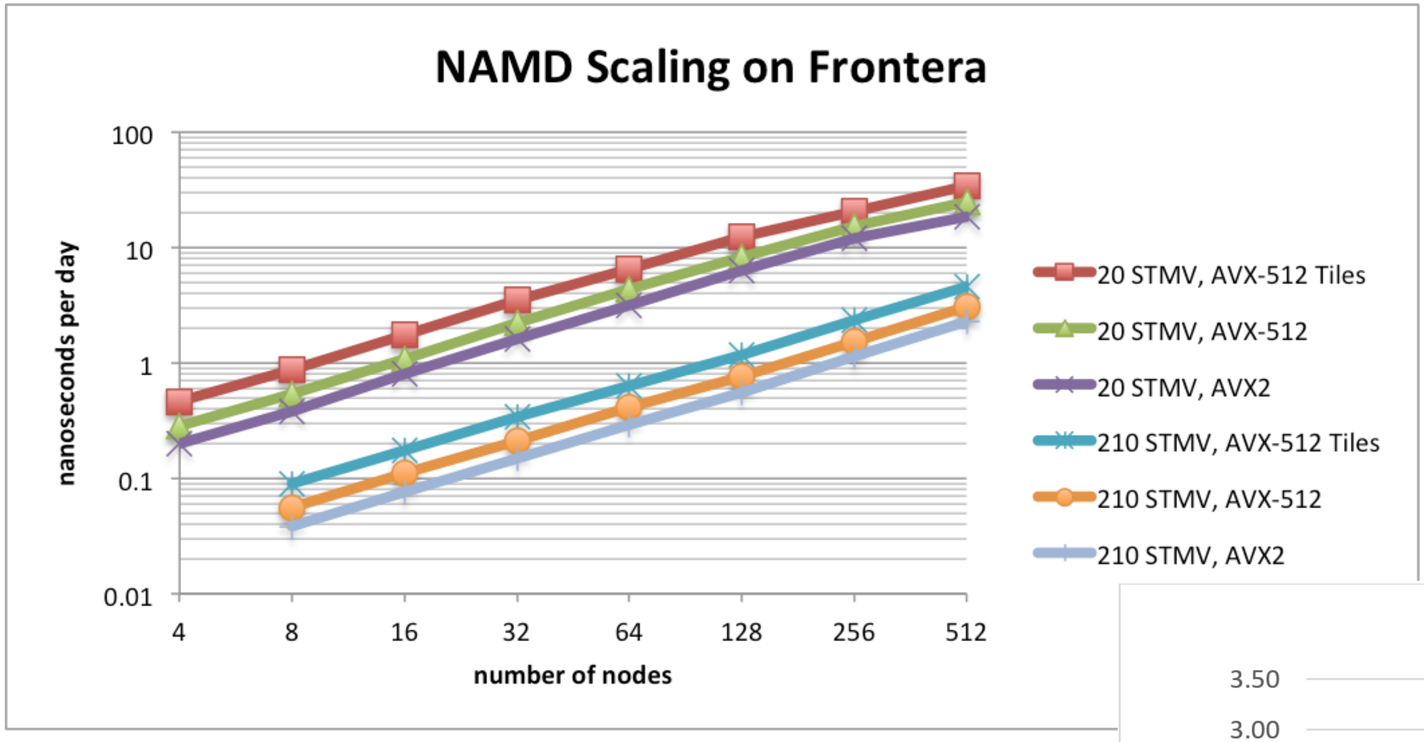

Porting of CUDA tiles kernels for non-bonded force (Mike Brown)

Benchmarks show
1.8x performance boost

Also runs on AMD Zen 4

Amaro Lab (UCSD) benefitted for their coronavirus spike simulations

Initially released in NAMD 2.15alpha

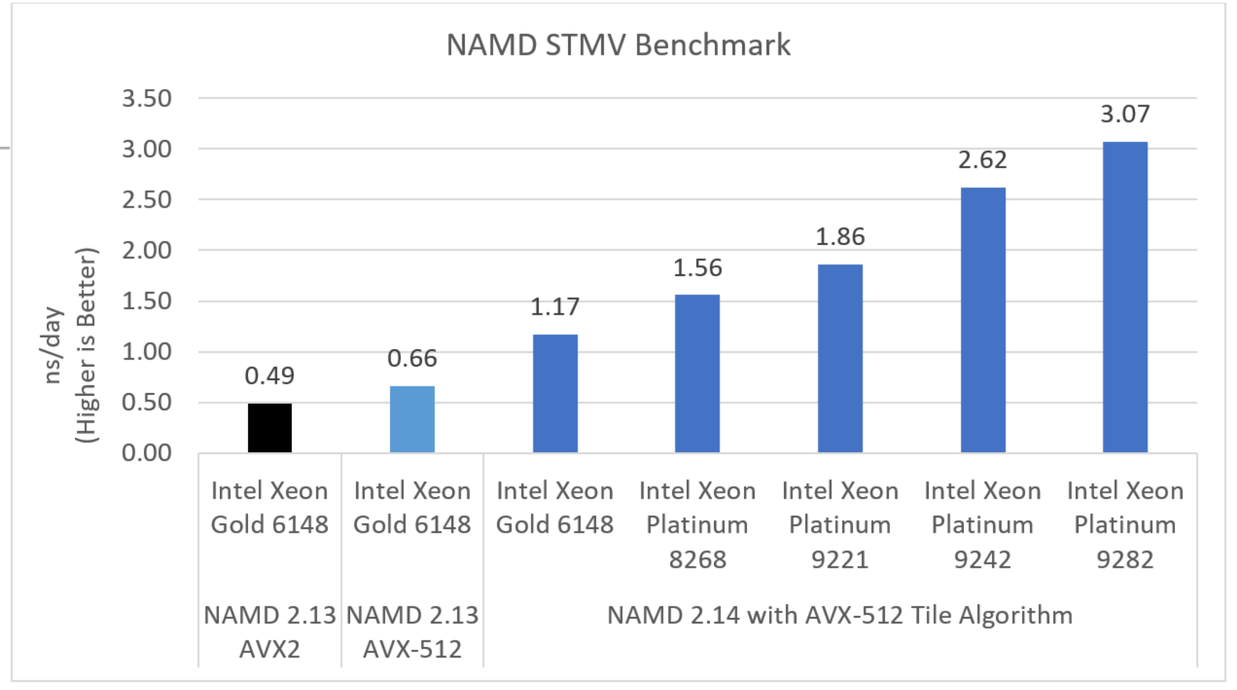

NIH National Institutes of Health

New GPU-resident Approach

Move integrator to GPU and maintain data between time steps

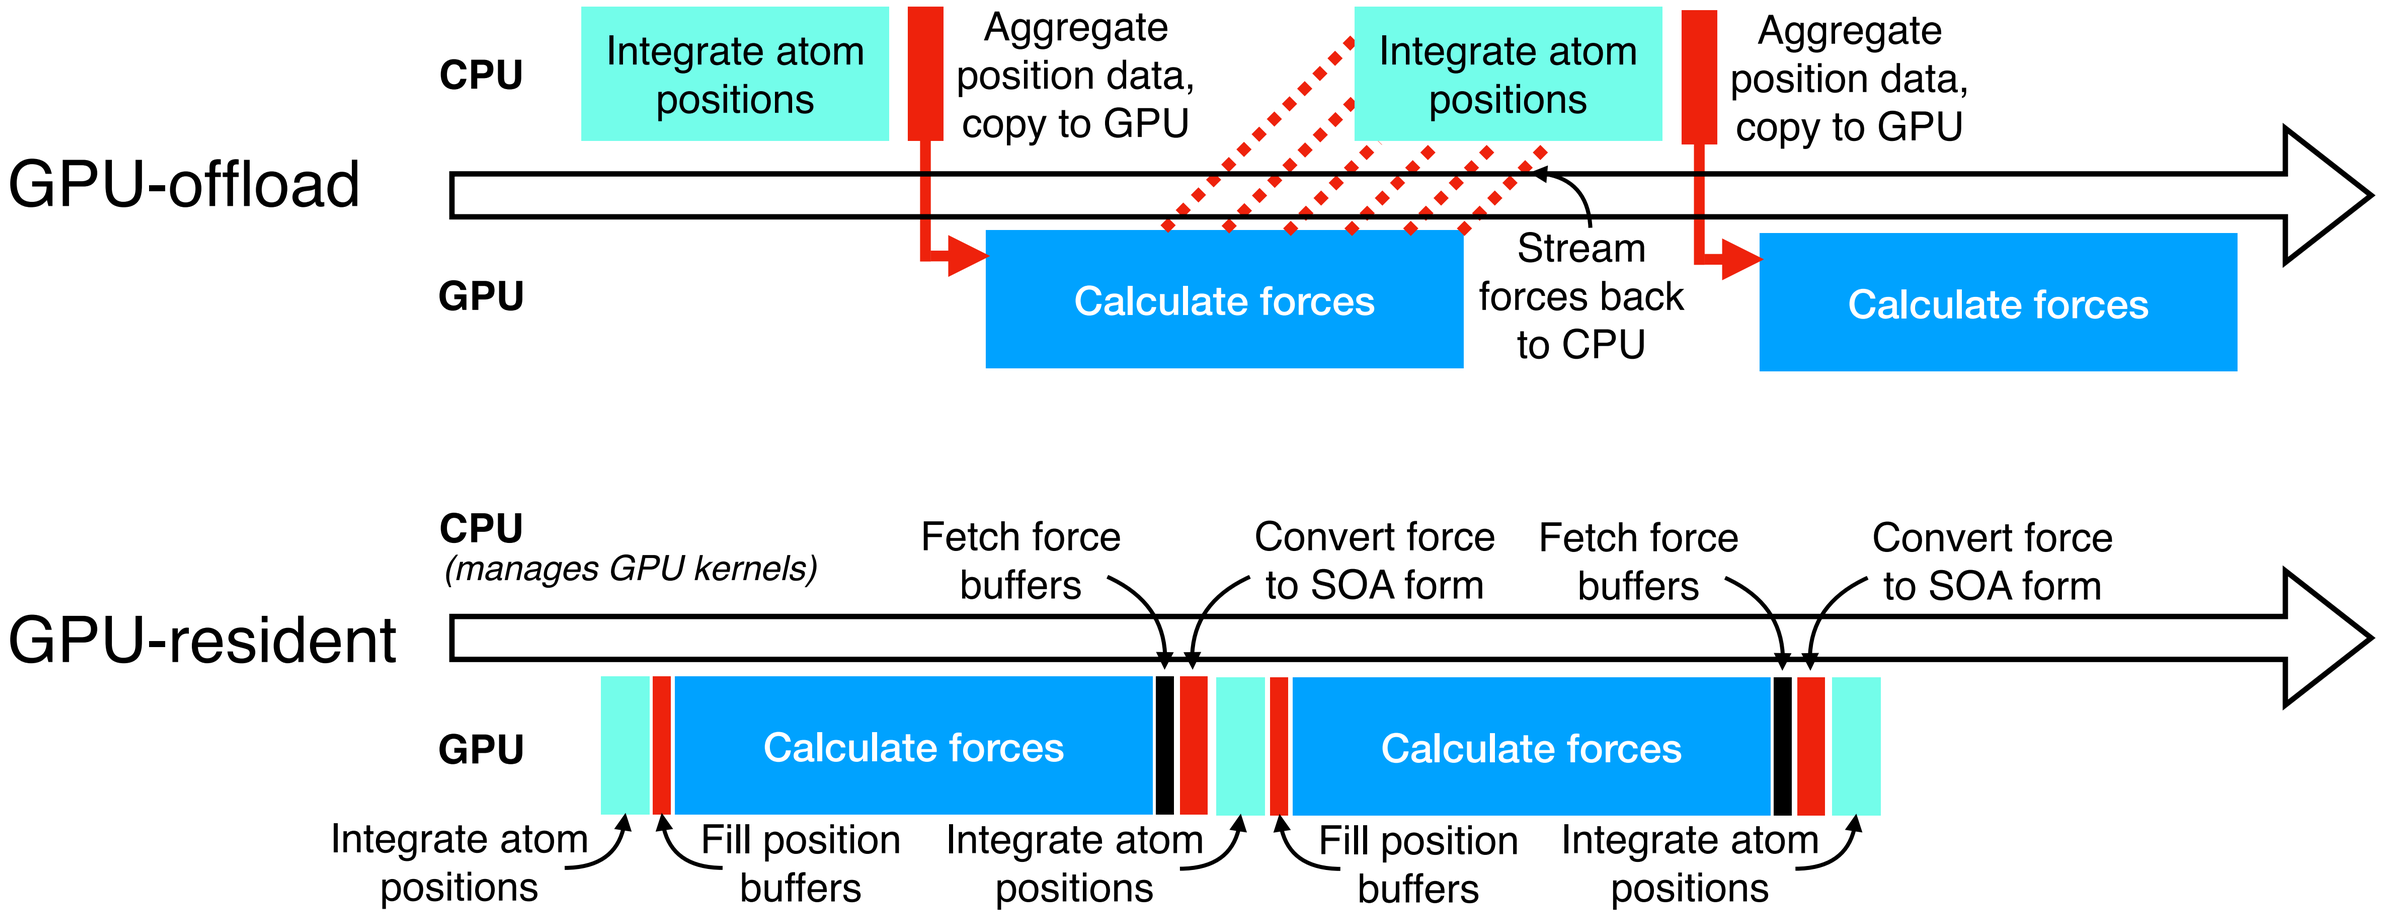

Adapting Parallel Scaling to GPU-resident Approach

Some communication required: multicasts and reductions

- Update atom positions in each patch during integration
- Perform **position multicast** into compute objects
- Compute new forces
- Perform **force reduction** back to patches
- GPUs need load-store memory access between different devices ***within every time step***, with data sizes on the order of 8KB per access

See past NVIDIA GTC talks for more details: s31529, s41378, s51693

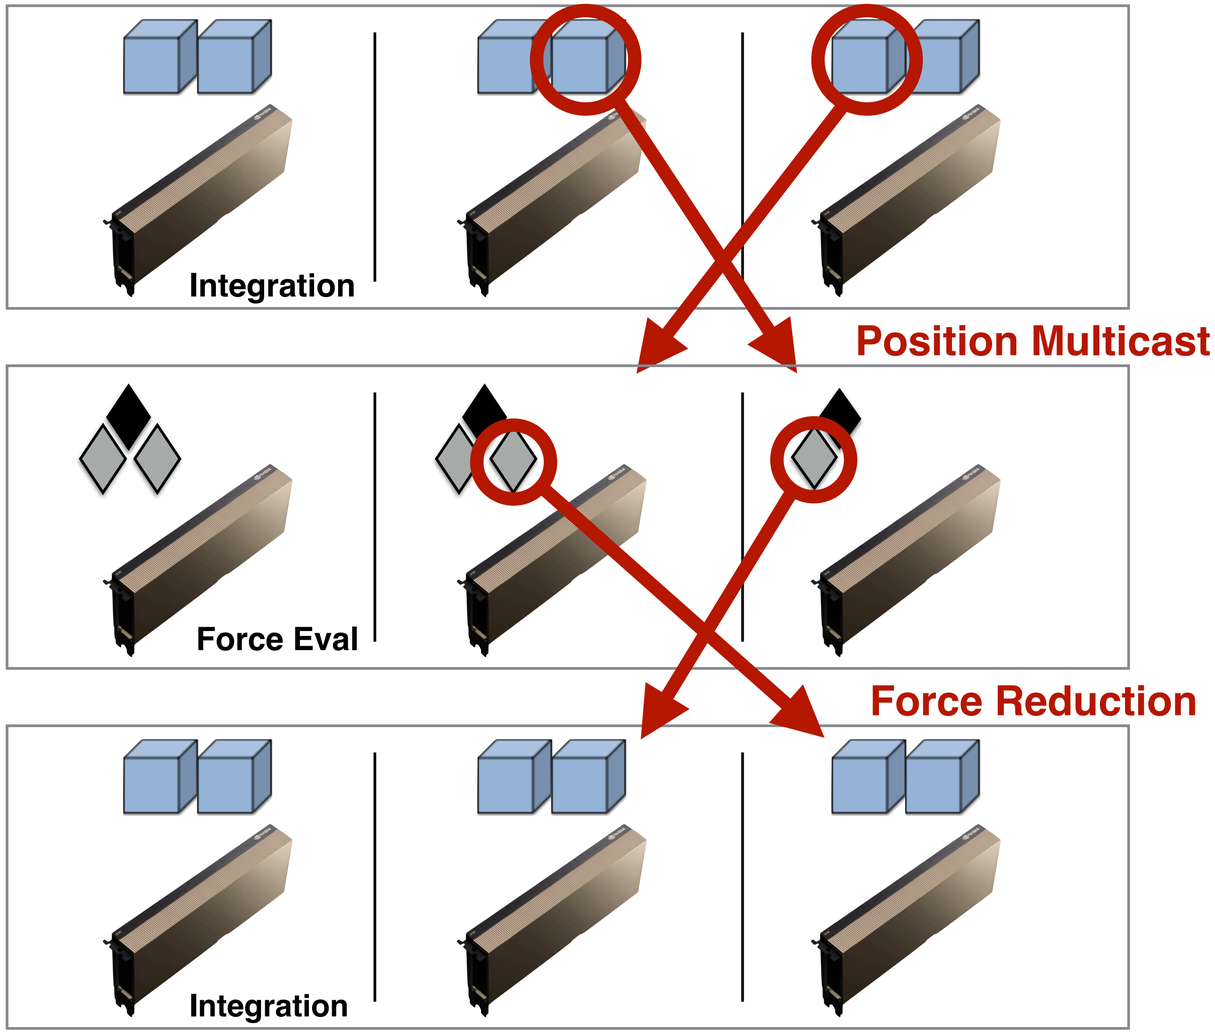

Single GPU Performance Improvements

A100

NVE simulation (constant energy):
- DHFR: AMBER-like force field (9 Å cutoff), HMR with 4 fs time step, PME, rigid bond constraints, "margin" 2 Å, two-away-Z.
- ApoA1: CHARMM force field (12 Å cutoff), multiple time stepping with 2 fs time step and 4 fs PME, rigid bond constraints.
- STMV: CHARMM force field (12 Å cutoff), multiple time stepping with 2 fs time step and 4 fs PME, rigid bond constraints.
- Spike ACE-2: CHARMM force field (12 Å cutoff), multiple time stepping with 2 fs time step and 4 fs PME, rigid bond constraints.
- Each measurement calculates the average ns/day running dynamics for 3 minutes of wall clock time.

Platform details:
- 1 GPU and CPU cores from HGX-A100 (4x A100-SXM4-40GB, NVLink, 2x AMD EPYC 74F3 (Milan) 24-core processor)
- GPU-offload performs best for each system using **all** 48 cores.
- GPU-resident: DHFR — 2 cores, ApoA1 — 4 cores, STMV — 8 cores, Spike ACE-2 — 8 cores.

	GPU-resident ns/day	GPU-offload ns/day	speedup
DHFR (23.6k)	**1174**	**330.4**	**3.55x**
ApoA1 (92.2k)	190.4	63.88	2.98x
STMV (1.06M)	16.64	7.547	2.20x
Spike ACE-2 (8.56M)	1.875	0.7711	2.43x

Comparing GPU-resident with GPU-offload

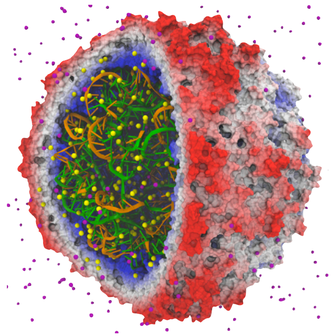

STMV
1.06M atoms

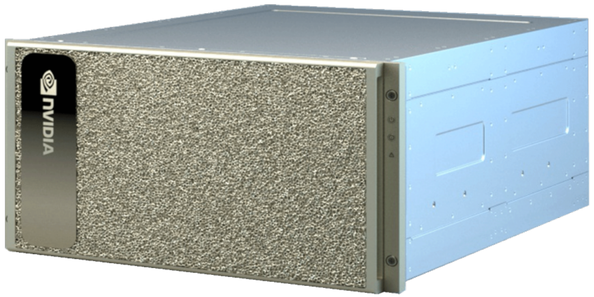

DGX-A100

Simulation details:
NVE, CHARMM force field, cutoff distance 12 Å,
MTS with 2 fs time step and 4 fs PME, rigid bond constraints.
GPU-resident mode sets performance tuning parameter "margin" to 4 Å for GPU-resident version, with PME PEs set to 8, 7, 5, 1 for numbers of GPUs 1, 2, 4, and 8, respectively.
GPU-offload mode scales CPU core usage, 8 cores per GPU.

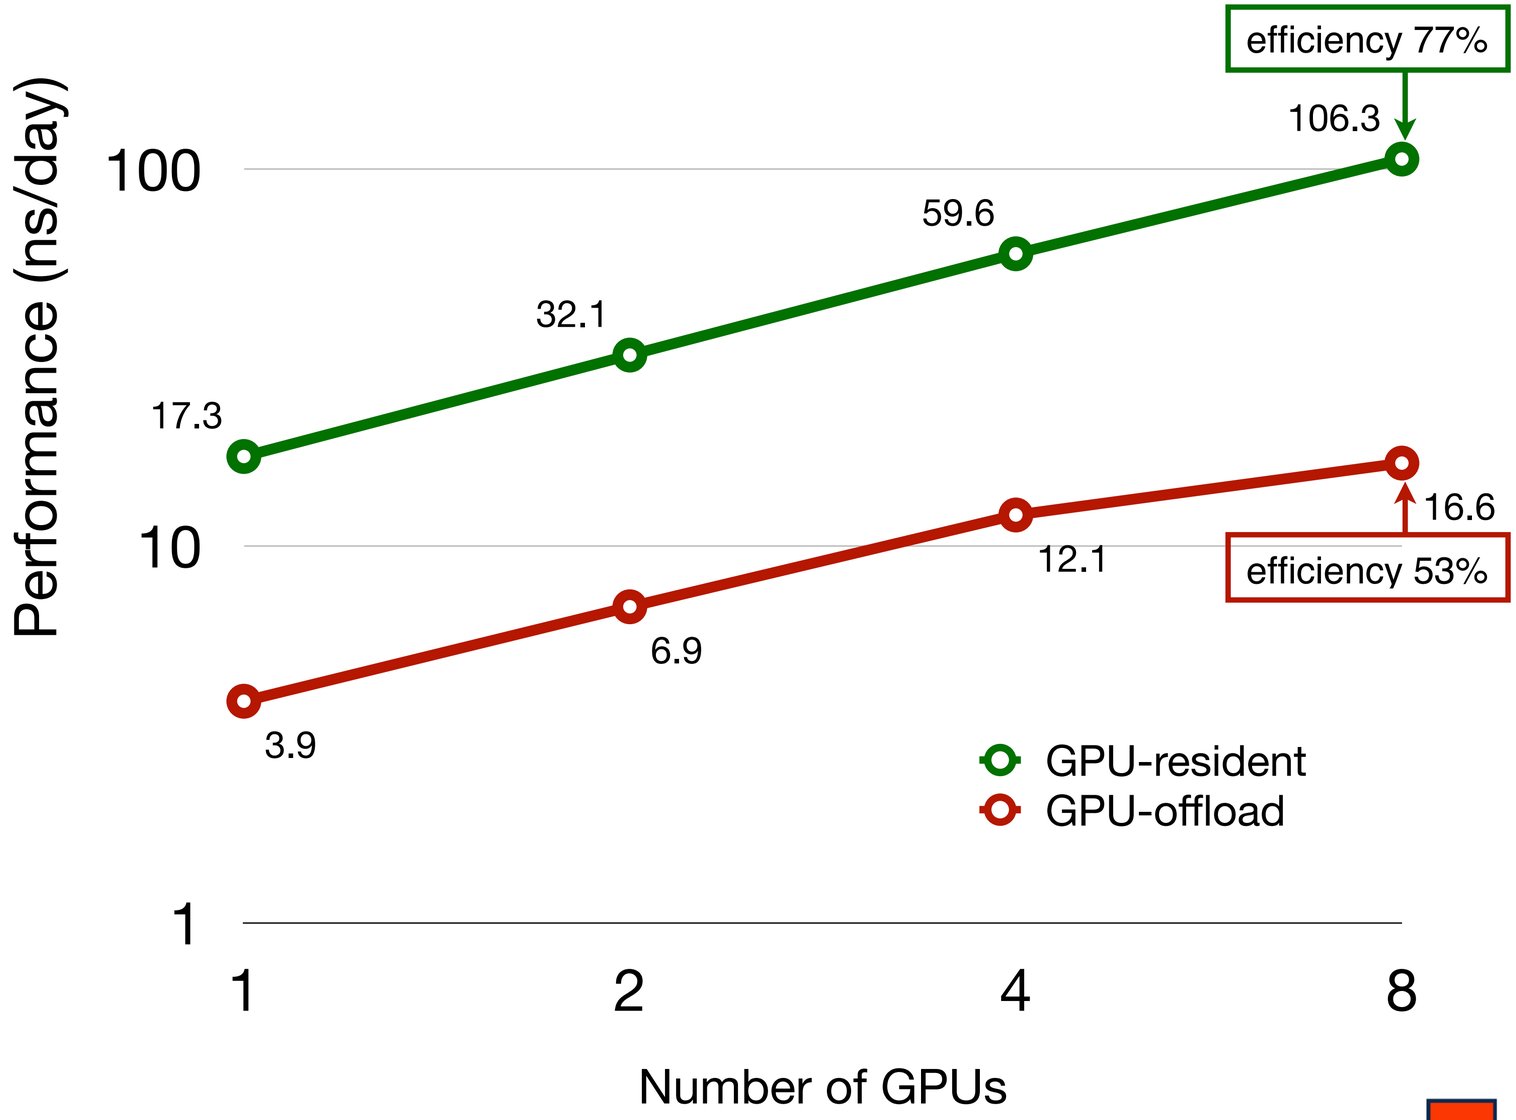

NIH National Institutes of Health

NAMD Simulating SARS-CoV-2 on Frontera and Summit

Collaboration with Amaro Lab at UCSD, images rendered by VMD

Winner of Gordon Bell Special Prize at SC20, project involved overall 1.13 Zettaflops of NAMD simulation

(A) Virion, (B) Spike, (C) Glycan shield conformations

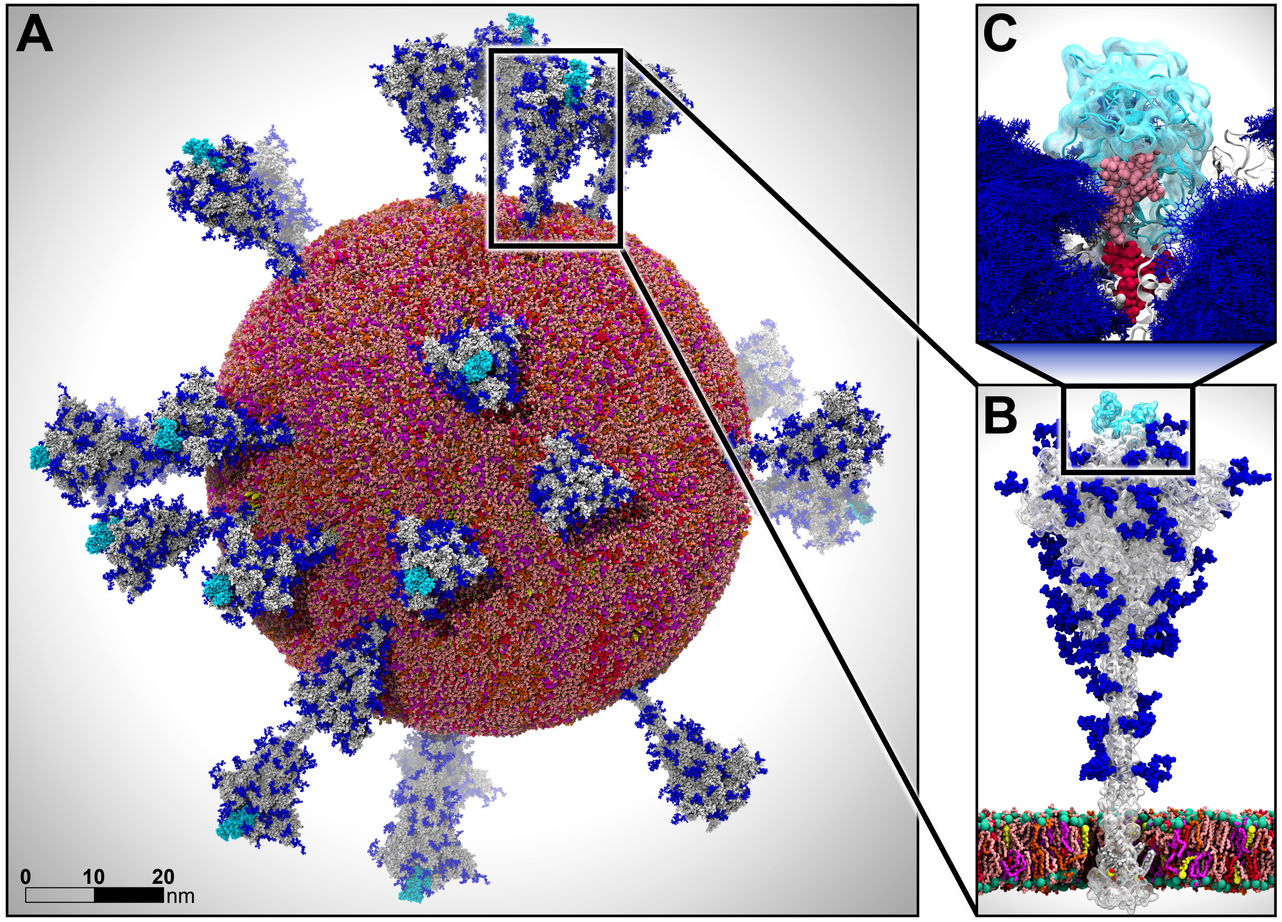

Scaling performance:

- ~305M atom virion
- ~8.5M atom spike

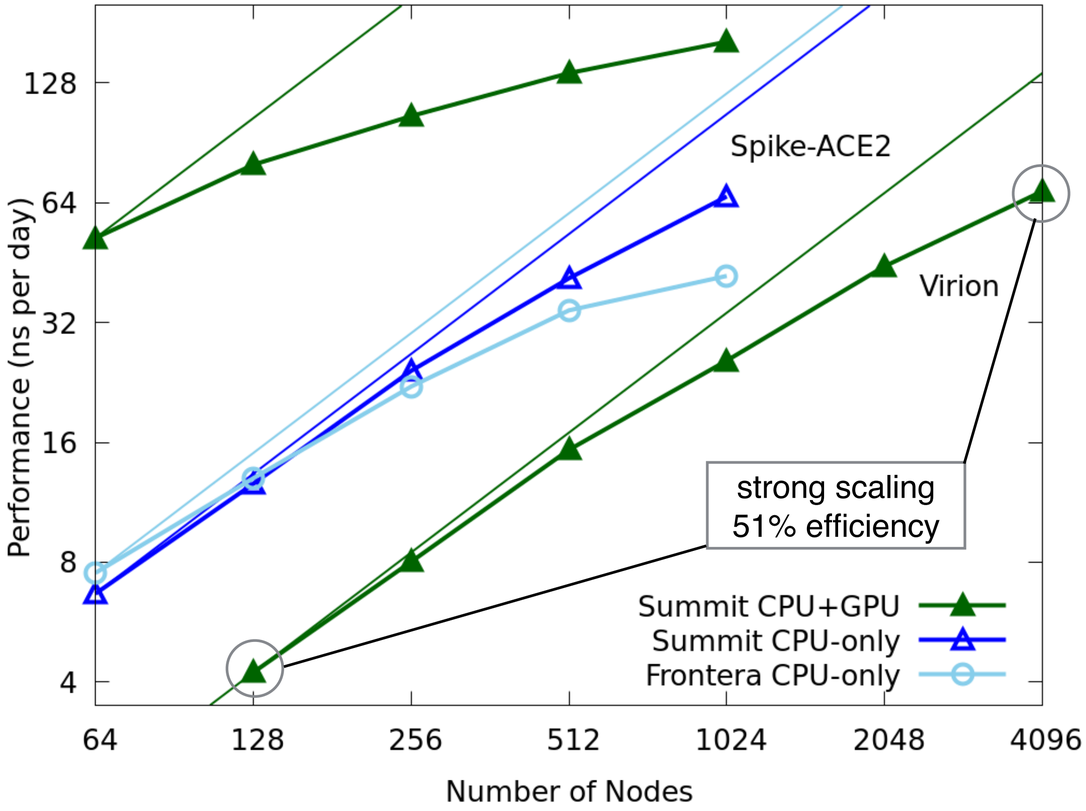

Casalino, et al. *bioRxiv* (2020) https://doi.org/10.1101/2020.11.19.390187

NIH National Institutes of Health

NAMD Simulating Aerosolized SARS-CoV-2 on Summit

Collaboration with Amaro Lab at UCSD, images rendered by VMD

Finalist for Gordon Bell Special Prize at SC21

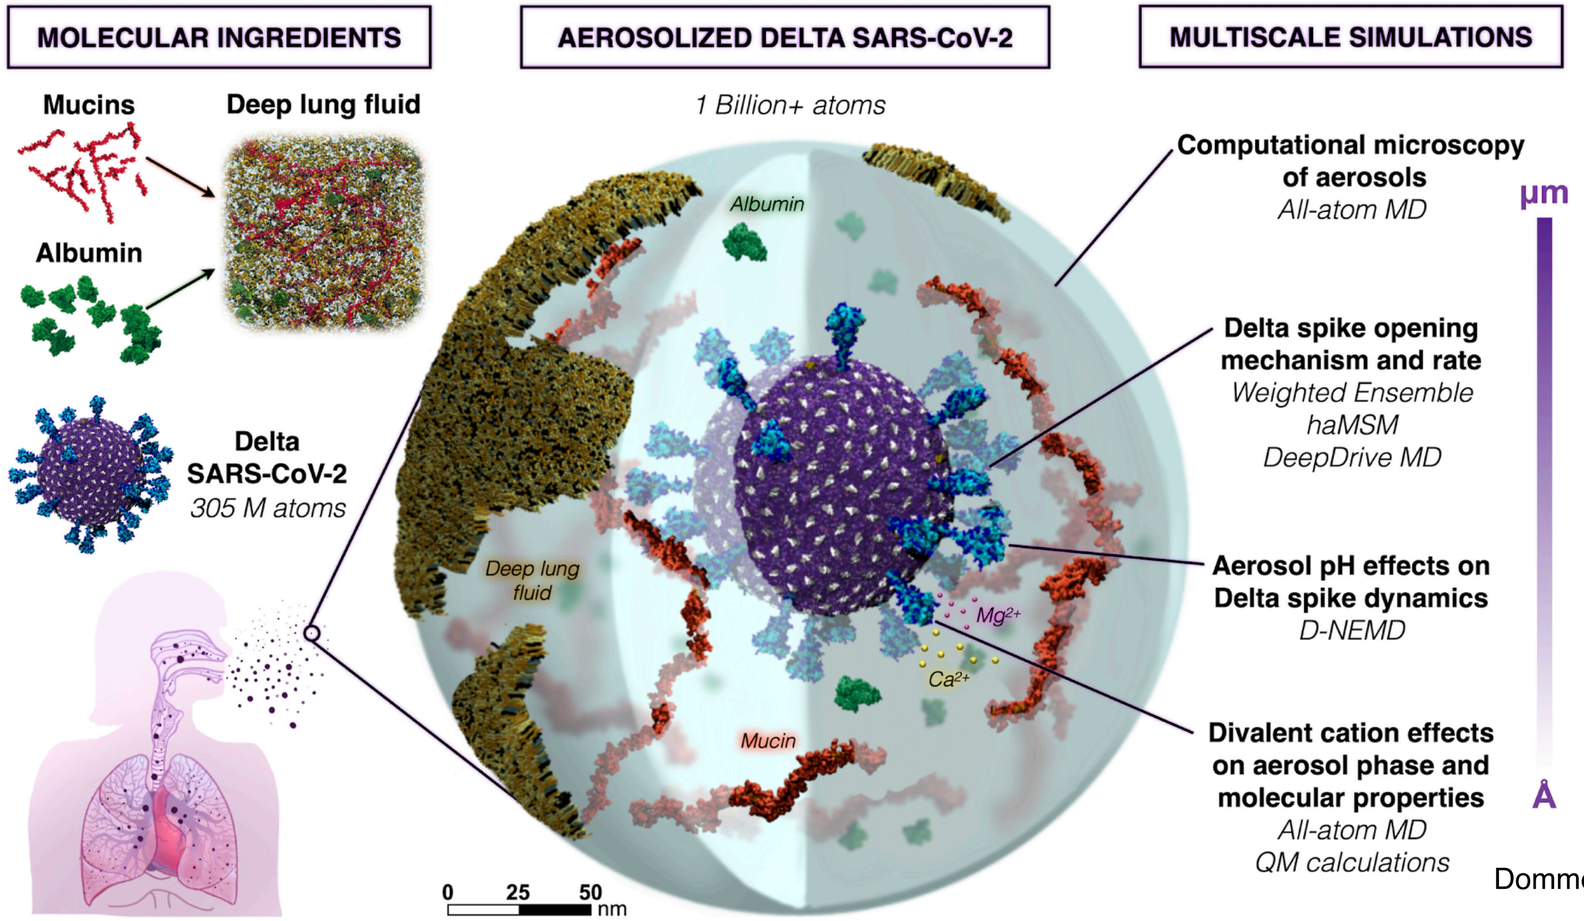

Scaling performance:

- ~1B atom aerosol

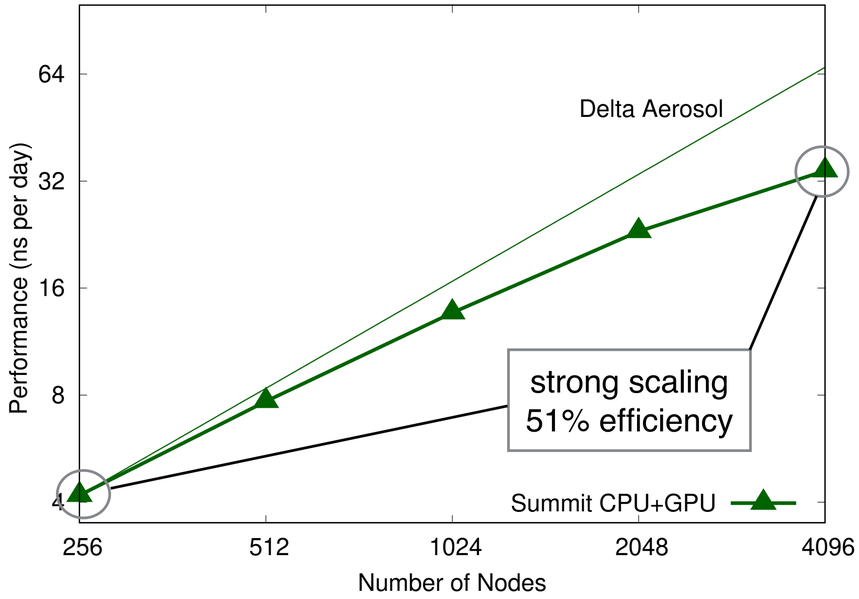

Dommer, et al. *bioRxiv* (2021) https://doi.org/10.1101/2021.11.12.468428

A Case Study: Excitatory Amino Acid Transporters (EAATs)

Work of **Ashkan Fakharzadeh** from Tajkhorshid Lab

- **EAATs** regulate concentration of **glutamate** in synaptic cleft of our neural cell.
- Dysregulation of EAATs → Parkinson, Alzheimer, depression
- EAATs transport 1glutamate, 3 Na^+, H^+, and counter-transport K^+ in a *complex process*!
- We study **transport cycle** through molecular dynamics simulations!

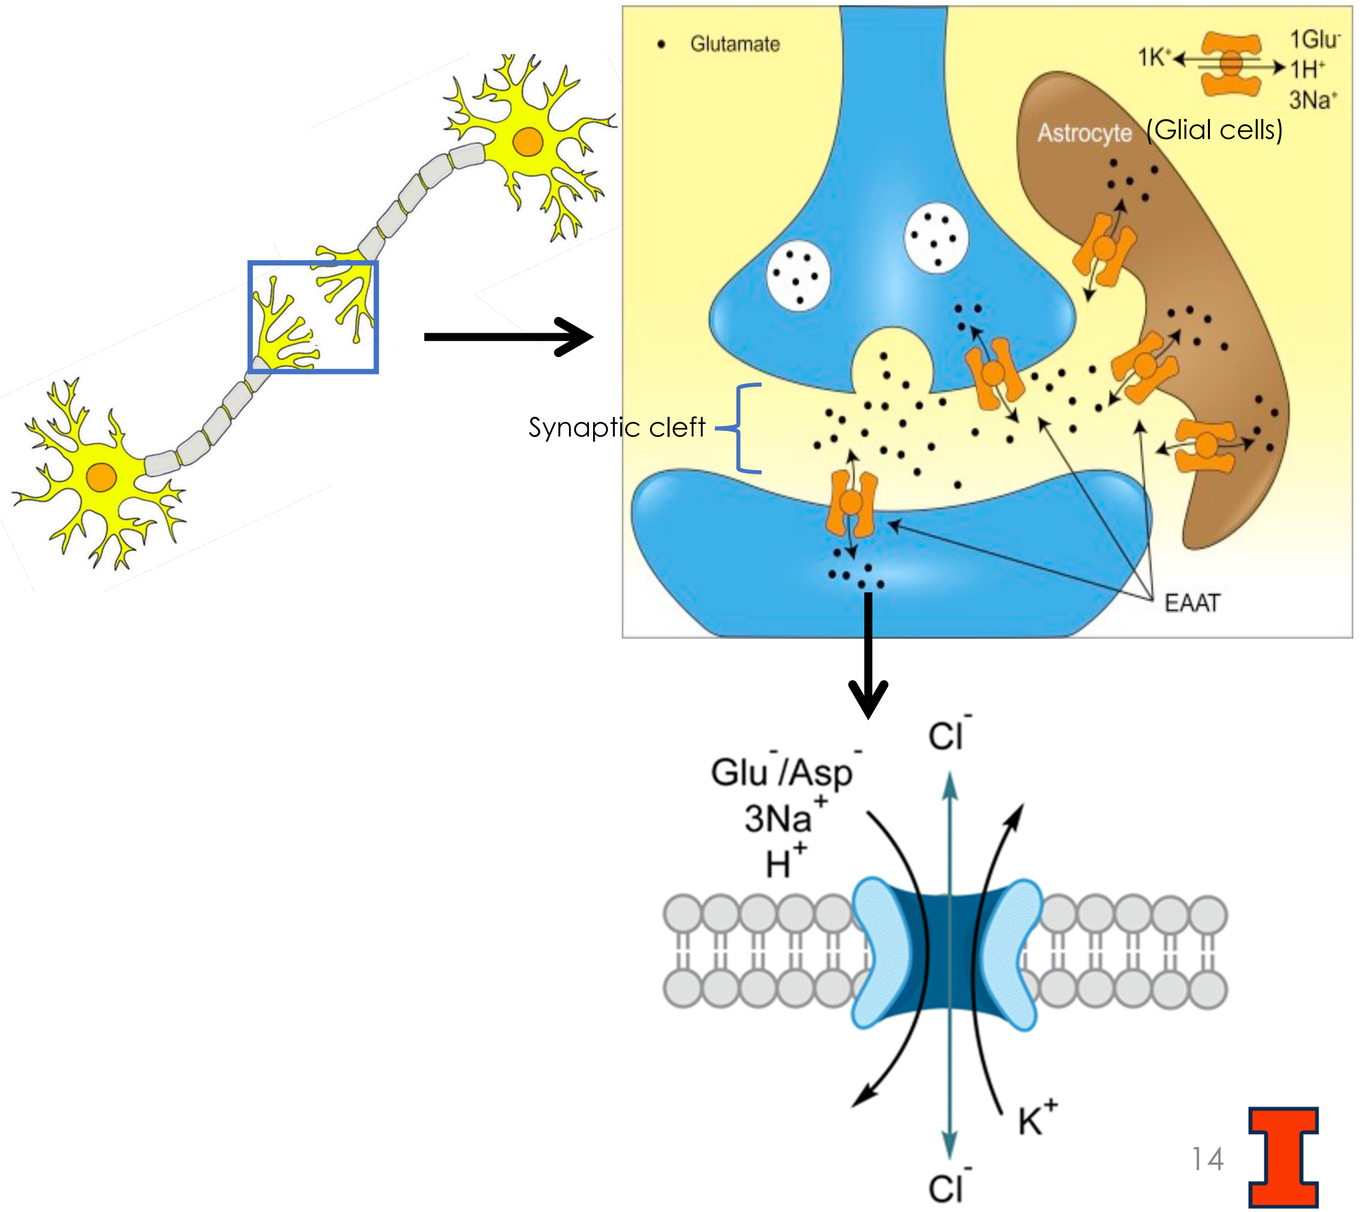

Simulations of dEAAT2

- Regular MD simulations were performed to study **local dynamics** of dEAAT2

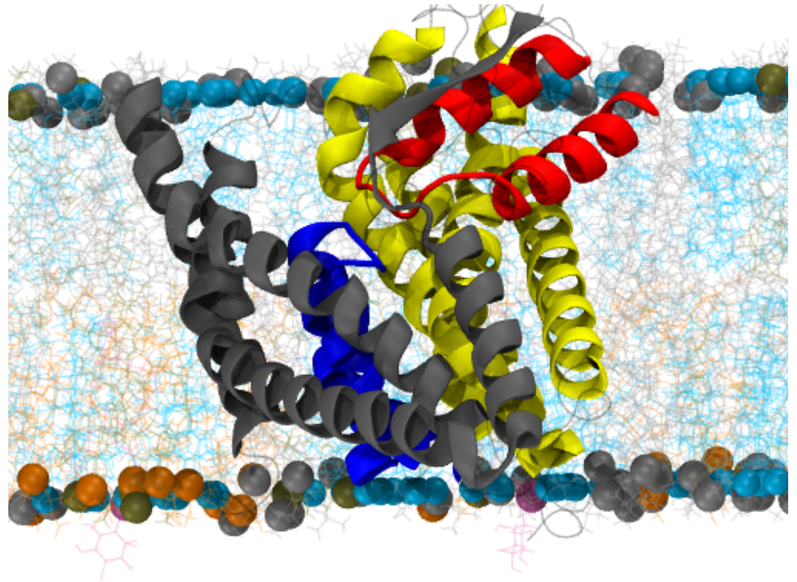

Simulation details:

Crystal structure embedded in an explicit glia-like membrane
Solvated with TIP3P water molecules. Neutralized with NaCl salt with an added 150 mM salt concentration.
Simulation size: ~200K atoms, box size: 123x123x117 $Å^3$
Temperature 310°K
Simulations were performed in NAMD3 with CHARMM36m force field
Electrostatic interactions: particle mesh Ewald method
Integration: 2 fs time step
Minimized and equilibrated for 10ns with backbone constrained; 200ns unconstrained equilibrium MD; repeated 1-3 times

- Advance features such as free energy perturbation and constant pH were used to study **ion binding sequence.**

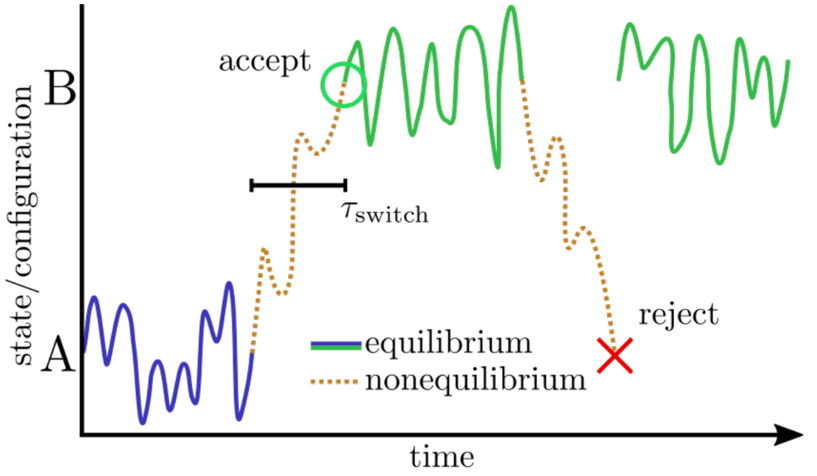

Constant pH MD: A series of cycles composed of an MD and nonequilibrium MD/Monte Carlo steps

Simulation details:

7-8 replica (pH 3.4, 4.4, …, 9.4)
Performed 4000 cycles of 10000 steps MD and 40000 steps nonequilibrium MD/MC
Upper bound of simulation time: 50ns

Performance of dEAAT2 Simulations on Frontera

- Run command examples:
  - **Single GPU node:**

```
ibrun namd3 +ppn 16 +pemap 0-7,16-23 +pmepes 7
+devices 0,1 CONFIG_FILE > LOG_FILE
```

7 PEs per PME device, and 9 PEs per non-PME device

  - **Multi-copy Multi-GPU:**

```
charmrun ++n NUM_REPLICA ++mpiexec ++remote-shell
ibrun namd3 ++ppn 7 +pemap 1-7,17-23,9-15,25-31
+commap 0,16,8,24 +replicas NUM_REPLICA
+devicesperreplica 1 CONFIG_FILE +stdout
LOG_FILE.%d
```

  - **Multi-copy Multi-CPU:**

```
ibrun namd3 +ppn 13 +pemap 4-55:2,5-55:2 +commap
0,2,1,3 CONFIG_FILE +replicas NUM_REPLICA +stdout
LOG_FILE.%d
```

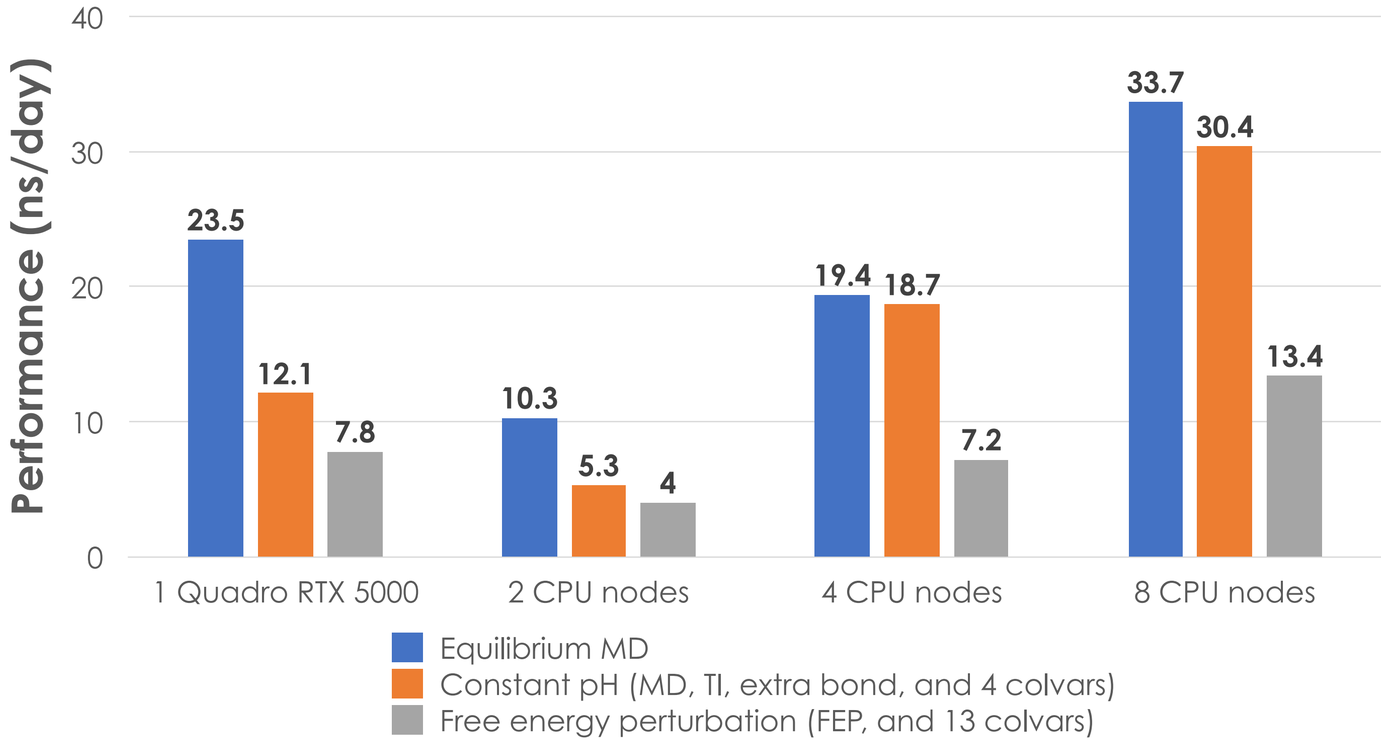

Performances are per replica.
For constant pH, the performance is the average of MD and nonequilibrium MD/MC.
Each GPU node can run up to 4 replicas with cost of 3 SU/hr.

NAMD Simulations on Grace Hopper

Preparing for Vista

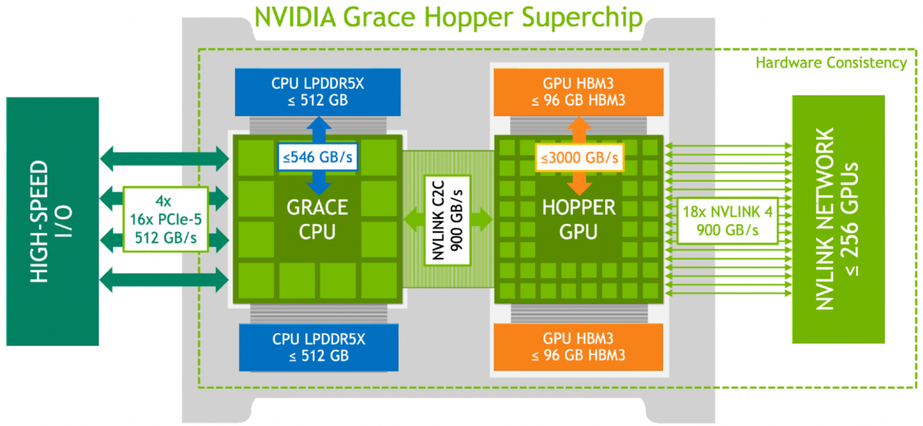

https://developer.nvidia.com/blog/nvidia-grace-hopper-superchip-architecture-in-depth/

- Enables fast, low-latency communication between CPU and GPU via NVLink
- Provides memory coherency between host and device
- Has much higher CPU memory bandwidth per GPU than x86
- Greatly reduces CPU-side bottlenecks, such as using Colvars with GPU-resident simulation

LaINDY ~200k atom simulation

Uses several collective variables:
spin, distance, orientation

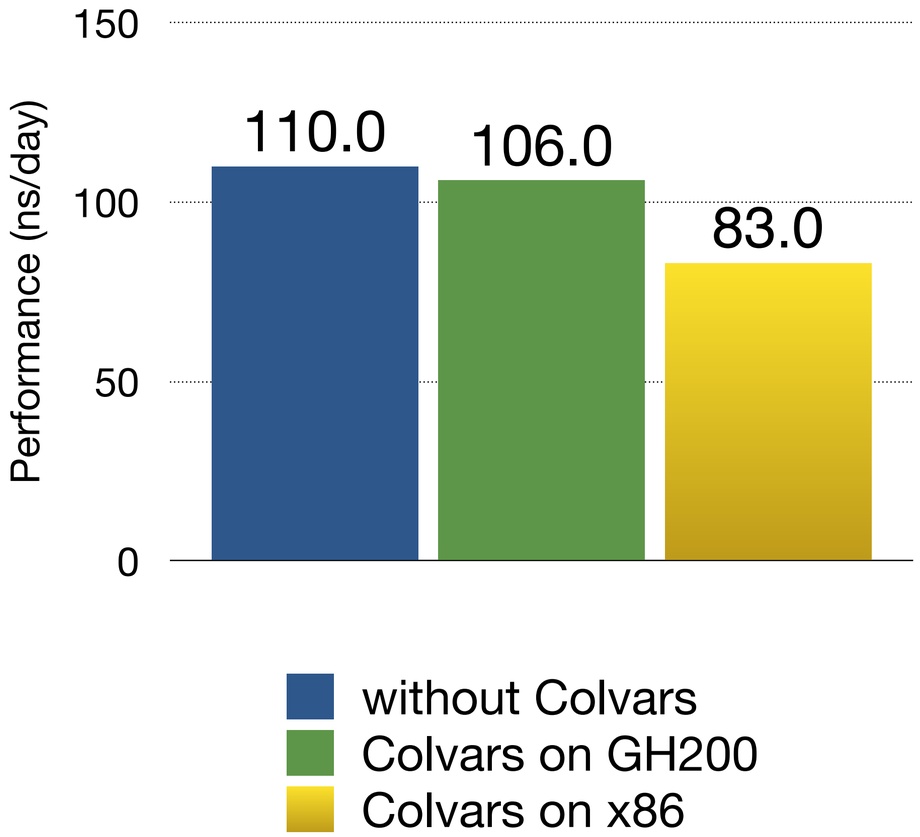

Sauer, Trebesch, et al. *eLife* (2020) https://doi.org/10.7554/eLife.61350

NIH National Institutes of Health

Acknowledgments

- Eric Bohm (UIUC), Mike Brown (NVIDIA, Intel), Haochuan Chen (UIUC), Antti-Pekka Hynninen (NVIDIA, ORNL), Julio Maia (AMD, UIUC), Tareq Malas (Meta, Intel), Jim Phillips (UIUC), Mohammad Soroush Barhaghi (Schrödinger, UIUC), John Stone (NVIDIA, UIUC), Peng Wang (NVIDIA), Ke Yue (Intel)
- NIH grants P41-GM104601 and R24-GM145965
- Intel oneAPI Academic Center of Excellence

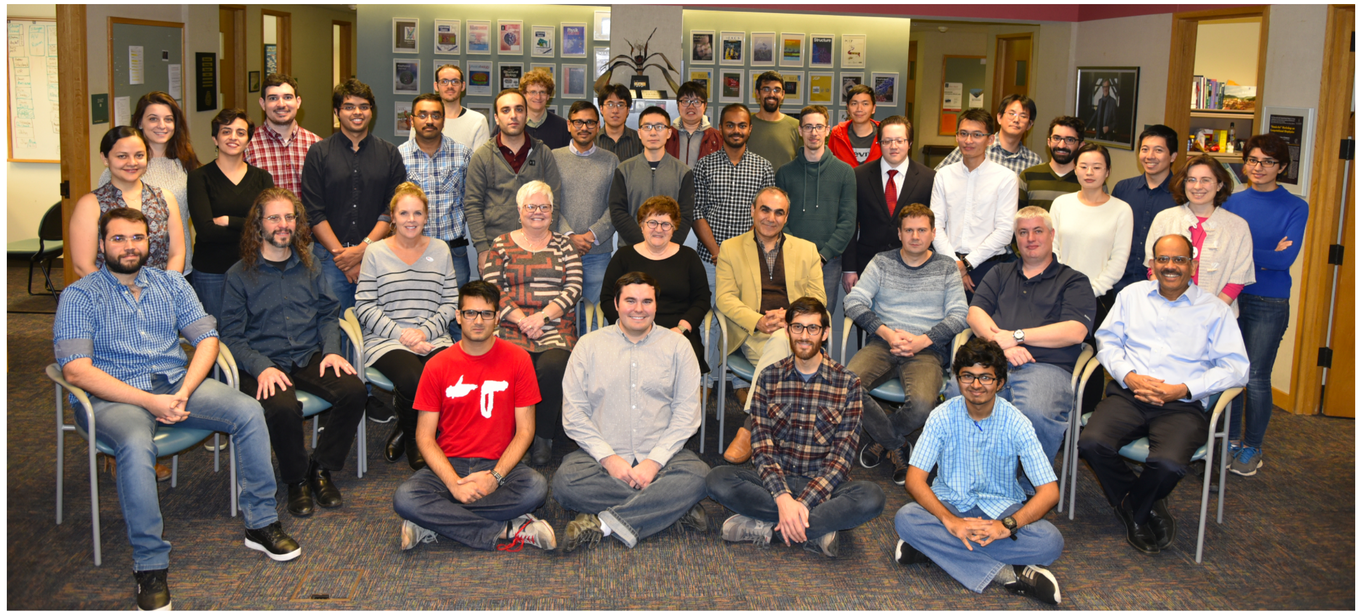

NIH Center for Macromolecular Modeling and Bioinformatics (2019)
Beckman Institute, University of Illinois at Urbana-Champaign