AMBER: The How, What and Why on an Intel® Xeon Phi™

Perri Needham

&

Ross Walker

(SDSC / UCSD)

What is Molecular Dynamics?

- ***A Method for simulating how a chemical system evolves with time.***

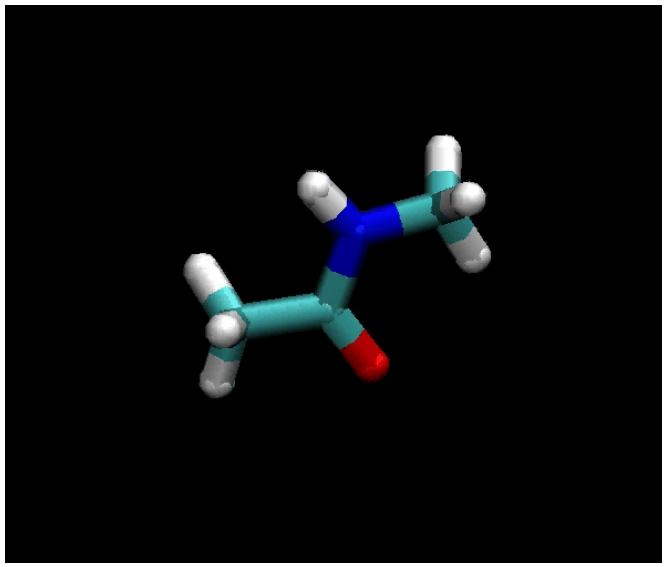

- ***Common Software: AMBER, CHARMM, NAMD, LAMMPS, GROMACS.***

- ***Used to simulate biological systems such as enzymes (Biological Factories)***

What can we Learn from Molecular Dynamics Simulations?

- **Key: Gives *time ordered* information.**
- **Allows simulation of how a biological system evolves with time.**
  - Reaction Rates
  - Activation Pathways
  - Structural Stability
  - Folding / unfolding pathways

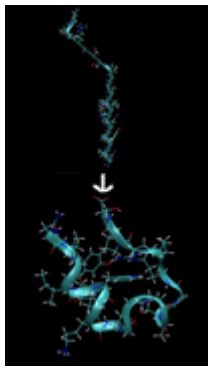

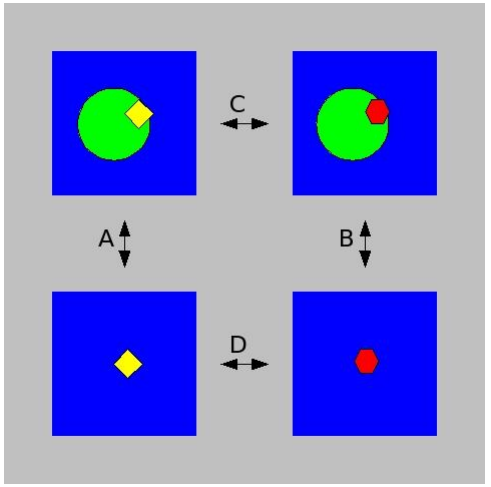

- **Can ‘experiment’ on a computer.**
  - Study mutation effects.
  - Environmental effects (temperature / pressure)
  - Study properties that cannot be measured experimentally.
  - ...

Why have we traditionally needed Supercomputers for MD? (Complex Equations)

$$U(R) = \sum_{bonds} K_r \left(r - r_{eq}\right)^2 + \sum_{angles} K_\theta \left(\theta - \theta_{eq}\right)^2$$

$$+ \sum_{dihedrals} \frac{V_n}{2}\left(1 + \cos\left[n\phi - \gamma\right]\right) + \sum_{i<j}^{atoms} \frac{A_{ij}}{R_{ij}^{12}} - \frac{B_{ij}}{R_{ij}^{6}}$$

$$+ \sum_{i<j}^{atoms} \frac{q_i q_j}{\varepsilon R_{ij}}$$

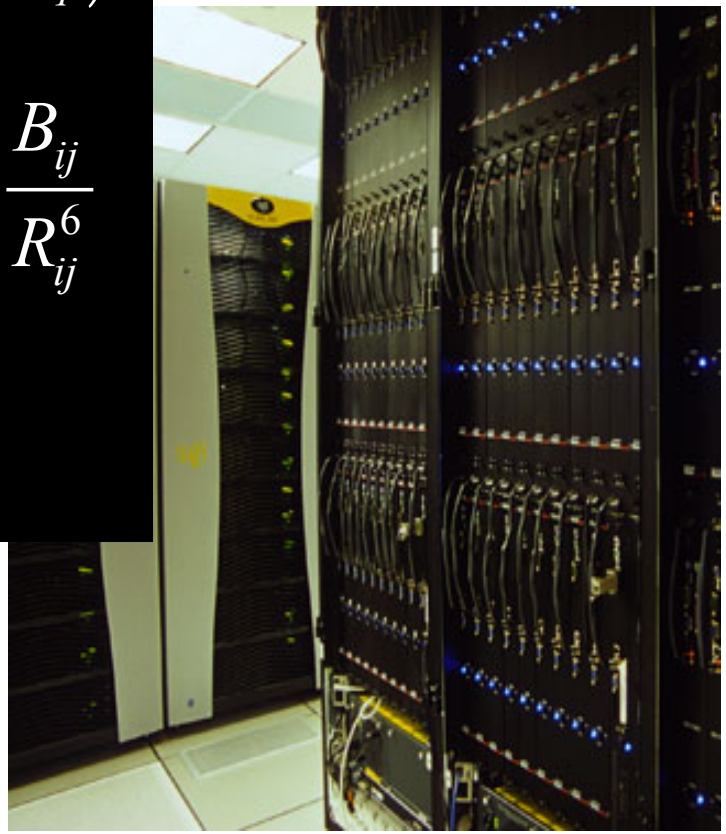

Why have we traditionally needed Supercomputers for MD? Lots of Atoms

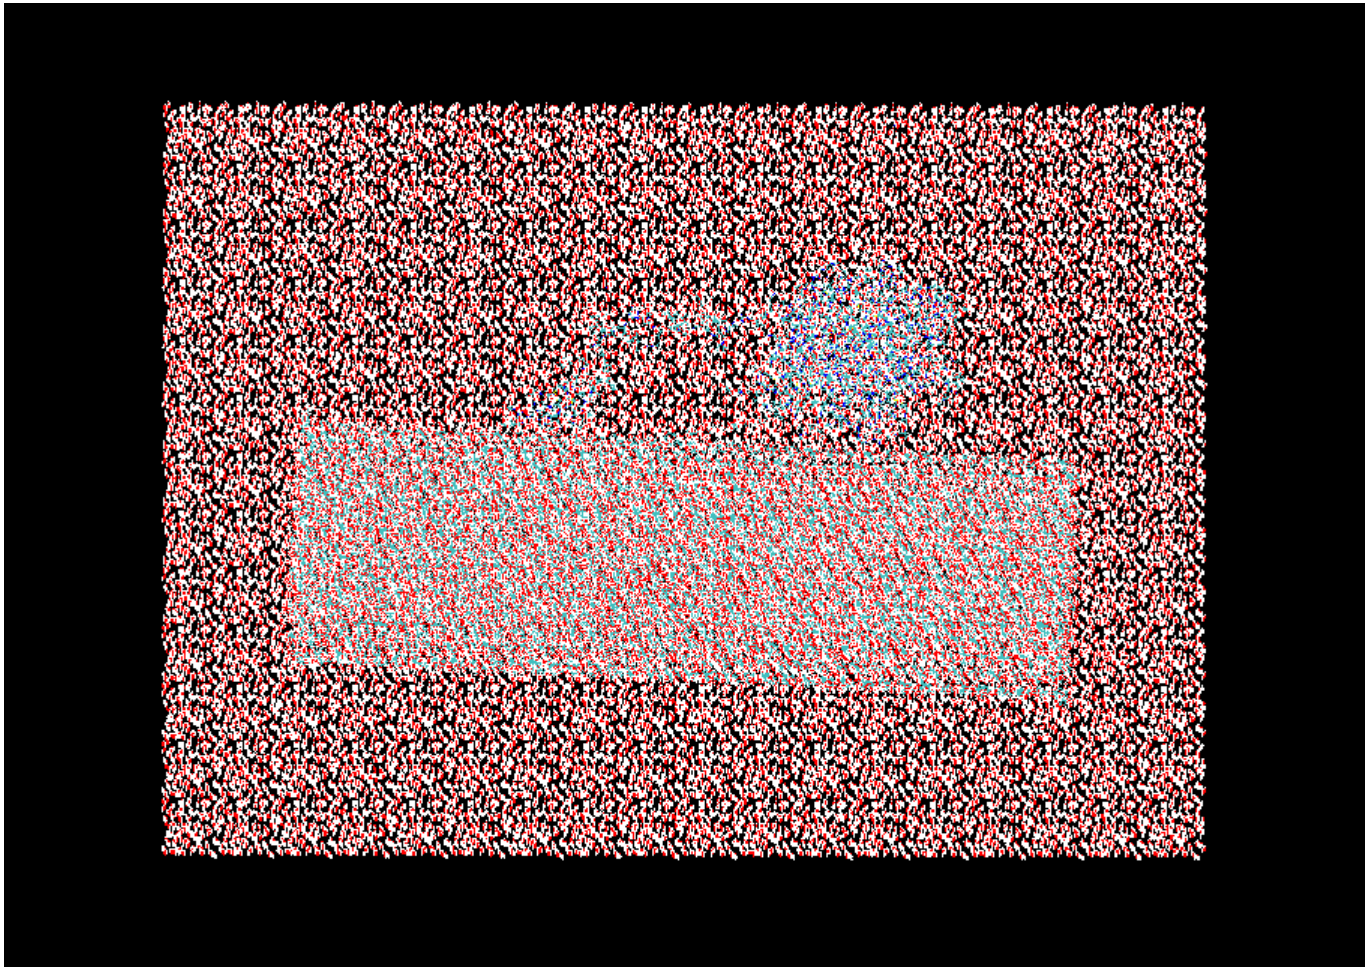

Why have we traditionally needed Supercomputers for MD? Lots of Time Steps

- ***Maximum time per step is limited by fastest motion in system (vibration of bonds).***
  = 2 femto seconds (0.000000000000002 seconds)(Light travels 0.006mm in 2 fs)

- ***Biological activity occurs on the nano-second to micro-second timescale.***
  1 micro second = 0.000001 seconds

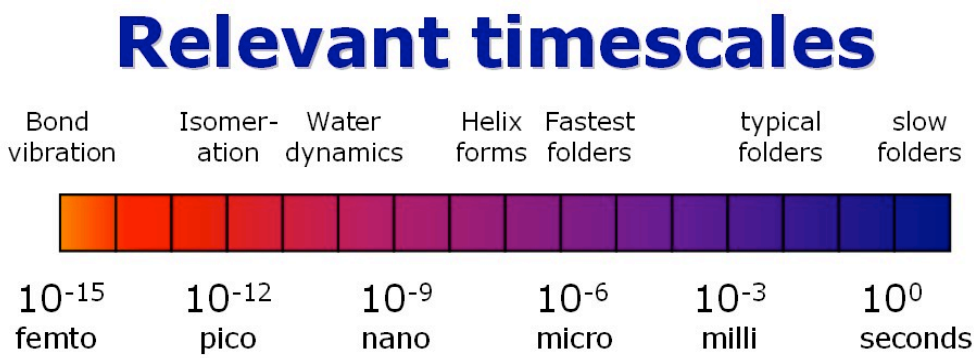

- **16 order of magnitude range**
  - Femtosecond timesteps
  - Need to simulate micro to milliseconds

SO WE NEED

500 million steps to reach 1 microsecond!!!

Why accelerate AMBER with an Intel® Xeon Phi™ coprocessor?

- ***Scaling to large machines does not work.***
- ***Too much communication at high core counts.***
- ***Coprocessors within a node reduce the communication overhead.***
- ***Intel® Xeon Phi™ coprocessors offer high computational power with a single node.***
- ***Intel® Xeon Phi™ coprocessors now supported for AMBER in offload and native mode.***

How does AMBER offload work on an Intel® Xeon Phi™ coprocessor?

CPU: Spatial decomposition

MIC: Thread-based decomposition

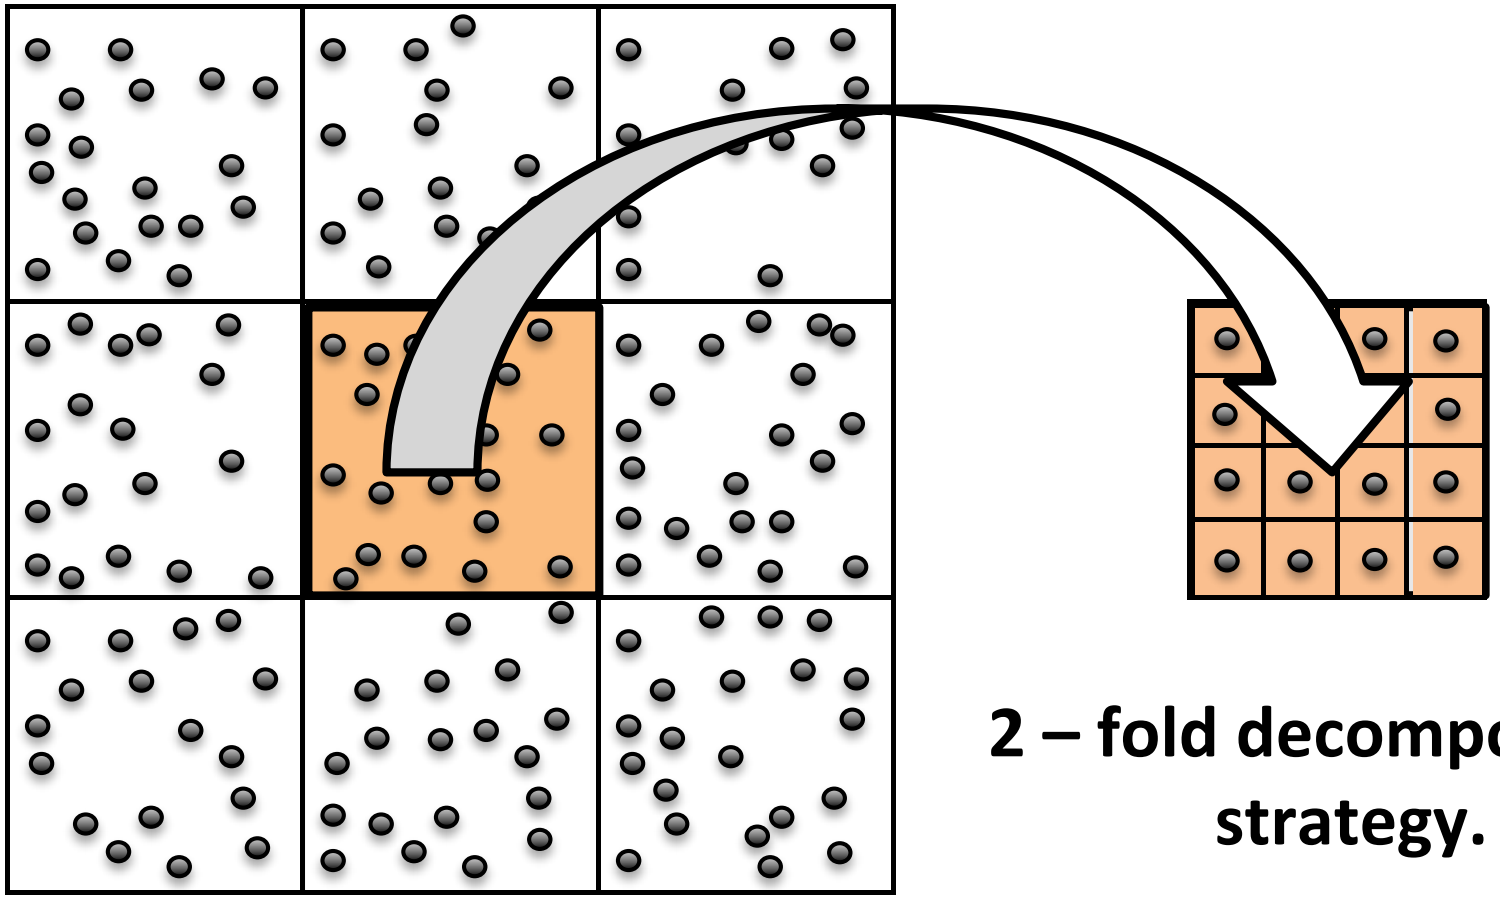

2 – fold decomposition strategy.

How does AMBER load-balance work in the offload algorithm?

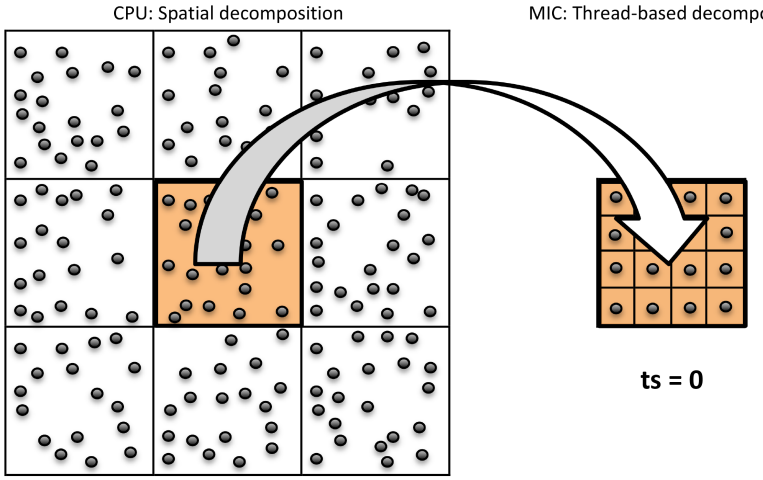

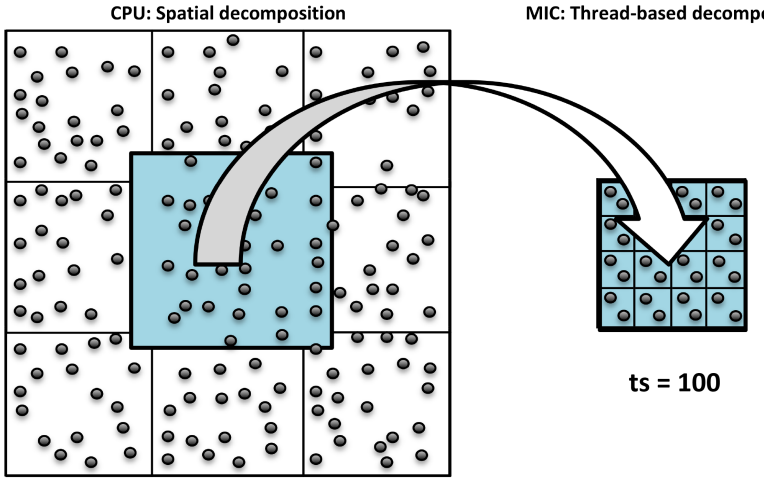

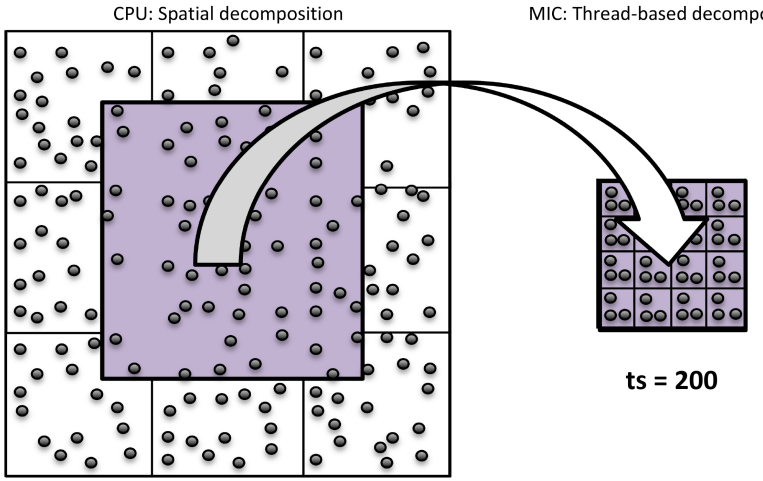

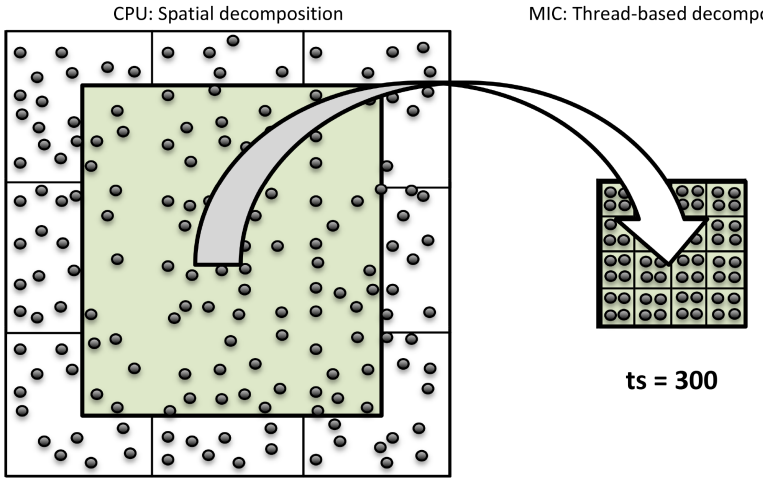

How does AMBER perform on Intel® Xeon® and Xeon Phi™ architectures?

CPU/MIC	Cellulose (408,609 atoms)	STMV (1,067,095 atoms)
unoptimized SNB (2 x 12 cores)	0.80	0.21
unoptimized IVB (2 x 12 cores)	1.40	0.34
optimized IVB (2 x 12 cores)	1.53	0.41
optimized IVB (2 x 12 cores) + Intel Xeon Phi 7120 (61 cores)	2.00	0.55

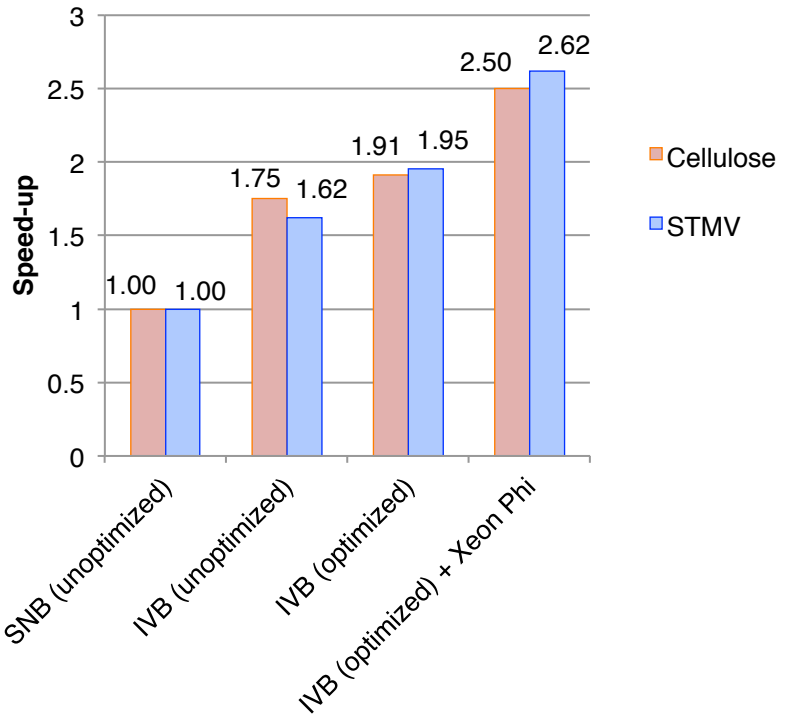

Performance of Cellulose (408,609 atoms) and STMV (1,067,095 atoms) measured in ns day^{-1} using 2 Intel Xeon E5 2660 v2 processor (2 x 12 cores) code name Sandy Bridge (SNB), 2 Intel Xeon 2697 v2 processor (2 x 12 cores) code name Ivy Bridge (IVB) and 2 Intel Xeon 2697 v2 processor (2 x 12 cores) code

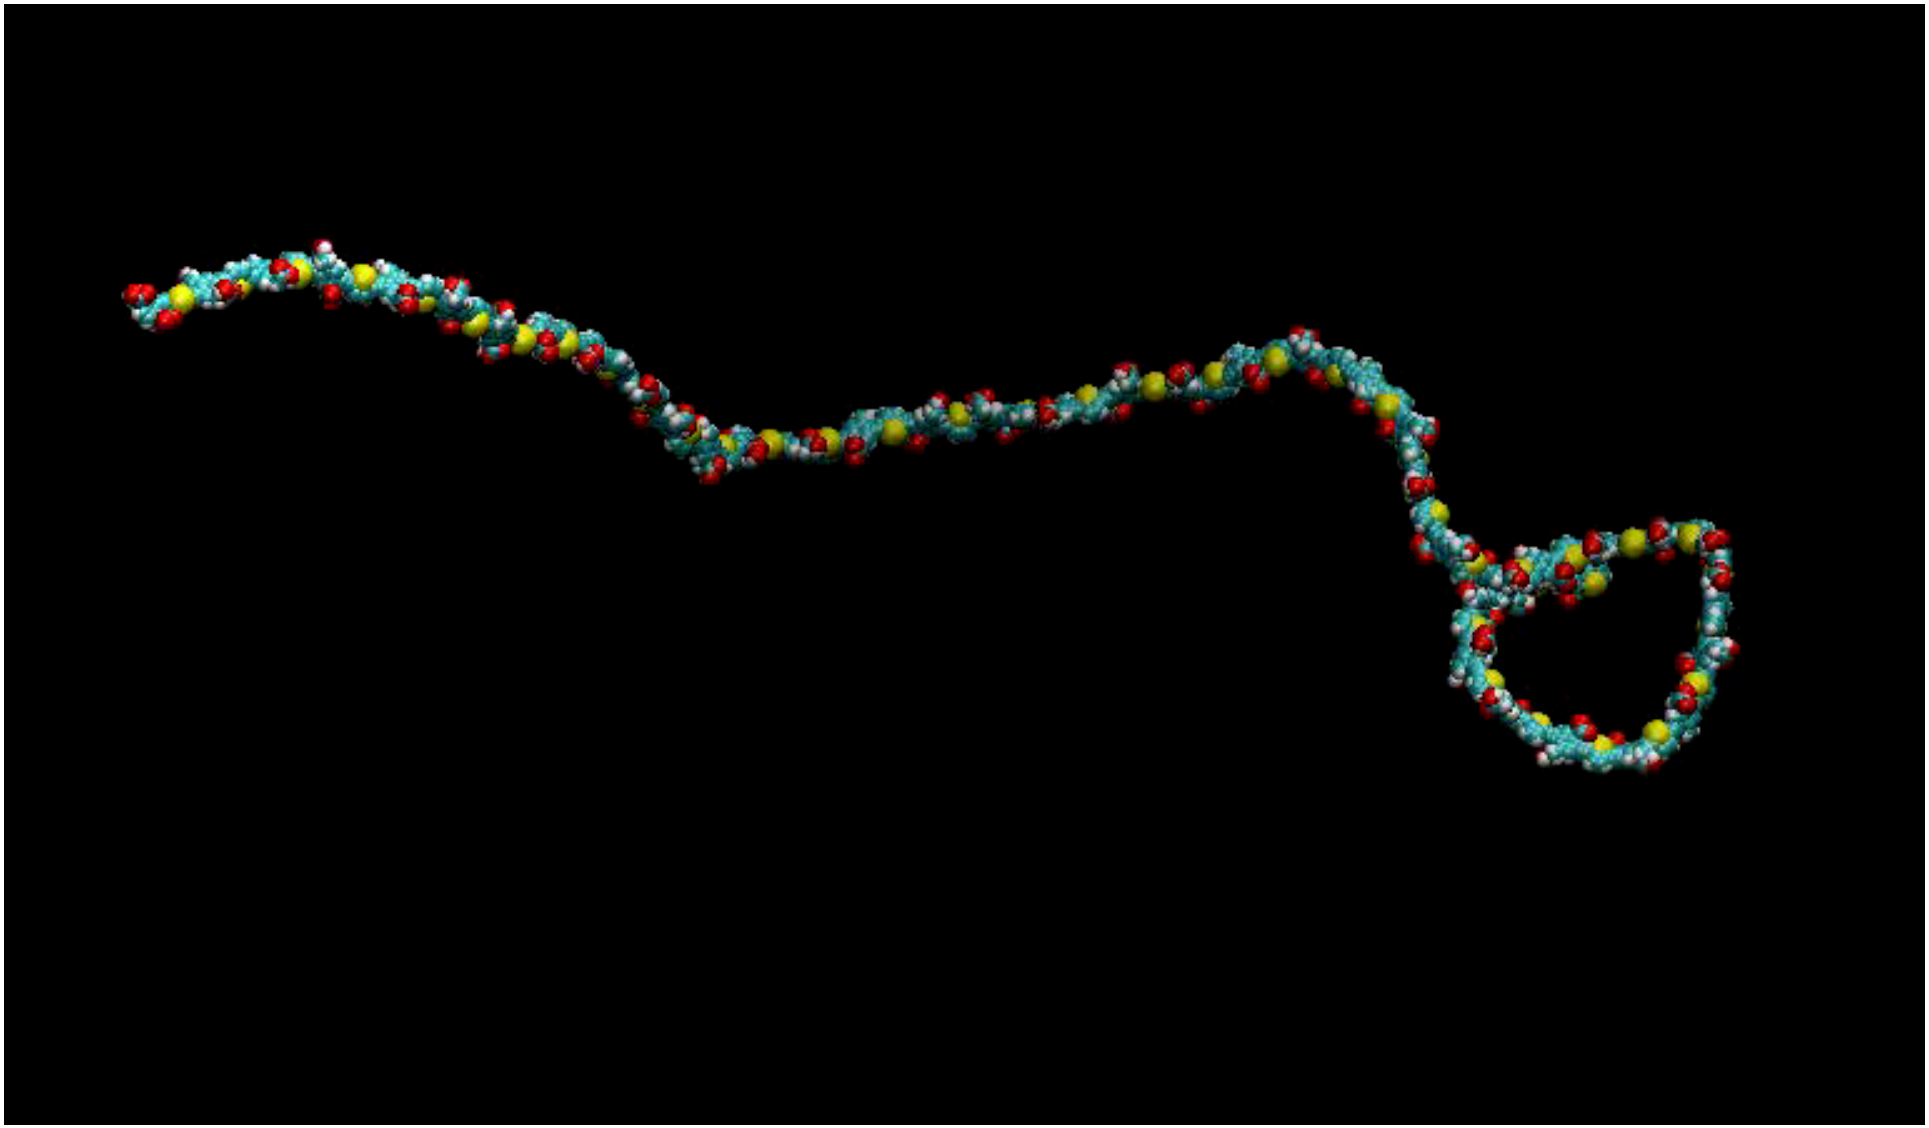

Conclusions and Long Term Vision

Conclusions

- **Partnership with Intel has been very successful and productive.**
- **Performance of latest offload code is currently > 2x baseline.**

Long Term Visions

- **Intel® Xeon® & Intel® Xeon Phi™ performance competitive with GPUs.**
- **Scaling to large node counts.**
- **Code extensively optimized but still clean and easily maintainable.**

Thank you for listening