


NAMD and Charm++ on Xeon Phi




James Phillips

Beckman Institute, University of Illinois


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

Computational Microscopy


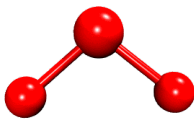
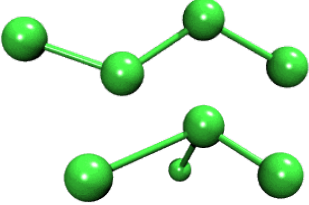
Ribosome: synthesizes proteins from genetic information, target for antibiotics



Silicon nanopore: bionanodevice for sequencing DNA efficiently



Molecular Mechanics Force Field

$$\begin{aligned}
 U(\vec{R}) = & \underbrace{\sum_{bonds} k_i^{bond} (r_i - r_0)^2}_{U_{bond}} + \underbrace{\sum_{angles} k_i^{angle} (\theta_i - \theta_0)^2}_{U_{angle}} + \\
 & \underbrace{\sum_{dihedrals} k_i^{dihedral} [1 + \cos(n_i \phi_i + \delta_i)]}_{U_{dihedral}} + \\
 & \underbrace{\sum_i \sum_{j \neq i} 4\epsilon_{ij} \left[\left(\frac{\sigma_{ij}}{r_{ij}} \right)^{12} - \left(\frac{\sigma_{ij}}{r_{ij}} \right)^6 \right] + \sum_i \sum_{j \neq i} \frac{q_i q_j}{\epsilon r_{ij}}}_{U_{nonbond}}
 \end{aligned}$$

Classical Molecular Dynamics

Energy function: $U(\vec{r}_1, \vec{r}_2, \dots \vec{r}_N) = U(\vec{R})$

used to determine the force on each atom:

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

Newton's equation represents a set of N second order differential equations which are solved numerically via the Verlet integrator at discrete time steps to determine the trajectory of each atom.


$$\vec{r}_i(t + \Delta t) = 2\vec{r}_i(t) - \vec{r}_i(t - \Delta t) + \frac{\Delta t^2}{m_i} \vec{F}_i(t)$$

Small terms added to control temperature and pressure.

nature

THE INTERNATIONAL WEEKLY JOURNAL OF SCIENCE

THE HIV-1 CAPSID





Atomic
structure
of the AIDS
pathogen's
protein coat

PAGE 643


2013 *HPCwire* Editors' Choice Award for Best Use of HPC in Life Sciences



HIV Infective Cycle



HIV Treatment



NIH Biomedical Technology Research Center for Macromolecular Modeling and Bioinformatics

Developers of the widely used computational biology software VMD and NAMD

250,000 registered VMD users
61,000 registered NAMD users

600 publications (since 1972)
over 54,000 citations

5 faculty members
8 developers
1 systems
administrator
17 postdocs
46 graduate students
3 administrative staff

*Renewed 2012-2017
with 10.0 score (NIH)*

research projects include: virus
capsids, ribosome, photosynthesis,
protein folding, membrane reshaping,
animal magnetoreception

Achievements Built on People



Tajkorshid, Luthey-Schulten, Stone, Schulten, Phillips, Kale, Mallon

NAMD Serves NIH Users and Goals

Practical Supercomputing for Biomedical Research

- 60,000 users can't all be computer experts.
 - 18% are NIH-funded; many in other countries.
 - 17,000 have downloaded more than one version.
 - 4000 citations of NAMD reference papers.
- One program available on all platforms.
 - Desktops and laptops – setup and testing
 - Linux clusters – affordable local workhorses
 - Supercomputers – free allocations on XSEDE
 - Blue Waters – sustained petaflop/s performance
 - GPUs/Xeon Phi – next-generation supercomputing
- User knowledge is preserved across platforms.
 - No change in input or output files.
 - Run any simulation on **any number of cores**.
- Available free of charge to all.




Hands-On Workshops



Oak Ridge TITAN

Structural data drives simulations




NIH Center Facilities Enable Petascale Biology

Over the past five years the Center has assembled all necessary hardware and infrastructure to prepare and analyze petascale molecular dynamics simulations, and ***makes these facilities available to visiting researchers.***



External Resources,
90% of our
Computer Power



Petascale Gateway Facility


Storage



Compute



Visualization



High-End Workstations
Accessible to Visitors

Virtual Facilities Enable Petascale Anywhere




High-end visualization and analysis workstations currently available only in person at the Beckman Institute must be *virtualized and embedded at supercomputer centers*.

Storage




Compute

Visualization



Compressed Video



1 Gigabit Network




Remote Visualization Now

- TACC Stampede supports this today
 - Includes nodes with 1TB memory
 - Not virtualized, allocate full dedicated node
 - New Maverick cluster recently added
- Blue Waters – no visualization resource
- Titan – new Rhea “viz” cluster drops GPUs

External Users Race Ahead

Influenza Virus, 100 nm diameter, lipid/protein
210 million atoms fully solvated


Using TACC Stampede as on-ramp to petascale.




Jacob Durrant
Fearless Post-Doc
Amaro Lab, UCSD

TACC Stampede

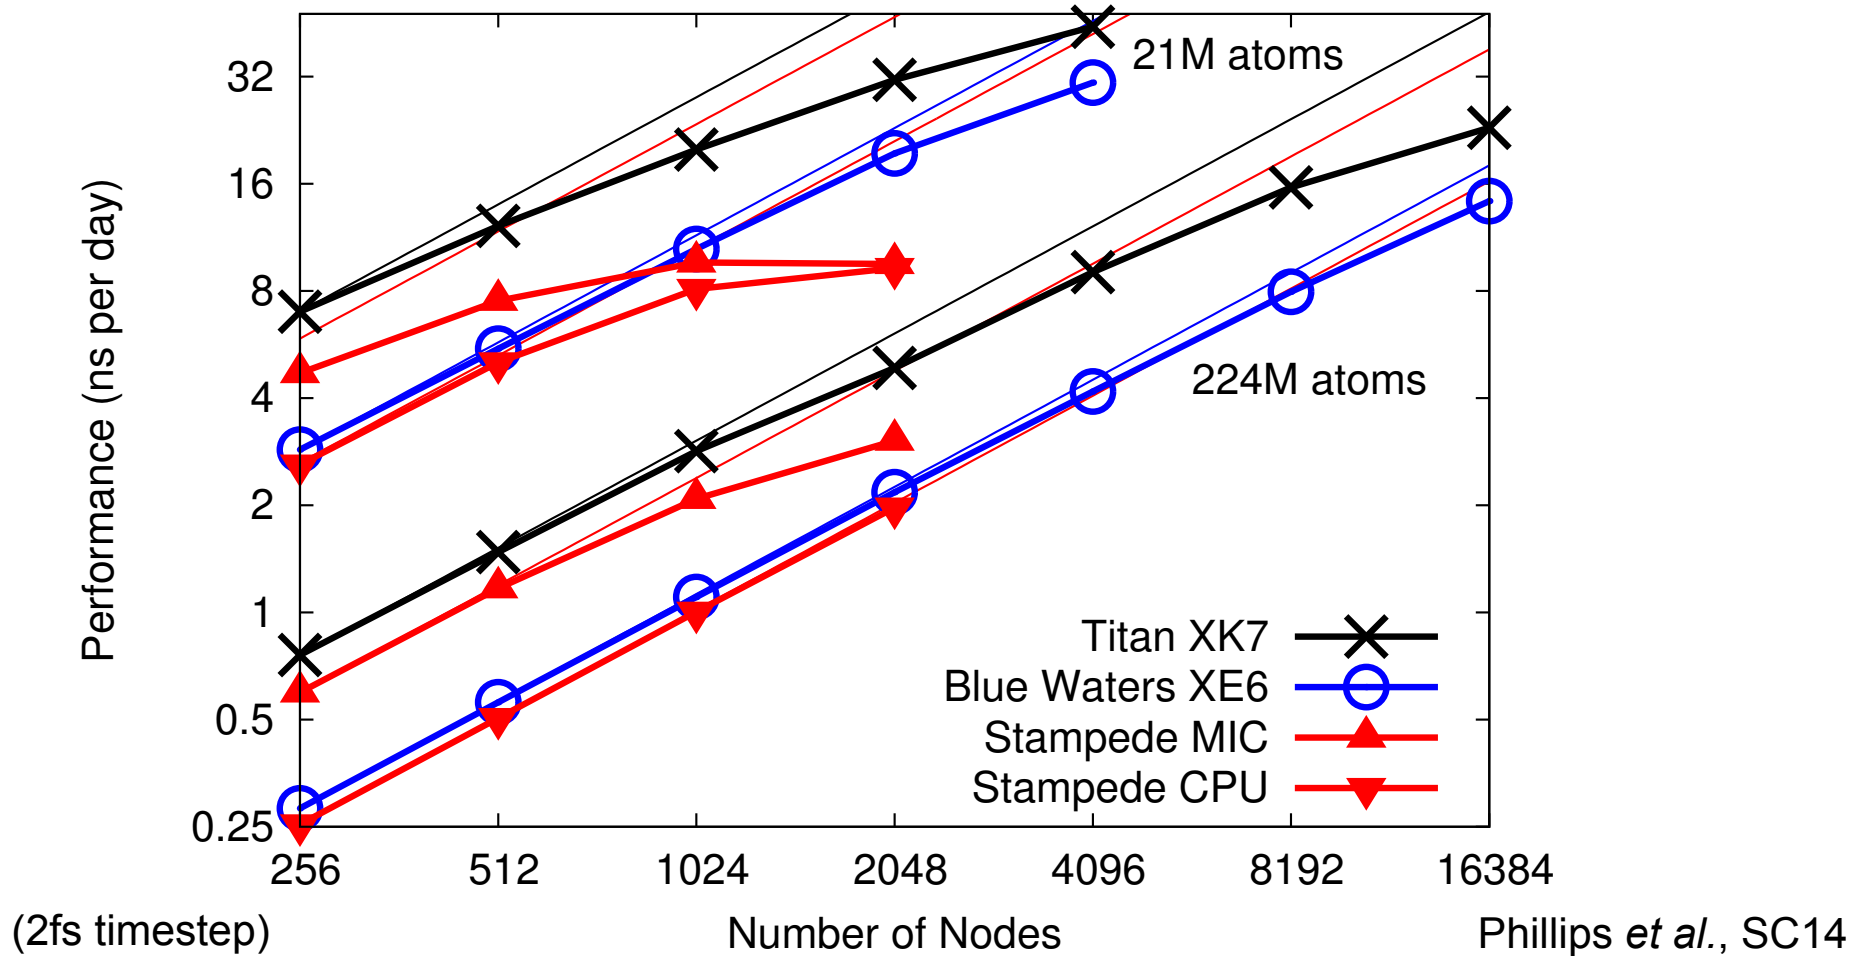
Remote visualization is also needed for performance analysis on petascale systems



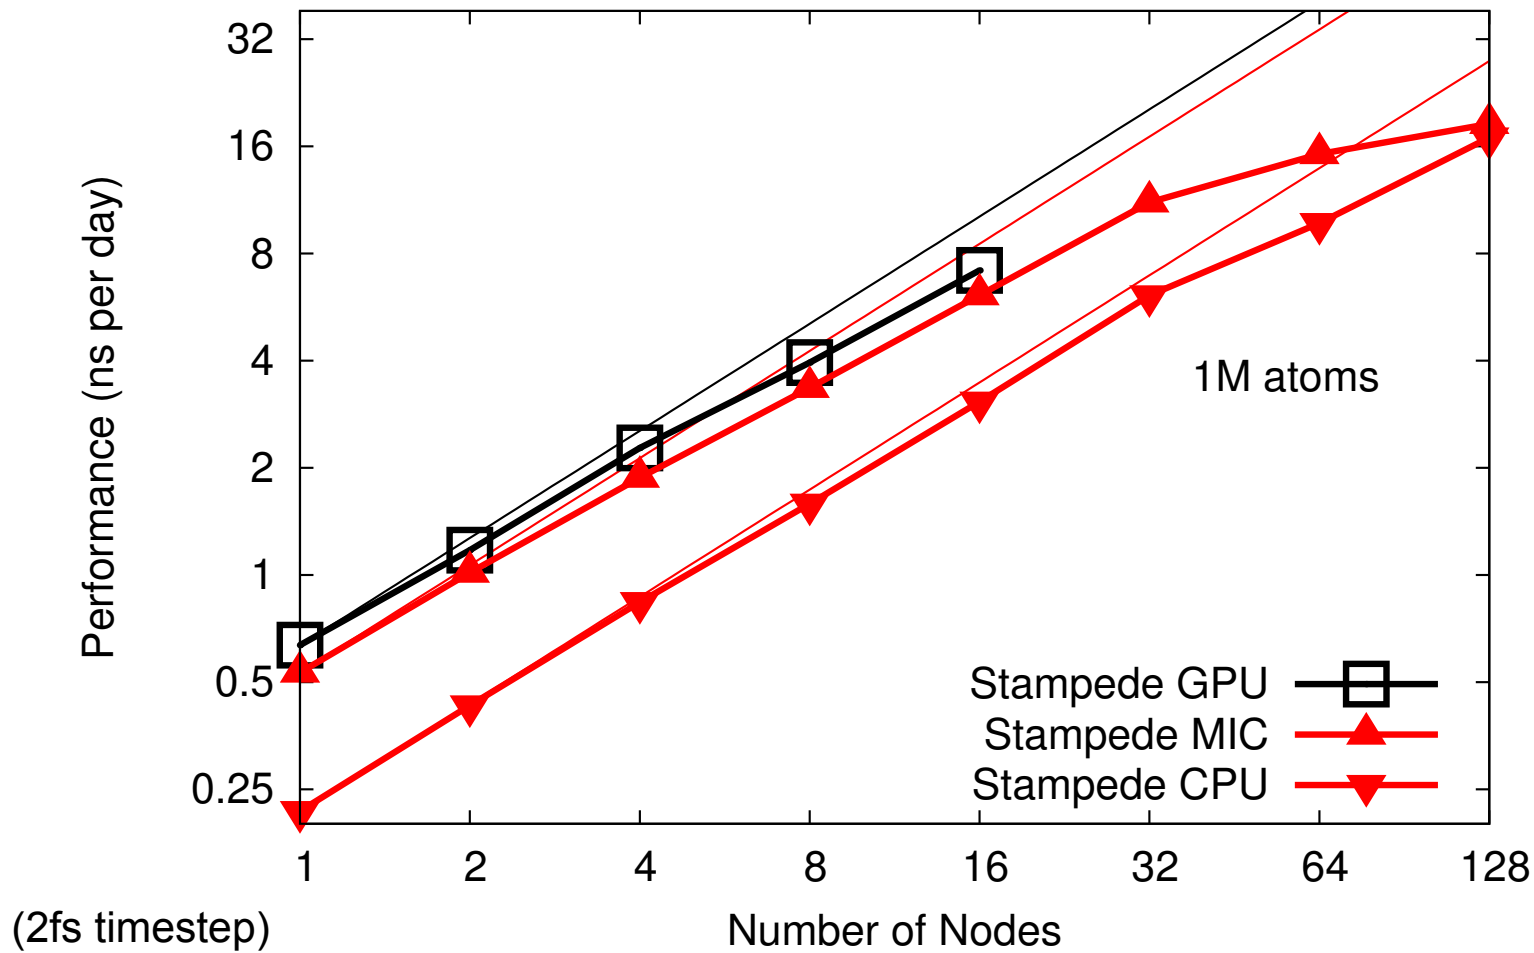
NAMD on Petascale Platforms



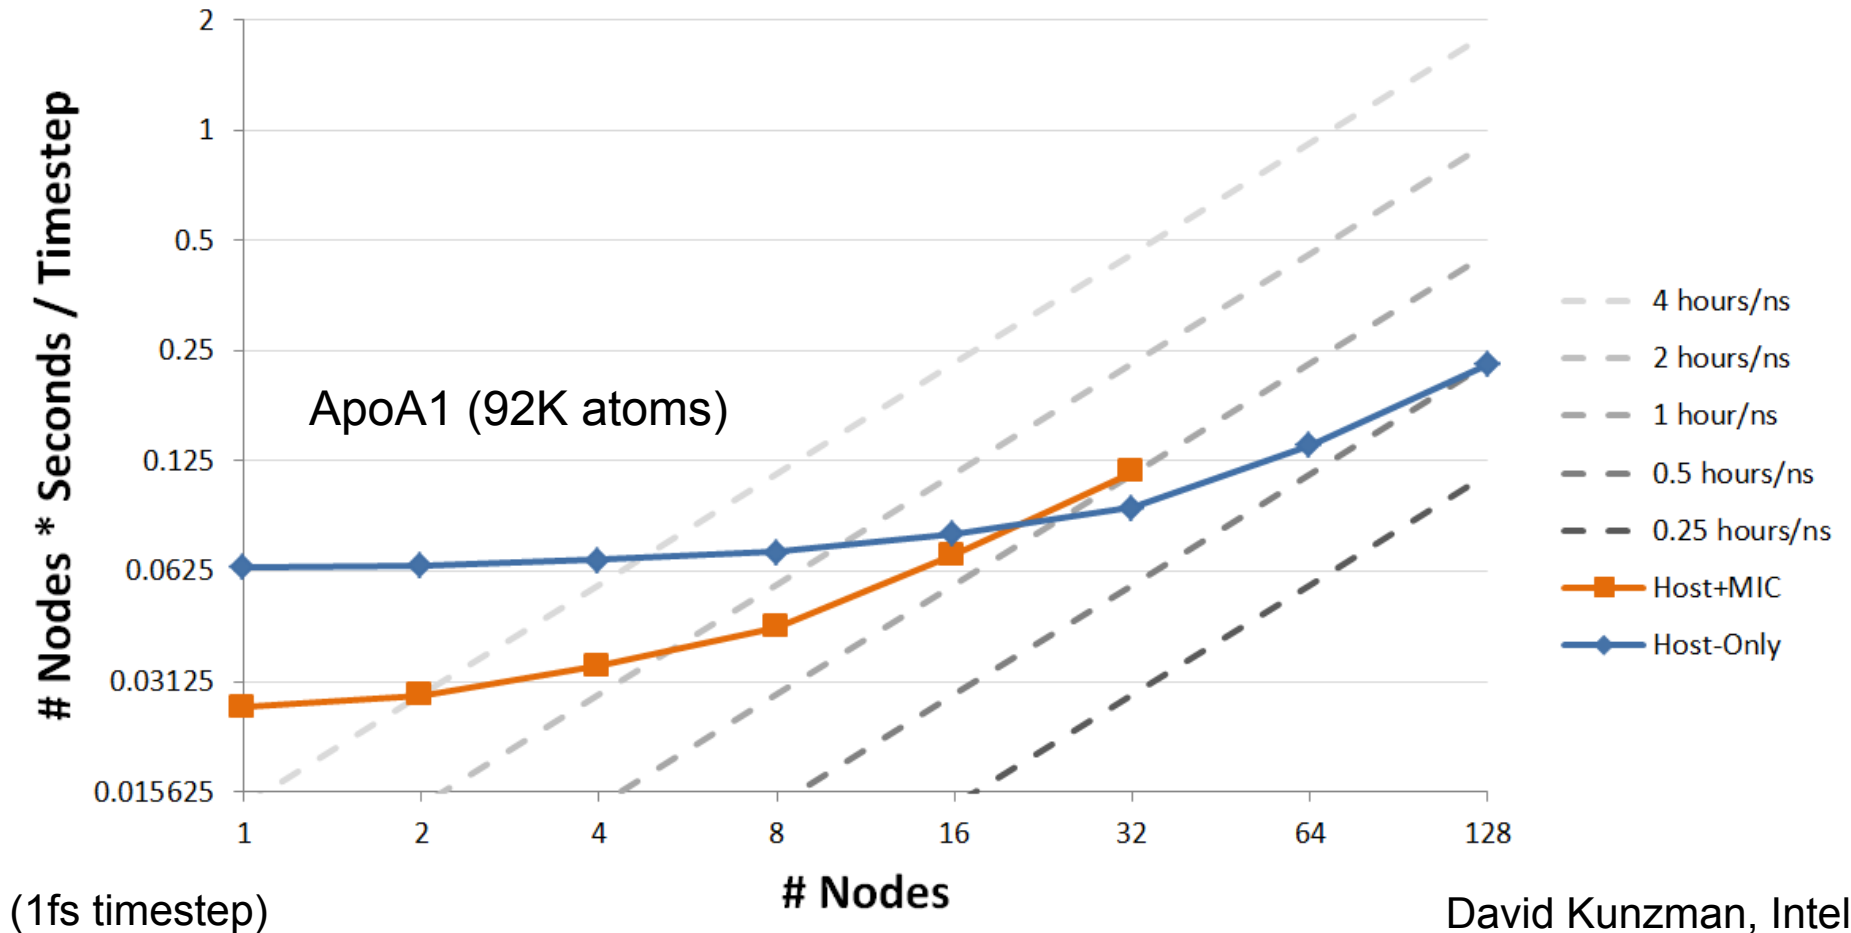
NAMD on Petascale Platforms




NAMD on TACC Stampede



NAMD on TACC Stampede



NAMD Benefits from Charm++




Charm++ Used by NAMD

- Parallel C++ with *data driven* objects.
- Asynchronous method invocation.
- Prioritized scheduling of messages/execution.
- Measurement-based load balancing.
- Portable messaging layer.

NAMD Hybrid Decomposition

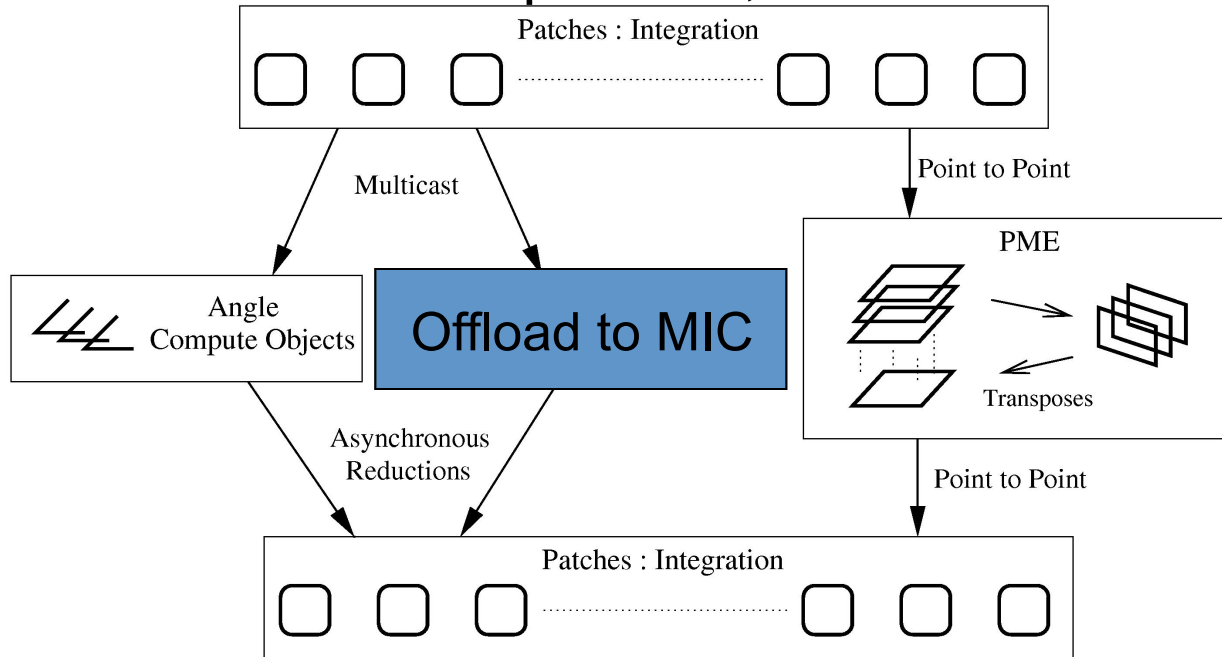
Kale *et al.*, *J. Comp. Phys.* 151:283-312, 1999.



- Spatially decompose data and communication.
- Separate but related work decomposition.
- “Compute objects” facilitate iterative, measurement-based load balancing system.

NAMD Overlapping Execution

Phillips *et al.*, SC2002.



Objects are assigned to processors and queued as data arrives.

Charm++ and MIC Options

- Naïve Offload
- Aggregated Offload
- Coprocessor Only
- Heterogeneous Cluster
- OpenMP Teams

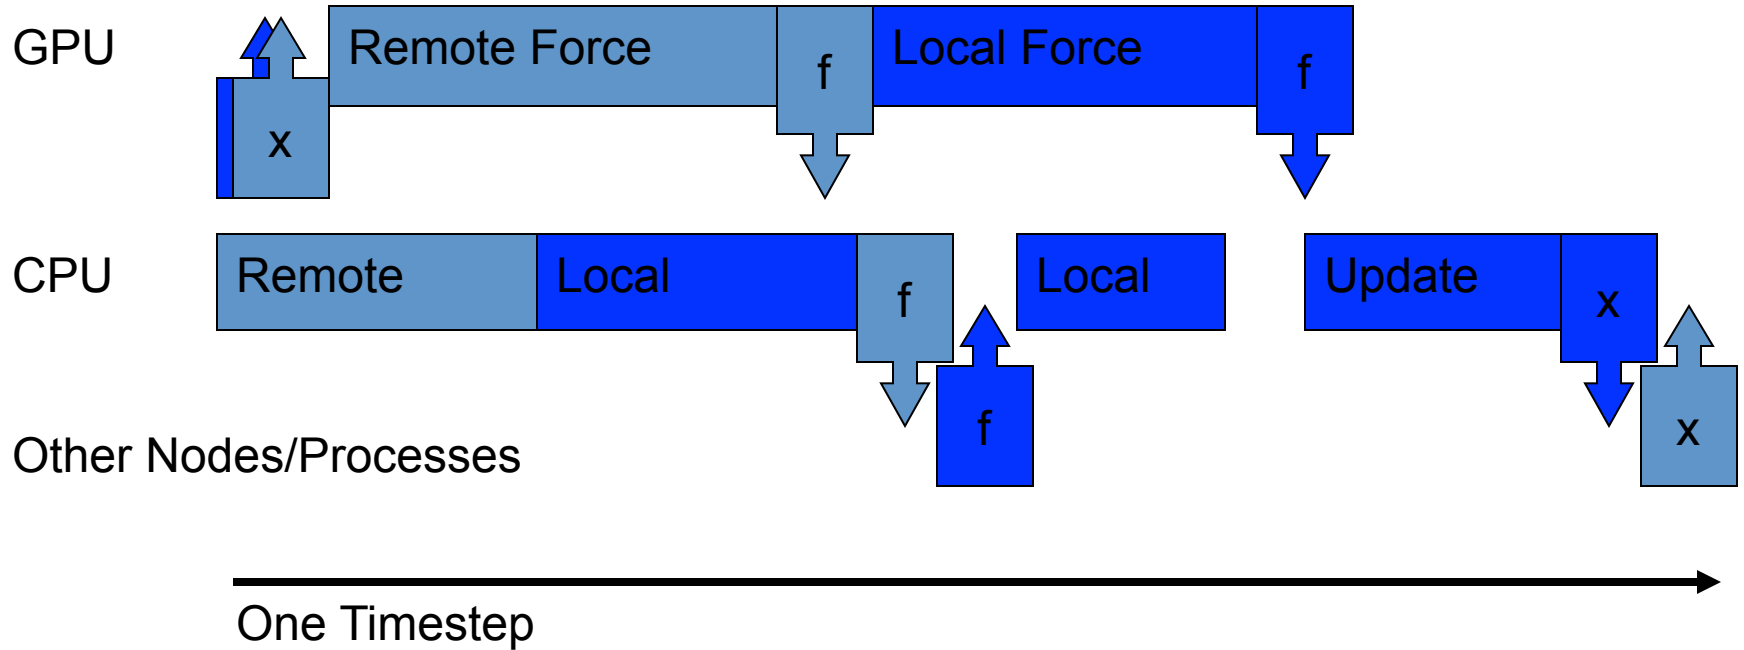
Charm++ and MIC Options

- Naïve Offload
 - Just add offload directives
 - Close your eyes and trust the runtime
 - How many host threads can offload at once?
 - How do they share the MIC cores?
 - NAMD work grainsize is too small for MIC

Charm++ and MIC Options

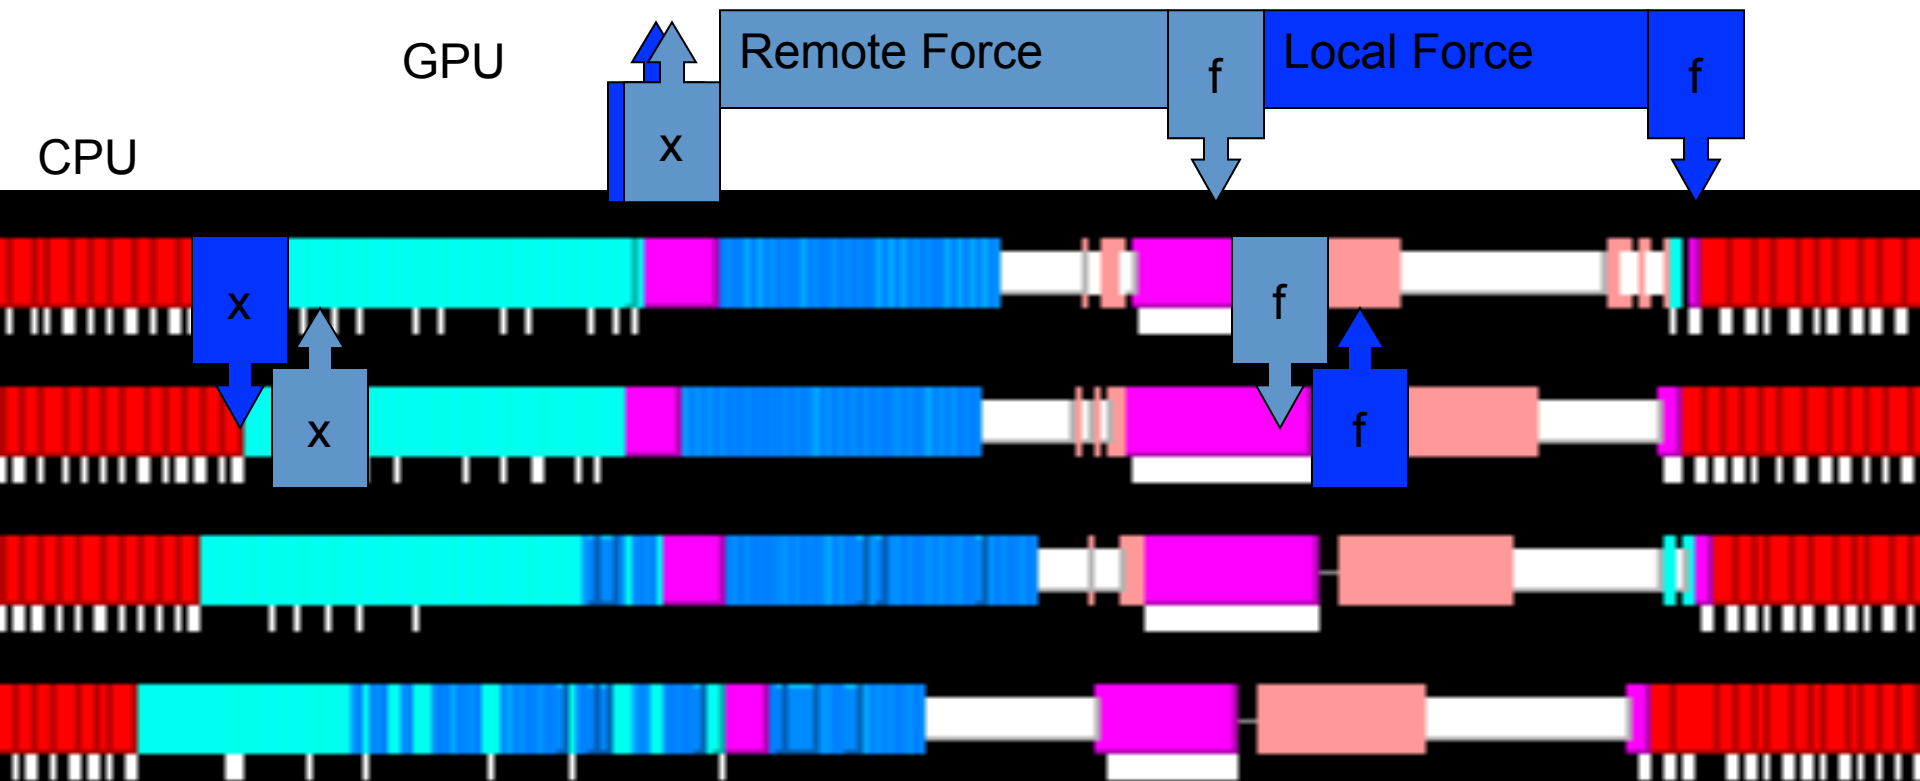
- Aggregated Offload
 - Keep NAMD work decomposition
 - Collect and bulk-copy data
 - Bulk-launch tasks in single offload
 - Method initially developed for CUDA
 - NAMD MIC offload is clone of CUDA offload

Overlapping GPU and CPU with Communication



Actual Timelines from NAMD

Generated using Charm++ tool “Projections” <http://charm.cs.uiuc.edu/>



Charm++ and MIC Options

- Coprocessor Only
 - Ignore host, run Charm++ scheduler on MIC
 - Requires all code to run acceptably on MIC
 - Not the case for KNC, hopefully for KNL
 - Fine for single card, maybe not for multi-node
 - Useful for preparing for KNL processor

Charm++ and MIC Options

- Heterogeneous Cluster
 - Treat host and MIC as two Charm++ nodes
 - Requires adapting to different core counts
 - Requires adapting to different core speeds
 - Performance ratio varies by function and **data**
 - E.g., self, face, and corner computes in NAMD
 - Full employment for computer scientists!

Charm++ and MIC Options

- OpenMP Thread Teams
 - Grainsize too large for single MIC thread
 - Grainsize too small for entire MIC
 - Let Charm++ control OpenMP thread teams
 - E.g, MIC = 15 Charm++ PEs
 - Each PE = 4 cores and 16 threads
 - Parallelize loops using OpenMP directives


MIC Vectorization Options

- Intrinsic – what we have
 - Written by David Kunzmann of Intel
 - Currently in production
- Compiler – what we want
 - `#pragma [omp] simd assert`
 - Currently ~20% slower with refactored kernel
 - “One missed optimization” but improving fast
- ISPC – our backup plan
 - Similar to CUDA, consider if compiler fails

Conclusions

- NAMD for Xeon Phi co-processors currently in production on TACC Stampede.
 - Available in nightly builds and upcoming 2.10 release.
 - Performance comparable to GPUs, scaling poor.
 - Implementation by David Kunzman of Intel (for details see his one-hour talk at last year's IXPUG meeting).
- Targeting KNL processors (e.g., NERSC Cori).
- Hoping to vectorize using `#pragma [omp] simd`.

Thanks to: NIH, NSF, DOE, NCSA, TACC, Intel(**David Kunzman**),
NVIDIA (Sarah Tariq, Patric Zhao, Sky Wu, Justin Luitjens, Nikolai Sakharlykh),
Cray (Sarah Anderson, Ryan Olson), NCSA (Robert Brunner),
PPL (Eric Bohm, Yanhua Sun, Gengbin Zheng, Nikhil Jain)
and 19 years of NAMD and Charm++ developers and users.



James Phillips

Beckman Institute, University of Illinois
<http://www.ks.uiuc.edu/Research/namd/>