


What is Charm++?

Charm++ is an asynchronous and message-driven paradigm developed by Parallel Programming Lab (PPL). Our goal is to develop technology that improves performance of parallel applications while also improving programmer productivity.

- ▶ Fine grain decomposition and processor virtualization
- ▶ Asynchronous message driven
- ▶ Adaptive and portable runtime system
- ▶ Automatic dynamic Load balancing
- ▶ Scalable fault tolerance up to thousands of processors
- ▶ Minimizing energy, power and execution time
- ▶ AMPI - Adaptive Message Passing Interface
- ▶ XCharJ: MultiParadigm Programming with compiler support



What is NAMD?

NAMD is a parallel, object-oriented molecular dynamics code designed for high-performance simulation of large biomolecular systems. It is the result of an interdisciplinary collaboration between Prof. Kale, computer science Prof. Robert D. Skeel, and physics Prof. Klaus J. Schulten at the Theoretical and Computational Biophysics Group (TCBG) of Beckman Institute.

- ▶ Due to atomic-level time and length scales, each time step is 1 fsec
- ▶ A meaningful simulation requires 1 microsecond or longer (billions of steps)
- ▶ Each timestep is carried out in milliseconds, it requires two months to complete a microsecond simulation
- ▶ How to achieve lower and lower timestep using more processors?

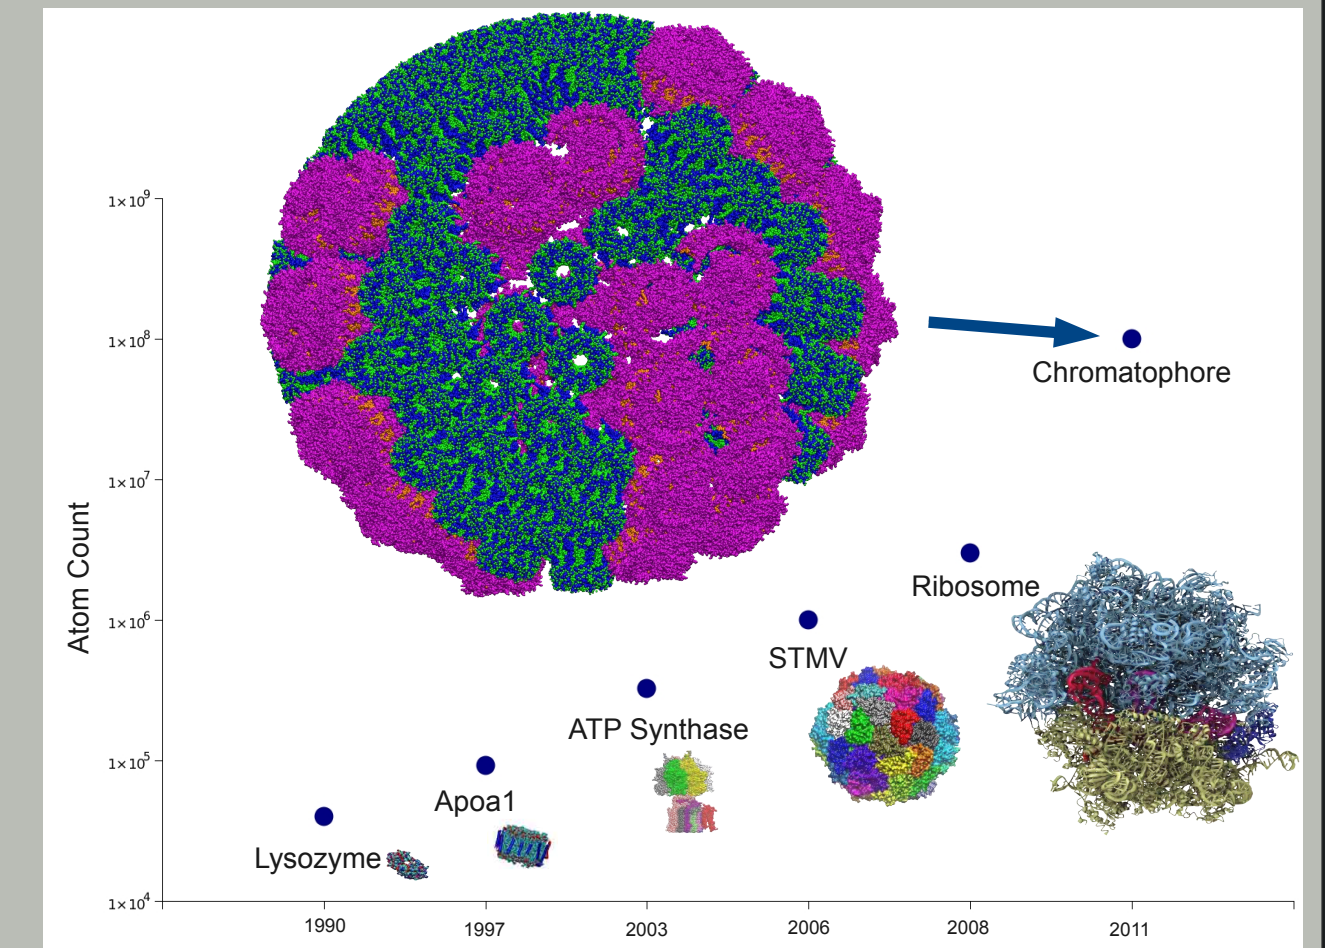
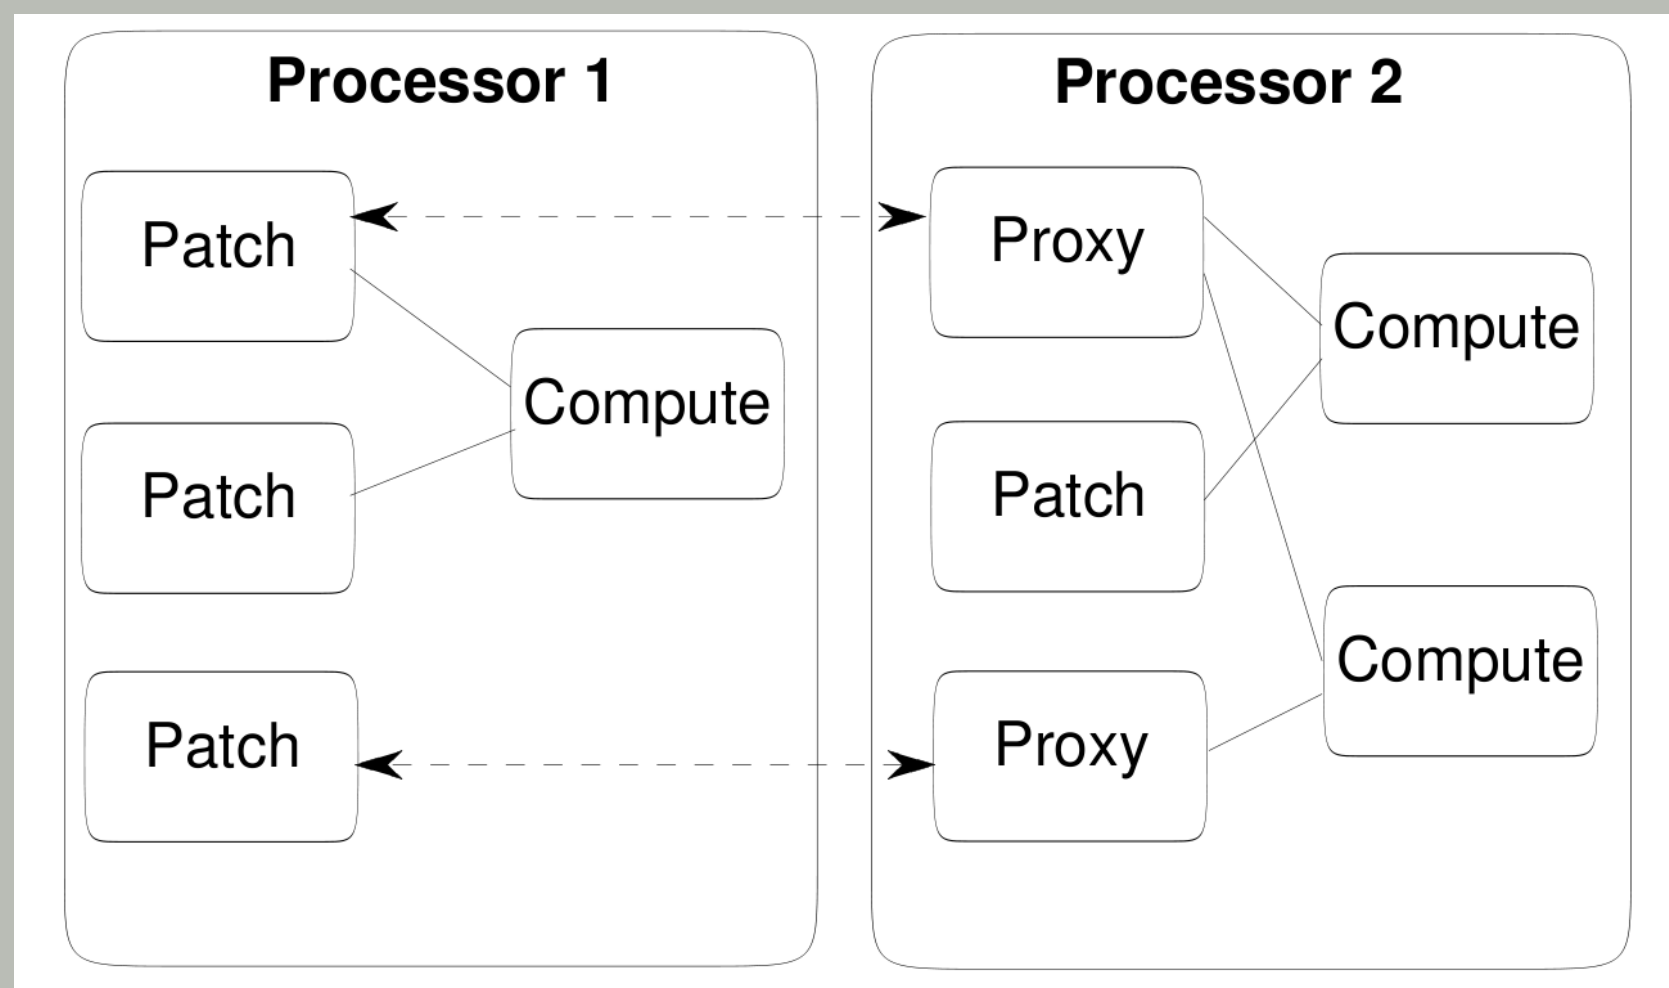



Figure: size of molecular systems that can be studied using all-atom molecular dynamics simulations has steadily increased from that of Lysozyme (40,000 atoms) in the 1990s to 100-million atoms as in the spherical chromatophore model shown above.

Parallelizing NAMD Using Charm++

- ▶ Reduce computation using short-range and long-range forces
- ▶ Non-bonded computation is intensive
- ▶ Spatial and force decomposition
- ▶ Overlap computation and communication
- ▶ Parallel I/O for large systems
- ▶ Load balancing
- ▶ Fine grain decomposition to increase parallelism



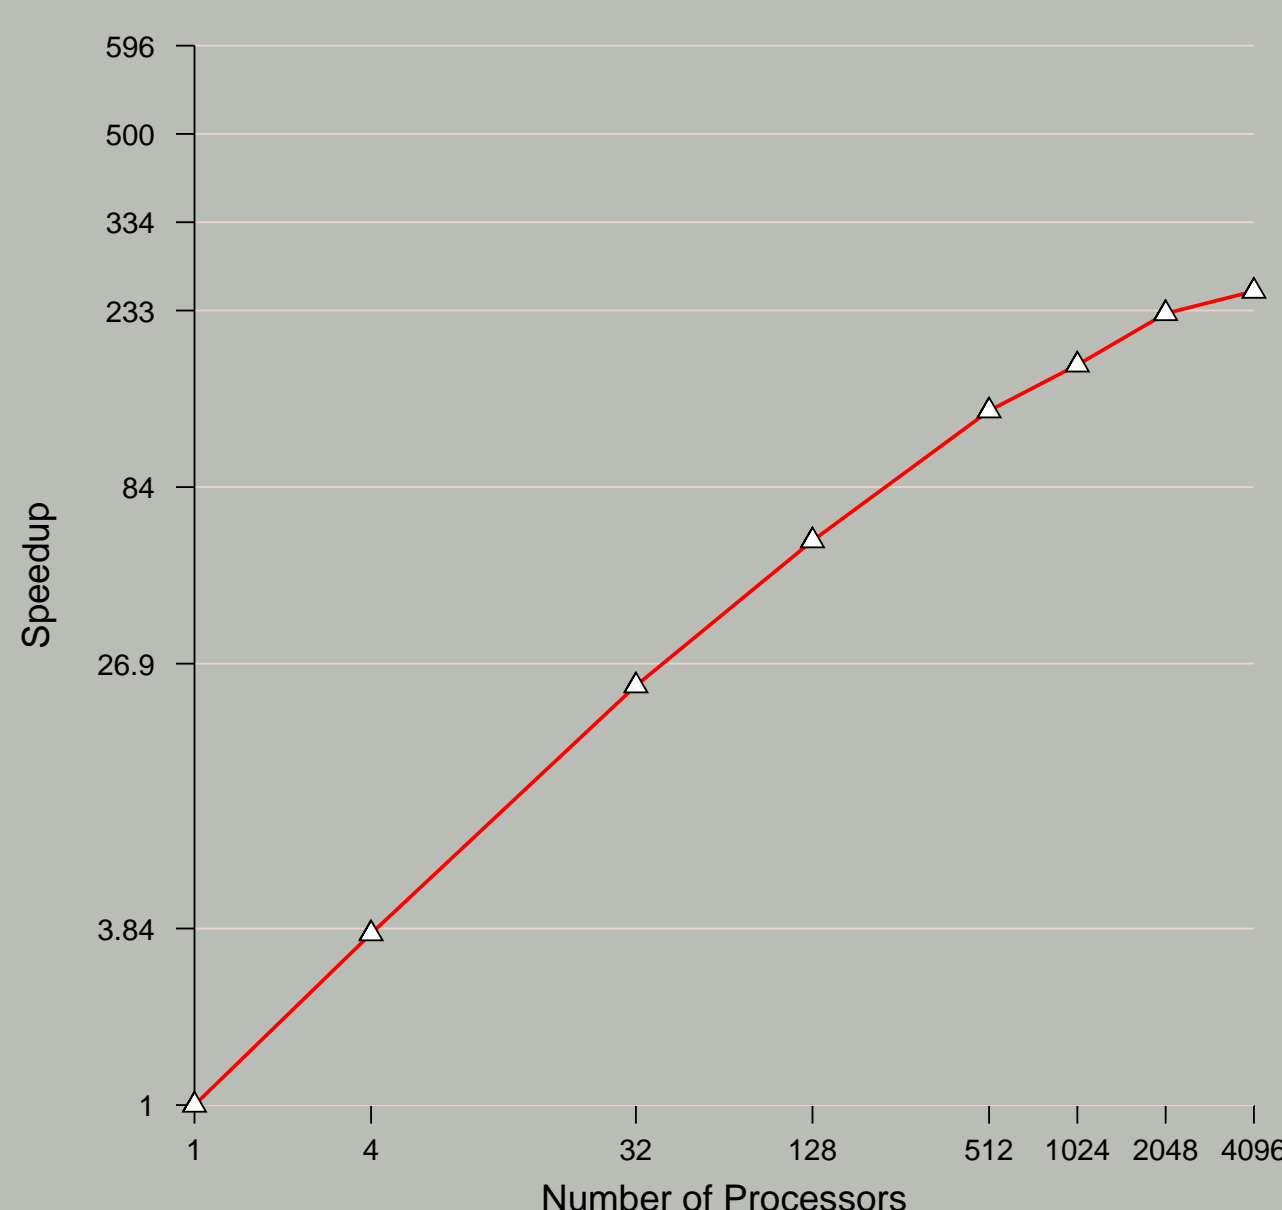
Performance Visualization and Analysis



- Projections framework
- ▶ Projections for performance tuning
 - ▶ Instrument data during run-time
 - ▶ Detect load balancing problem
 - ▶ Communication problems
 - ▶ Using Projections to find load balancing problem and fix it

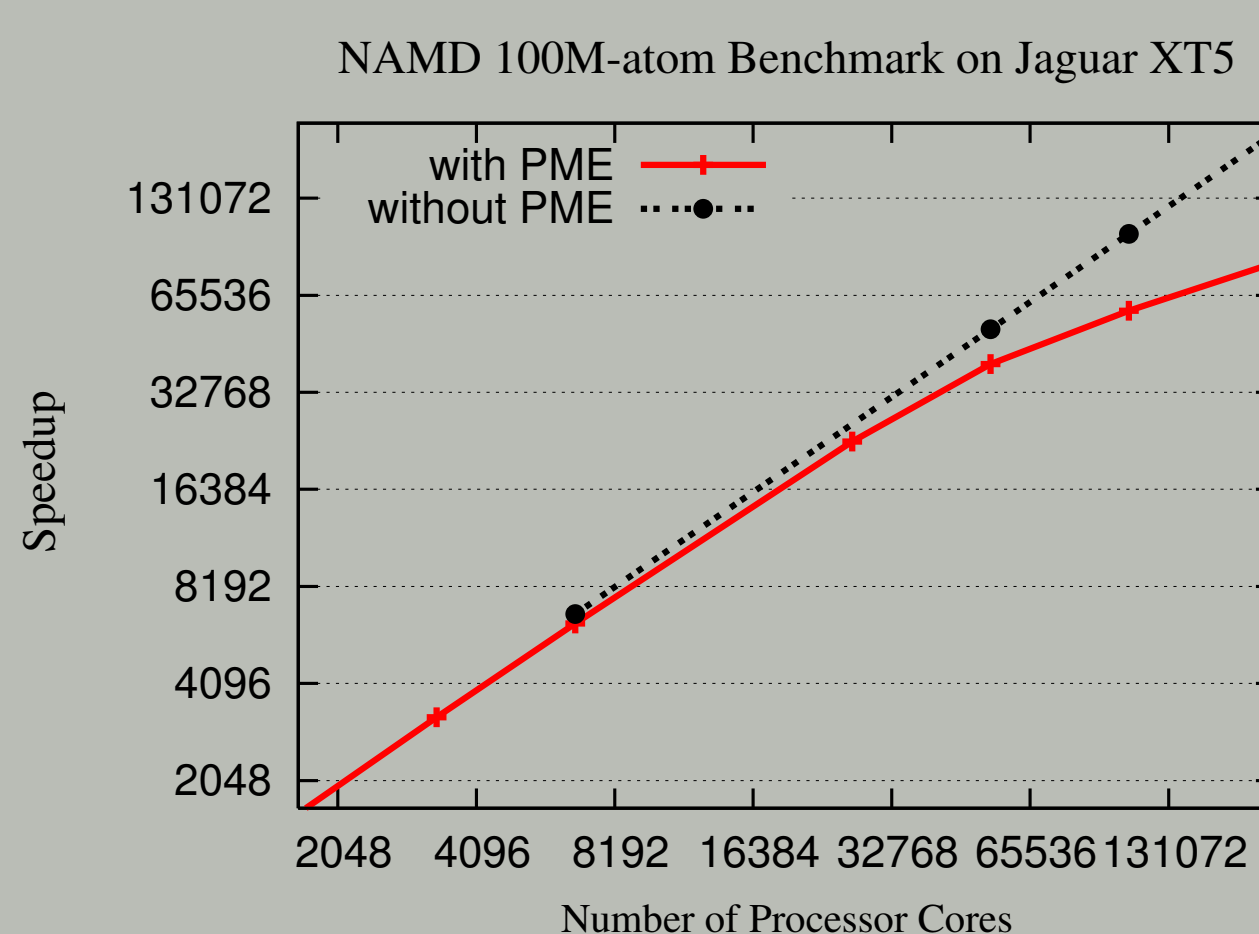
Strong Scaling for Small System

- ▶ DHFR - 23,558 atoms system
- ▶ Blue Gene/P at Argonne National Lab
 - ▶ 40,960 quad-core processors
 - ▶ Peak performance: 557 teraflops



100M-atoms Simulation

- ▶ Jaguar at Oak Ridge National Lab
 - ▶ 224,256 AMD Opteron cores
 - ▶ Peak 2.3Petaflops
- ▶ 9ms/step on full machine
- ▶ 93% parallel efficiency (vs 6720 cores)
- ▶ Long-range force calculation (PME) does not scale well
- ▶ Optimizing all-to-all is challenging



Blue Waters @ Illinois

- ▶ Cray Inc and NCSA
- ▶ 49,000 AMD Opteron 6200 Processors
- ▶ 3,000 NVIDIA GPUs
- ▶ 1.5PB system memory, 25 PB storage
- ▶ 11.5 Petaflops peak performance
- ▶ Simulation of 100M-atom system benchmark on Blue Waters aims at 4ms/step

