

NAMD Team:

Abhinav Bhatele

Eric Bohm

Sameer Kumar

David Kunzman

Chao Mei

Chee Wai Lee


Kumaresh P.

James Phillips

Gengbin Zheng

Laxmikant Kale

Klaus Schulten



Scaling Challenges in NAMD: Past and Future



ILLINOIS


UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

PARALLEL
PROGRAMMING LAB

DEPT. OF COMPUTER SCIENCE, UNIVERSITY OF ILLINOIS




Outline



- NAMD: An Introduction
- Past Scaling Challenges
 - Conflicting Adaptive Runtime Techniques
 - PME Computation
 - Memory Requirements
- Performance Results
- Comparison with other MD codes
- Future Challenges:
 - Load Balancing
 - Parallel I/O
 - Fine-grained Parallelization

What is NAMD ?

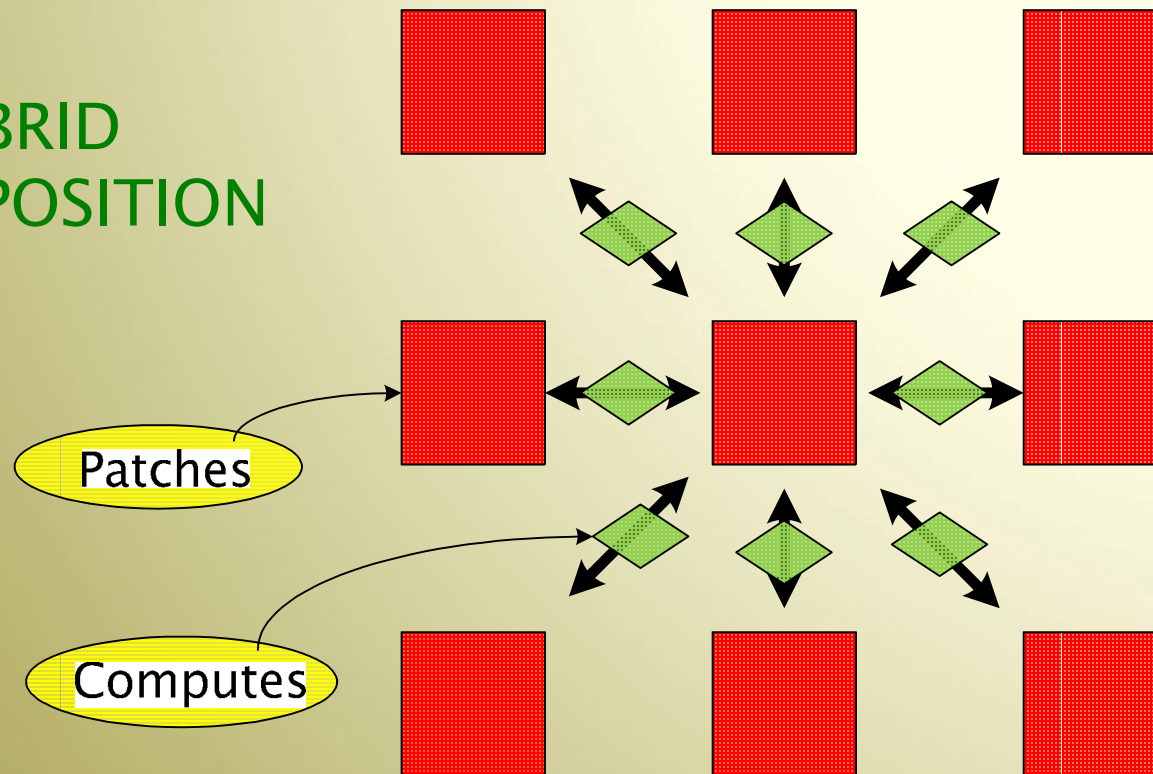


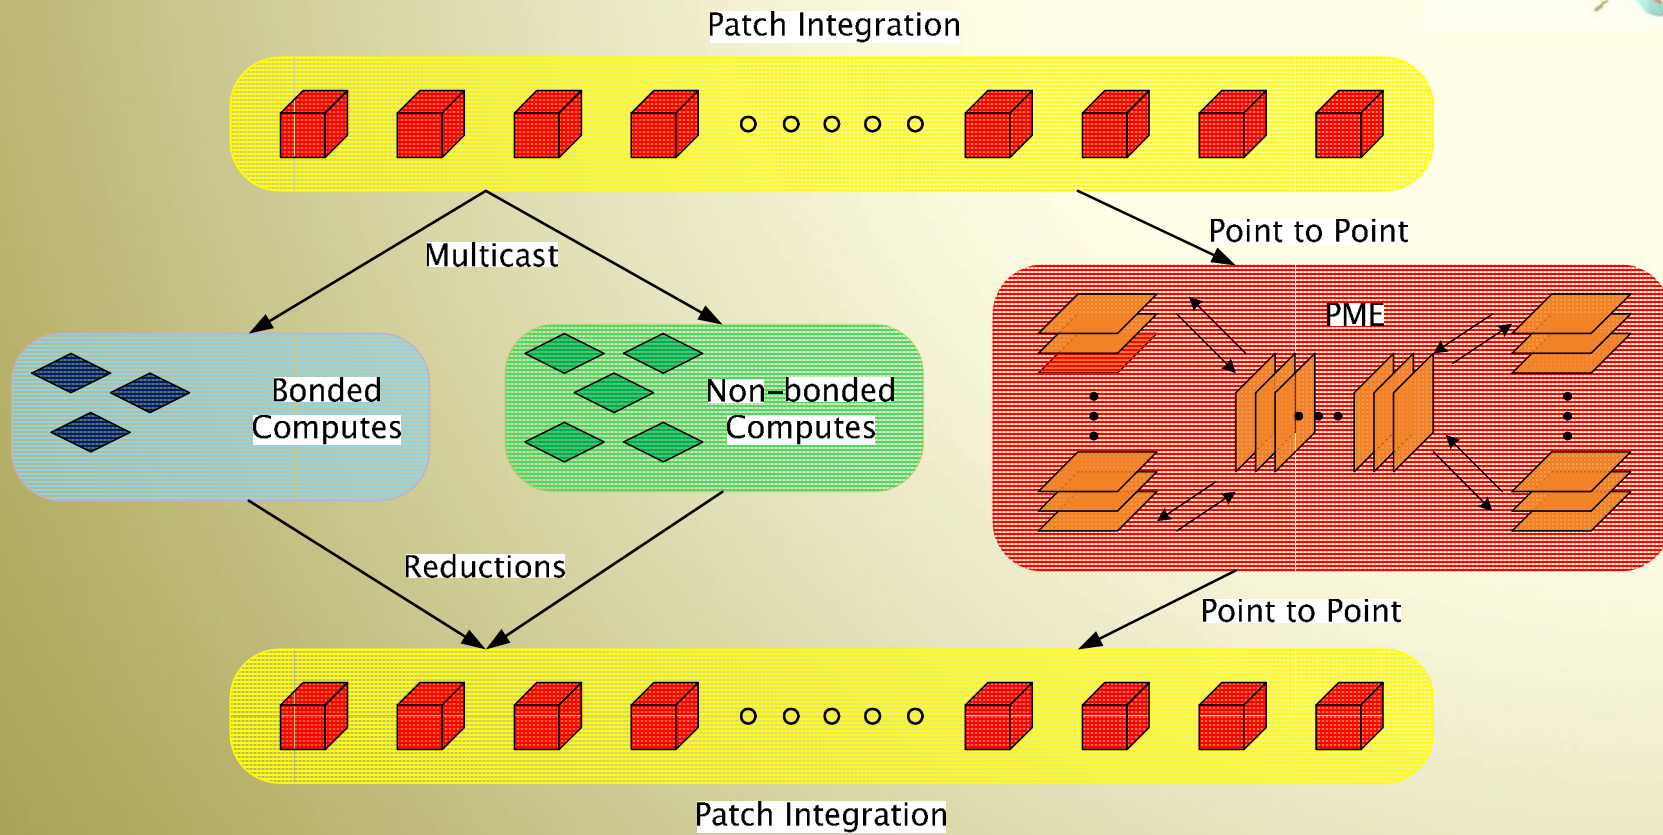

- A parallel molecular dynamics application
- Simulate the life of a bio-molecule
- How is the simulation performed ?
 - Simulation window broken down into a large number of time steps (typically 1 fs each)
 - Forces on every atom calculated every time step
 - Velocities and positions updated and atoms migrated to their new positions

How is NAMD parallelized?




HYBRID
DECOMPOSITION

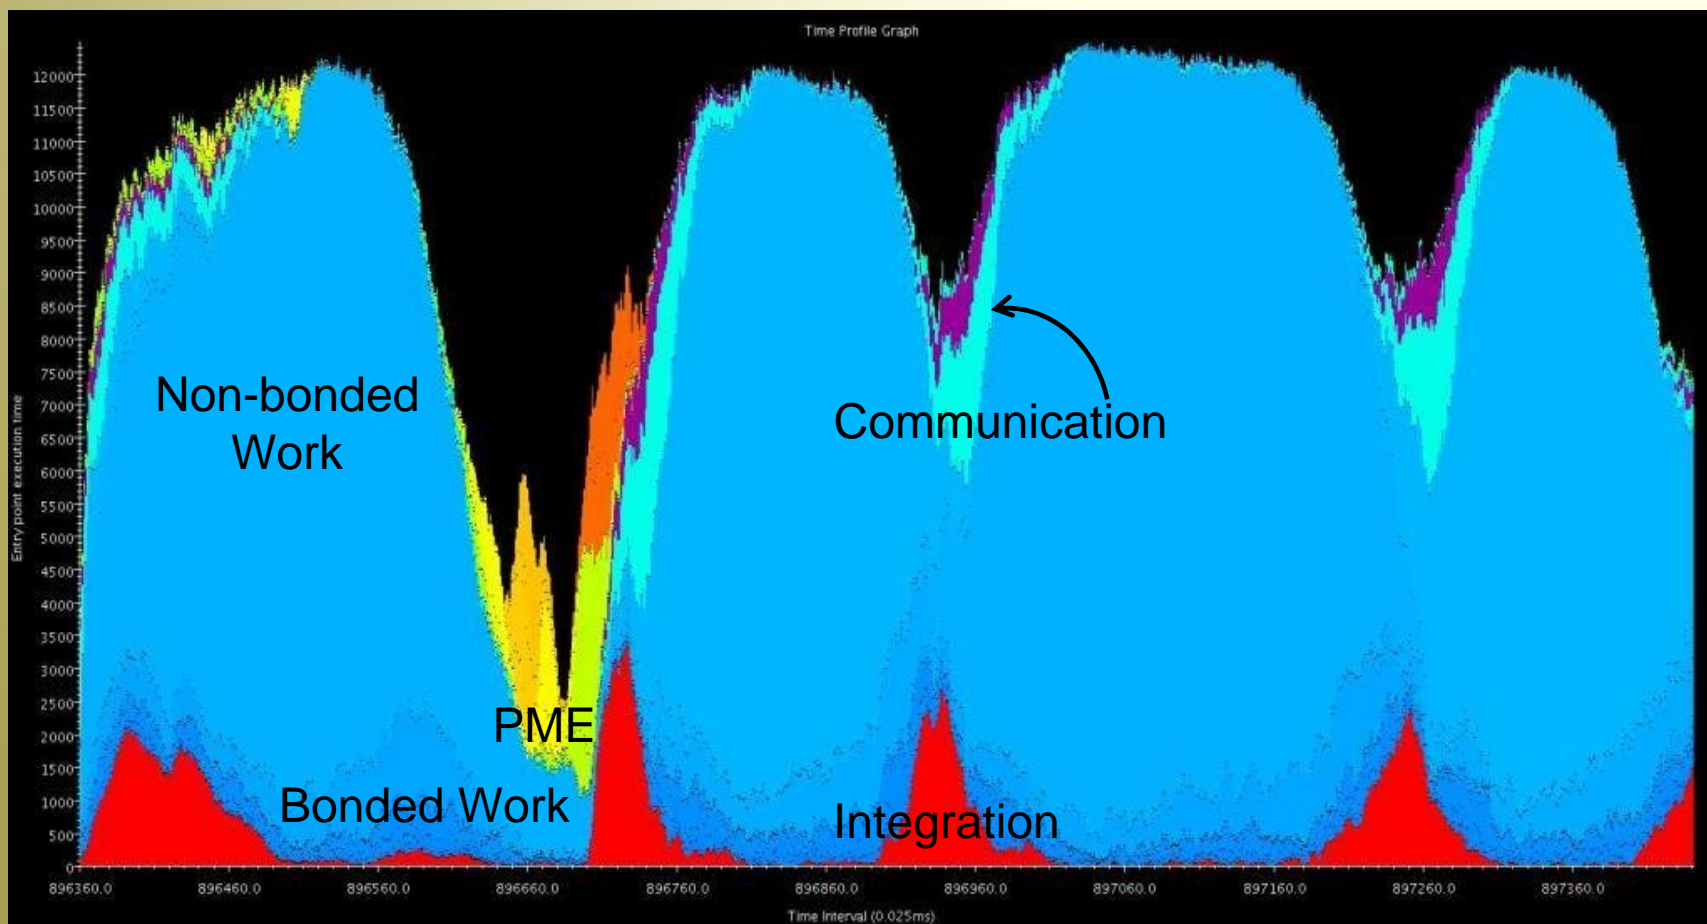

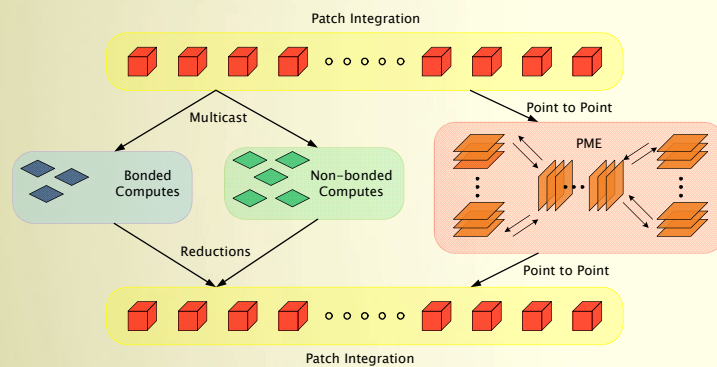




What makes NAMD efficient?



- Charm++ runtime support
 - Asynchronous message-driven model
 - Adaptive overlap of communication and computation




What makes NAMD efficient ?

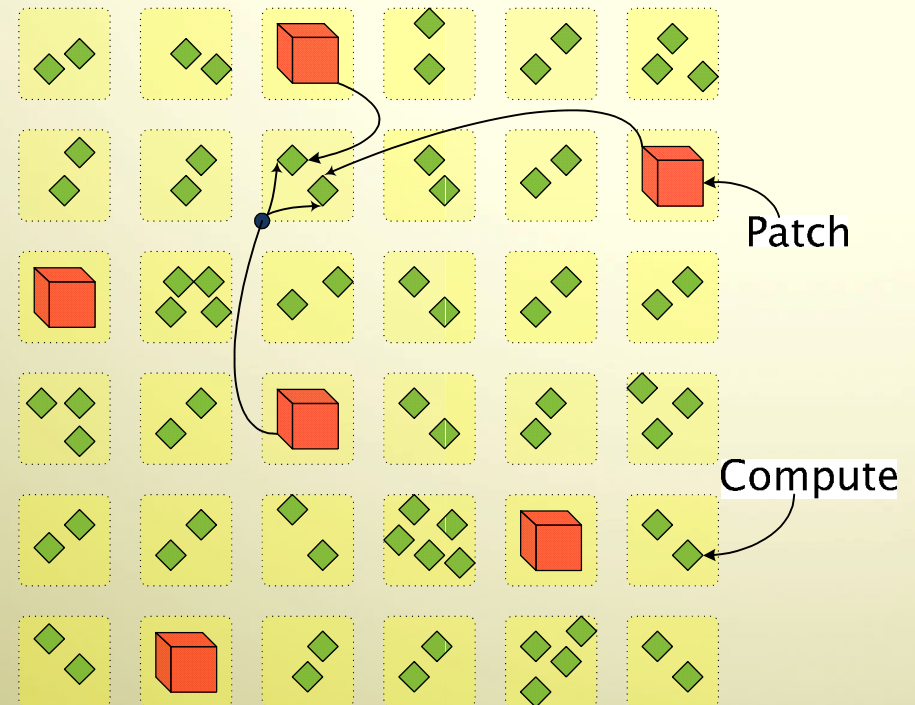
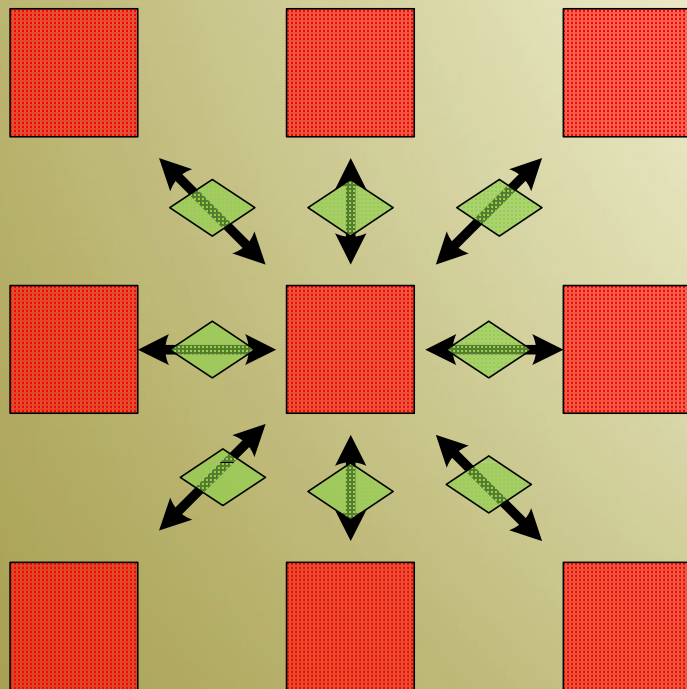


- Charm++ runtime support
 - Asynchronous message-driven model
 - Adaptive overlap of communication and computation
- Load balancing support
 - **Difficult problem: balancing heterogeneous computation**
 - **Measurement-based load balancing**


What makes NAMD highly scalable ?



- Hybrid decomposition scheme
- Variants of this hybrid scheme used by Blue Matter and Desmond




Scaling Challenges

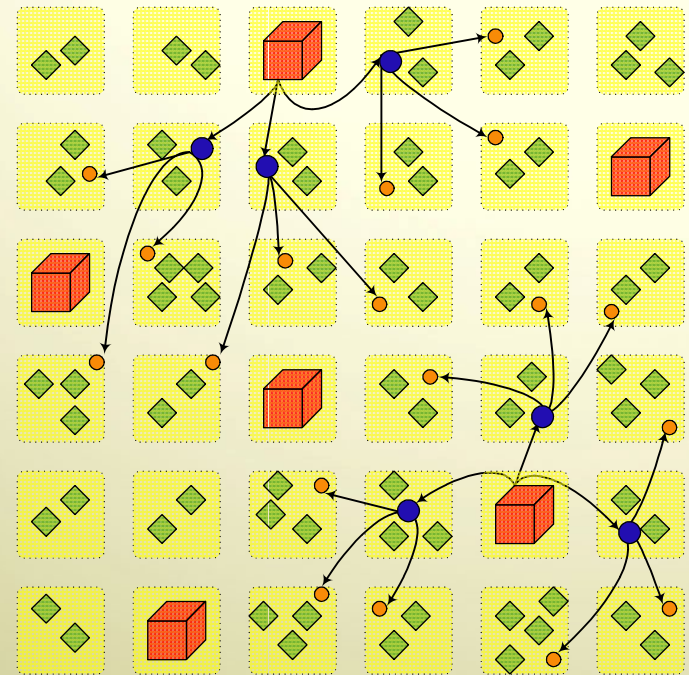




- Scaling a few thousand atom simulations to tens of thousands of processors
 - Interaction of adaptive runtime techniques
 - Optimizing the PME implementation
- Running multi-million atom simulations on machines with limited memory
 - Memory Optimizations



Conflicting Adaptive Runtime Techniques




- Patches multicast data to computes
- At load balancing step, computes re-assigned to processors
- Tree re-built after computes have migrated











- Solution
 - Persistent spanning trees
 - Centralized spanning tree creation
- Unifying the two techniques

PME Calculation




- Particle Mesh Ewald (PME) method used for long range interactions
 - 1D decomposition of the FFT grid
- PME is a small portion of the total computation
 - Better than the 2D decomposition for small number of processors
- On larger partitions
 - Use a 2D decomposition
 - More parallelism and better overlap

Automatic Runtime Decisions




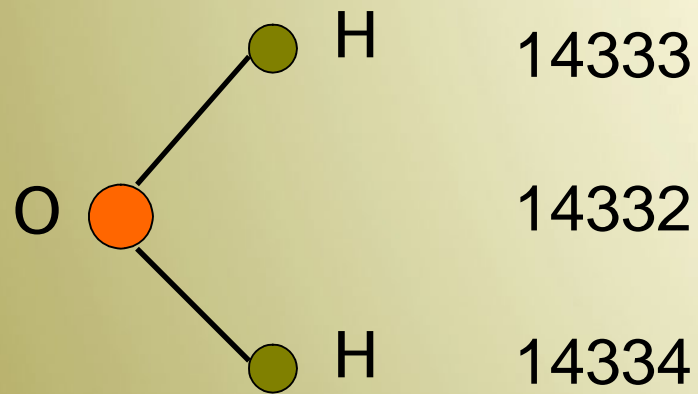



- Use of 1D or 2D algorithm for PME
- Use of spanning trees for multicast
- Splitting of patches for fine-grained parallelism
- Depend on:
 - Characteristics of the machine
 - No. of processors
 - No. of atoms in the simulation


Reducing the memory footprint




- Exploit the fact that building blocks for a biomolecule have common structures
- Store information about a particular kind of atom only once



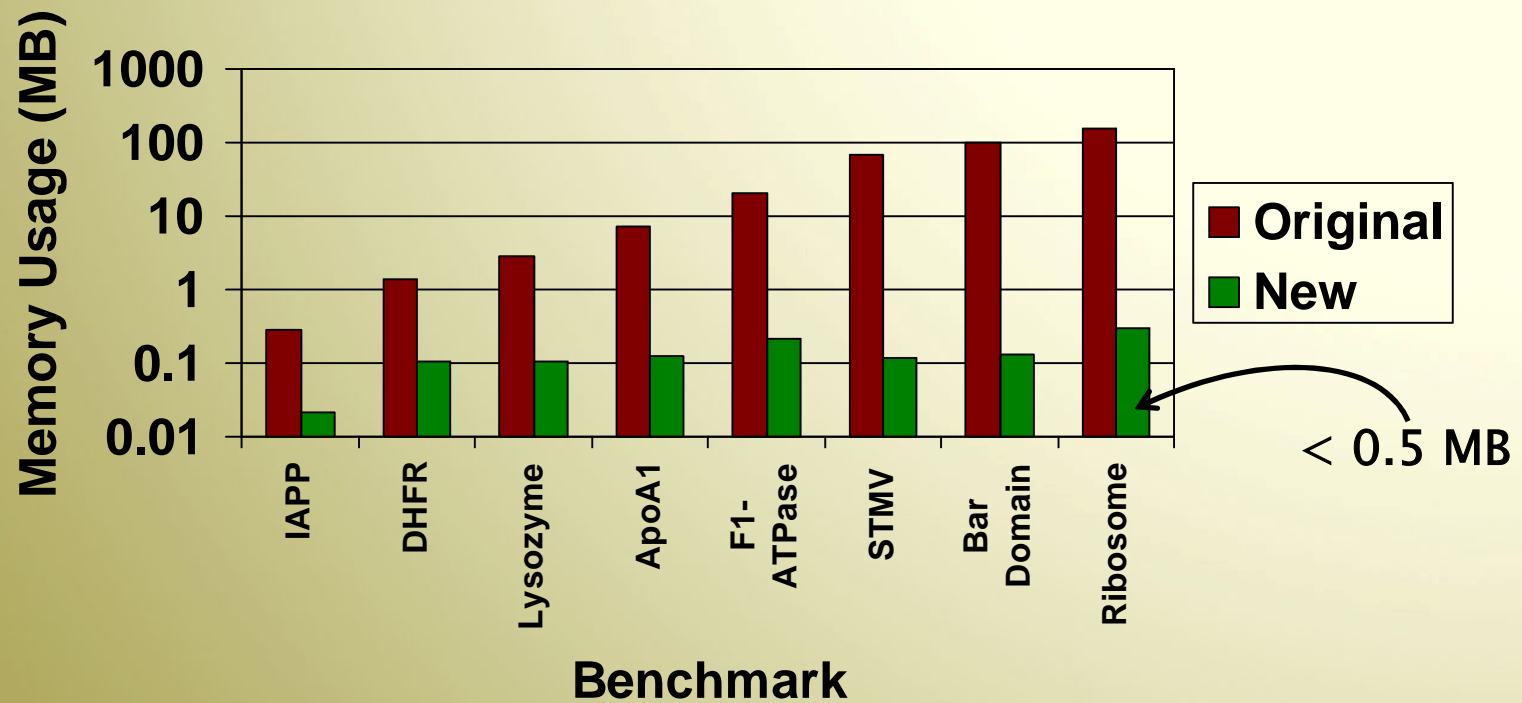
Reducing the memory footprint





- Exploit the fact that building blocks for a biomolecule have common structures
- Store information about a particular kind of atom only once
- Static atom information increases only with the addition of unique proteins in the simulation
- Allows simulation of 2.8 M Ribosome on Blue Gene/L



Memory Reduction





NAMD on Blue Gene/L




1 million atom simulation on 64K processors (LLNL BG/L)

NAMD on Cray XT3/XT4



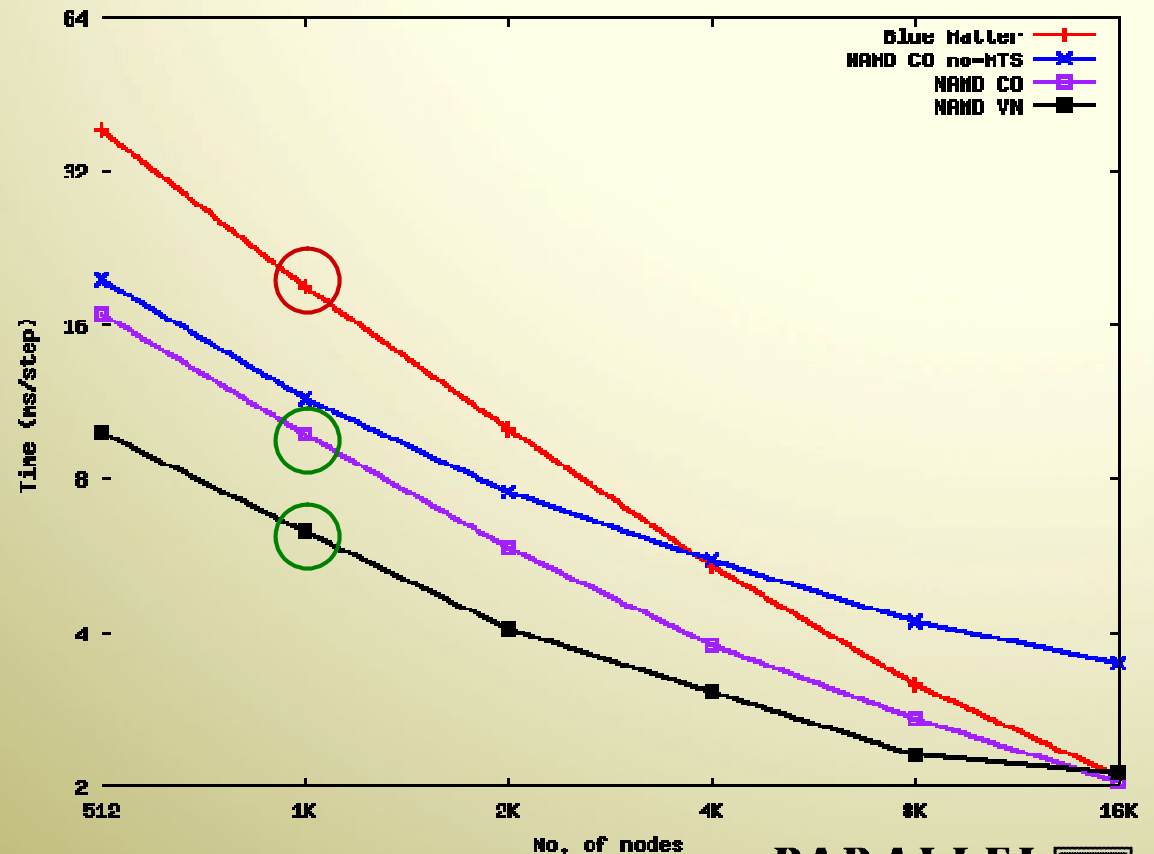
5570 atom
simulation on
512 processors
at 1.2 ms/step


Comparison with Blue Matter



- Blue Matter developed specifically for Blue Gene/L
- NAMD running on 4K cores of XT3 is comparable to BM running on 32K cores of BG/L


Time for ApoA1 (ms/step)






Number of Nodes	512	1024	2048	4096	8192	16384
<i>Blue Matter (2 pes/node)</i>	38.42	18.95	9.97	5.39	3.14	2.09
<i>NAMD CO mode (1 pe/node)</i>	16.83	9.73	5.8	3.78	2.71	2.04
<i>NAMD VN mode (2 pes/node)</i>	9.82	6.26	4.06	3.06	2.29	2.11
<i>NAMD CO mode (No MTS)</i>	19.59	11.42	7.48	5.52	4.2	3.46
<i>NAMD VN mode (No MTS)</i>	11.99	9.99	5.62	5.3	3.7	-


Comparison with Desmond



- Desmond is a proprietary MD program
- Uses single precision and exploits SSE instructions
- Low-level infiniband primitives tuned for MD



Time (ms/step) for Desmond on 2.4 GHz Opterons and NAMD on 2.6 GHz Xeons






Number of Cores	8	16	32	64	128	256	512	1024	2048
<i>Desmond ApoA1</i>	256.8	126.8	64.3	33.5	18.2	9.4	5.2	3.0	2.0
<i>NAMD ApoA1</i>	199.3	104.9	50.7	26.5	13.4	7.1	4.2	2.5	1.9
<i>Desmond DHFR</i>	41.4	21.0	11.5	6.3	3.7	2.0	1.4	-	-
<i>NAMD DHFR</i>	27.3	14.9	8.09	4.3	2.4	1.5	1.1	1.0	

NAMD on Blue Gene/P




Future Work




- Optimizing PME computation
 - Use of one-sided puts between FFTs
- Reducing communication and other overheads with increasing fine-grained parallelism
- Running NAMD on Blue Waters
 - Improved distributed load balancers
 - Parallel Input/Output

Summary



- NAMD is a highly scalable and portable MD program
 - Runs on a variety of architectures
 - Available free of cost on machines at most supercomputing centers
 - Supports a range of sizes of molecular systems
- Uses adaptive runtime techniques for high scalability
- Automatic selection of algorithms at runtime best suited for the scenario
- With new optimizations, NAMD is ready for the next generation of parallel machines



Questions ?



ILLINOIS

UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

PARALLEL
PROGRAMMING LAB

DEPT. OF COMPUTER SCIENCE, UNIVERSITY OF ILLINOIS

