MD Study of Ice Nucleation

Multi-pronged Approach to Accelerating Results by >1000x

GE Global Research
Advanced Computing Lab
Masako Yamada

imagination at work

Opportunities in Cold-Climate Win

Wind energy production > 285 GW/year and grov

- Cold regions favorable
  - Lower human population
  - Good wind conditions
  - >50 GW opportunity in next decade ~$2million/MW inst
- Technical need
  - Anti-icing surfaces
    - 3-10% energy losses due to icing
    - Shut-downs
    - Active heating expensive

imagination at work

VTT Technical Research Centre of Finland
http://www.vtt.fi/news/2013/28052013_wind_energy.jsp?lang=en

MD to probe freezing on surfaces

Diverse demands

- Large simulations > 1 million molecule droplet
- Long simulations > 1 microsecond
- Many simulations > 1000's independent drop

Need for leadership-class resource

- 40M CPU-h on Jaguar at ORNL (2011)
- 40M CPU/GPU-h on Titan at ORNL (2013)

hybrid

Billion-fold acceleration in 20 year

Year	Software/Language	# of Molecules	Hardware
1995	Pascal	Few	Desktop M
2000	C, Fortran90	Hundreds	IBM SP, SG
2010	NAMD, LAMMPS	1000's	Linux HPC
Present	GPU-enabled LAMMPS	Millions	Titan

1995

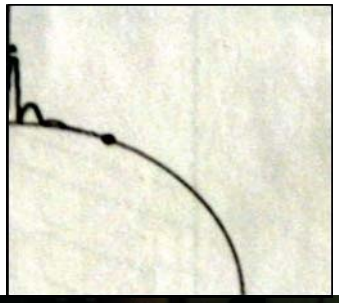

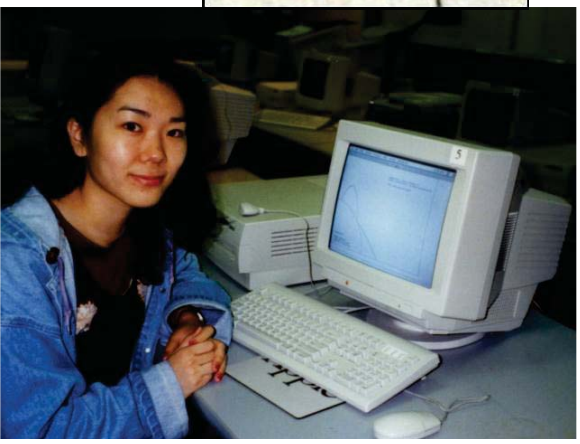

2000

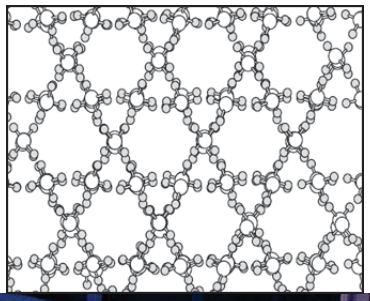

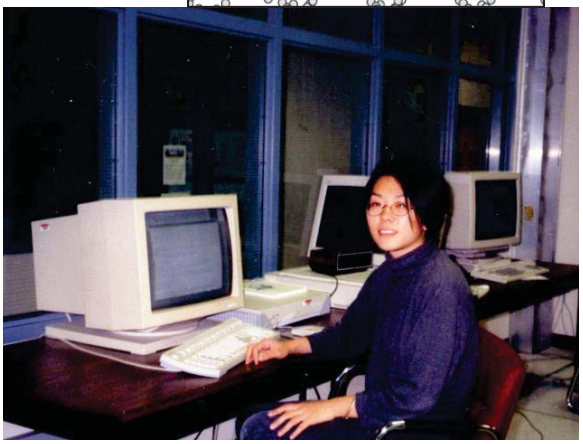

2013

Progress in HPC = Progress in MD

- Faster processors
- More processors
- Larger memory
- Faster interconnects
- Algorithms/Optimization

>1000x overall speedup since 201

keeping system size and # of nodes f

1. New water model – 40-50x
2. Dynamic load balancing – 2-3x
3. GPU acceleration – 5x
4. Parallel Replica Method – up to 1

"Who can I add to the team?

imagination at work

Overview of LAMMPS

- Popular open-source molecular dynami
  code developed by Sandia Nat'l Lab
- Pre-populated with many popular pair-
  and many-body potentials
- One of 6 CAAR (early acceptance)
  applications on Titan

1) New water model – 40-50x

- mW water introduced in 2009; Nature paper
- One 3-body particle ➔ one water molecule
- Properties comparable/better than existing r
- Much faster than point-charge models

SPC/E on Jaguar (2011)

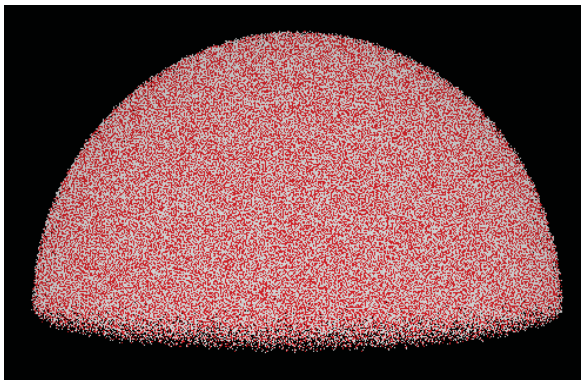

mW on Titan (2013)

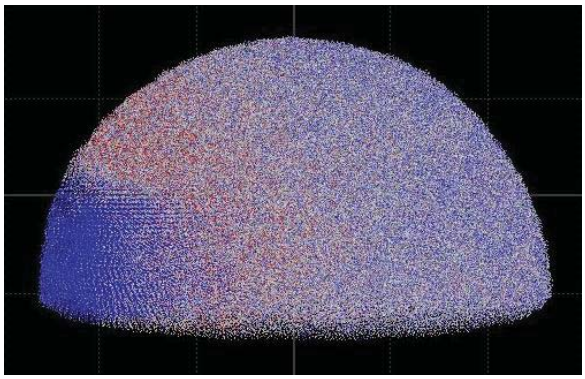

imagination at work

mW-surface interaction potential

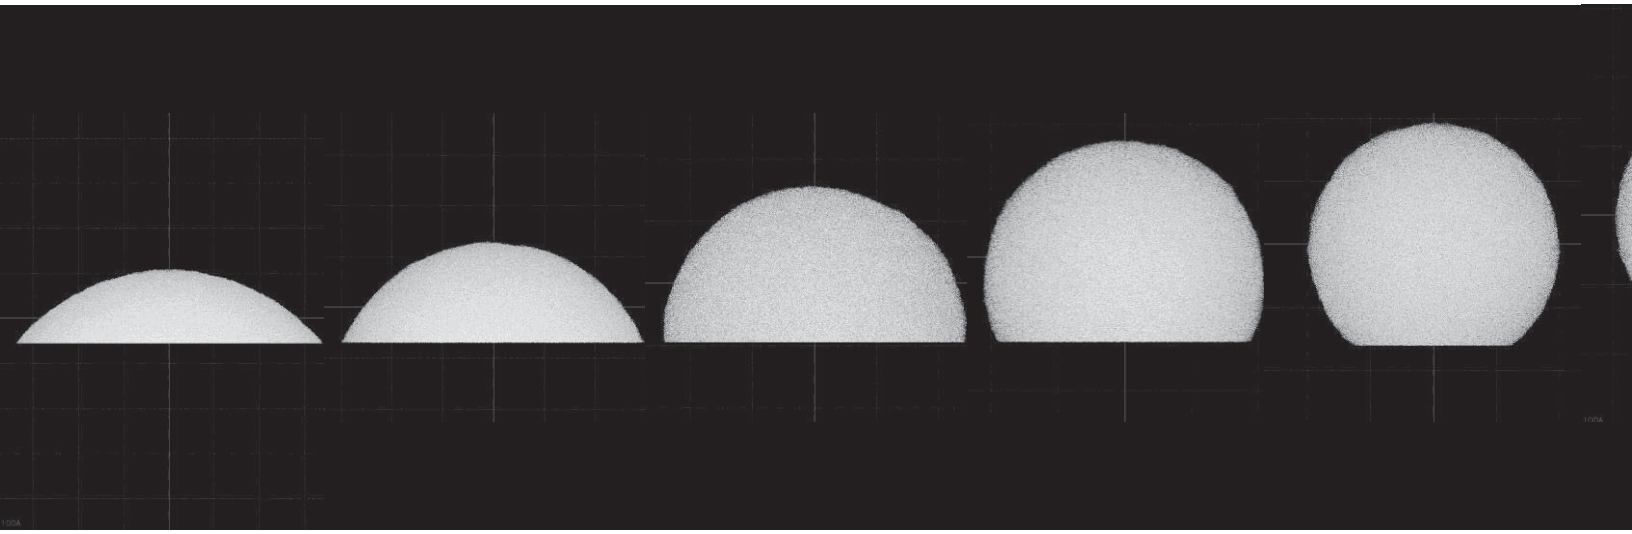

Interaction potential developed at GE Global Resear

imagination at work

2) Dynamic load balancing – 2-3x

Sandia developed in 2012; we did some beta-t

Adjusts size of processor sub-domains to equa number of particles

Up to 2-3x speedup for our production droplet

Needs some user-specified processor mappin

No load balancing

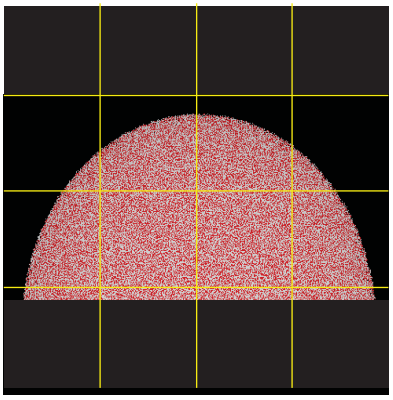

Default load balancing

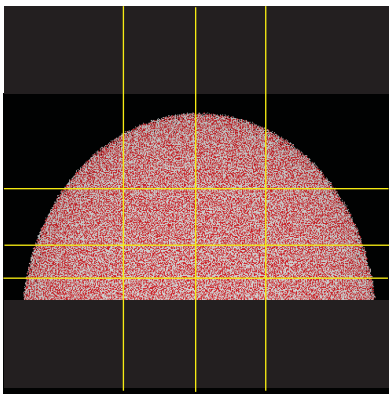

User-specified m

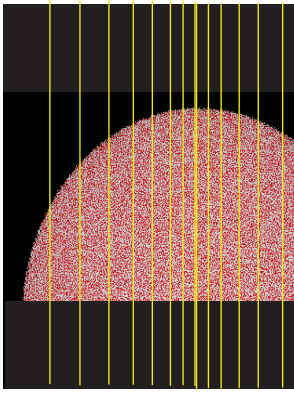

imagination at work

Our production droplets run on 64 node

3) GPU acceleration – 5x

Host	Accelerato
Time integration	3-body potent
Thermostat/barostat	Neighbor-list
Bond/angle calculations	
Statistics	

Generic 3-body potential

$$U = \begin{cases} \sum_i \sum_{j \neq i} \sum_{k>j} \phi(\boldsymbol{p}_i, \boldsymbol{p}_j, \boldsymbol{p}_k) & r_{ij} < r_c, r_{ik} < \\ 0 & \text{otherwise} \end{cases}$$

Good candidate for GPU

1. Occupies majority of computational time
2. Can be decomposed into independent kernels/work-items

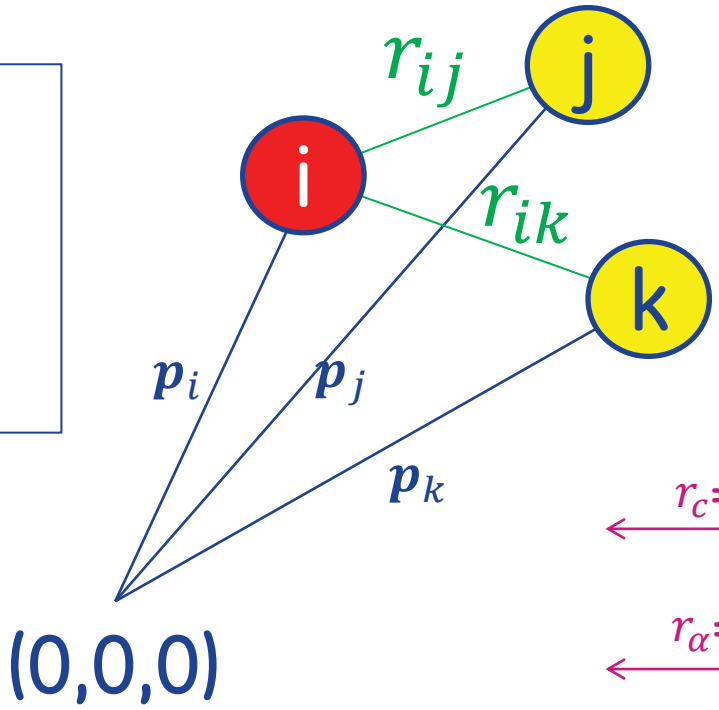

Stillinger-Weber
MEAM
Tersoff
REBO/AIREBO
Bond-order…

imagination at work

Stillinger-Weber 3-body potential

$$U = \sum_i \sum_{j<i} \phi_2(r_{ij}) + \sum_i \sum_{j \neq i} \sum_{k>j} \phi_3(r_{ij}, r_{ik}, \theta_j$$

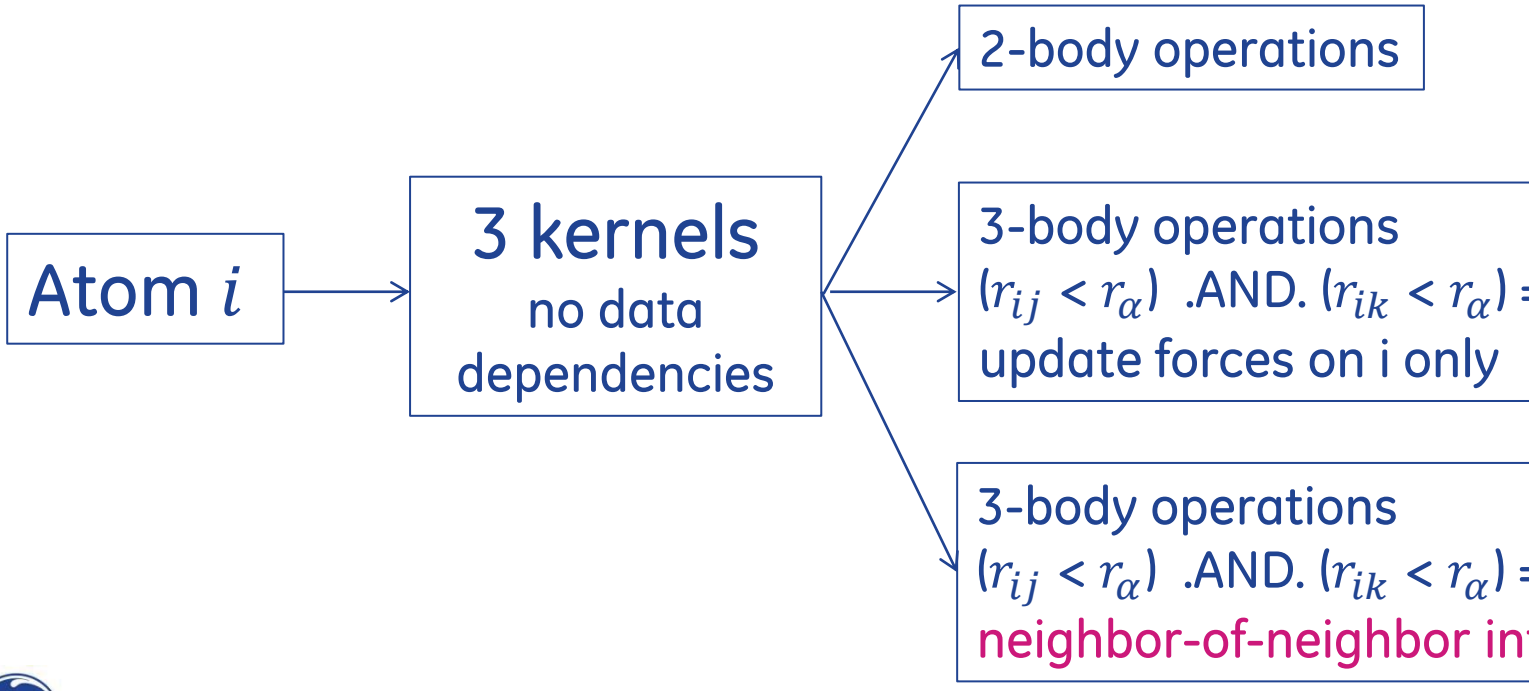

imagination at work

Neighbor List on GPU

- 3-body force-decomposition approach invol
  neighbor-of-neighbor operations
- Requires additional overhead
  - increase in border size shared by two processes
  - neighbor list for ghost atoms “straddling” across
- GPU not necessarily faster than CPU but less
  spent in host-accelerator data transfer

4) Parallel Replica Method – up to

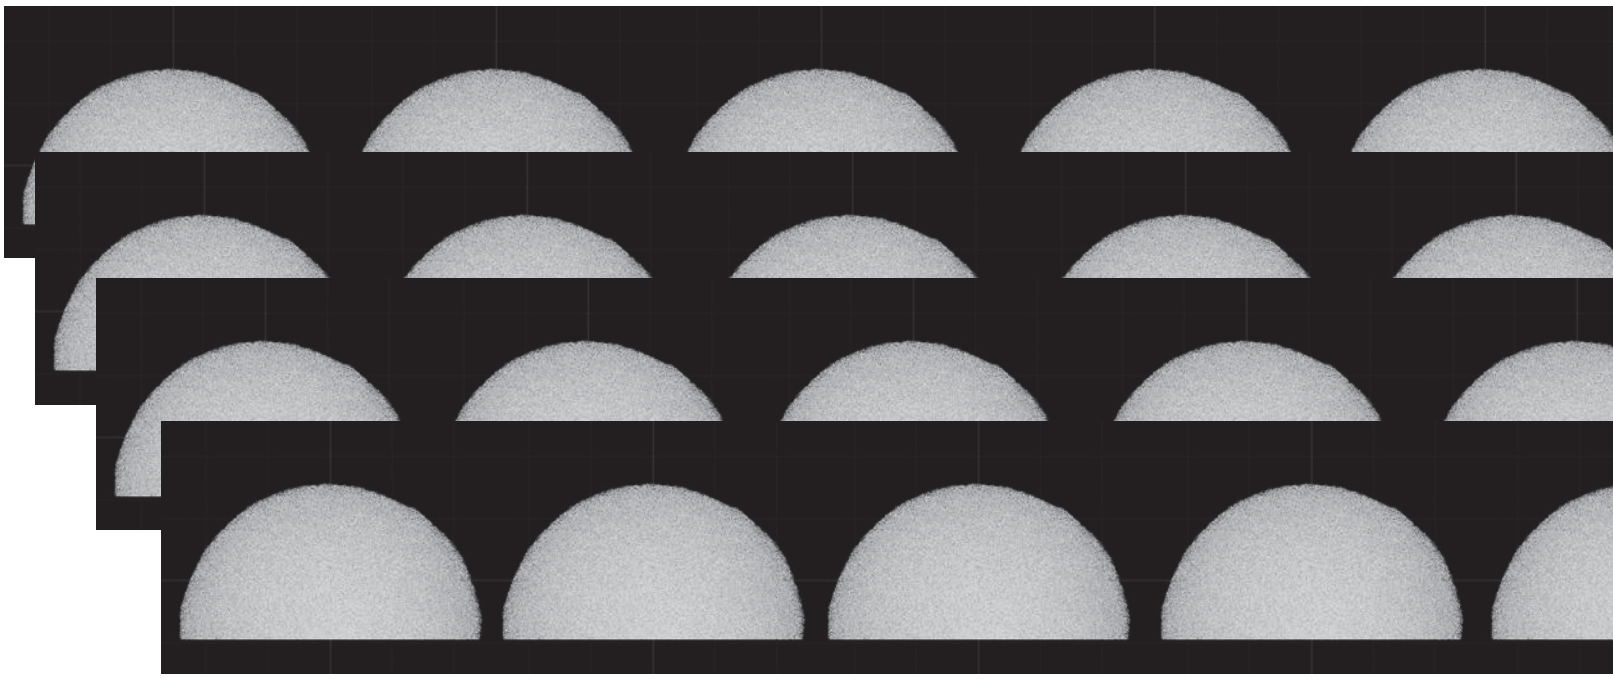

Launch N identical replicates simultaneou
varying only the random seed (N ~ large)

Quasi-time-parallelization

$$t_{\text{first droplet freezes}} \sim N \times t_{\text{avg droplet freezes}}$$

- PR method accelerates observation of rare, exponentially distributed events
- Titan can accommodate up to N= 200
- Can stop entire run after first drop free

Fun fact if N = 200

$$t_{\text{last droplet freezes}} \sim 6 \times t_{\text{avg dro}}$$

imagination at work

Also - Post-processing and Viz

Big Data – Total 50TB

- 1 million molecules per snapshot
- Dozens of snapshots per file
- 10,000's files

Big Compute – NOT simple search/sort

- Execute three-body calculation again
- Subtle pattern-matching of intra-molecular po
- Post-processing is a EOS/Rhea job!

Big Visualization – need dedicated viz resou

Examples

Perturbation

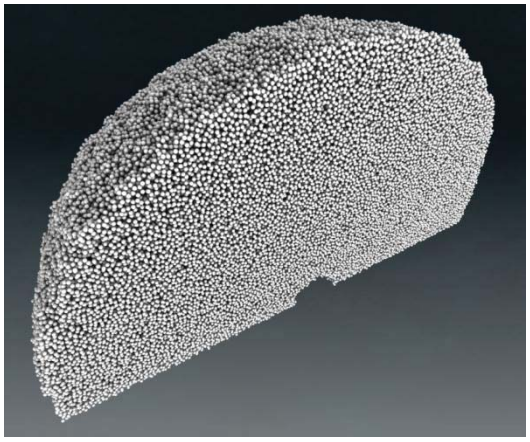

Matheson, ORNL

Nucleation

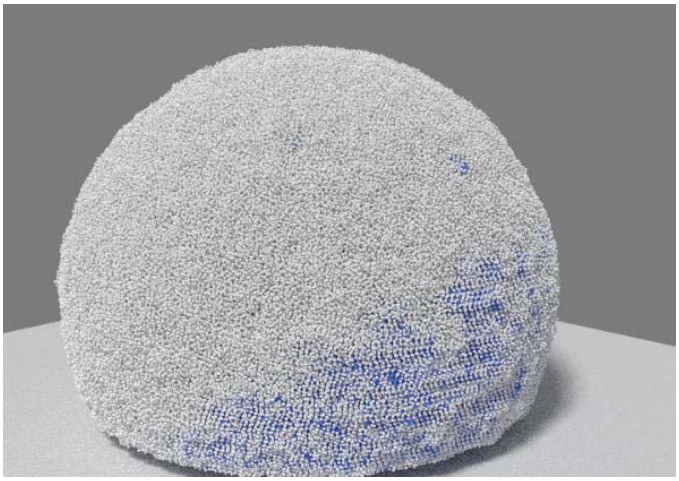

Matheson, ORNL

Grai

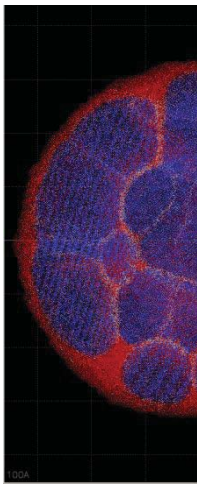

Yar

Thanks to Lens, Rhea, EOS and M. Mathes

imagination at work

Credits

- Mike Brown (ORNL) – GPU acceleration
- Paul Crozier (Sandia) – dynamic load balancing
- Valeria Molinero (Utah) – mW potential
- Aaron Keyes (Umich, Berkeley) – Steinhardt-Nelson order parameters
- Art Voter/Danny Perez (LANL) – Parallel Replica method
- Mike Matheson (ORNL) -- Visualization
- Jack Wells, Suzy Tichenor (ORNL) – General
- Azar Alizadeh, Branden Moore, Rick Arthur, Margaret Blohm (GE Global R

This research was conducted in part under the auspices of the Office of Advanced Scientific Computing Research, Office of Science, U.S. Department of Energy under Contract No. DEAC05-00OR22725 with UT-Battelle, LLC. This research was also conducted in part under the auspices of the GE Global Research High Performance Computing program.

Backup

Overview of MD

Atom-by-atom modeling of materials

- N-body problem
- Discrete, numerical integration

Biology, chemistry requires good water mo

- Dozens of potentials available
- Most use pair-wise interactions
- Most non-polarizable/rigid

MD *always* on the forefront of HPC

Relevant GPU acceleration activity

Pair-wise potentials

- LAMMPS already GPU-enabled

Three-body potentials

- Impressive acceleration… but for crystal soli

Present work

- >5x acceleration demonstrated using LAMM
- Works for liquids, glass, vapor

Redundant Computation Approac

Atom-decomposition

- 1 atom → 1 computational kernel only
- fewest operations (and effective parallelization) but
  - shared memory access a bottleneck

Force-decomposition

- 1 atom → 3 computational kernels required
- redundant computations but
  - reduced shared memory issues
  - many work-items = more effective use of cores

Load 1 million molecules on Host/

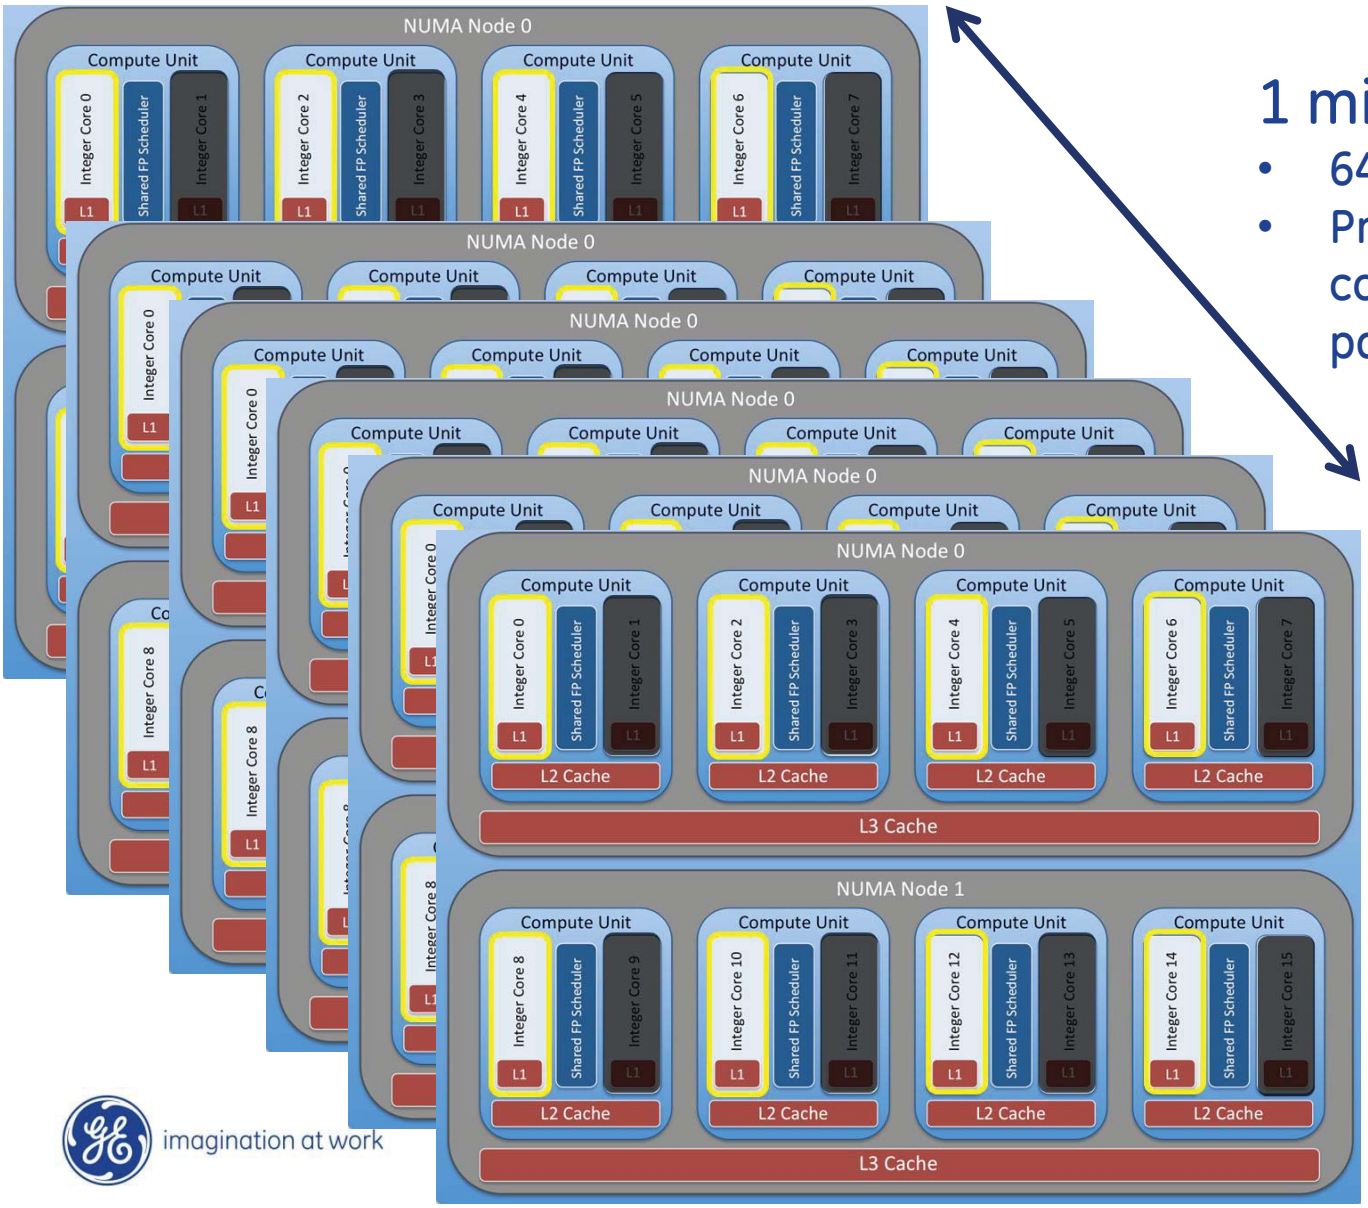

1 million mole

- 64 nodes
- Processor sub
  correspond to
  partitioning of

- 8 MPI ta
- 1 core/p

imagination at work

Per node ~ 15,000 molecules

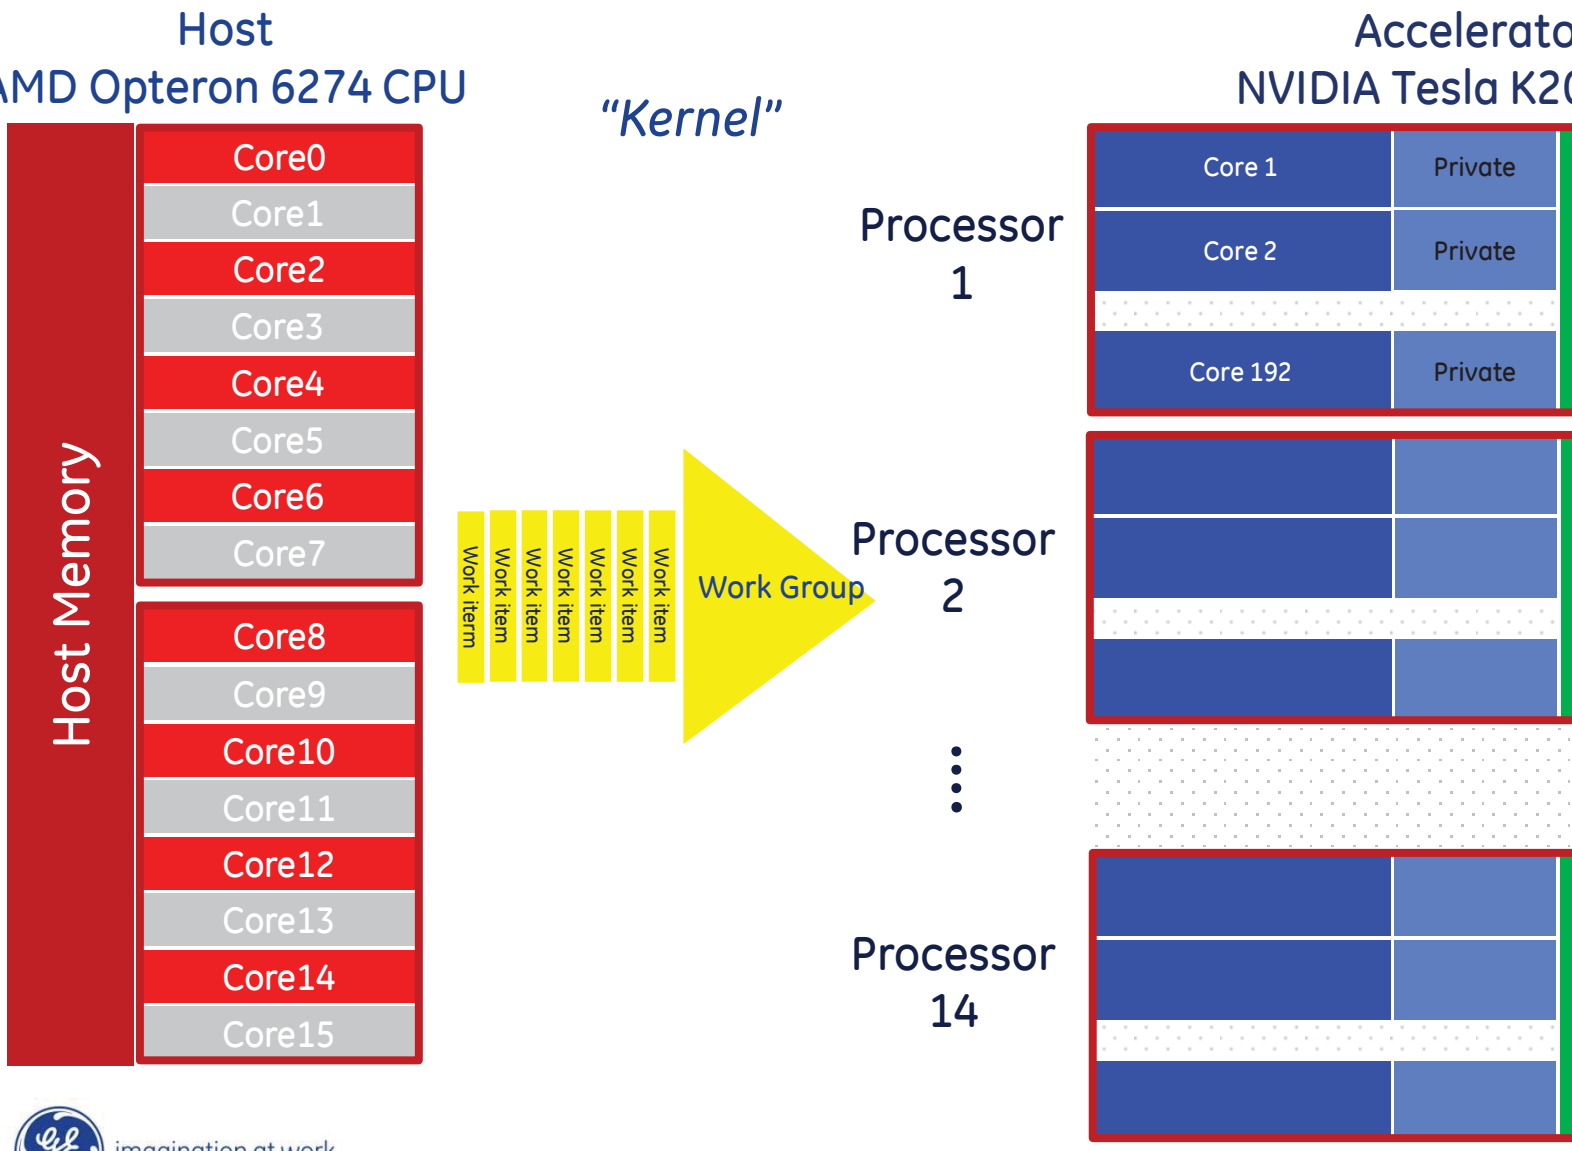

imagination at work

Work item = fundamental unit of activity

References

- http://www.vtt.fi/news/2013/28052013_wind_energy.jsp?lang=en
- W. Michael Brown, W. M and Yamada, M. Implementing Molecular Dynamics on Hybrid High Performance Compute Potentials. Computer Physics Communications. 2013.Computer Physics Communications, (2013)
- C. Hou, J. Xu, P. Wang, W. Huang, X. Wang, Computer Physics Communications (2013)
- Shi, B. and Dhir, V. K. Molecular dynamics simulation of the contact angle of liquids on solid surfaces. The Journal of 130, 3 (01/21/ 2009), 034705-034705; Sergi, D., Scocchi, G. and Ortona, A. Molecular dynamics simulations of the co between water droplets and graphite surfaces. Fluid Phase Equilibria, 332, 0 (10/25/ 2012), 173-177.
- Oxtoby, D. W. Homogeneous nucleation: theory and experiment. Journal of Physics: Condensed Matter, 4, 38 1992),
- Plimpton, S. Fast Parallel Algorithms for Short-Range Molecular Dynamics. Journal of Computational Physics, 117, 1
- Humphrey, W., Dalke, A. and Schulten, K. VMD: Visual molecular dynamics. Journal of Molecular Graphics, 14, 1 (2//
- Keys, A. S. Shape Matching Analysis Code. University of Michigan, City, 2011; Keys, A. S., Iacovella, C. R. and Glotzer, Characterizing Structure Through Shape Matching and Applications to Self-Assembly. Annual Review of Condensed 1 (2011/03/01 2011), 263-285; Steinhardt, P. J., Nelson, D. R. and Ronchetti, M. Bond-orientational order in liquids an Physical Review B, 28, 2 (07/15/ 1983), 784-805.
- Stillinger, F. H. and Weber, T. A. Computer simulation of local order in condensed phases of silicon. Physical Review 1985), 5262-5271.
- Berendsen, H. J. C., Grigera, J. R. and Straatsma, T. P. The missing term in effective pair potentials. The Journal of Ph 91, 24 (1987/11/01 1987), 6269-6271.
- Molinero, V. and Moore, E. B. Water Modeled As an Intermediate Element between Carbon and Silicon†. The Journal Chemistry B, 113, 13 (2009/04/02 2008), 4008-4016; Moore, E. B. and Molinero, V. Structural transformation in super controls the crystallization rate of ice. Nature, 479, 7374 (11/24/print 2011), 506-508.
- Yamada, M., Mossa, S., Stanley, H. E. and Sciortino, F. Interplay between Time-Temperature Transformation and the Phase Transition in Water. Physical Review Letters, 88, 19 (04/26/ 2002), 195701.
- Brown, W. M., Wang, P., Plimpton, S. J. and Tharrington, A. N. Implementing molecular dynamics on hybrid high per computers – short range forces. Computer Physics Communications, 182, 4 (4// 2011), 898-911.
- Voter, A. F. Parallel replica method for dynamics of infrequent events. Physical Review B, 57, 22 (06/01/ 1998), R1398