

## TRILLION-ATOM MOLECULAR DYNAMICS BECOMES A REALITY

TIMOTHY C. GERMANN\* and KAI KADAU†

*Theoretical Division, Los Alamos National Laboratory  
Los Alamos, New Mexico 87545, USA*

\* *tcg@lanl.gov*

† *kkadau@lanl.gov*

Received 5 September 2008

Accepted 6 September 2008

By utilizing the molecular dynamics code *SPaSM* on Livermore's BlueGene/L architecture, consisting of 212 992 IBM PowerPC440 700 MHz processors, a molecular dynamics simulation was run with one trillion atoms. To demonstrate the practicality and future potential of such ultra large-scale simulations, the onset of the mechanical shear instability occurring in a system of Lennard-Jones particles arranged in a simple cubic lattice was simulated. The evolution of the instability was analyzed on-the-fly using the in-house developed massively parallel graphical object-rendering code *MD\_render*.

*Keywords:* Molecular dynamics; BlueGene/L; high performance computing; SPaSM; large-scale; trillion-atom; visualization.

PACS Nos.: 07.05.Tp, 83.10.Rp, 71.15.Pd.

### 1. Introduction

In the fall of 2007 Lawrence Livermore National Laboratory (LLNL)'s BlueGene/L machine was extended from its prior configuration of 131 072 IBM PowerPC 440 processors (700 MHz clock speed, providing 2.8 GFlop/s theoretical peak performance, and 256 MB memory per processor) by 81 920 additional processors, identical to the others but with twice the memory (512 MB), giving a total of 212 992 CPUs with an aggregate memory of 72 TB. As the additional CPUs have twice as much memory, only 52 TB can effectively be used when employing the whole machine in a load-balanced fashion, i.e., in a molecular dynamics (MD) simulation where each processor is assigned roughly the same number of particles. A trillion particles with the position, velocity, and force vectors represented in single precision (12 bytes for each three-dimensional vector) and two additional integers (4 bytes each) to store the particle type and a user-defined tag consume 44 bytes per particle — or about 40 TB for  $10^{12}$  particles. Even after accounting for some additional memory consumption for the Eulerian cells in which particles are organized, on-the-fly

analysis and visualization, and density inhomogeneities on the order of 20%,<sup>1–3</sup> we demonstrate the first-ever trillion-particle MD simulation on the current Blue-Gene/L configuration at LLNL.

## 2. The Simulation

We set up a system containing  $10\,000 \times 10\,000 \times 10\,000$  Lennard-Jones (LJ) particles in a simple cubic lattice. The lattice constant was chosen to be at the minimum of the LJ 6–12 potential (i.e.,  $0.891\sigma$ ), and the initial temperature was set to 0.001 by assigning the particles' velocity components according to a random Gaussian distribution (for a description of LJ units, see Refs. 1 and 4). With this trillion atom simulation we wanted to honor the *father* of the *SPaSM* (Scalable Parallel Short-range Molecular dynamics) code,<sup>5–7</sup> Peter S. Lomdahl, who together with his student Dave Beazley first initiated and spearheaded *SPaSM* in the early 1990s, first for the Connection-Machine 5 (CM-5) and subsequently for a variety of platforms, where the *SPaSM* code very successfully simulated millions of atoms addressing impact and crack propagation phenomena in solids. Using a gray-scale digital image of Peter in front of a Cray T3D ( $1066 \times 1066$  pixel resolution), we added — according to the gray-scale value of the pixel of the picture — an additional  $+x$  velocity component ranging from 0 to 6 (in Lennard-Jones units). This was done for all  $z$ -planes ( $z$  being the direction orthogonal to the paper-plane in Fig. 1) of the crystal. In Fig. 1 the spheres are shaded according to their velocity, which initially consisted of a superposition of a thermal part and the directional velocity assigned by the image. The rendering of the whole system with *MD\_render* (with a resolution of  $1200 \times 1200$  pixels) took only about 25 s, while the zoomed-in picture took about twice as long, as fewer processors had to calculate the scene.<sup>2</sup> The simulation




Fig. 1. 2D slice of the initial configuration of the 3D trillion atom simulation. The gray value indicates the absolute velocity of the atoms (0 = white, 6 = black in Lennard-Jones velocity units). By zooming-in to about 0.5% of the system length near the tip of the letter “S”, the pixelized nature of the image used to assign the local nonthermal velocities can be seen. The fuzzy gray area represents regions where only thermal velocities are present. By using the nearest neighbor distance for Cu (0.255 nm), the edge length of the simulated crystal is  $2.55 \mu\text{m}$ .

was run for 40 timesteps with a timestep of 0.001, using the analytical LJ 6–12 potential truncated at  $2.5\sigma$ . This small timestep was needed in order to properly resolve the strong local shocks originating from the initial velocity pattern and the mechanical shear instability of the simple cubic LJ system (for less “violent” simulations, a timestep between 0.005–0.01 is usually adequate). The wall-clock time




Fig. 2. Evolution of the trillion-atom simulation for timesteps 10, 30, and 40. The three different zoom-ins for timestep 40 (the first zoom-in shows 2% of the edge length of the system) show the nucleation pattern of the local shear instability, as well as the atomistic nature of the crystal (gray-scale indicates the velocity of the atoms (0 = white, 6 = black)).

for one iteration was between 45–49 s, which falls in line with the earlier timings on the 131 072 processor BlueGene/L cluster after accounting for the lower density of the present simple-cubic lattice.<sup>3</sup> Images rendered on-the-fly using *MD\_render* are shown in Fig. 2; the overall appearance does not change much up to 20 timesteps, but at timestep 30 the picture appears fuzzier due to local shock gradients and associated shear effects. The gray area (initially consisting of only thermal velocities) also deforms by nucleating thin long (on the order of 10 atoms in length) regions aligned to  $\langle 100 \rangle$  directions that have increased velocities. As this work is a continuation of a series of record-breaking MD simulations reported in the journal *Int. J. Mod. Phys. C*,<sup>8,2,3</sup> and was only meant to demonstrate the potential and practicality of ultra large-scale MD simulations, we stopped the simulation after 40 timesteps.

### 3. Conclusions

We have demonstrated that on today's high-performance architectures a 3D cubic crystal with an edge length of about  $2.5 \mu\text{m}$  can be simulated; a 10 ps simulation could be performed within a couple of days. Assuming that the current trend of computational power doubling about every 18 months continues, we can expect to describe metallographic and meso-turbulence scales (for example, 0.25 mm for 0.1  $\mu\text{s}$ ) in 20 and 50 years by atomistic methods in 2D and 3D, respectively. Furthermore, one can expect that engineering scales (for example, 2.5 cm for 1 ms) can be fully resolved by atomistic methods in 60 and 100 years in 2D and 3D, respectively. As we have shown earlier, important progress can readily be made with today's computational capacity in the areas of material science, turbulence, and agent-based modeling;<sup>4,9–15</sup> these future advances will enable other studies that are currently impractical, such as polycrystalline materials with  $\mu\text{m}$  grain sizes, direct *ab initio* calculation of forces in large-scale MD, or uncertainty propagation and quantification using ensembles of simulations. By efficiently combining atomistic methods with more coarse-grained methods, or with techniques for accelerating infrequent event processes,<sup>16</sup> significant improvement should be possible much earlier. Assuming further progress is made in these algorithmic areas, alongside the expected advances in hardware capabilities, atomistic methods such as MD could potentially revolutionize the field of engineering. This could lead to a more predictive and therefore very efficient and reliable approach in the science and engineering of materials, including challenges involving complex turbulent flows.

### Acknowledgments

The authors would like to thank Peter S. Lomdahl and David Beazley for initiating and pursuing the *SPaSM* project and for many interesting discussions and support. This work was carried out under the auspices of the National Nuclear Security Administration of the US Department of Energy at Los Alamos National Laboratory

under contract No. DE-AC52-06NA25396, with funding by the Advanced Simulation and Computing (ASC) program.

## References

1. T. C. Germann, K. Kadau and P. S. Lomdahl, in *Supercomputing '05*, SC05 IEEE ACM 1-59593-061-2/05/0011 (2005).
2. K. Kadau, T. C. Germann and P. S. Lomdahl, *Int. J. Mod. Phys. C* **15**, 193 (2004).
3. K. Kadau, T. C. Germann and P. S. Lomdahl, *Int. J. Mod. Phys. C* **17**, 1755 (2006).
4. K. Kadau, T. C. Germann, N. G. Hadjiconstantinou, P. S. Lomdahl, G. Dimonte, B. L. Holian and B. J. Alder, *Proc. Natl. Acad. Sci.* **101**, 5851 (2004).
5. P. S. Lomdahl, P. Tamayo, N. Grønbech-Jensen and D. M. Beazley, in *Proc. Supercomputing 93*, ed. G. S. Ansell (IEEE Computer Society Press, Los Alamitos, CA, 1993), p. 520.
6. D. M. Beazley and P. S. Lomdahl, *Parallel Computing* **20**, 173 (1994).
7. D. M. Beazley and P. S. Lomdahl, *Comput. Phys.* **11**, 230 (1997).
8. J. Roth, F. Gähler and H.-R. Trebin, *Int. J. Mod. Phys. C* **11**, 317 (2000).
9. B. L. Holian and P. S. Lomdahl, *Science* **280**, 2085 (1998).
10. T. C. Germann, B. L. Holian, P. S. Lomdahl and R. Ravelo, *Phys. Rev. Lett.* **84**, 5351 (2000).
11. K. Kadau, T. C. Germann, P. S. Lomdahl and B. L. Holian, *Science* **296**, 1681 (2002).
12. T. C. Germann, K. Kadau, I. M. Longini, Jr. and C. A. Macken, *Proc. Natl. Acad. Sci.* **103**, 5935 (2006).
13. K. Kadau, C. Rosenblatt, J. L. Barber, T. C. Barber, Z. Huang, P. Carles and B. J. Alder, *Proc. Natl. Acad. Sci.* **104**, 7741 (2007).
14. K. Kadau, T. C. Germann, P. S. Lomdahl, R. C. Albers, J. S. Wark, A. Higginbotham and B. L. Holian, *Phys. Rev. Lett.* **98**, 135701 (2007).
15. D. H. Kalantar, J. F. Belak, G. W. Collins, J. D. Colvin, H. M. Davies, J. H. Eggert, T. C. Germann, J. Hawreliak, B. L. Holian, K. Kadau *et al.*, *Phys. Rev. Lett.* **95**, 075502 (2005).
16. A. F. Voter, F. Montalenti and T. C. Germann, *Ann. Rev. Mater. Res.* **32**, 321 (2002).