Molecular Dynamics Comes of Age: 320-Billion-Atom Simulation on BlueGene/L

Kai Kadau, Peter S. Lomdahl, T-14; Timothy C. Germann, X-1-SMMP

As the computational power is increasing, molecular dynamics (MD) simulations are becoming more important in physics, and science in general. We demonstrate weak and strong scaling of our MD code SPaSM on Lawrence Livermore National Laboratory's (LLNL's) BlueGene/L (BGL) architecture containing 131,072 IBM PowerPC440 processors. A maximum of 320 billion atoms have been simulated in double precision, which corresponds to a cubic piece of solid copper with an edge length of 1.56 μm.

Since the invention of MD in 1957 by Berni Alder and Tom Wainwright [1], the computational power has increased from about 1000 floating point operations per second (Flops) of early vacuum tube machines, like the UNIVAC at LLNL or the MANIAC at Los Alamos National Laboratory, to 360 TFlops on architectures like the BlueGene/L at LLNL [2]. With this increase of computational power the early 100-particle hard sphere system has grown into more sophisticated particle interactions—like smooth pair and many body potentials, as well as *ab initio* methods calculating the interatomic forces based on quantum mechanic principles—and the number of simulated particles has been increased to multiple billions. The importance of MD in science has evolved with this increase in capabilities as can be illustrated by the MD-related scientific literature fraction evolution (see Fig. 1). Today, 7 out of 1000 published scientific articles are related to MD; looking at a specific physics journal such as *Physical Review Letters,* this number is about 10 times larger, with the growth still in the exponential regime.

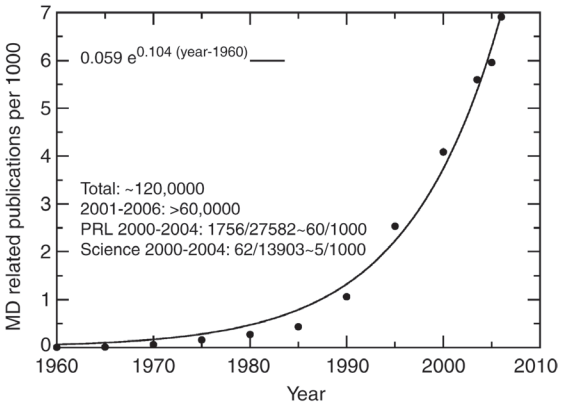

Fig. 1.
Evolution of the literature fraction of MD-related publications (i.e., molecular dynamics is a keyword or in the abstract). The literature search was performed using the following databases: BIOSIS, Engineering Index, Inspec, ISI Proceedings, ISI SciSearch, ISI Social SciSearch. An interdisciplinary journal such as* Science *follows this general trend, while specific physics journals such as* Physical Review Letters *have a significant larger fraction of MD-related publications. An exponential fit represents the general trend of the data.

With the introduced SPaSM [3] algorithm, multiple short performance runs were carried out on the BGL system using 4096-131,072 CPUs and 1-320 billion particles arranged in a face-centered cubic lattice and interacting via a Lennard-Jones potential with $r_{cut} = 2.5\ \sigma$ and with $r_{cut} = 5\ \sigma$. (See Fig. 2.)

The calculations were all performed in double precision, which means that each particle structure consisted of 88 bytes [position, velocity, and force vectors (24 bytes each), integer for type (4 bytes) and a tag (4 bytes), and a double to analyze/characterize the atoms (8 bytes)]. Hence, 320 billion atoms consume 26226 GB for the particle structure information alone. Additional memory is needed for the cell structure, various buffers, and the executable and operating system on each node. This additional memory only amounts to less than 20% of the whole memory consumption (LLNL's BGL has

32768 GB main memory). These numbers suggest that SpaSM, with its efficient memory management, would be able to perform with almost a trillion atoms in single precision on ASC Purple, which has a main memory of 48832 GB.

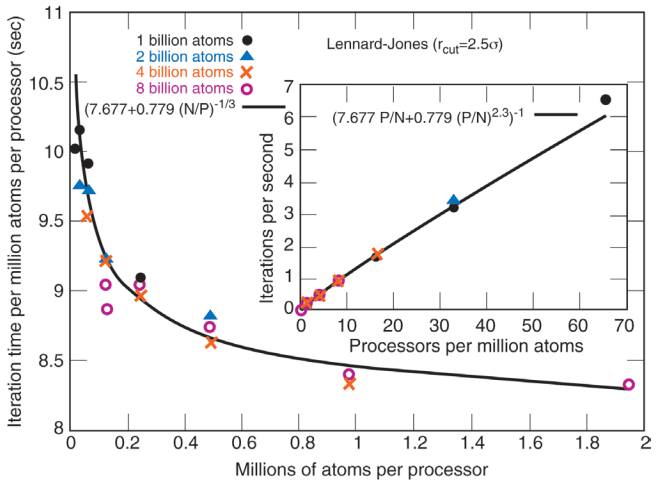

Fig. 2.
Strong scaling for 1–8 billion atoms distributed on 4096-65,536 processors of the IBM BGL machine at LLNL. For a small number of atoms per processor the communication/calculation ratio gets larger and the iteration time per atom per processor increases. The inset shows the speed-up as the number of processors per million atom increases.

Sophisticated use of the evergrowing computational power makes it possible to investigate challenges in physics and related science on the fundamental atomistic level. The 320 billion atoms simulated on BGL could represent a cubic piece of metal with an edge length well over a μm (Fig. 3). With one iteration taking only slightly more than 20 s, a physical simulation on the microscale in 3-D is now, in principle, possible in a matter of weeks. One of the biggest challenges to make these ultralarge-scale production runs feasible is the analysis and the input/output (I/O) of the enormous data produced on long-term storage, such as huge parallel file systems. The I/O is also very important for restart capabilities since time to failure of these enormous multiprocessor machines is usually only on the order of weeks if at all.

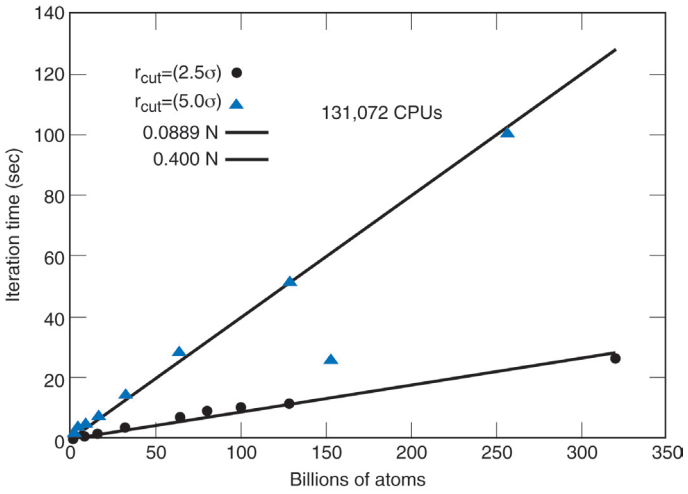

Fig. 3.
Iteration time for up to 320 billion atoms on 131,072 processors of the BGL architecture at LLNL. The achieved performance for 320 billion atoms simulated by a Lennard-Jones potential with a cutoff of 2.5 σ was 27.2 TFlops, and 48.1 TFlops for 256 billion atoms interacting via a Lennard-Jones potential with a larger cutoff of 5.0 σ.

For more information contact Kai Kadau at kkadau@lanl.gov.

[1] B. J. Alder and T. E. Wainwright, *J. Chem. Phys.* **27**, 1208 (1957).
[2] T.C. Germann, et al., *Supercomputing '05*, SC05 IEEE ACM 1-59593-061-2/05/0011 (2005).
[3] D.M. Beazley and P.S. Lomdahl, *Parallel Computing* **20**, 173–195 (1994).

Funding Acknowledgments
NNSA's Advanced Simulation and Computing (ASC).