

Grand Challenge: Millisecond-scale Molecular Dynamics Simulations

By Eduardo D'Azevedo

David E. Shaw, chief scientist of D.E. Shaw Research and a senior research fellow at the Center for Computational Biology and Bioinformatics, Columbia University, opened the SIAM Conference on Parallel Processing for Scientific Computing with a plenary talk on his group's exciting new research: designing massively parallel machine architectures and algorithms for the grand challenge of millisecond-scale molecular dynamics simulations of biological macromolecules.

D.E. Shaw Research (DESRES) conducts research in computational biochemistry to elucidate structures and dynamic behaviors of proteins. Molecular dynamics simulation is an important tool for the modeling of protein-size systems (25,000–50,000 atoms in water); applications include the development of new drugs. The grand challenge is to simulate such systems as long trajectories, in the millisecond time scale, where biologically interesting phenomena occur. Among these phenomena are the folding of proteins, the binding of drugs to molecular targets, interactions between proteins, and the dynamics of conformational changes in macromolecules. To put this molecular dynamics challenge into perspective, a single processor can simulate about one nanosecond in a day, and a massively parallel code might be able to simulate about one hundred nanoseconds per day. Meeting this grand challenge will thus require close to a hundred-fold speedup, which in turn will require new massively parallel architectures and innovative algorithms.

A traditional MD algorithm performs a discrete time integration of particles by computing the interaction forces among the particles, and uses Newton's law to advance the velocities and positions of the particles. The particles in a biological system are acted on by both "bonded" forces, such as the stretching, bending, and twisting of molecular bonds, and "non-bonded" forces, such as long-range electrostatic and short-range van der Waals interactions. The Lennard–Jones potential (also known as the 6–12 potential) is a simple mathematical model commonly used to approximate van der Waals interactions—long-range attraction and short-range repulsion. The 6–12 potential is of the form $V(r) = 4\epsilon[(\sigma/r)^{12} - (\sigma/r)^6]$. The non-bonded interactions account for most of the computation. Distant interactions can be computed efficiently via Ewald summation, the fast Fourier transform, the fast multipole method, or multigrid solvers. What remains is computing the range-limited interactions with neighboring particles within a cut-off distance R .

Spatial domain decomposition is commonly used to assign particles to boxes in a background rectangular grid, which is partitioned among processors (see Figure 1). A key aspect of a successful algorithm for computing the range-limited N -body problem is efficient communication. On a distributed-memory machine, each particle needs to have access to data for neighboring particles within the cut-off distance R . Two parallel algorithms are commonly used for computing the range-limited N -body problem: the traditional "home territory" method, in which two particles always interact within a processor that contains at least one of them, and the "neutral territory," or NT, method [3], in which two particles can interact in a processor to which neither is assigned. Surprisingly, the NT algorithm is often faster in massively parallel environments. Traditional spatial decomposition methods require interprocessor communication proportional to R^3 . This is also true for some neutral territory methods, such as the mid-point method [2].

For other NT methods, the interprocessor communication required is only $O(R^{3/2})$. The communication volume approaches zero as the number of processors is increased. Figure 2 shows the difference in communication patterns between the traditional and NT methods.

DESRES has developed a new MD code, called Desmond, that uses the novel NT parallel algorithms and numerical techniques to achieve very high performance levels on hardware ranging from a conventional commodity cluster to high-performance machines. A version of Desmond will be available at no cost for non-commercial use at universities and other not-for-profit research institutions. A commercial version of Desmond will also be available later this year.

Another project under way at DESRES is the development and implementation of Anton [4], a specialized massively parallel supercomput-

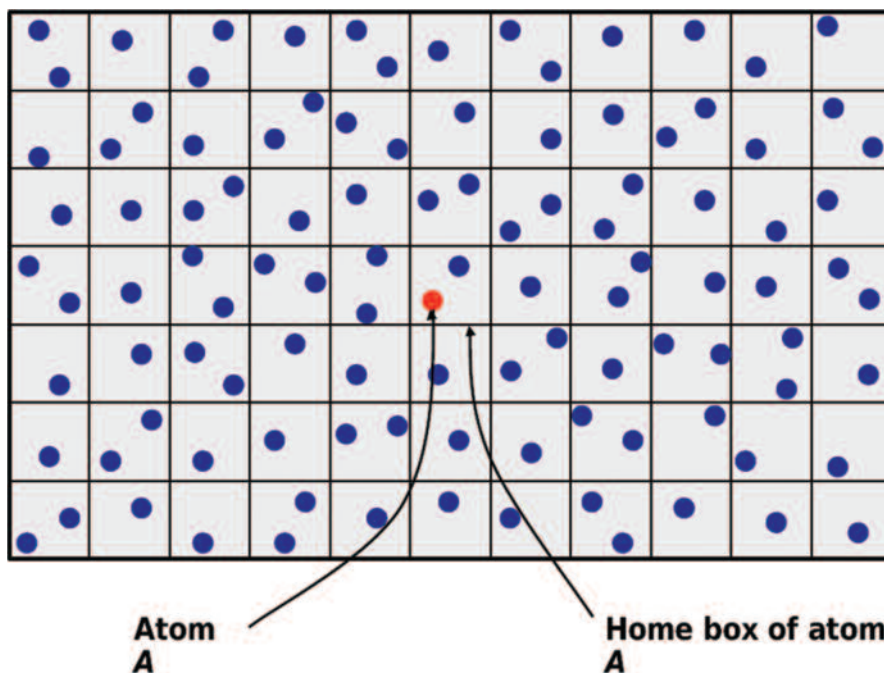


Figure 1. Spatial decomposition of particles into a background grid. Figures courtesy of D.E. Shaw.

er designed to execute MD simulations hundreds of times faster than has been possible to date. Anton achieves tremendous speedups through judicious use of “arithmetic specialization”: Flexibility or programmability is available only where needed; elsewhere, hardware is tailored for high speeds. The basic algorithms for computing electrostatic forces and van der Waals interactions are an example; the algorithms, although unlikely to change, can be adapted to different problems through the use of parameters and tables.

One segment of Anton consists of 512 application-specific integrated circuits (ASICs), connected as an $8 \times 8 \times 8$ three-dimensional torus.

The 3D torus reflects the physical space being simulated, with its nearest-neighbor connections and periodic boundary conditions. Communication on Anton is carefully choreographed so that data flows to a processor only when needed. Communication overhead is reduced, and off-chip memory accesses are almost never needed. Each ASIC achieves about 500 times the performance of a general-purpose microprocessor, but uses the same amount of power.

Each of Anton’s chips contains 32 particle-interaction pipelines—specialized fixed hardware that can efficiently compute the energy, potential, and force vectors. Anton uses more flexible programmable hardware for algorithms that are less regular or more likely to change, as for calculations of bonded interactions, bond-length constraints, and other alternative integration techniques. The flexible subsystem exploits three forms of parallelism: multicore parallelism, with four Tensilica cores (custom instructions) and eight geometry cores (instruction-level parallelism), and single instruction, multiple data (SIMD) parallelism, with 3D vectors manipulated as a single operation. Each geometry core is implemented as a very long instruction word (VLIW) processor. Test runs on a real (prototype) Anton processor have produced very promising performance results.

The MD simulation has yielded scientific insight into the mechanism of “NhaA antiporting,” in the selective transport of H^+ and Na^+ in cytoplasm [1]. In another case, a model of the “DFG flip” in Abl kinase, the simulation predictions for effects of pH variations and effects of mutations are consistent with experimental results from John Kuriyan’s laboratory at the Howard Hughes Medical Institute at the University of California, Berkeley.

By the end of 2008, with the expected completion of the full Anton machine and the possibility of regular millisecond-scale MD simulations, we can look forward to exciting new discoveries.

References

- [1] I.T. Arkin et al., *Mechanism of Na^+/H^+ antiporting*, *Science*, 317 (2007), 799.
- [2] K.J. Bowers, E. Chow, H. Xu, R.O. Dror, M.P. Eastwood, B.A. Gregersen, J.L. Klepeis, I. Kolossvary, M.A. Moraes, F.D. Sacerdoti, J.K. Salmon, Y. Shan, and D.E. Shaw, *Scalable algorithms for molecular dynamics simulations on commodity clusters*, Proceedings of the ACM/IEEE Conference on Supercomputing, Tampa, Florida, November 2006, <http://sc06.supercomputing.org/schedule/pdf/pap259.pdf>.
- [3] K.J. Bowers, R.O. Dror, and D.E. Shaw, *Overview of neutral territory methods for the parallel evaluation of pairwise particle interactions*, *J. Physics: Conference Series*, 16 (2005), 300–304.
- [4] D.E. Shaw et al., *Anton, a special-purpose machine for molecular dynamics simulation*, *ACM SIGARCH Comput. Architec. News*, 35:2 (May 2007).

Eduardo D’Azevedo is a member of the Computer Science and Mathematics Division at Oak Ridge National Laboratory.




Figure 2. Meeting on neutral territory. Communication patterns in traditional (left) and NT (right) methods.