Anton: A Specialized Machine for Millisecond-Scale Molecular Dynamics Simulations of Proteins

David E. Shaw
D. E. Shaw Research and
Center for Computational Biology and Bioinformatics, Columbia University

The ability to perform long, accurate molecular dynamics (MD) simulations involving proteins and other biological macromolecules could in principle lead to important scientific advances and provide a powerful new tool for drug discovery. A wide range of biologically important processes, however, occur over time scales on the order of a millisecond -- several orders of magnitude beyond the duration of the longest previous MD simulations. Our research group has recently completed a specialized, massively parallel machine called Anton, which is capable of calculating millisecond-scale molecular trajectories at an atomic level of detail. The machine has greatly extended the power of simulation as a tool for understanding the structure and dynamics of proteins, and has already allowed us to observe and analyze important biological phenomena that have not previously been accessible to either computational or experimental study.

Anton's dramatically accelerated execution of MD simulations is attributable in large part to specialized logic for the high-speed calculation of pairwise interactions between particles and/or gridpoints separated by no more than some specified cutoff radius -- a "range-limited" version of the classic *N*-body problem in physics.

In particular, each of Anton's 512 ASICs, which are implemented using 90-nm technology, includes a "high-throughput interaction subsystem" (HTIS) incorporating 32 highly specialized pipelines running at 800 MHz. During every cycle, each of these pipelines produces a pairwise-interaction result that would require approximately 50 arithmetic operations to calculate on a general-purpose processor. Novel algorithms and architectural features are used to greatly reduce the requirements for intra- and inter-chip communication, allowing Anton to feed these pipelines and collect their results at a speed sufficient to take advantage of the machine's computational power. The ASIC also includes a "flexible subsystem" based on eight programmable "geometry cores," each containing eight arithmetic pipelines.

This talk will describe the architecture of the Anton machine, with special attention to the manner in which the HTIS achieves unusually high throughput in calculating the effects of range-limited *N*-body interactions.

1063-6889/09 $25.00 © 2009 IEEE
DOI 10.1109/ARITH.2009.33