

“Anton, a Special-Purpose Machine for Molecular Dynamics Simulation”



DAVID E. SHAW, MARTIN M. DENEROFF, ET AL
D. E. SHAW RESEARCH
ISCA, 2007

PRESENTED BY ALEX EDELSBURG



Introduction



"Do one thing and do it well"

- Special-purpose logic designed for MD calculations
- Parallel algorithms tuned for the machine to minimize comms

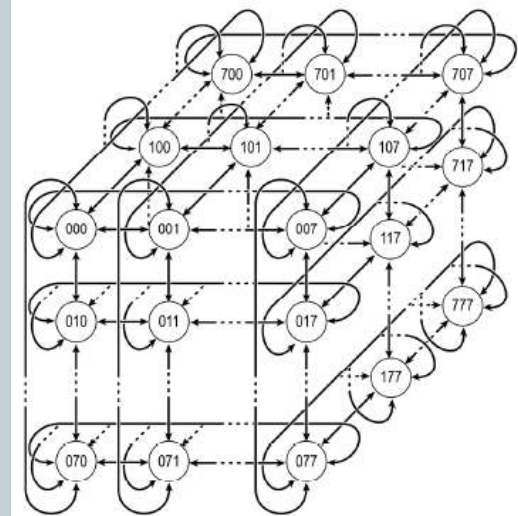
Motivation

- Simulate biochem phenomena that can't be observed in the lab
- Phenomena occur at millisecond timescale
- Best simulations do microsecond timescale

Hardware Overview



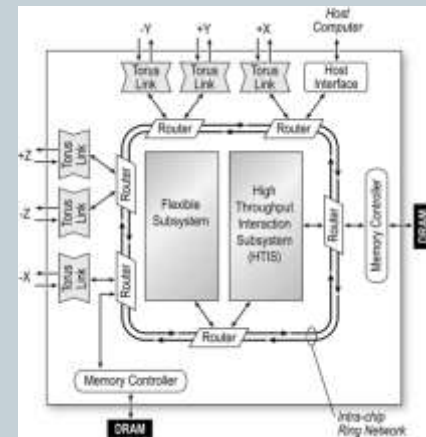
- 512 identical MD-specific ASICs
- Specialized high-speed comms
 - High-speed 3D torus
- Programmable core that directs ASICs
- Balanced system – keep MD ASICs fed
 - Weakness of past MD-specific systems



MD Calculations



- Motion of collection of atoms using classical laws of physics
 - Atoms of target molecule
 - Atoms of solvent (usu water)
- Force of every particle on every other
- Integration of all forces to yield new velocity & position



Parallelization



- Break cell up into a grid, much like ocean
 - Each grid box lives on an ASIC
 - Each particle has a home node (ASIC)
- NT method for pairwise non-bonded interactions
 - Each ASIC responsible for its own grid + part of neighboring grids
 - Can calculate the interactions somewhere other than home node – "neutral territory"
 - This allows for more even spread of calculations

MD ASICs



- Don't need large caches
- Devote more chip area to computations
- Pay attention to Amdahl's Law when speeding up particular operations

- Systolic array
 - Each PPIM in the grid does the following: receive, compute, transmit
 - Atoms "flow" through the array and intersect at the various PPIMs

Conclusions



- **Special hardware is a big win**
 - 80-100x faster than competition
 - Hardware design led to algorithm insights
- **...if you can justify it**
 - Failure of LISP Machines, database accelerators
 - Must beat Moore's Law
- **It works!**
 - They built it

Questions/Concerns



- When does a problem domain warrant specialized hardware?
 - They explained the Moore's Law justification for specialized hardware
- Are GPGPUs a more cost-effective option?
 - GPGPUs have become more prevalent since the publication of this paper.
- Would a hierarchy of interconnection structures improve communications?
 - Small collection of nodes = snoopy
 - Large collection of meta-nodes = directory
- Could Anton be used for other classes of problems with similar computational needs?