Accelerating Biomolecular Simulation using a Scalable Network of Reconfigurable Hardware

Arun Patel[1], Christopher Madill[2,3], Manuel Saldaña[1], Christopher Comis[1], Dave Chui[1], Sam Lee[1], Régis Pomès[2,3], Paul Chow[1].

1 Department of Electrical and Computer Engineering, University of Toronto. 2 Department of Structural Biology and Biochemistry, The Hospital for Sick Children. 3 Department of Biochemistry, University of Toronto.

Motivation

- Computer simulations of biomolecules is increasingly playing an important role in medical research.
- Understanding the balance of physical forces governing atomic-level interactions is a central challenge in modern Biochemistry.
- Computer simulations have been successfully applied to study biophysical phenomena including:
  - i) Cell membrane transport.
  - ii) Molecular conformational equilibria.
  - iii) Protein / Ligand docking.
  - iv) Time-dependent molecular motion.
- Current technological limitations restrict the size and length of simulations.
- The primary objective of this study is to engineer a scalable Molecular Dynamics simulator which is capable of outperforming supercomputers and computing clusters.
- This project is a collaboration between the Department of Electrical and Computer Engineering, and the Department of Biochemistry, University of Toronto and the Department of Structural Biology and Biochemistry, The Hospital for Sick Children.

Force Calculations ($\vec{F} = -\nabla E$)

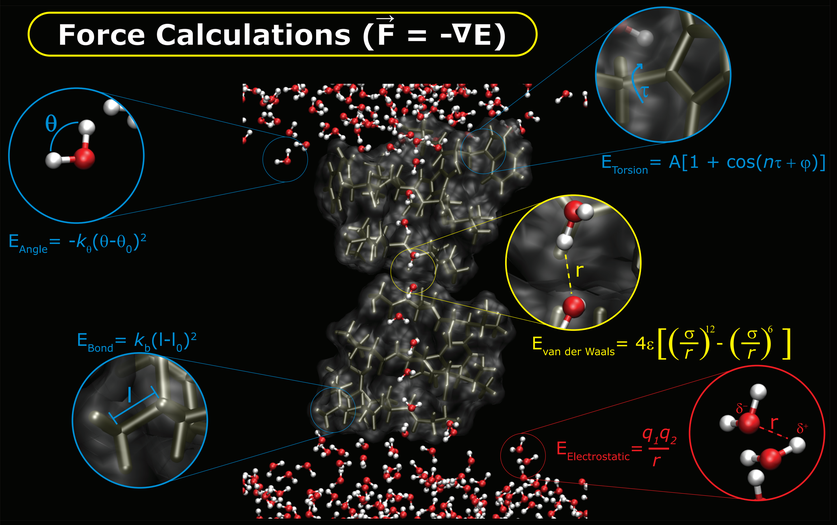

Molecular Dynamics Architecture

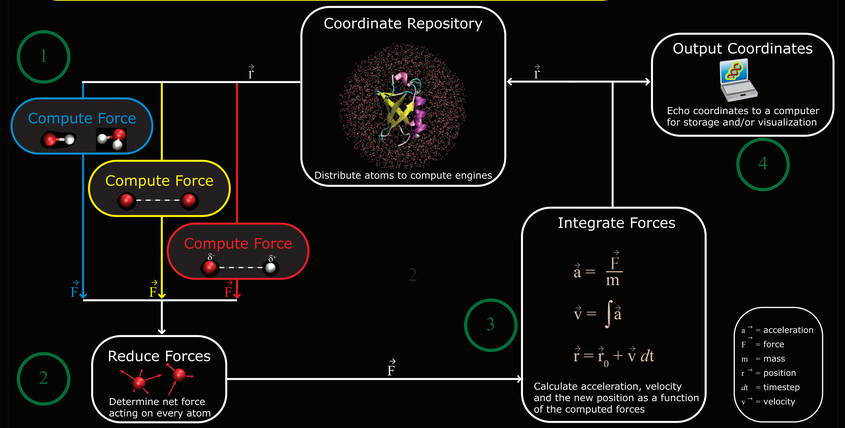

Conventional vs. Proposed Implementations

- NAMD, a state-of-the-art MD program, employs many techniques to improve simulation performance including:
  - Spatial decomposition to enable parallel force evaluations.
  - PM-Ewald algorithm to improve electrostatic calculation efficiency.
  - Nonbonded cutoff to reduce the number of pairwise calculations.
- Despite these enhancements, a protein folding reaction which occurs in 10^{-5} seconds would require about 30 years of CPU time to simulate.

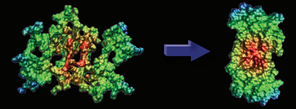

- Consequently, MD simulations are often run on supercomputers or large clusters.

3000-node Molecular Dynamics Cluster at the Pittsburgh Supercomputing Center

- Our proposed implementation is not another attempt at building a supercomputer. This is an interdisciplinary effort to design a machine that performs MD simulations roughly 10^3x faster than the current supercomputer-based approach.
- MD is a highly-parallel problem with a large computation-to-data transfer ratio.
  - Computation throughput is improved vastly by performing time-consuming computation kernels with hardware accelerators.
  - Using a network of interconnected FPGAs, we achieve greater parallelization and higher integration than typical supercomputers.
  - The fully-interconnected network topology, developed by profiling NAMD network activity, is ideal for transferring small data quickly.
- A single rack of FPGA-based hardware can out-perform a supercomputer in MD simulation!

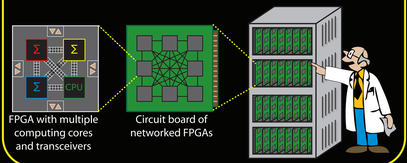

Hardware Components

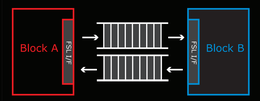

- Fast Simplex Links abstract on-chip hardware communication
- Standard interface enables rapid integration of system modules

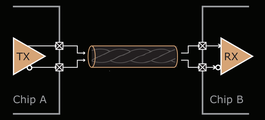

- Multi-Gigabit Transceivers allow gigabit-rate communication
- Serial Protocol enables a fully-interconnected network topology

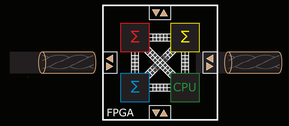

- FPGAs combine computation and communication hardware on-chip
- Reconfigurability = Flexibility
- High integration density possible

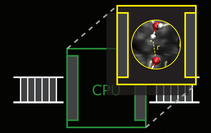

- Programming model allows mix of hardware and software elements
- Soft-processors can be used to emulate hardware functionality

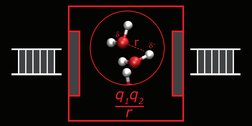

- Ewald Electrostatic Force Engine
- Computes electrostatic forces in hardware with N·log(N) algorithm
- Other engines use N^2 algorithm

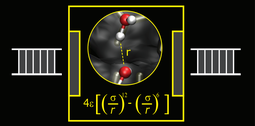

- Lennard-Jones Potential Engine
- Computes van der Waals forces
- Outperforms existing software implementations by up to 88x

Hardware Prototypes

First Generation Prototype

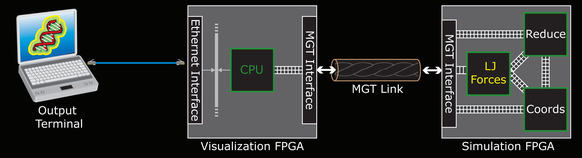

- Goal of initial prototype is to develop programming model
- Uses standardized Message-Passing Interface (MPI) software for communication
- Accelerators are emulated using soft-processors
- Control processor separated from computation units
- Lennard-Jones Forces only

Second Generation Prototype

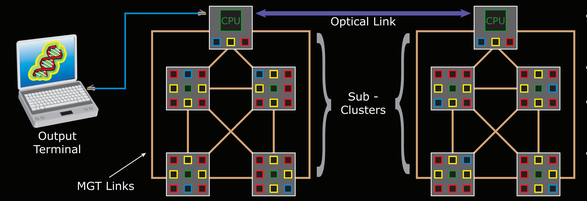

- Second prototype divided into fully-interconnected sub-clusters of five FPGAs
- Optical links used to connect sub-clusters together
- Embedded soft-processors handle control flow and perform scheduling duties
- Each FPGA "node" contains heterogeneous mixture of processors and accelerators