

An Investigation of the Effects of Error-Correcting Code on GPU-accelerated Molecular Dynamics Simulations

Ross C. Walker

San Diego Supercomputer Center
Dept of Chemistry & Biochemistry
UC San Diego

Robin M. Betz

San Diego Supercomputer Center


What is Molecular Dynamics?

- > Simulation of dynamic properties of condensed phase physical / biological systems
 - > Enzymes/Proteins
 - > Drug Molecules
 - > Biological Catalysts
- > Classical energy function
 - > Parameterized force fields (bonds, angles, dihedrals, VDW, charges...)
 - > Integration of Newton's equations of motion
 - > Atoms considered as points, electrons are implicit

$$V(x) = \sum_{bonds} K_b (r - r_0)^2 + \sum_{angles} K_a (\phi - \phi_0)^2 + \sum_{dihedrals} \frac{K_d}{2} [1 + \cos(n\vartheta + \gamma)] + \sum_{i,j < i} \left(\frac{A}{r^{12}} - \frac{C}{r^6} + \frac{q_i q_j}{r} \right)$$

What can we do with Molecular Dynamics?

- > Simulation of time dependent properties
 - > Protein domain motions
 - > Small protein folds
 - > Spectroscopic properties
- > Simulation of ensemble properties
 - > Binding free energies
 - > Reaction pathways
 - > Free energy surfaces



Binding of Tamiflu to influenza neuraminidase

AMBER GPU-Acceleration


- > Collaboration between NVIDIA and AMBER development team
- > SPFP model for calculation accuracy and speed
- > Focus on accuracy
 - > Passes all AMBER development tests
 - > Bitwise identical trajectories for runs with same random seed
 - > Conserves energy comparably to double-precision CPU code



NVIDIA®



AMBER GPU-Acceleration

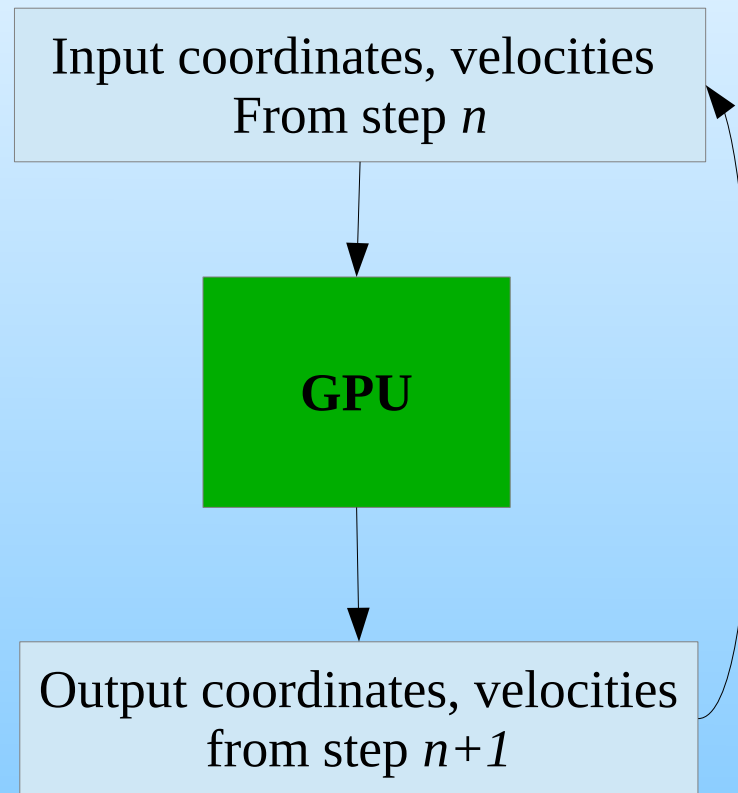


JAC (DHFR) benchmark: 23,559 atoms PME

Error Correcting Code (ECC)

- > Random bit-flip events (ECC errors) occur in memory
- > ECC using Hamming codes to check memory is unaltered
- > Can detect and correct single bit-flip errors
- > Can detect double bit-flip errors
- > Costs about 10% of memory and speed in GPU simulations

Error Correcting Code (ECC)




Error Correcting Code (ECC)




Normal
Operation

Error Correcting Code (ECC)




Normal
Operation

Error Correcting Code (ECC)



Error Correcting Code (ECC)



Error Correcting Code (ECC)




Error Correcting Code (ECC)


1001




Error Correcting Code (ECC)



Error Correcting Code (ECC)



Error Correcting Code (ECC)




Error Correcting Code (ECC)




Error Correcting Code (ECC)


1001




Error Correcting Code (ECC)




Error Correcting Code (ECC)



Error Correcting Code (ECC)




Error Correcting Code (ECC)





How Frequent Are ECC Events?

- > “DRAM Errors in the Wild” Google study
 - > 25,000-75,000 errors/billion device hours/MBit
 - > ~8% of DIMMs affected by errors each year
 - > Majority of errors are “hard errors” => physical defect
 - > Soft errors perhaps 2% of all errors
- > Unknown error frequency for GPUs

Simulation




Satellite Tobacco Mosaic
Virus




Simulation

- > Satellite Tobacco Mosaic Virus (STMV)
- > 1,067,095 atoms in explicit solvent
- > Uses 48% (2.6 GB) of GPU memory
- > NPT simulation at 300K with 2fs timestep
- > Start from equilibrated trajectory



Supercomputer

- > Keeneland GPU supercomputer at Georgia Tech
- > 240 available nodes for our test
- > 3 M2090 GPUs per node
- > Run approximately 10 hours on all nodes with ECC on, then 10 hours with ECC off




Results

- > ECC on:
 - > 3 runs failed due to node I/O error
 - > 717 runs completed successfully
- > ECC off:
 - > 1 run hung during middle of simulation
 - > 719 runs completed successfully

No ECC events detected


Results

- > Mean time for 0.3 ns of simulation:
 - > ON: 9.951 hours
 - > OFF: 9.096 hours
- > Standard deviation:
 - > ON: 0.079 hour
 - > OFF: 0.006 hour



Results

- > Mean time for 0.3 ns of simulation:
 - > ON: 9.951 hours
 - > OFF: 9.096 hours
- > Standard deviation:
 - > ON: 0.079 hour
 - > OFF: 0.006 hour



Conclusion

- > ECC soft error events are very rare
- > Increased likelihood of failure with longer wall clock times
- > Turning ECC off may benefit AMBER users

Conclusion

- > ECC soft error events are very rare
- > Increased likelihood of failure with longer wall clock times
- > Turning ECC off may benefit AMBER users

More research is needed into ECC soft error events

Acknowledgments

- > Dr. Ross Walker
- > Keeneland
- > NVIDIA
- > NSF



Keeneland