GPU-Enabled Monte Carlo Simulation Makes Study of Cardiac Arrhythmia Possible

Tuan M. Hoang-Trong[1], Aman Ullah[1], W. J. Lederer[2], M. Saleet Jafri[1,2]
[1]Krasnow Institute for Advanced Study, George Mason University, Fairfax VA 22030
[2]Center for Biomedican Engineering and Technology, University of Maryland School of Medicine, Baltimore MD 21201

ABSTRACT

Heart disease is the leading cause of death in the developed world. Many of these deaths occur through fatal arrhythmia that may develop from dysfunction at cellular and subcellular levels. Multi-scale stochastic computational models that integrate the function of individual transmembrane proteins called ion channels that number in the millions in the cardiac myocyte are critical to understand the complex dynamics of the myriad components that comprise the heart and thus, is needed to understand arrhythmia. Simulation of such computational complexity is now possible due to our patented Ultrafast Monte Carlo Algorithm and its implementation on GPUs. are required to integrate data.

INTRODUCTION

Multi-scale computational models are required to integrate data to understand the complex dynamics of biological systems and to predict how changes to the system can lead to disease. These models integrate components that can span spatial and temporal scales of several orders of magnitude. In our work, we apply these principles to the understanding of heart disease which is the leading cause of death in the developed world. Many of these deaths occur through fatal arrhythmia. These arrhythmias can be initiated by the dysfunction of a specific protein molecule resulting from a genetic defect or some perturbation of the physiological system. In particular, we are studying how defects in proteins (calcium channels) that govern calcium regulation in the heart can initiate an arrhythmia. Our simulations model the stochastic behavior of these calcium channels and span up to the tissue level. In a single heart muscle cell there are over 1,000,000 of these calcium channels. In the heart, there are 2-3 billion of these cells.

When we started this work we tried simulating this system using standard Monte Carlo simulation methods on a CPU or using MPI however, these simulations took a month to simulate one second of simulation time. Traditionally to make such a computationally expensive tractable, model reduction methods and simplification of the system were used. We even developed ensemble density methods that could replace the stochastic system by its average behavior. However, these removed necessary detail of the system or removed the stochasticity needed to trigger an arrhythmia. We then developed our new patented Ultra-fast Monte Carlo Method that builds upon the theory of stochastic automata networks. Our improvements offered a 100-fold speed up over these methods.

We ported our method to the first generation C1060 and achieved another 20-30x speed up. This allowed us to reduce the 31 days or run time to 45 seconds. With each ensuing generation of GPU architecture we have achieved a 2x speedup.

The poster will describe the challenges that we encountered in our multi-scale systems biology study of cardiac arrhythmia and how GPU and algorithmic development helped us to overcome these limitations.

THE MODEL

Schematic of SR Ca^{2+} leak model and release site schematic. **(A)** Model compartments and Ca^{2+} fluxes (solid arrows). **(B)** Transition state-diagram for the two-state Markov chain describing a single RyR. **(C)** Transition-state diagram for the Markov chain representing the RyR cluster where each state indicates the number of open RyRs (No) in the CRU (e.g., 0, 1, 2, 48, 49). **(D)** L-type channel model from Sun. C1 and C6 are closed states with C6 the resting state. O2 and O3 are the open states. C4 is Ca2+-dependent inactivated state. C5 is the voltage-dependent inactivated state.

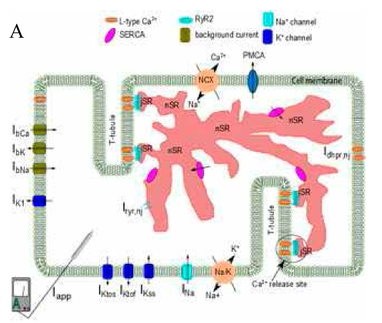

- Rat ventricular myocyte: 20,000 CRUs/cell

Each CRU:

- 7 L-type Ca^{2+} channels (LCCs) + 49 RyR2
- Novel LCC model using Markov chain with both V_m- dependent activation/inactivation and Ca^{2+}-dependent inactivation.
- Model RyR2 as 2-state with both cytosolic Ca^{2+} sensitivity and luminal Ca^{2+} regulation.
- RyR gating incorporate energetic coupling formulation.

COMPUTATIONAL CHALLENGES

Stochastic Simulation is very computationally expensive. Various approaches have been suggested to address these.

Reduction methods

- Reduction methods make assumptions about the system to make the reduction possible. These might not be valid under all relevant conditions.
- Sometimes the reduction require simplification of the system by reducing model complexity. This limits the details that can be included.

Monte Carlo Simulation of simpler system

- Other attempts have simplified the dynamics of ryanodine receptor gating or the system.
- Using slower kinetics, fewer channels, release units, omitting physiological/biophysical detail reduced veracity of the model.

We have developed our patented Ultrafast Monte Carlo Method that makes the computation possible.

ULTRA FAST MONTE CARLO ALGORITHM

Process

Create Infinitesimal generator matrix (Q-matrix) that contains transition rate constants for single cluster (Markov-based model)

Use Kronecker Product to form transition matrix for the heterogeneous clusters. These can depend upon the system conditions.

Create a state space matrix describing the state of all combinations of random variables.

Perform Monte Carlo Step which multiplies the transition matrix for the cluster with the state space matrix.

CUDA kernel 1

Integrate differential equations that depend upon the random variables and will determine the system conditions.

CUDA kernel 2

update

update

Implementation Detail

This can result in large matrices (>10Gb) so we use compact representations of the matrices.

We have developed an algorithm to calculate the Kronecker Product of the compact matrices without needing to un-compact them. Here, all matrix-matrix and matrix-vector operations are avoided to greatly reduce computational expense.

This reduces state space as we count the number of random variable in each state rather than the state of each random variable.

Create a probability transition matrix, based on the state space matrix describing the state of all combinations of random variables.

BENCHMARKS (C1060)

Method	Traditional Monte Carlo On CPU	Ultra-fastMonte Carlo On CPU	Ultra-fast Monte Carlo On GPU
Runtime	11000 min	20 min	45 sec
Speedup	1x	550x	15000x

Note that an additional 4x speed-up has been achieved with Fermi and Kepler over CPU alone.

APPLICATION TO CARDIAC ARRHYTHMIA

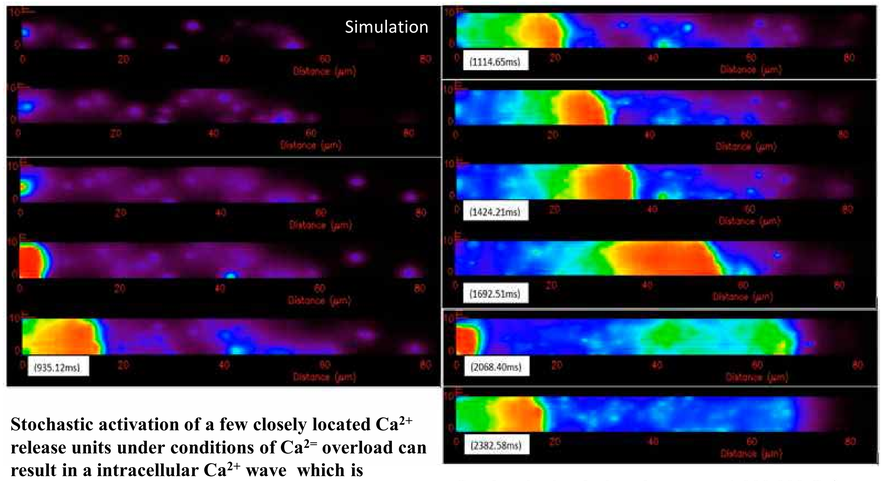

Stochastic activation of a few closely located Ca^{2+} release units under conditions of $Ca^{2=}$ overload can result in a intracellular Ca^{2+} wave which is thought to be sufficient to initiate an arrhythmic event on the cellular level. Simulations are show above and experiment below for comparison.

Stochastic simulations have over 1,000,000 Ca^{2+} release channels located in the myocyte which has ~ 3,000,000 grid points are easily solved on a single GPU in a few hours. Large simulations are planned that incorporate several cells that will leverage multi-GPU calculation.

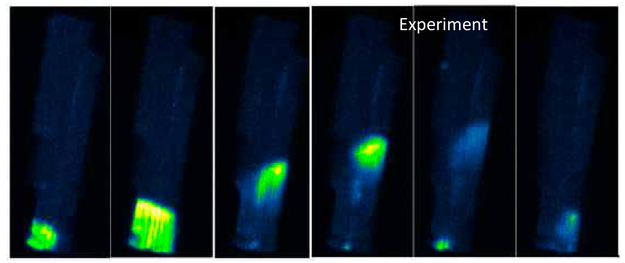

We simulate ~40,000 myocytes represented by compartmental models that have stochastic Ca ²+ release. The geometry of the fine structures of the heart lowers the number of cells needed to initiate an arrhythmia from a few hundred cells to 12-64 cells. The heart has many fine structures such as Purkinje fibers and trabeculae (lining the chambers of the heart) that are in effect 1D structures). We use simulated trabeculae that interfaces with the heart wall.

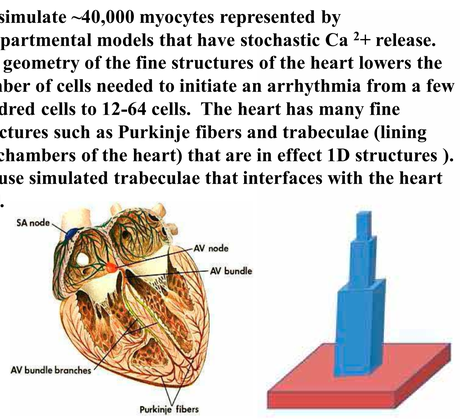

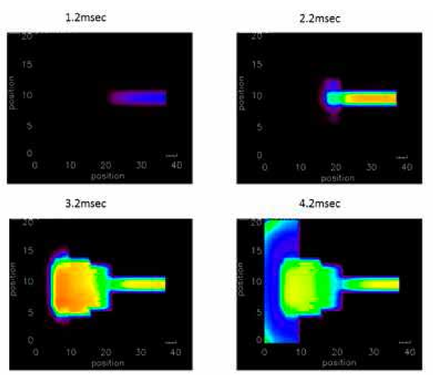

CONCLUSIONS

In conclusion, we presented here a brief description of our Ultrafast Monte Carlo Algorithm and how we have applied it to modeling cardiac arrhythmia.

- The method offers a significant speed up that is augmented significantly by use of GPU.
- These simulations using this method have provided novel insights into mechanisms of cardiac arrhythmias.

ACKNOWLEDGMENTS

This work is supported by the National Institutes of Health Grants 5R01HL105239 and 5R01AR057348, NSF 0443843, and a grant from the NVIDIA Corporation. The research leading to the results has received funding from the European Community's Seventh Framework Programme FP7/2007-2013 under grant agreement No HEALTH-F2-2009-241526, EUTrigTreat.