

# Heterogeneous Computing for Cardiac Electrophysiology

Xing Cai

Simula Research Laboratory & University of Oslo, Norway

April 26, 2019

- Introduction
- Two computational problems in cardiac electrophysiology
  - Tissue-scale electrical activity coupled with cell-level calcium handling
  - Action potential wave propagation in the entire heart
- “A peek into the future”, through the **eX<sup>3</sup>** infrastructure

# Heterogeneous computing

# Top500 Linpack ranking, November 2018

Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	<b>Summit</b> - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband , IBM DOE/SC/Oak Ridge National Laboratory United States	2,397,824	143,500.0	200,794.9	9,783
2	<b>Sierra</b> - IBM Power System S922LC, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband , IBM / NVIDIA / Mellanox DOE/NNSA/LLNL United States	1,572,480	94,640.0	125,712.0	7,438
3	<b>Sunway TaihuLight</b> - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway , NRCPC National Supercomputing Center in Wuxi China	10,649,600	93,014.6	125,435.9	15,371
4	<b>Tianhe-2A</b> - TH-IVB-FEP Cluster, Intel Xeon E5-2692v2 12C 2.2GHz, TH Express-2, Matrix-2000 , NUDT National Super Computer Center in Guangzhou China	4,981,760	61,444.5	100,678.7	18,482
5	<b>Piz Daint</b> - Cray XC50, Xeon E5-2690v3 12C 2.6GHz, Aries interconnect , NVIDIA Tesla P100 , Cray Inc. Swiss National Supercomputing Centre (CSCS) Switzerland	387,872	21,230.0	27,154.3	2,384

# Hardware accelerators are powerful and energy-efficient

Hardware	Peak DP rate	Peak memory BW
Intel Xeon 8180 28-core CPU	2.06 TFlop/s	128 GB/s
NVIDIA V100 GPU	7.80 TFlop/s	900 GB/s
Intel Xeon Phi 5110P coprocessor	1.01 TFlop/s	320 GB/s




Accelerators have tremendous computing power, but their usage is non-trivial.

- Programming a hardware accelerator is different from programming a CPU
  - Unifying programming languages/standards/tools for heterogeneous computing do exist, but they have performance/functionality limitations
- To achieve overall performance of heterogeneous computing:
  - Single-accelerator performance
  - Single-node performance: CPU+accelerator(s)
  - Multi-node performance


# Cardiac electrophysiology

# The heart and its cells


An adult human heart has about 7.5 billion ventricular cells.



Longitudinal view



Cross-sectional view



Ventricular cells

# The “overall” picture

- The heart needs to be electrically stimulated so that it can contract during every heartbeat
- Proper **calcium handling** is crucial for normal heart function
- Several diseases can impair calcium handling
- Simulations can shed light on the complex interactions
- Multiscale models and high computational resolutions are needed

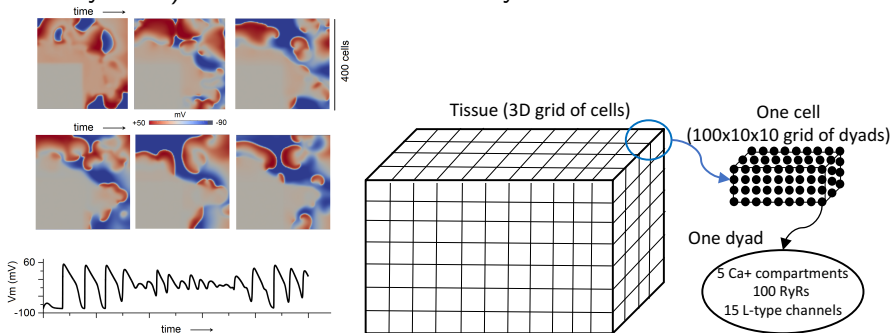
**Many challenging computational problems with biomedical importance**

## Two computational problems in cardiac electrophysiology

- 1 Tissue-scale electrical activity coupled with cell-level calcium handling
- 2 Action potential wave propagation in the entire heart

# Computational problem 1: Tissue-scale electrical activity coupled with cell-level calcium handling

Objective: To connect the tissue-scale abnormalities (such as arrhythmia) with subcellular calcium dynamics



Physiological realism in the true number of calcium release units per cell, while coupling a large number of cells together

A cardiac issue is considered as a 3D grid of cardiac cells.

On the tissue level, the monodomain equation describes the inter-cell electrical coupling:

$$\frac{\partial V_m}{\partial t} = \frac{-I_{\text{ion}}}{C_m} + D_x \frac{\partial^2 V_m}{\partial x^2} + D_y \frac{\partial^2 V_m}{\partial y^2} + D_z \frac{\partial^2 V_m}{\partial z^2}.$$

The ionic current  $I_{\text{ion}}$  per cell is the sum of various ionic currents determined by an underlying multiscale cell model of detailed calcium handling.

Each cell incorporates detailed stochastic calcium handling.

- There are 10,000 calcium release units, also called **dyads**, per cell.
- Each dyad consists of five calcium compartments:
  - myoplasm  $Ca_{myo}$ ,
  - submembrane space  $Ca_{ss}$ ,
  - network sarcoplasmic reticulum  $Ca_{NSR}$ ,
  - junctional sarcoplasmic reticulum  $Ca_{JSR}$ ,
  - dyadic space  $Ca_{ds}$ .
- The five calcium concentrations are modeled by a system of nonlinear reaction-diffusion equations.

- Finite differences for discretizing the differential equations
- Explicit time integration
- On the tissue level: solving the monodomain equation, each cardiac cell being a mesh point on the 3D global computational grid
- Within each cell: solving three 3D reaction-diffusion equations and two ODEs over a  $100 \times 10 \times 10$  intra-cell grid, where each dyad is a mesh point
- Within each dyad, 100 RyRs and 15 L-type channels follow a Markov model
- Memory requirement: about 3 MB per cell


- Computation per time step:

```
for (int k = 1; k <= num_subd_cells; k++) {  
  for (int j = 1; j <= num_dyads; j++) {  
    L-type probability calculation  
    L-type opening  
    RyR probability calculation  
    RyR opening  
    Ca concentration computation  
  }  
  diffusion between dyads  
}  
difusion between cells  
MPI communication
```

- Important performance optimizations: removal of redundant calculations, pre-calculations, code vectorization, economic numerical algorithms

## Using binomial distribution

There are 100 RyRs per dyad. Each RyR has four states, the 8 possible transitions between the states follow a Markov model.

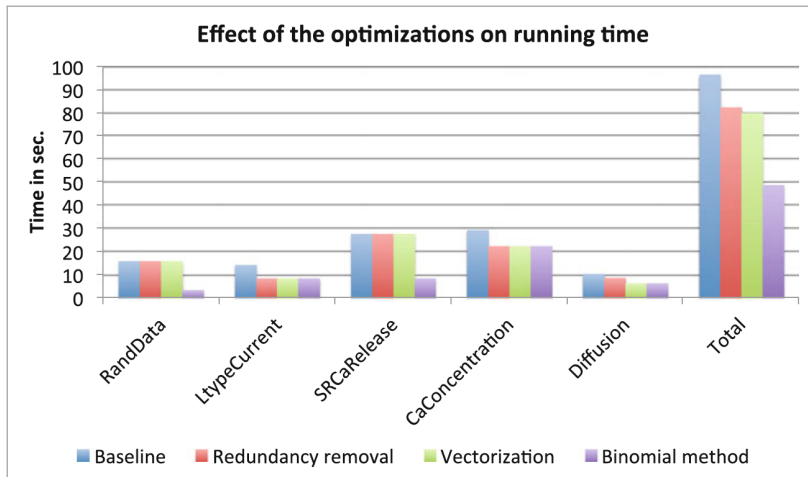


Instead of individually simulating the random state transitions for each RyR, we use 8 random samples from a binomial distribution. (Only the number of open RyRs per dyad is of interest.)


The cumulative probability of having up to  $k$  successes in  $n$  trials is

$$F(k, n, p) = \Pr(X \leq k) = \sum_{i=0}^k \binom{n}{i} p^i (1-p)^{n-i}$$


# Impact of performance optimizations



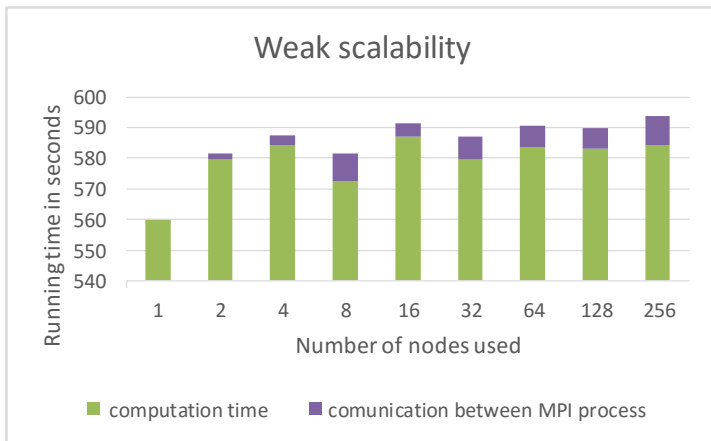
# Using heterogeneous supercomputer: Tianhe-2





- Tianhe-2: former No. 1 supercomputer (2013-2015)
- 16,000 compute nodes (before 2018)
  - Each node has 3 Xeon Phi Knights Corner 57-core coprocessors



Partitioning of the cells mostly with respect to the computing capability of the 3 KNCs and the CPU host




# Simulation snapshots



## Computational problem 2: Action potential wave propagation in the entire heart

Objective: To use realistic geometry and conductivity properties of the entire heart




Due to a spatially varying conductivity tensor field  $\mathbf{D}(x, y, z)$ :

$$\frac{\partial V_m}{\partial t} = \frac{-I_{\text{ion}}}{C_m} + \nabla \cdot (\mathbf{D} \nabla V_m)$$

Also, we have to adopt an unstructured 3D tetrahedral mesh.

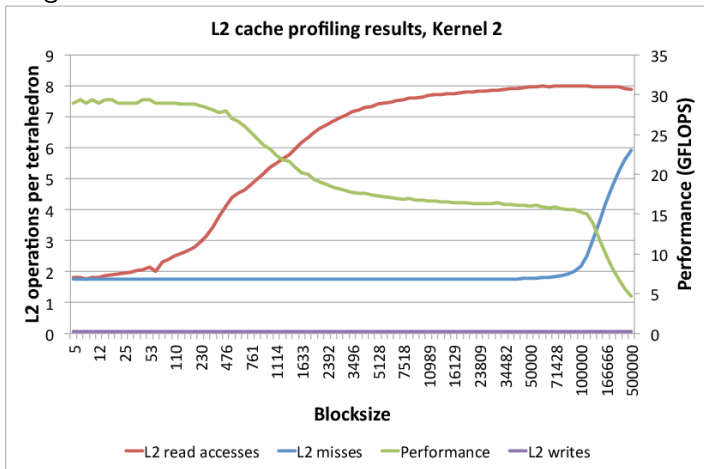
# Challenge with unstructured computational mesh



- Unstructured mesh  $\Rightarrow$  irregular & indirect data accesses
- Lower computational intensity (more data traffic per FLOP)

# Using a single GPU

Performance study of sparse matrix-vector multiplication (SpMV) on a single GPU:



Optimization consideration: ordering of the tetrahedrons in a “memory-friendly” fashion

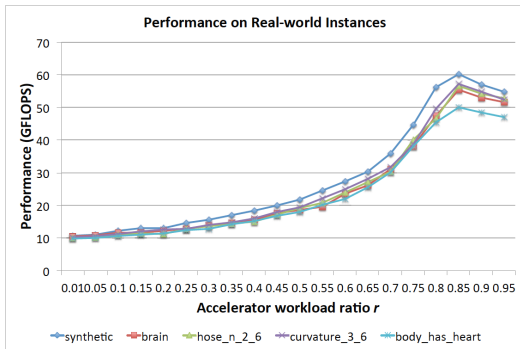
Example case: unstructured 3D tetrahedral mesh & cell-centered computation (4 neighbors per tetrahedron)

- Hardware: dual socket E5-2650 CPUs + two K20m GPUs
- Theoretical peak performance on the CPU side:  $\frac{11}{64} \times 102.4$  GB/s=17.6 GFLOPs
- Realistic performance on the CPU side: 11 GFLOPs
- Theoretical peak performance from the two GPUs:  
 $2 \times \frac{11}{64} \times 208$  GB/s=71.5 GFLOPs
- Realistic performance from the two GPUs: 54 GFLOPs

So, what should be the percentage of work given to the CPU and GPU sides?


# Work division between CPUs and GPUs

Theoretical performance estimates  $\rightarrow$  the GPU side should get  $\frac{71.5}{17.6+71.5} \approx 80\%$ , a realistic GPU workload ratio is  $\frac{54}{11+54} \approx 83\%$ .



Experiments: The best GPU workload ratio is indeed around 83%. However, if the CPU workload fits in L3, then the balance of workload division will be very different!


# CPU+GPU unstructured tetrahedral mesh computation





(a) Stampede (b) Wilkes

J. Langguth, M. Sourouri, G.T. Lines, S.B. Baden, X. Cai. *Scalable heterogeneous CPU-GPU computations for unstructured tetrahedral meshes*. IEEE Micro 35, pp. 6–15, 2015

# Snapshot of entire-heart monodomain simulation



# 3D mesh close-ups



- Realistic simulations of cardiac electrophysiology require multiscale mathematical models and high computational resolutions
- Supercomputers are essential hardware resources
- Modern/future supercomputers will have increasingly advanced features (aggressive SIMD vectorization, hardware accelerators, etc.)
  - Care is needed when choosing the numerical strategy
  - Elaborate parallel programming is required
  - New issues (e.g. heterogeneous partitioning) will need to be handled

“A peek into the future”



An Experimental Infrastructure for Exploration of Exascale Computing

## The eX<sup>3</sup> infrastructure will provide the Norwegian HPC community and collaborators with a foundation for exascale preparations



With funding from  
**The Research  
Council of Norway**

“The primary objective of the eX<sup>3</sup> project is to provide an experimental, heterogeneous computational cluster as a national infrastructure that will help Norwegian HPC researchers, HPC/Big Data users, data center management, and HPC industry prepare for the coming exascale era of computing.”

### Subgoals:

Access to bleeding-edge technology

Educate students, researchers and practitioners

Develop powerful programming models, best practices, and software tools

Provide objective measurements and analyses

**2018 – 2022**

By its nature, itself being the target for research, eX<sup>3</sup> is a different infrastructure

## The eX<sup>3</sup> consortium has partners representing the full HPC landscape

### HPC Researchers

**simula**



UNIVERSITY OF BERGEN



**NTNU**

Norwegian University of  
Science and Technology

### HPC Management

**UNINETT** **figma2**

### HPC Users

 **BigInsight**



**SIRIUS**



CENTER FOR CARDIOLOGICAL INNOVATION

### HPC Industry

  
**Dolphin**

**NUMA SCALE**


 **fabriscale**



**simula**

# eX<sup>3</sup> Phase 1 installation is almost done

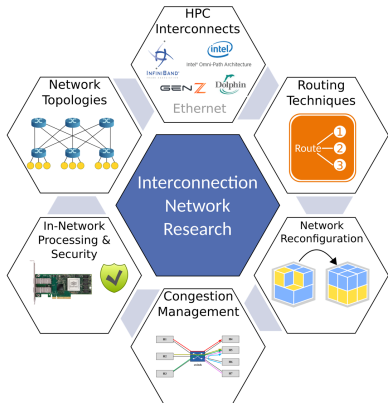
- CPU/GPU nodes with strong heterogeneity
  - Three types of top-bin CPUs: AMD Epyc 7601, Cavium ThunderX2 CN9980, Intel Scalable Platinum 8176M
  - DGX-2 system consisting of 16 NVIDIA V100 processors
- Interconnect: Mellanox HDR InfiniBand capable of 200 Gb/s
- Additional interconnect: Dolphin PCIe Gen3 x4 network
- High-performance memory and secondary storage



**The eX<sup>3</sup> infrastructure is heterogeneous on several levels, promoting new topics of research in scientific computing**

- **Hardware-aware heterogeneous work and data partitioning**
  - **Different compute nodes have different computing capabilities**
  - **Different data transfer speeds (intra-chip, intra-socket, intra-node, inter-node)**
  - **CPU hosts and hardware accelerators have different tasks & capabilities**
    - **Combination of numerical calculations and ML/DL**
  - **Co-existence of DRAM and NVMe**
    - **Enhanced possibilities for in-situ data analytics**
- **Co-existence of precision levels (double-, single-, half-, quarter-precision)**
  - **Mixed-precision calculations**

## Interconnection network research is full of challenges – eX<sup>3</sup> enables conducting experiments with disruptive technologies at a decent scale



- Multiple Interconnect solutions
  - Benchmarking
  - Interfacing
- Topologies and Routing
  - Full-scale HPC
  - Experimental enough to do reconfigurations
  - Fault-Tolerance
- Heterogeneous workloads
- Recent Advances
  - Smart-NICs
    - Security, Intrusion Detection
  - In-switch collectives

In summary, eX3 will be an extremely heterogeneous infrastructure bringing unprecedented research opportunities

