Towards detailed Organ-Scale Simulations in Cardiac Electrophysiology

Johannes Langguth (Simula research Laboratory)

joint work with

Namit Gaur, Hermenegild Arevalo, Chad Jarvis, Neringa Altanaite, Qiang Lan, and Xing Cai

Study Arrhythmia in the Human Heart via Electrophysiological Simulations

Simulations using highly detailed cell models can further the understanding of the processes that lead to cardiac arrythima and sudden cardiac death. Modeling of cardiac cells is an active field of research. Detailed cell models can have thousands of parameters and are thus challenging to compute at the organ scale.

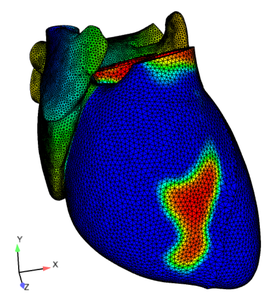

Tetrahedral mesh for whole heart simulation

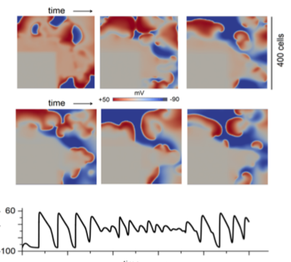

Simulation of diseased tissue

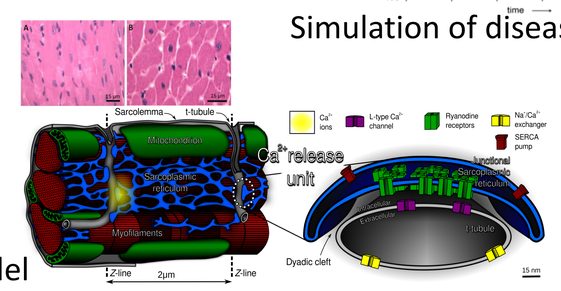

Detailed cell model

Stratifying the Risk of Sudden Cardiac Death to select Patients for ICD

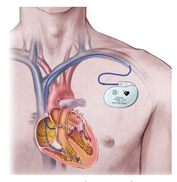

An **implantable cardioverter-defibrillator(ICD)** is vital for patients with high risk of sudden cardiac death. However, the procedure is expensive, invasive, and seriously impacts quality of life. Therefore, selecting only high-risk patients is vital.

Current clinical procedures for risk stratification are also invasive, expensive, and produce unsatisfactory results. In-silico simulations based on medical imaging have been shown to provide better prediction of patient outcomes.

Virtual Heart Model Creation

LGE MRI

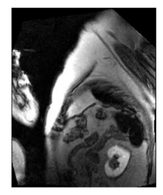

Ventricular Segmentation

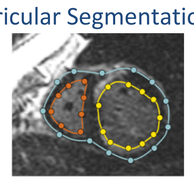

Infarct Segmentation

Normal Myocardium
Gray Zone (GZ)
Scar

Rule-Based Fiber Orientation

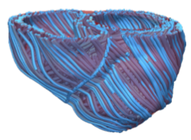

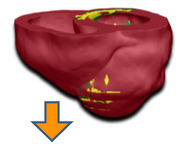

Virtual heart Arrhythmia Risk Predictor (VARP)

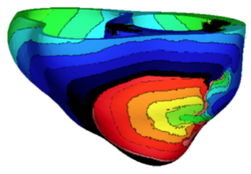

Non-Invasive
Safe
Effective

In-Silico Stimulation to induce Arrythmia

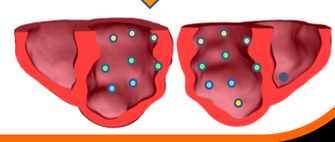

Exascale Problem

O'Hara-Rudy Model

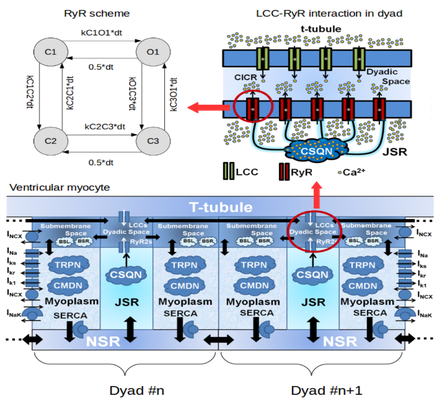

Each cell is modeled as a grid of calcium units, each having numerous ion channels that perform stochastic transitions. The models shows the emergence of regular heartbeats from stochastic cell behavior.

Organ scale simulation scope

- $2 * 10^9$ Cells in the heart
- 10^4 Dyads per cell
- 10^2 Ryanodine Receptors (RyRs) per dyad
- 10^4 Time steps per heartbeat

10^{19} possible state transitions

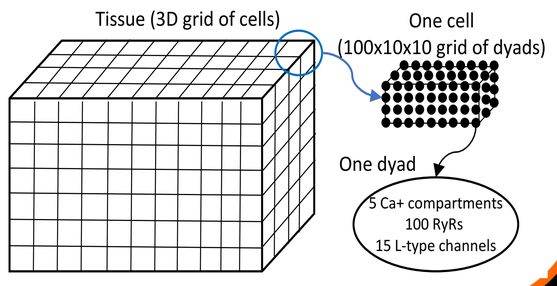

Implementation

Compute Kernels

1. Compute L-type opening probabilities
2. Simulate L-type opening
3. Compute RyR opening probabilities
4. Simulate RyR opening
5. Compute calcium concentrations
6. Dyad diffusion

Six compute kernels are used to simulate the change in calcium concentrations in each dyad. Kernels vary widely in Computational requirements.

Vectorized binomial sampling

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

State transitions are simulated efficiently via sampling from a binomial distribution. An optimized implementation is the key to high performance here.

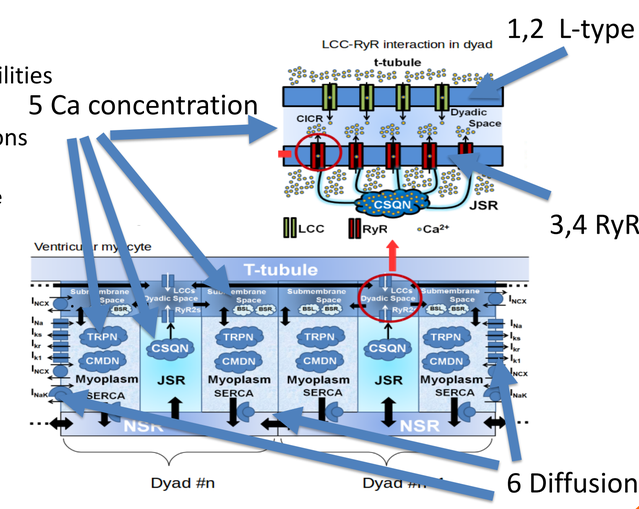

Xeon Phi: KNL shows massive improvements over KNC

Tianhe-2 Simulation runs

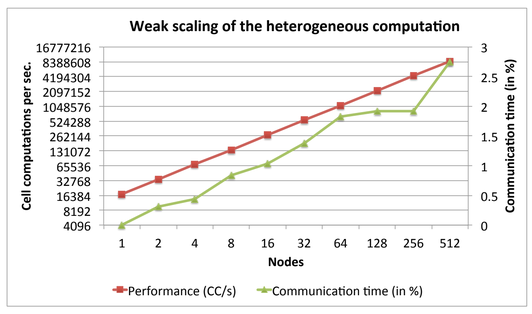

Due to the heavy computation, communication has almost no impact on performance. Load balancing under memory constraints is the more pressing issue.

4 Levels of Parallelism

- SIMD among dyads
- OpenMP among cells
- Device level subdomains
- Node level cuboid subdomains

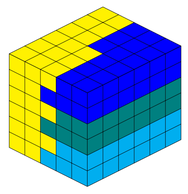

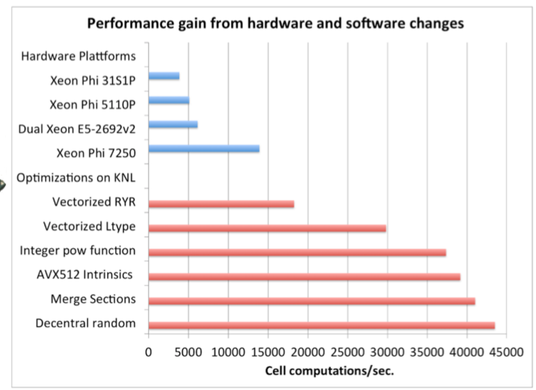

GPU Results

The GPU implementation provides higher performance than x86 – based computations. The bandwidth analysis shows that the higher memory bandwidth of the GPUs is crucial for performance. Correct mapping of the calcium units to the GPU threads was necessary for GPU performance.

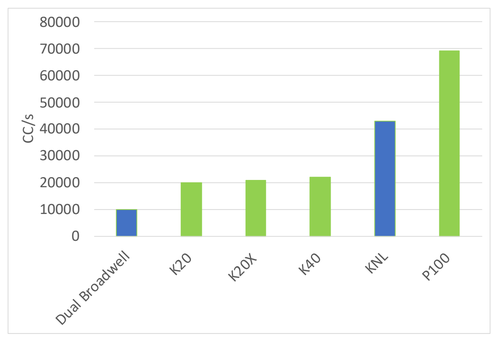

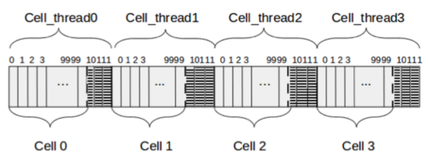

Dyad to thread assignment

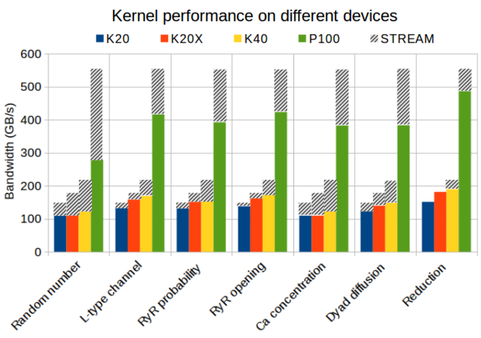

Upcoming Work

LYNX: Strong Scaling for Clinical Arrhythmia Risk Predictions

The goal of the LYNX project is to accelerate the computations for the Virtual heart Arrhythmia Risk Predictor (VARP) in clinical use and make them scalable on GPU based systems. Doing so requires strong scaling of the computation, with the ultimate objective being real-time simulations. Naturally, this significantly increases communication pressure, especially on modern heterogeneous systems with multi-GPU nodes. Furthermore, the irregular mesh complicates communication routines. Significant time was spent on maximizing computation-communication overlap while avoiding bandwidth contention.

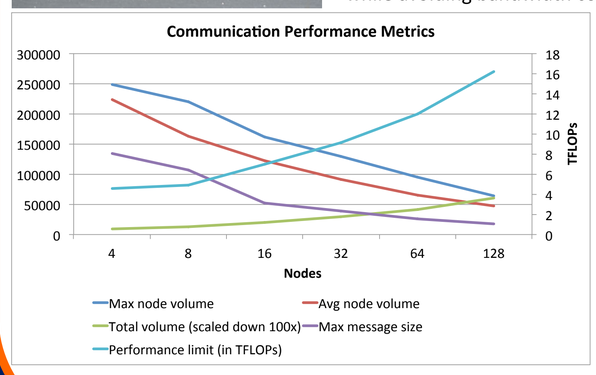

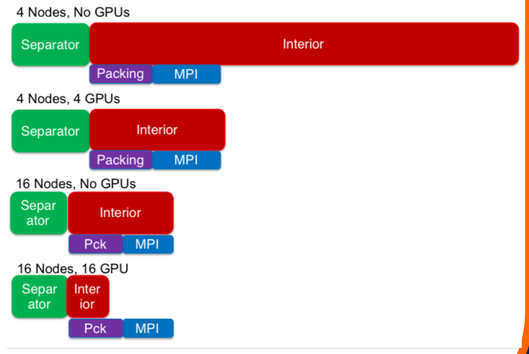

Are Detailed Organ-Scale Simulations Feasible ?

ORNL Summit

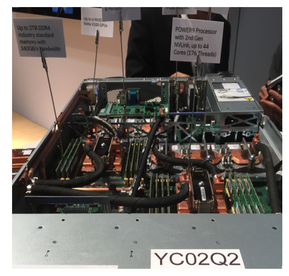

- Performance per GPU = 100.000 CC/s
- Performance per node = 600.000+ CC/s
- System Performance with perfect scaling = 2.750.000.000 CC/s
- Time to simulate heartbeat: < **3 hours**

- Memory for 2 billion cells: 6 PB
- Total GPU memory: 441 TB
- Total CPU memory: 2.355 PB
- Total NVRAM: 7.360 PB

Memory is the limiting factor for GPU computations. If swapping data into and out of the GPU is feasible, simulating the entire heart at this level of detail would be feasible on Summit already. However, even more detailed cell models exist. Their simulation at organ scale would require an exascale system.