


Physiologically realistic study of subcellular calcium dynamics with nanometer resolution



Kengo Nakajima
(The University of Tokyo, Japan)

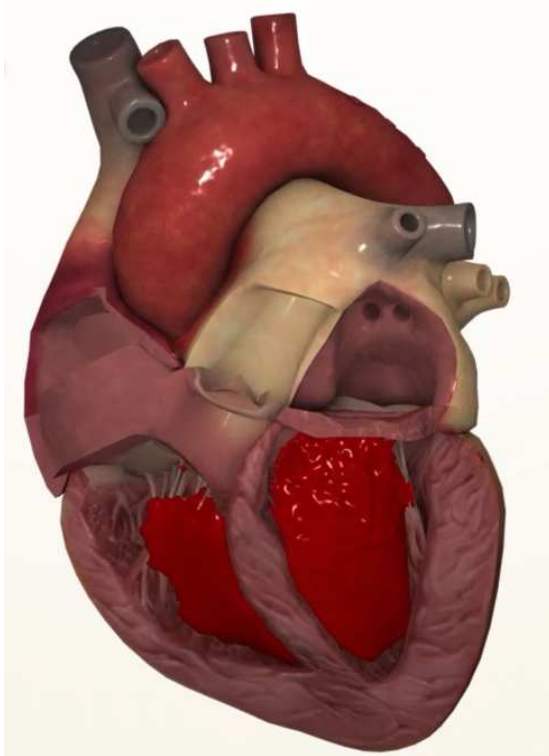
Xing Cai, Glenn Terje Lines
(Simula Research Laboratory, Norway)

C. Jarvis, J. van den Brink, J. Langguth (Simula Research Laboratory, Norway)


A. Ida, T. Hanawa, T. Hoshino, M. Matsumoto (The University of Tokyo, Japan)

M. Kawai (The University of Tokyo, Japan)

Background




- Synchronous and stable **calcium** handling is vital for normal **heart contraction**
- There are about 10,000 calcium release units in each heart cell, with structural details down to the **nanometer scale**
- Computer simulations of **subcellular calcium dynamics**, by solving elaborately coupled differential equations, require very high resolutions and thus huge computations



A schematic overview of a sarcomere (1/50 of a cardiac cell) and a calcium release unit

Governing equations & numerical scheme



- The entire solution domain, a 3D cube, represents a segment of a cardiac cell.
- It consists of several types of irregularly-shaped and intricately-connected **physiological domains**.
- Multiple calcium species are studied, the concentration of each species $u_s(x, y, z, t)$ is modeled by a **3D reaction-diffusion equation**:


$$\frac{\partial u_s}{\partial t}(x, y, z, t) = \nabla \cdot (\sigma_s(x, y, z) \nabla u_s(x, y, z, t)) + \sum_{i \neq s} R_{s,i}(u_s, u_i)$$

$$R_{s,i}(u_s, u_i) = k_{\text{on}}^{s,i} u_s (B_{\text{tot}}^i - u_i) - k_{\text{off}}^{s,i} u_i,$$

- The 3D cube is discretized by a uniform mesh of box-shaped computational voxels.
- The irregular physiological domains are resolved with the voxels.
- Explicit time integration plus 2nd-order finite differencing in space is adopted.
- Operator splitting is used to separate diffusion and reaction computations.
 - Different species can also use different interior time steps for diffusion comp.
- During every unified “outer” time step, all the increments due to diffusion and reaction are accumulated, which are added to the previous-outer-step solution at the end.

Project overview

- We aim to enable subcellular calcium dynamics simulations with **physiological realism**
 - using data of 3D super-resolution microscopy
 - **simulating a large number of calcium release units together (not done previously)**
- We need
 - hardware-compatible **optimizations** of an **old 3D simulator** of subcellular calcium dynamics
 - **calibration and validation** of the mathematical model and the parameters
- Importance:
 - Building a physiologically accurate description of healthy and pathological calcium releases
 - New knowledge about coding multiple inter-tangled stencil computations for the Knights Landing architecture (**Oakforest-PACS**)



FY2019 Research Activity 1


- Further optimization of the parallel simulator of subcellular calcium handling, with the following studied topics:
 - Reducing OpenMP overhead
 - Appropriate choice of OpenMP threads per MPI process
 - Improving overlap of MPI communication with computation
 - Tests of the simulator also on Xeon Skylake processors

Table 1: Time measurements (obtained on a 4-socket Skylake server) of five MPI implementations of the subcellular simulator; computational mesh: $672 \times 672 \times 168$, time steps: 1000

Version	CA_auto		LUT_auto		CA_man		LUT_man1		LUT_man2	
MPI procs	T_R	T_D	T_R	T_D	T_R	T_D	T_R	T_D	T_R	T_D
$2 \times 2 \times 2 = 8$	41.0	164.6	40.7	166.3	17.7	131.0	17.7	85.2	17.8	63.2
$4 \times 2 \times 2 = 16$	20.5	83.8	20.4	84.6	10.0	67.9	10.1	43.9	10.1	33.5
$4 \times 4 \times 2 = 32$	10.4	49.5	10.4	45.2	7.5	45.2	7.5	27.9	7.5	23.2
$4 \times 4 \times 4 = 64$	6.1	51.3	6.2	38.8	5.8	50.9	6.1	36.8	6.1	28.1
$8 \times 4 \times 4 = 128$	6.4	52.5	6.8	37.8	6.2	53.5	6.7	35.6	6.6	32.4

FY2019 Research Activity 2

More simulations involving multiple calcium release units (CRUs)



Some snapshots (2D cross sections) of a 3D simulation of multiple CRUs

FY2019 Research Activity 3

To prepare for implicit time integration (allowing larger time steps), effort was under way to investigate an implicit time-marching scheme

- A large-scale linear system needs to be solved per time step
- Preconditioned Krylov iterative solvers must be used
- Scalability of standard Krylov subspace methods suffers from the costly global synchronization steps (dot-products)
- Global-synchronization-free variants are thus preferred

```
1:  $r_0 := b - Ax_0$ ;  $u_0 := M^{-1}r_0$ ;  $p_0 := u_0$ 
2: for  $i=0 \dots$  do
3:    $s := Ap_i$ 
4:    $\alpha := (r_i, u_i) / (s, p_i)$ 
5:    $x_{i+1} := x_i + \alpha p_i$ 
6:    $r_{i+1} := r_i - \alpha s$ 
7:    $u_{i+1} := M^{-1}r_{i+1}$ 
8:    $\beta := (r_{i+1}, u_{i+1}) / (r_i, u_i)$ 
9:    $p_{i+1} := u_{i+1} + \beta p_i$ 
10: end do
```

Example: Results of dot-products are used just after they are calculated, as shown in lines 4-5 (a), and lines 8-9 (b) in the original Conjugate Gradient (CG) method. These dot-products are implemented by calls to MPI_Allreduce. The performance of the original CG method is thus significantly affected by the associated communication overhead.

FY2019 Research Activity 3 (cont'd)

Gropp's CG [Ghysels 2014]

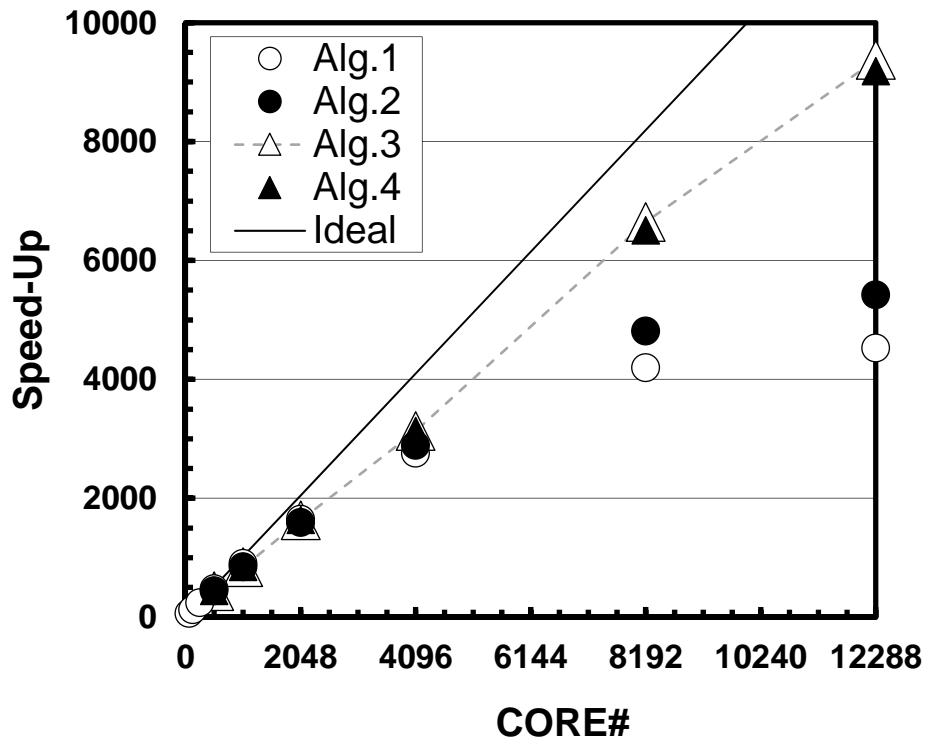
```

1:  $r_0 := b - Ax_0$ ;  $u_0 := M^{-1}r_0$ ;  $p_0 := u_0$ ;  $s_0 := Ap_0$ ;
 $\gamma_0 := (r_0, u_0)$ 
2: for  $i=0 \dots$  do
3:    $\delta := (p_i, s_i)$ 
4:    $q_i := M^{-1}s_i$ 
5:    $\alpha_i := \gamma_i / \delta$ 
6:    $x_{i+1} := x_i + \alpha_i p_i$ 
7:    $r_{i+1} := r_i - \alpha_i s_i$ 
8:    $u_{i+1} := u_i - \alpha_i q_i$ 
9:    $\gamma_{i+1} := (r_{i+1}, u_{i+1})$ 
10:   $w_{i+1} := Au_{i+1}$ 
11:   $\beta_{i+1} := \gamma_{i+1} / \gamma_i$ 
12:   $p_{i+1} := u_{i+1} + \beta_{i+1} p_i$ 
13:   $s_{i+1} := w_{i+1} + \beta_{i+1} s_i$ 
14: end for
    
```

Pipelined CG [Ghysels 2014]

```

1:  $r_0 := b - Ax_0$ ;  $u_0 := M^{-1}r_0$ ;  $w_0 := Au_0$ 
2: for  $i=0 \dots$  do
3:    $\gamma_i := (r_i, u_i)$ 
4:    $\delta := (w_i, u_i)$ 
5:    $m_i := M^{-1}w_i$ 
6:    $n_i := Am_i$ 
7:   if  $i > 0$  then
8:      $\beta_i := \gamma_i / \gamma_{i-1}$ ;  $\alpha_i := \gamma_i / (\delta - \beta_i \gamma_i / \alpha_{i-1})$ 
9:   else
10:     $\beta_i := 0$ ;  $\alpha_i := \gamma_i / \delta$ 
11:   end if
12:    $z_i := n_i + \beta_i z_{i-1}$ 
13:    $q_i := m_i + \beta_i q_{i-1}$ 
14:    $s_i := w_i + \beta_i s_{i-1}$ 
15:    $p_i := u_i + \beta_i p_{i-1}$ 
16:    $x_{i+1} := x_i + \alpha_i p_i$ 
17:    $r_{i+1} := r_i - \alpha_i s_i$ 
18:    $u_{i+1} := u_i - \alpha_i q_i$ 
19:    $w_{i+1} := w_i - \alpha_i z_i$ 
20: end for
    
```




Results for Strong Scaling on Reedbush-U System using up to 384 nodes (12,288 cores), Alg.1: Original CG, Alg.2: Chronopoulos/Gear, Alg.3: Pipelined CG, Alg.4: Gropp's CG [Hanawa 2016] These results in were obtained in July 2016 using Intel MPI Library 2015 or 2016.

FY2019 Research Activity 3 (cont'd)

In Intel MPI Library 2019, performance of Asynchronous Progress Thread has been significantly improved, where efficient asynchronous collective communications are supported by software. Therefore, Communication-Computation Overlapping will be improved by Asynchronous Progress.

During FY.2019, we have re-evaluated performance of pipelined CG methods using up to 4,096 nodes of OFP system with the improved capability of Asynchronous Progress Thread



Results for Strong Scaling on Oakforest-PACS (OFP) System using up to 4,096 nodes (262,144 cores), Alg.1-AR: Original CG, Alg.3-AR/3-IAR: Pipelined CG, Alg.4-AR/4-IAR: Gropp's CG, AR: MPI_Allreduce, IAR: MPI_lallreduce [5]

Node #	Core #	Alg.3-IAR	Alg.4-IAR
128	8,192	1.03	1.06
256	16,384	1.13	1.21
512	32,768	1.05	1.20
1,024	65,536	1.06	1.25
2,048	131,072	1.08	1.34
4,096	262,144	1.19	1.38

Results for Strong Scaling on Oakforest-PACS (OFP) System using up to 4,096 nodes (262,144 cores), Speed-up of Alg.3-IAR and Alg.4-IAR against Alg.1-AR at each core/node number

Publications/Presentations in FY2019

- **Proceedings of international conferences (Refereed)**

Chad Jarvis, Glenn Terje Lines, Johannes Langguth, Kengo Nakajima, Xing Cai. Combining algorithmic rethinking and AVX-512 intrinsics for efficient simulation of subcellular calcium signaling. Proceedings of ICCS 2019 Conference, 2019

Johannes Langguth, Hermenegild Arevalo, Kristian Hustad, Xing Cai. *Towards detailed real-time simulations of cardiac arrhythmia*. Proceedings of the International Conference in Computing in Cardiology, 2019.

- **1 international conference paper (Non-refereed)**

Kengo Nakajima. Parallel Multigrid with Adaptive Multilevel hCGA on Manycore Clusters, Extreme-Scale/Exascale Applications China, Japan, World, ISC High Performance 2019, Frankfurt, Germany, 2019 (Invited Talk)

- **2 presentations at domestic conference (Non-refereed)**