31st Annual International Conference of the IEEE EMBS
Minneapolis, Minnesota, USA, September 2-6, 2009

Strong scaling and speedup to 16,384 processors in cardiac electro – mechanical simulations

Matthias Reumann, *Member, IEEE,* Blake G. Fitch, *Member, IEEE*, Aleksandr Rayshubskiy, David U. J. Keller, *Member IEEE*, Gunnar Seemann, Olaf Dössel, *Member, IEEE,* Michael C. Pitman, John J. Rice

***Abstract*— High performance computing is required to make feasible simulations of whole organ models of the heart with biophysically detailed cellular models in a clinical setting. Increasing model detail by simulating electrophysiology and mechanical models increases computation demands. We present scaling results of an electro – mechanical cardiac model of two ventricles and compare them to our previously published results using an electrophysiological model only. The anatomical data-set was given by both ventricles of the Visible Female data-set in a 0.2 mm resolution. Fiber orientation was included. Data decomposition for the distribution onto the distributed memory system was carried out by orthogonal recursive bisection. Load weight ratios for non – tissue vs. tissue elements used in the data decomposition were 1:1, 1:2, 1:5, 1:10, 1:25, 1:38.85, 1:50 and 1:100. The ten Tusscher et al. (2004) electrophysiological cell model was used and the Rice et al. (1999) model for the computation of the calcium transient dependent force. Scaling results for 512, 1024, 2048, 4096, 8192 and 16,384 processors were obtained for 1 ms simulation time. The simulations were carried out on an IBM Blue Gene/L supercomputer. The results show linear scaling from 512 to 16,384 processors with speedup factors between 1.82 and 2.14 between partitions. The most optimal load ratio was 1:25 for on all partitions. However, a shift towards load ratios with higher weight for the tissue elements can be recognized as can be expected when adding computational complexity to the model while keeping the same communication setup. This work demonstrates that it is potentially possible to run simulations of 0.5 s using the presented electro-mechanical cardiac model within 1.5 hours.**

***Index Terms*— Multi-physical heart models, computational biology, parallel supercomputer, orthogonal recursive bisection**

I. Introduction

Full scale simulations of the electrophysiology of the heart required the computation power of high performance computing (HPC) systems [1 - 3]. The electrophysiological model needs to be augmented by mechanical, contraction deformation models and models of blood flow to model the heart in its full complexity for personalized medicine [4]. Increasing model complexity will increase the computational demands. Also, the difference in computational load for tissue vs. non – tissue elements will become even more significant which highlights the need for optimal load balancing measures.

We have developed the orthogonal recursive bisection algorithm for cardiac simulations previously and investigated its effectiveness with respect to different cardiac cell models [5] and load ratios [1] in simulations of cardiac electrophysiology on the IBM Blue Gene/L supercomputer [6].

In this article we increase model complexity by simulating a two ventricular electro–mechanical model. We will present load balancing and scaling results to 16,384 processors and compare the results with our previous simulations.

II. Methods

A. Cardiac model

We chose the ten Tusscher et al. (TT) model [7] as a representation of cellular electrophysiology. The calcium transient of the TT model was used as input to the Rice et al. model [8] to compute the force generated per cell. A limitation of our electro – mechanical model is that we did not include electro – mechanical feedback. However, since we wanted to investigate computation performance, the electro – mechanical feedback plays a minor role and can be disregarded at this stage of investigation.

The diffusion terms to model electrophysiological excitation can be determined by the mono- or bidomain equations [9]. In this study we have used the monodomain equation [1] to have only one communication cycle per time step. The monodomain equation is given by

$$\frac{\lambda}{1-\lambda}\nabla(M_i\nabla v) = \chi C_m \frac{\partial v}{\partial t} + \chi I_{ion} \qquad (1)$$

with the constant scalar λ that defines an equal anisotropy ratio between the extracellular and intracellular conductivity matricies $M_e = \lambda M_i$, v is the transmembrane voltage, χ is the membrane area to volume ratio, C_m the membrane capacitance and I_{ion} the ionic current. For boundary conditions the monodomain model yields

Manuscript received April 15, 2008.

Matthias Reumann is with the Computational Biology Center, IBM TJ Watson Research Center, Yorktown Heights, 1101 Kitchawan Road, Route 134, NY 10598, USA (corresponding author; phone: +49-177-6727731; fax: +1-914-945-4217; e-mail: mreumann@ieee.org).

Gunnar Seemann (e-mail: Gunnar.Seemann@ibt.uni-karlsruhe.de), David Keller (David.Keller@ibt.uni-karlsruhe.de), and Olaf Doessel (e-mail: Olaf.Doessel@ibt.uni-karlsruhe.de) are with the Institute of Biomedical Engineering, KIT, Kaiserstrasse 12, 76131 Karlsruhe, Germany.

Blake Fitch (e-mail: bgf@us.ibm.com), Aleksandr Rayshubskiy (e-mail: arayshu@us.ibm.com), Mike Pitman (e-mail: pitman@us.ibm.com) and Jeremy Rice (e-mail: johnrice@us.ibm.com) are with the Computational Biology Center, IBM TJ Watson Research Center, Yorktown Heights, 1101 Kitchawan Road, Route 134, NY 10598, USA.

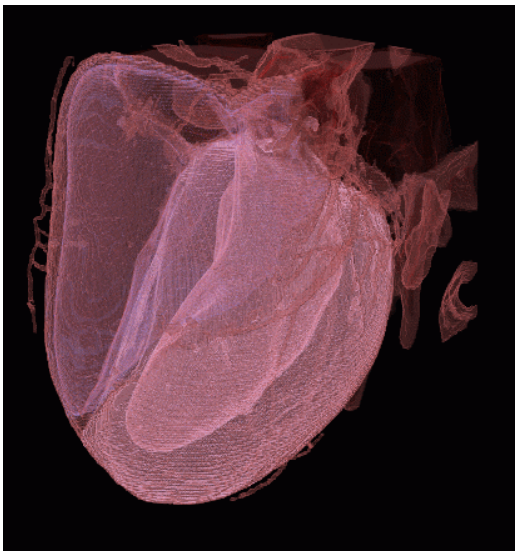

Fig. 1. Visible female data – set of the National Library of Medicine, Bethesda, Maryland, USA, showing the ventricles as used in our scaling simulations. Also visible are parts other cardiac structures like the atria and the cardio – vascular system that has been set to bath medium in the simulations and thus did not contribute to the computational complexity.

$$n\nabla(M_i\nabla v) = 0 \qquad (2)$$

with the outward unit normal vector n of the hearts boundary. The model includes fiber orientation account for heterogeneous anisotropic conduction.

The monodomain reaction – diffusion system was implemented using an operator splitting approach and the finite difference method [3, 9, 10] based on the regular Cartesian grid of the anatomical model.

B. Data decomposition and communication framework

The Laplacian operator in equations 1 – 2 require only neighboring elements for the computation of the diffusion term. Thus, we can decompose the data set into subvolumes that are distributed one to each processor. Adding a ghost layer around the subvolume where the values of the adjacent elements are stored enables the computation of the diffusion term in a distributed memory system. In this study we use the orthogonal recursive bisection (ORB) algorithm first introduced in N-body problems in molecular dynamics simulations [11, 12] adjusted to our cardiac model. A detailed description of the ORB algorithm used and the corresponding communication setup can be found in [1]. Here we will only outline the underlying idea..

A list of x – y – z co-ordinates and the respective computational load of the associated element is created based on the anatomical data of the Visible Female ventricles (National Library of Medicine, Bethesda, Maryland, USA) with resolution of 0.2 mm cubic elements (Fig. 1). The model consists of over 128.9 million elements in an array of 537 by 492 by 488 with approximately 32.5 million active ventricular elements. A number of lists for the ORB input was created to account for different load ratios of non – tissue vs tissue elements. Tissue elements here refer to elements in the data set that are associated with ventricular tissue for which the TT and Rice et al. model was computed. The cell models were not computed for non – tissue elements. To compare the results of this study with our previous simulations [1] we decided to test load ratio 1:1, 1:2, 1:5, 1:10, 1:25, 1:38.85, 1:50 and 1:100.

The volume is first cut into two halves along the x axis by first establishing the computational load of the volume then counting the load along the x axis. The x value that determines half of the total load gives the x value for the decomposition of the volume. The created subvolumes are cut similarly along y and z axis in a recursive fashion to create a binary tree representation. The so called ORB tree has n levels with 2^n subvolumes. The ORB output is a list of the lowest level subvolumes as well as the physical location of the computational node to which the subvolume of data will be distributed to. Since the Blue Gene supercomputer is organized into 2^n partitions, the ORB tree can be conveniently be mapped onto the physical system. The ORB algorithm maps the anatomical subvolumes in x – y – z co-ordinates to the N_x – N_y – N_z processor partition.

The subvolumes are stored in a list of bounding boxes associated with a processor rank. By comparing the bounding box borders of the subvolumes, a communication list is created for each processor so that it knows its communication partners and the respective overlap in data to be communicated [1, 5]. The communication framework is implemented using standard non-blocking Message Passing Interface (MPI) functions. At each communication phase, each processor goes through the list of communication partners, sends and received the data for the transmembrane voltage and stores it in the ghost layer.

C. Scaling analysis

Strong scaling results were obtained for 1 ms simulation time, i. e. 100 time steps, on 512, 1024, 2048, 4096, 8192 and 16,384 processors. A stimulus was set in the left ventricular apex region for the length of the simulation. The simulations were carried on the distributed memory IBM Blue Gene/L supercomputer [6].

We measured the computation and communication time separately as well as the total run time. We define a load balance value by determining the average computation time over all processors and time steps divided by the maximum computation time over all processors and time steps. Similarly we compute the load balance for the communication phase.

The speedup factor is typically defined as

$$Speedup\ factor = t_{N1}/t_{Nx} \qquad (3)$$

with t_{N1} being the fastest run time on one processor and t_{Nx} the fastest run time on the partition with $N = x$ processors. The smallest processor partition used in this study was 512. Theoretically, the speedup with respect to one processor is 512 times for a partition with 512 processors. We adjust this value by multiplying it with the load balance value for the fastest computation phase over all load ratios on partition N = 512 and take it as reference for the speedup to 16,384 processors.

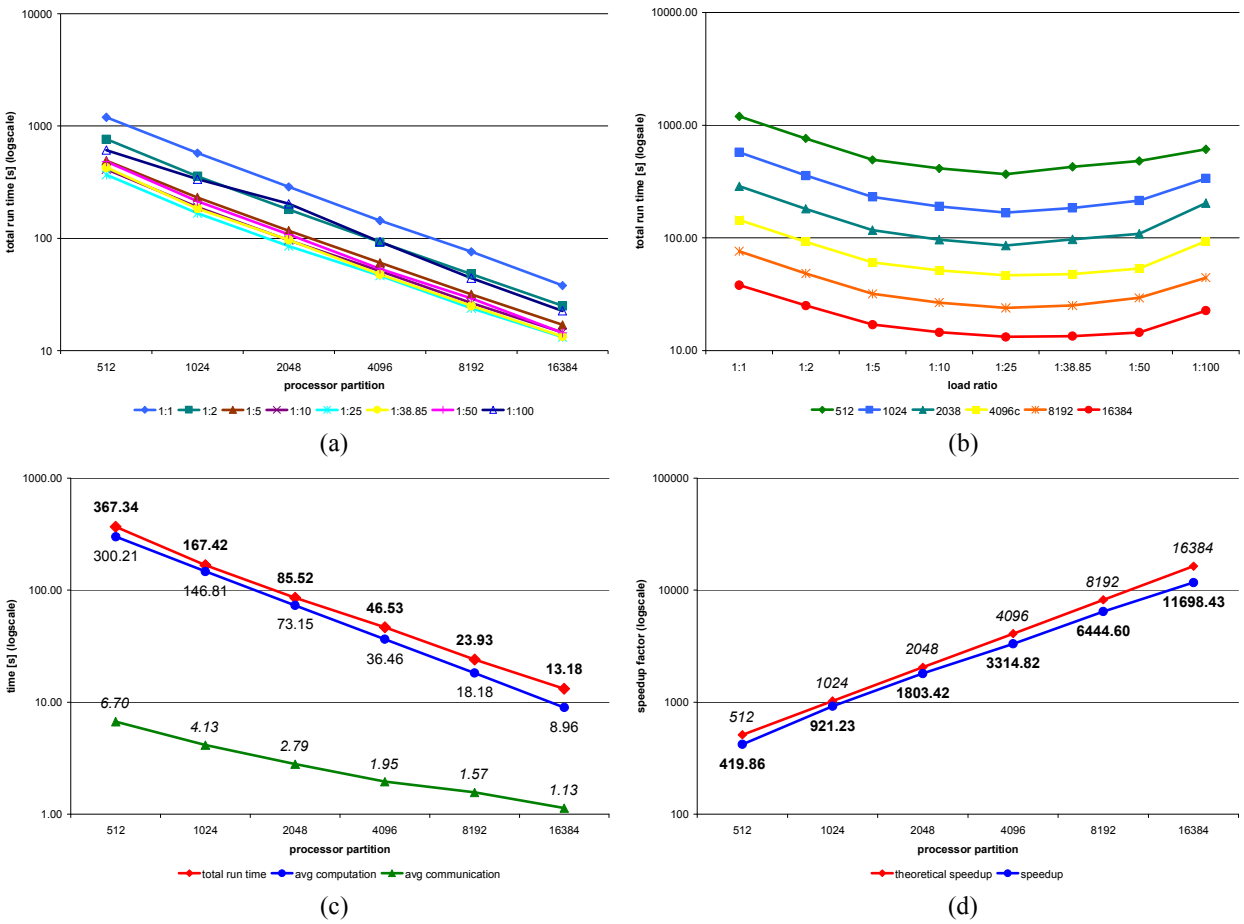

Fig. 2. Scaling results for different number of processors (N). (a) The graph shows the computation time on all partitions for all load ratios. Load ratio 1:25 achieved the best results on all partitions. (b) The total run time versus load ratio on all processor partitions is shown. This figure illustrates that very similar speedup factors are achieved for all load ratios. (c) Displayed are the average communication time (green triangles) and computation time (blue circles) as well as the total run time (red diamonds) for partitions N = 512 to N = 16,384. The data points are joined by linear interpolation for visualization purposes. Linear scaling in both communication and computation can be recognized. (d) Shown is the adjusted speedup factors for the simulations with respect to the theoretical speedup

.

III. Results

A. Data decomposition and communication

The results for the data decomposition in this study are the same as for our previously published simulations [1] where only an electrophysiological model was computed because the same load values were used for the same anatomical data set. Thus, the results can be summarized by highlighting that small load ratios like 1:1 will yield subvolumes of equal or close to equal size but with large differences in the number of tissue elements per subvolume that determine the computational load. Increasing the load for tissue elements leads to a greater distribution of subvolume sizes and tissue elements per subvolume are more evenly spread.

When increasing the processor partition by factor 2, each subvolume in the smaller partition is divided by factor 2. Thus, the subvolume sizes are half as large including on average half as many tissue elements. The results confirm this behavior. The distribution of tissue elements vs. subvolume sizes keeps its shape for all load ratios but the scale of the distribution is reduced by about factor 2 when increasing the number of processors by 2.

B. Scaling results and speedup

Figure 2 and table 1 summarize the results of the strong scaling simulations. The simulations scale linearly for all load ratios (fig. 2.a). The load ratio with the fastest run times is 1:25 for all partitions (fig. 2.b). The fastest run times (fig. 2.c) are 367.34 s, 167.42 s, 85.52 s, 46.53s, 23.93 s and 13.18 s on 512, 1024, 2048, 4096, 8192 and 16,384 processors respectively. However, while the run times for load ratio 1:38.85 are over 10 % slower on partition 512 – 2048, they are only 2.7, 4.7 and 1.6 % slower for partitions with 4096, 8192 and 16,384 processors, respectively.

The adjusted speedup factor is below the theoretical speedup (fig. 2.d). However, the speedup between a processor partition and the next larger partition is 2.19, 1.96, 1.84, 1.94 and 1.82. The simulation times for the average

TABLE I
LOAD BALANCE AND SPEEDUP FACTORS FOR THE ELECTRO-MECHANICAL SIMULATIONS

Partition size	N_x	N_y	N_z	Total run time [s]	Load balance computation	Load balance communication	Speedup	Speedup reference N = 512	Speedup between partitions
512	8	8	8	367.34	0.8200	0.1488	419.86	1.00	1.00
1024	8	8	16	167.43	0.7850	0.1307	921.22	2.19	2.19
2048	8	16	16	85.52	0.7382	0.1422	1803.42	4.30	1.96
4096	16	16	16	46.53	0.6503	0.1270	3314.82	7.90	1.84
8192	16	32	16	23.93	0.5664	0.1321	6444.60	15.35	1.94
16,384	32	32	16	13.18	0.4391	0.1070	11698.43	27.86	1.82

computation phase are close to the total run time. The computation phase scales better than the overall run time and the communication time as shown in fig. 2.c. Disregarding the communication overhead the speedup achieved is even higher and just above the theoretical value of 2. While the load balance value for the computation is over 0.8 on 512 processors, it deteriorates to 0.44 on 16,384 processors (tab. 1). The load balance value for the communication phase does not reach above 0.15 for all partitions.

IV. DISCUSSION

Given the results of the electro–mechanical simulations the first impression is that the difference in simulation times compared with the simulation times of just the electrophysiological model is small and hardly noticeable. The same load ratio 1:25 yields the fastest simulations times. However, a closer look confirms that the computation load is indeed increased. The total run time is increased by 12, 9, 7, 11, 14 and 18 % for partitions 512 to 16,384, respectively, compared with the corresponding simulations of the electrophysiological model. Also, the speedup factors are below slightly smaller. While the performance on 512 processors is similar comparing electrophysiological and electro – mechanical simulations, the performance on larger processor partitions is less good for the electro – mechanical model. Performance and thus scaling deteriorates the larger the processor partition. However, an interesting result can be recognized in figure 2.b. The simulation times for load ratio 1:38.85 are very close and hardly distinguishable to those of load ratio 1:25. They are even below the run times load ratio 1:10 for the three larger partitions. Previously in the electrophysiological simulations, load ratio 1:10 was closer to the fastest simulations times. This clearly indicates a shift of computational load towards higher load ratios when using more complex cardiac models. This clearly needs to be taken into account when optimizing the data decomposition strategy in the future.

V. CONCLUSIONS

This work confirms the hypothesis that more complex cardiac models require a higher computational load that needs to be taken into account in load balancing and data decomposition. However, our simulation results also indicate that a simulation of 0.5 s using a detailed electro - mechanical model on a detailed anatomical model could be carried out in less than two hours.

ACKNOWLEDGEMENT

We thank Fred Mintzer who has generously provided us the time on the IBM T. J. Watson Blue Gene/L system. In addition we would like to thank Dave Singer for his hardware support. Finally we would like to thank Robert Germain who has provided us with the tools to analyze the time trace data for all simulation runs

REFERENCES

[1] M. Reumann, B. G. Fitch, M. C. Pitman, J. J. Rice. Orthogonal Recursive Bisection as Data Decomposition Strategy for Cardiac Modeling in High Performance Computer Simulations on the IBM Blue Gene/L supercomputer. BMC Systems Biology 2009 (submitted)

[2] M. Potse, A. Vinet. Large-scale integrative modeling of the human heart. 22. International Symposium on High Performance Computing Systems and Applications (HPCS 2008). 2008. Quebec, Canada.

[3] J. Pitt-Francis, A. Garny, D. Gavaghan. Enabling computer models of the heart for high-performance computers and the grid. Philos Transact A Math Phys Eng Sci 2006;364:1501-16

[4] M. Reumann, V. Gurev, J. J. Rice. Cardiac models for personalized medicine. Personalized Medicine. 2009;6(1):45-66

[5] M. Reumann, B. G. Fitch, A. Rayshubskiy et. al. Large scale cardiac modeling on the Blue Gene supercomputer. Conf Proc IEEE Eng Med Biol Soc. 2008;2008;577-580

[6] Gara, A, MA Blumrich, D Chen et al. Overview of the Blue Gene/L system architecture. IBM J. Res. & Dev. 2005;49, 195-212.

[7] K. H. W. J. ten Tusscher, D. Noble, P. J. Noble and A. V. Panfilov. A model for human ventricular tissue. Am J Physiol. 2004;286:H1573-H1598.

[8] J. J. Rice, R. L. Winslow and W. C. Hunter. Comparison of putative cooperative mechanisms in cardiac muscle: Length dependence and synamic responses. Am. J. Physiol. 1999;276:H1734-H1754

[9] M. Potse, B. Dubé, J. Richer, A. Vinet and R. M. Gulrajani. A Comparison of Monodomain and Bidomain Reaction-Diffusion Models for Action Potential Propagation in the Human Heart. IEEE Trans Biomed Eng. 2006;53(12):2425-2435

[10] H. I. Saleheen and K. T. Ng. New Finite Difference Formulations for General Inhomogeneous Anisotropic Bioelectric Problems. IEEE Trans Biomed Eng. 1997;44(9):800-809R. A. FitzHugh. Impulses and physiological states in theoretical models of nerve membrane. Biophysical Journal. 1961;1:445

[11] L. Nyland, J. Prins, R.H. Yun, J. Hermans, H. -C. Kum, and L. Wang. Achieving scalable parallel molecular dynamics using dynamic spatial decomposition techniques. Journal of Parallel and Distributed Computing, 1997;47(2): 125-138

[12] B. G. Fitch, A. Rayshubskiy, M. Eleftheriou, T. J. C. Ward, M. Giampapa, M. C. Pitman and R. S. Germain. Blue Matter: Approching the limits of Concurrency for classical molecular dynamics. In Proc. Supercomputing 2006