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

Orthogonal recursive bisection data decomposition for high performance computing in cardiac model simulations: Dependence on anatomical geometry

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*—Orthogonal recursive bisection (ORB) algorithm can be used as data decomposition strategy to distribute a large data set of a cardiac model to a distributed memory supercomputer. It has been shown previously that good scaling results can be achieved using the ORB algorithm for data decomposition. However, the ORB algorithm depends on the distribution of computational load of each element in the data set. In this work we investigated the dependence of data decomposition and load balancing on different rotations of the anatomical data set to achieve optimization in load balancing. 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. The data set was rotated by 90 degrees around x, y and z axis, respectively. By either translating or by simply taking the magnitude of the resulting negative coordinates we were able to create 14 data sets of the same anatomy with different orientation and position in the overall volume. Computation load 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 to investigate the effect of different load ratios on the data decomposition. The ten Tusscher et al. (2004) electrophysiological cell model was used in monodomain simulations of 1 ms simulation time to compare performance using the different data sets and orientations. The simulations were carried out for load ratio 1:10, 1:25 and 1:38.85 on a 512 processor partition of the IBM Blue Gene/L supercomputer. The results show that the data decomposition does depend on the orientation and position of the anatomy in the global volume. The difference in total run time between the data sets is 10 s for a simulation time of 1 ms. This yields a difference of about 28 h for a simulation of 10 s simulation time. However, given larger processor partitions, the difference in run time decreases and becomes less significant. Depending on the processor partition size, future work will have to consider the orientation of the anatomy in the global volume for longer simulation runs.**

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

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 (email: 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.

I. Introduction

THE need for high performance computing (HPC) in the simulation of detailed whole organ cardiac models has become more pronounced in the past few years [1]. Previously, cardiac models were reported to run on HPC systems with over 750 processors [2] using a simple data decomposition strategy of dividing the data volume in equally sized cubes along the Cartesian coordinates with respect to the number of processors. We have shown that the more sophisticated orthogonal recursive bisection algorithm [3, 4] applicable to large scale cardiac simulations on up to 16,384 processors [5 – 7] on the parallel, distributed memory IBM Blue Gene/L supercomputer [8].

Load balancing and scaling has so far only been applied to a single biventricular anatomical data set. In this work we investigated how different rotations of the anatomical data set influence data decomposition as well as load balancing.

II. Methods

A. Cardiac model

We chose the ten Tusscher et al. (TT) model [9] as a representation of cellular electrophysiology. The diffusion terms were represented by the monodomain equation [5 – 7] so that we had only one communication cycle per time step. The monodomain equation is given by

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

with the constant scalar λ that defines an equal anisotropy ratio between the extracellular and intracellular conductivity matrices $M_e = \lambda M_i$, 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

$$n\nabla(M_i\nabla v) = 0 \qquad \text{(equation 2)}$$

with the outward unit normal vector n of the hearts boundary.

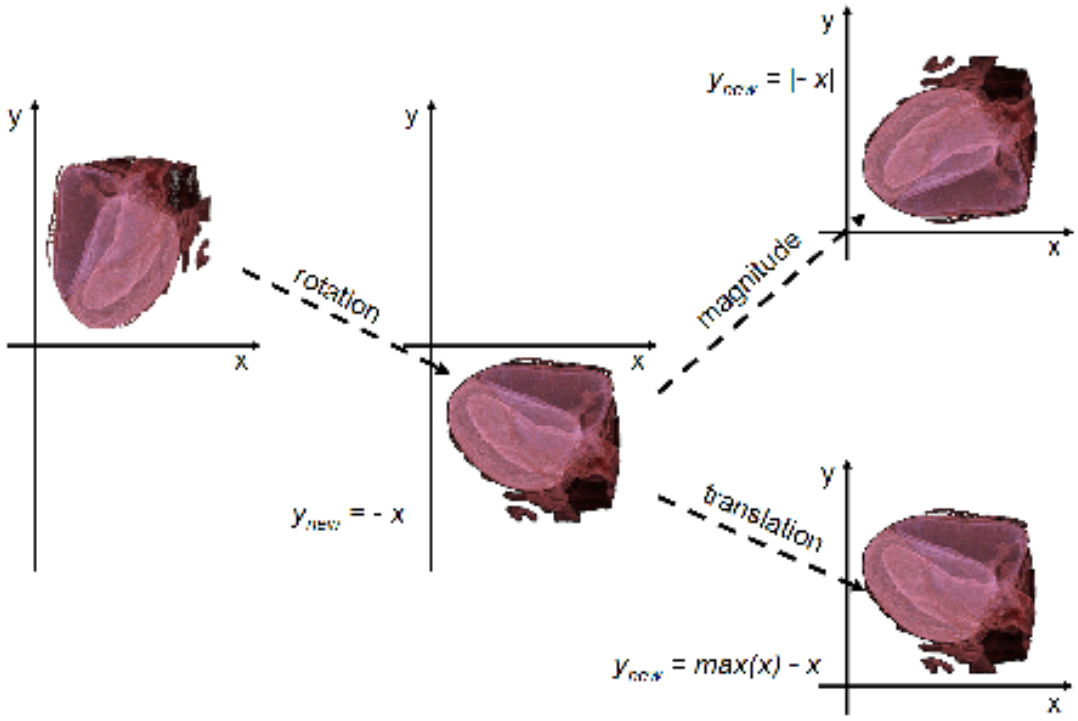

Fig. 1. The figure illustrates how the differently oriented anatomical data sets were created. A projection of the ventricles of the Visible Female data set, National Library of Medicine, Bethesda, Maryland, USA, is shown on the left in x – y coordinates. A rotation of the data set by 90 degrees around the z – axis (not shown in schematic) is achieved by simply swapping x and y coordinates of each point (middle). However, the new y value y_{new} is the negative x value due to rotation matrix. For the computation we only store positive y values. This can be achieved by taking the magnitude of x (top right), ie. flipping the data set across the x axis or by translating the data set along the y axis so that all elements of the data set are positive (bottom right).

The model includes fiber orientation to take heterogeneous anisotropic conduction into account. The monodomain reaction – diffusion system was implemented using an operator splitting approach and the finite difference method [5, 10 – 12] based on the regular Cartesian grid of the anatomical model.

B. Anatomical data sets

We created 14 anatomical data sets with different orientation of the Visible Female ventricles (National Library of Medicine, Bethesda, Maryland, USA). The original data set consists of over 128.9 M elements in an array of 537 by 492 by 488 voxels with approximately 32.5 M active ventricular elements. We mapped this data set to a 544 by 544 by 544 volume, ie. the total volume was increased to approximately 161 M voxels. On the one hand, we simply padded the original data set by keeping the original data coordinates and filling the remaining volume with bath medium (data set 1). On the other hand, we centered the original volume inside the new volume (data set 2) to have an equal amount of bath medium around the original data set.

We then rotated both data sets by $\alpha = 90$ degrees around the x, y and z axis, respectively. The rotation matrix for a rotation around the z axis is given by

$$\begin{pmatrix} x_{rot} \\ y_{rot} \\ z_{rot} \end{pmatrix} = \begin{pmatrix} \cos(\alpha) & -\sin(\alpha) & 0 \\ \sin(\alpha) & \cos(\alpha) & 0 \\ 0 & 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} x \\ y \\ z \end{pmatrix} \qquad \text{(equation 3)}$$

Thus, negative coordinates result from a rotation of $\alpha = 90$ degrees. To eliminate this, we took both the magnitude of the new coordinates (denoted by *MagRot* from hereon) or we translated the rotated data set along the respective axis by 544 (denoted by *TransRot* from hereon). Figure 1 illustrates the two new data sets created by the *MagRot* and *TransRot* operation for the rotation around the z axis. Thus, by applying these two operations on both data sets in the 544^3 volume yields six anatomical data sets with different orientation in the volume. We denote these data sets with *noRot*, *MagRotX*, *MagRotY*, *MagRotZ*, *TransRotX*, *TransRotY* and *TransRotZ* to indicate the operation performed on the original data set (*noRot*) and the axis around which the rotation was performed. For the data set which was centered, the same nomenclature applies with the difference that we included the term *Center* in the notation, e. g. *CenterMagRotX* and *CenterTransRotX*.

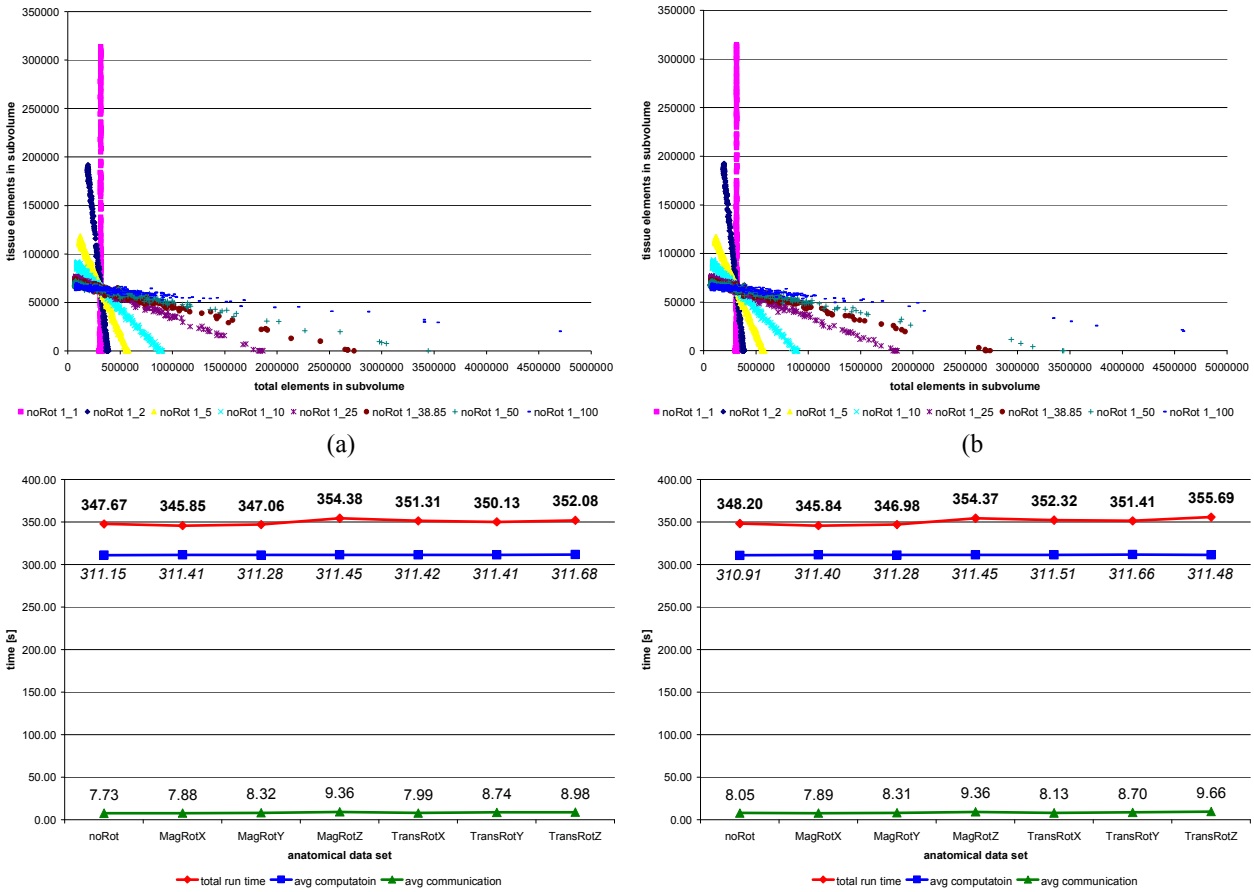

Fig. 2. This figure shows the data decomposition with respect to tissue elements vs. total elements in subvolume for data set 1 (a) and data set 2 (b) with the original anatomical orientation. Differences in distributions are hardly noticeable. However differences exist as can be seen especially for larger subvolume sizes when comparing data sets 1 and 2 and their respective rotations. The difference in run times is less than 10 s for data set 1 and the respective rotations (c) and just above 10 s for data set 2 and its rotations (d). The average communication and computation times hardly differ for any anatomical orientation in the data sets.

C. Data decomposition and communication

The Laplacian operator in equations 1 – 2 require only neighboring elements for the computation of the diffusion term, which enables us to decompose the data sets and distribute the subvolumes onto a distributed memory parallel supercomputer. A list of x, y and z coordinates and the respective computational load value is the input for the ORB. We use load ratios 1:1, 1:2, 1:5, 1:10, 1:25, 1:38.85, 1:50 and 1:100 for non – tissue vs. tissue element loads. Tissue elements are associated with a cell model represent ventricular tissue. Non – tissue elements do not represent cardiac tissue and no cell model is computed at these points. Since our previous work [5, 7] shows the best performance for load ratio 1:25 we use only 1:10, 1:25, and 1:38.85 for the simulations to cover the adjacent load ratios also. The ORB algorithm creates a binary tree with n levels and thus 2^n subvolumes. The anatomical subvolumes in x – y – z coordinates can be conveniently mapped onto the $N_x – N_y – N_z$ Blue Gene/L processor partition. Communication between adjacent subvolumes is carried out by standard non – blocking MPI functions where the communicated values at the boundaries of the subvolumes are stored in a ghost layer surrounding each subvolume. The data decomposition strategy is applied to all 14 data sets. The simulations are performed as described in [5, 7]. Timing data for computation, communication and total run time is gained by running 1 ms simulations on a 512 processor Blue Gene/L partition.

III. RESULTS

A. Data decomposition

The decomposition results for all data sets show a similar pattern with respect to tissue elements vs. total elements per subvolume (fig. 2.a – b). Load ratio 1:1 will result in equal subvolume sizes with large differences in tissue elements per subvolume. In contrast, load ratio 1:100 will have the smallest distribution of tissue elements per subvolume but the widest spectrum in subvolume sizes. However, the decomposition is not equal for all data sets. This can be recognized best when looking at the larger subvolume sizes

where the distribution is not so dense (fig. 2.a – b).

In general, the greater the differences in model weights for computational load of tissue and non – tissue elements, the wider the spectrum of subvolume sizes and the smaller the spectrum of tissue elements per subvolume.

B. Timing data

Figure 2.c – d shows the average computation and communication times as well as the maximum run time for all simulations with load ratio 1:25 which was the fastest in all simulations. However, we note that the simulations with load ratio 1:38.85 were faster than the simulations for load ratio 1:10. In our previous simulations [5] with an overall smaller data set, the simulations of load ratio 1:10 performed better. The timing data shows that the total run time for all simulations is between 345.85 s and 354.38 s for data set 1 and its rotations and between 345.84 s and 355.69 s for data set 2 and the respective rotations. Thus, the difference between the fastest simulations on all data sets is about 10 s. It can also be noted that the difference between data sets and rotations of average computation time is within 1 s for all simulations. Also, the average computation time is within 1.5 s for all simulations.

IV. Discussion

The data decomposition shows that the ORB algorithm output does depend on the orientation of the anatomical data set in the data volume. The differences in data decomposition lead to differences in run time. While the average computation and communication times hardly differ between data sets and rotation, the total run time shows differences up to 3%. This difference seems small but might become significant for longer simulation times. If a simulation with 10 s simulation time were to be carried out, the difference in run time can be projected to be 100,000 s, ie. almost 28 h. In simulations of cardiac arrhythmia, the aim is to run simulations of minutes and even hours to understand the progression of the arrhythmia eventually. Hence, the orientation of the anatomy in the data set become significant when determining the ideal data decomposition for optimal load balancing and run times. In contrast, if larger processor partitions are used in the simulations, the orientation of the anatomy might become less significant. A larger number of subvolumes are created on larger processor partitions. E. g. the ORB algorithms divides the subvolumes for a 512 porcessor partition in half to create 1024 subvolumes for the 1024 processor partition. So computational load differences are reduced also. However, the exact influence on the communication overhead remains unclear. More investigation is needed here.

In the presented study, the data decomposition is not carried out during run time. A change in the simulation code could include it at the beginning of the simulation. This will add a fixed one time overhead to the simulation. An exception is a data decomposition optimization that adjusts the load balancing during run time, i.e. the data decomposition is carried out at the beginning of the simulation, several time steps are computed and timing data of the simulation is used as new input to adjust the data decomposition to yield an optimized load balance. In this case, the overhead of computing the data decomposition for load balancing increases. However, often simulations are based on the same decomposition which allows the data decomposition to be carried out only once offline.

V. Conclusions

This work confirms the hypothesis that the data decomposition using the ORB algorithm is dependent on the anatomy and its orientation in the data volume. The difference in run time on a 512 processor partition will have a greater impact on long simulation runs. Thus, depending on the processor partition size, future work will have to consider the orientation of the anatomy in the global volume for longer simulation runs. However, since the difference in run time for short simulation times is small comparing all orientations in the presented work, it can be neglected in the data decomposition strategy for short simulation times.

References

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

[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] 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

[4] 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

[5] 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)

[6] 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

[7] M. Reumann, B. G. Fitch, A. Rayshubskiy et. al. Strong scaling and speedup to 16,384 processors in cardiac electro – mechanical simulations. Conf Proc IEEE Eng Med Biol Soc. 2009 (submitted)

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

[9] 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.

[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] 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

[12] 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