

Improving the Efficiency of Parallel FEM Simulations on Voxel Domains

N. Kosturski, S. Margenov, and Y. Vutov

Institute of Information and Communication Technologies,
Bulgarian Academy of Sciences

Abstract. In this work, we consider large-scale finite element modeling on voxel grids. We are targeting the IBM Blue Gene/P computer, which features a 3D torus interconnect. Our previous parallelization approach was to divide the domain in one spatial direction only, which lead to limited parallelism. Here, we extend it to all three spatial directions in order to match the interconnect topology.

As a sample problem, we consider the simulation of the thermal and electrical processes, involved in the radio-frequency (RF) ablation procedure. RF ablation is a low invasive technique for the treatment of hepatic tumors, utilizing AC current to destroy the tumor cells by heating. A 3D voxel approach is used for finite element method (FEM) approximation of the involved partial differential equations. After the space discretization, the backward Euler scheme is used for the time stepping.

We study the impact of the domain partitioning on the performance of a parallel preconditioned conjugate gradient (PCG) solver for the arising large linear systems. As a preconditioner, we use BoomerAMG – a parallel algebraic multigrid implementation from the package Hypre, developed in LLNL, Livermore. The implementation is tested on the IBM Blue Gene/P massively parallel computer.

1 Introduction

This work is motivated by the need to improve the parallel efficiency of our supercomputer simulation of RF hepatic tumor ablation on the IBM Blue Gene/P massively parallel computer [5]. This simulation is based on a cubical computational domain, represented by a structured voxel grid. Here, different parallel partitioning strategies (illustrated on Figure 1) are given special attention.

Our previous implementation allowed only 1D partitioning of the computational domain among processors. The biggest limitation of this approach is that the maximum number of processors that can be utilized can never exceed the voxel image resolution in one direction. For higher voxel resolutions, the subdomains assigned to each processor may easily require more than the available amount of memory.

There is another disadvantage of the 1D partitioning strategy that is specific to the IBM Blue Gene/P computer. The communication patterns associated with the 1D partitioning cannot fully utilize the available hardware interconnect, which in this case is either a 3D mesh or a 3D torus, depending on the




Fig. 1. Examples of Domain Partitioning Strategies

number of processors used. With a 3D partitioning, on the other hand, the communication patterns can be naturally mapped to the hardware interconnect topology. Therefore, this strategy is expected to provide the best scalability. The presented parallel tests fully confirm this.

2 Radio-Frequency Tumor Ablation

Let us turn our attention to the considered numerical simulation. RF ablation is an alternative, low invasive technique for the treatment of hepatic tumors, utilizing AC current to destroy the tumor cells by heating ([6, 7]). The destruction of the cells occurs at temperatures of 45°C – 50°C . The procedure is relatively safe, as it does not require open surgery.




Fig. 2. The Structure of a Fully Deployed RF Probe

The considered RF probe is illustrated on Figure 2. It consists of a stainless steel trocar with four nickel-titanium retractable electrodes. Polyurethane is used to insulate the trocar. The RF ablation procedure starts by placing the straight RF probe inside the tumor. The surgeon performs this under computer tomography (CT) or ultrasound guidance. Once the probe is in place, the electrodes are deployed and RF current is initiated. Both the surfaces areas of the uninsulated part of the trocar and the electrodes conduct RF current.

The human liver has a very complex structure, composed of materials with unique thermal and electrical properties. There are three types of blood vessels with different sizes and flow velocities. Here, we consider a simplified test problem, where the liver consists of homogeneous hepatic tissue and blood vessels.

The RF ablation procedure destroys the unwanted tissue by heating, arising when the energy dissipated by the electric current flowing through a conductor is converted to heat. The bio-heat time-dependent partial differential equation [6, 7]

$$\rho c \frac{\partial T}{\partial t} = \nabla \cdot k \nabla T + J \cdot E - h_{bl}(T - T_{bl}) \quad (1)$$

is used to model the heating process during the RF ablation. The term $J \cdot E$ in (1) represents the thermal energy arising from the current flow and the term $h_{bl}(T - T_{bl})$ accounts for the heat loss due to blood perfusion.

The following initial and boundary conditions are applied

$$\begin{aligned} T &= 37^\circ\text{C} \text{ when } t = 0 \text{ at } \Omega, \\ T &= 37^\circ\text{C} \text{ when } t \geq 0 \text{ at } \partial\Omega. \end{aligned} \quad (2)$$

The following notations are used in (1) and (2):

- Ω – the entire domain of the model;
- $\partial\Omega$ – the boundary of the domain;
- ρ – density (kg/m^3);
- c – specific heat (J/kg K);
- k – thermal conductivity (W/m K);
- J – current density (A/m);
- E – electric field intensity (V/m);
- T_{bl} – blood temperature (37°C);
- w_{bl} – blood perfusion ($1/\text{s}$);
- $h_{bl} = \rho_{bl} c_{bl} w_{bl}$ – convective heat transfer coefficient accounting for the blood perfusion in the model.

The bio-heat problem is solved in two steps. The first step is finding the potential distribution V of the current flow. With the considered RF probe design, the current is flowing from the conducting electrodes to a dispersive electrode on the patient's body. The electrical flow is modeled by the Laplace equation

$$\nabla \cdot \sigma \nabla V = 0, \quad (3)$$

with boundary conditions

$$\begin{aligned} V &= 0 \text{ at } \partial\Omega, \\ V &= V_0 \text{ at } \partial\Omega_{el}. \end{aligned}$$

The following notations are used in the above equations:

- V – potential distribution in Ω ;
- σ – electric conductivity (S/m);
- V_0 – applied RF voltage;
- $\partial\Omega_{el}$ – surface of the conducting part of the RF probe.

After determining the potential distribution, the electric field intensity can be computed from

$$E = -\nabla V,$$

and the current density from

$$J = \sigma E.$$

The second step is to solve the heat transfer equation (1) using the heat source $J \cdot E$ obtained in the first step.

For the numerical solution of both of the above discussed steps of the simulation the Finite Element Method (FEM) in space is used ([2]). Linear conforming elements are chosen in this study. To apply the linear FEM discretization to the voxel domain, each voxel is split into six tetrahedra. To solve the bio-heat equation, after the space discretization, the time derivative is discretized via finite differences and the backward Euler scheme is used ([3]).

Let us denote with K^* the stiffness matrix coming from the FEM discretization of the Laplace equation (3). It can be written in the form

$$K^* = \left[\int_{\Omega} \sigma \nabla \Phi_i \cdot \nabla \Phi_j d\mathbf{x} \right]_{i,j=1}^N,$$

where $\{\Phi_i\}_{i=1}^N$ are the FEM basis functions.

The system of linear algebraic equations

$$K^* X = 0 \tag{4}$$

is to be solved to find the nodal values X of the potential distribution.

The electric field intensity and the current density are then expressed by the partial derivatives of the potential distribution in each finite element. This way, the nodal values F for the thermal energy $E \cdot J$ arising from the current flow are obtained.

Let us now turn our attention to the discrete formulation of the bio-heat equation. Let us denote with K and M the stiffness and mass matrices from the finite element discretization of (1). They can be written as

$$K = \left[\int_{\Omega} k \nabla \Phi_i \cdot \nabla \Phi_j d\mathbf{x} \right]_{i,j=1}^N,$$

$$M = \left[\int_{\Omega} \rho c \Phi_i \Phi_j d\mathbf{x} \right]_{i,j=1}^N.$$

Let us also denote with Ω_{bl} the subdomain of Ω occupied by blood vessels and with M_{bl} the matrix

$$M_{bl} = \left[\int_{\Omega} \delta_{bl} h_{bl} \Phi_i \Phi_j d\mathbf{x} \right]_{i,j=1}^N,$$

where

$$\delta_{bl}(x) = \begin{cases} 1 & \text{for } x \in \Omega_{bl}, \\ 0 & \text{for } x \in \Omega \setminus \Omega_{bl}. \end{cases}$$

Then, the parabolic equation (1) can be written in matrix form as:

$$M \frac{\partial T}{\partial t} + (K + M_{bl})T = F + M_{bl}T_{bl}. \quad (5)$$

If we denote with τ the time-step, with T^{n+1} the solution at the current time level, and with T^n the solution at the previous time level and approximate the time derivative in (5) we obtain the following system of linear algebraic equations for the nodal values of T^{n+1}

$$(M + \tau(K + M_{bl}))T^{n+1} = MT^n + \tau(F + M_{bl}T_{bl}). \quad (6)$$

The matrices of the linear systems (4) and (6) are ill-conditioned and very large, having around 10^8 rows. Since they are symmetric and positive definite, we use the PCG [1] method, which is the most efficient solution method in this case.

A parallel algebraic multigrid implementation is used as a preconditioner. The matrix $A = M + \tau(K + M_{bl})$ from (6) is assembled only once on the first time step and not varied after that. The corresponding AMG preconditioner is also constructed only on the first time step.

3 Parallel Tests

The results presented in this section are based on a high-resolution voxel-based representation of the computational domain. The domain consists of liver and tumor tissues, a large bifurcated blood vessel and the RF ablation probe (see Figure 3). We use three different domains (with sizes $127 \times 127 \times 127$, $255 \times 255 \times 255$, and $511 \times 511 \times 511$) to compare the performance and the weak scalability of the simulation using the three considered partitioning strategies.

Table 1 lists the thermal and electrical properties of the materials, which are taken from [6] as well as the blood perfusion coefficient $w_{bl} = 6.4 \times 10^{-3}$ 1/s. For the test simulations, a RF voltage of 10 V is applied for a duration of 8 minutes. A time step of $\tau = 5$ s is used.

Large-scale systems of linear algebraic equations arise from the FEM discretization of the considered problem, requiring an efficient parallel implementation. A parallel PCG solver is used here. The selected preconditioner is Boomer-AMG [4, 8] – a state of the art parallel preconditioner of optimal complexity. A relative PCG stopping criterion in the form

$$\mathbf{r}_k^T C^{-1} \mathbf{r}_k \leq \varepsilon^2 \mathbf{r}_0^T C^{-1} \mathbf{r}_0, \quad \varepsilon = 10^{-6},$$




Fig. 3. High-Resolution 3D Voxel Representation of the Computational Domain

where \mathbf{r}_k stands for the residual at the k -th step of the PCG method, is used.

The settings for the BoomerAMG preconditioner were carefully tuned for maximum scalability in time. The selected coarsening algorithm is *Falgout-CLJP. Modified classical interpolation* is applied. The selected relaxation method is *hybrid symmetric Gauss-Seidel or SSOR*. To decrease the operator and grid complexities two levels of aggressive coarsening are used and the maximum number of elements per row for the interpolation is restricted to six. Smaller operator and grid complexities lead to faster iterations and reduced memory requirements, but can also affect the convergence rate of the solver. Thus, the values of the last two parameters must be carefully chosen to provide the best balance. With the above described setup, the solutions of the linear systems on each time step required 1–3 PCG iterations.

Table 1. Thermal and Electrical Properties of the Materials

Material	ρ (kg/m ³)	c (J/kg K)	k (W/m K)	σ (S/m)
Ni-Ti	6 450	840	18	1×10^8
Stainless steel	21 500	132	71	4×10^8
Liver	1 060	3 600	0.512	0.333
Blood	1 000	4 180	0.543	0.667
Polyurethane	70	1 045	0.026	10^{-5}

The presented parallel tests are performed on the IBM Blue Gene/P machine at the Bulgarian Supercomputing Center (see <http://www.scc.acad.bg/>). This supercomputer consists of 2048 PowerPC 450 based compute nodes, each with four cores running at 850 MHz and 2 GB RAM. It is equipped with a torus network for the point to point communications capable of 5.1 GB/s and a tree network for global communications with a bandwidth of 1.7 GB/s. Our software is implemented in C++, using MPI for the parallelization. It is compiled using the IBM XL C++ compiler with the following options: “-05 -qstrict”.

Table 2. Parallel Times and Weak Scaling for the Complete Simulation

Domain size	$N_p = P_x \times P_y \times P_z$	Unknowns	N_{it}	Time	Weak scaling
$127 \times 127 \times 127$	$8 = 8 \times 1 \times 1$	2 097 152	161	1 225.00 s	
$255 \times 255 \times 255$	$64 = 64 \times 1 \times 1$	16 777 216	128	5 951.08 s	21 %
$511 \times 511 \times 511$	$512 = 512 \times 1 \times 1$	134 217 728	—	> 24 h	< 2 %
$127 \times 127 \times 127$	$8 = 4 \times 2 \times 1$	2 097 152	167	1 137.83 s	
$255 \times 255 \times 255$	$64 = 8 \times 8 \times 1$	16 777 216	129	1 203.29 s	95 %
$511 \times 511 \times 511$	$512 = 32 \times 16 \times 1$	134 217 728	114	1 581.13 s	72 %
$127 \times 127 \times 127$	$8 = 2 \times 2 \times 2$	2 097 152	167	1 137.91 s	
$255 \times 255 \times 255$	$64 = 4 \times 4 \times 4$	16 777 216	128	1 062.30 s	107 %
$511 \times 511 \times 511$	$512 = 8 \times 8 \times 8$	134 217 728	114	1 155.08 s	99 %

The parallel times for the whole numerical simulation are presented in Table 2. The three parts of the table correspond to the three considered partitioning strategies. Here, N_p is the number of processors, P_x , P_y , P_z are the number of partitions in direction x , y , z respectively, and N_{it} is the number of PCG iterations performed during the simulation. The simulation in the 1D partitioning case for the largest domain could not finish in 24 hours, which is the hard limit for problems of this size on the Blue Gene/P machine in the Bulgarian Supercomputing Center. The weak scaling with respect to the smallest simulation domain is provided for each partitioning strategy. A big advantage of the 3D partitioning strategy over the other two is observed. The 1D partitioning is the least scalable of the considered three.

Equivalent computations are performed with each partitioning strategy, as each processor holds the same number of unknowns and the number of PCG iterations is almost the same. The communications, however, are quite different. In the case of 3D partitioning, the partition allocated to each processor has the same size ($64 \times 64 \times 64$ unknowns) for all domains. Typical communications involve the transfer of values for all the unknowns on the interfaces between neighboring partitions. Therefore, in the 3D partitioning case, the communication time should be independent of the domain size. Moreover, by mapping the partitioning to the underlying interconnect topology, we ensure that the commu-

nications in each direction can be performed in parallel. With the 1D partitioning strategy the size of the interfaces increases four times with each next domain size, which is the reason for the much lower scalability. The scalability of above 100 % can be explained by the smaller number of PCG iterations in this case.

4 Concluding remarks

Parallel tests of a large-scale, time-dependent, voxel-based simulation of the RF ablation procedure are presented. Such simulations require very efficient use of supercomputer resources. Three approaches to the partitioning of the computational domain are presented and compared. Using a 3D domain partitioning leads to a substantial scalability improvement over our previous work, as clearly demonstrated by the parallel test results.

Our future plans include enhancing the model in order to simulate a more complex type of RF ablation probe, involving fluid injection during the procedure. Overcoming the previous restriction on the number of processors in our implementation as well as improving its scalability were crucial for enabling us to take the next step.

Acknowledgments

This work is partly supported by the Bulgarian NSF Grants DCVP 02/1 and DPRP7RP02/13. We also kindly acknowledge the support of the Bulgarian Supercomputing Center for the access to the IBM Blue Gene/P supercomputer.

References

1. O. Axelsson, *Iterative Solution Methods*, Cambridge University Press, 1996.
2. S. Brenner, L. Scott, *The mathematical theory of finite element methods*, Texts in applied mathematics, 15, Springer-Verlag, 1994.
3. E. Hairer, S.P. Norsett, G. Wanner *Solving ordinary differential equations I, II*, Springer Series in Comp. Math., 2000, 2002
4. V.E. Henson, U.M. Yang, *BoomerAMG: A parallel algebraic multigrid solver and preconditioner*, Applied Numerical Mathematics 41 (1), Elsevier, 2002, 155–177.
5. N. Kosturski, S. Margenov, *Supercomputer Simulation of Radio-Frequency Hepatic Tumor Ablation*, AMiTaNS'10 Proceedings, AIP CP vol. 1301, pp. 486–493.
6. S. Tungjitkusolmun, S.T. Staelin, D. Haemmerich, J.Z. Tsai, H. Cao, J.G. Webster, F.T. Lee, D.M. Mahvi, V.R. Vorperian, *Three-dimensional finite-element analyses for radio-frequency hepatic tumor ablation*, IEEE transactions on biomedical engineering 49 (1), 2002, 3–9.
7. S. Tungjitkusolmun, E.J. Woo, H. Cao, J.Z. Tsai, V.R. Vorperian, J.G. Webster, *Thermal-electrical finite element modelling for radio frequency cardiac ablation: Effects of changes in myocardial properties*, Medical and Biological Engineering and Computing 38 (5), 2000. 562–568.
8. Lawrence Livermore National Laboratory, *Scalable Linear Solvers Project*, http://www.llnl.gov/CASC/linear_solvers/.