AVOIDING THE PROBLEM OF FE MESHING: A PARALLEL ALGEBRAIC MULTIGRID WITH MULTIPLE RIGHT-HAND SIDE TREATMENT FOR AN EFFICIENT AND MEMORY-ECONOMICAL COMPUTATION OF HIGH RESOLUTION EEG AND MEG LEAD FIELD BASES

C.H.Wolters[1,2,3], A. Anwander[3], S. Reitzinger[4] and G. Haase[5]

(1) Scientific Computing and Imaging Institute, University of Utah, Salt Lake City, USA (2) Max Planck Institute for Mathematics in the Sciences, Leipzig, Germany (3) Max Planck Institute for Human Cognitive and Brain Sciences, Leipzig, Germany (4) CST GmbH, Darmstadt, Germany (5) Inst. for Mathematics and Computational Sciences, University of Graz, Austria.

ABSTRACT

When using the Finite Element (FE) method for volume conductor modeling in EEG/MEG source reconstruction, one problem consists of generating an appropriate FE mesh. If a software package were able to efficiently perform FE computations with even some million unknowns, then we only could transform each voxel from a segmented MR or CT image into a finite element and would avoid the meshing problem. On this poster we will present a new efficient strategy, the parallel algebraic multigrid preconditioned conjugate gradient method with simultaneous treatment of multiple right-hand sides for the computation of EEG and MEG lead field bases. We will show that this method is memory-economical and furthermore exploits a much higher cache hit rate to speed the computation by about a factor of 2 compared to former approaches. Together with the concept of the EEG and MEG lead field bases, the complexity of realistic high resolution anisotropic FE forward modeling within the EEG/MEG inverse problem is significantly reduced. For FE-meshes with some few hundred thousand unknowns, the computations can now be performed in some few minutes on a single processor PC. Our parallel approach furthermore offers the possibility to efficiently treat FE-meshes with some million unknowns so that in the future, the problem of FE meshing can be omitted.

KEY WORDS

EEG/MEG Source Reconstruction, Finite Element Method, Finite Element meshing, Lead Field Bases, Algebraic MultiGrid, Preconditioned Conjugate Gradient Methods, Treatment of Multiple Right-Hand Sides, Cache Algorithms, Parallelization.

INTRODUCTION

When choosing the FE method for volume conductor modeling within the EEG/MEG inverse problem, the construction of the lead field bases requires "number of EEG/MEG sensors " many solutions of large sparse FE systems of linear equations in a setup phase for each individual head model [1,2]. Therefore, preconditioning techniques for the iterative solution process are important to speed the computation and keep the necessary memory amount in reasonable areas. It was previously shown that the Algebraic MultiGrid preconditioned Conjugate Gradient (AMG-CG) method is a very efficient solver for inhomogeneous anisotropic high resolution FE forward modeling [3,4]. Even if Multigrid-techniques have proven to be of optimal order with respect to memory requirements and computational costs, a further distribution onto multiple computational nodes and memories gets necessary if the resolution exceeds a limit of some hundred thousand nodes. For those purposes, we recently presented a parallel AMG-CG iterative solver technique for EEG/MEG source localization [5]. Here, we will discuss a new strategy for a further speedup, the simultaneous treatment of multiple right-hand sides. We will first present performance results for the computation of EEG and MEG lead field bases for a realistic anisotropic FE head model on a single processor machine. We then show the efficiency of the new approach in a parallel computation.

METHODS

In a first step, compared to the AMG-CG presented in [3,5], general algorithmical improvements were implemented for the new Multiple Right-Hand Side AMG-CG (MultiRHS-AMG-CG) [6].

The old memory management for the stiffness and interpolation matrices was replaced by the classical Compact Row Storage (CRS) format in order to decrease the number of cache misses. Within the AMG algorithm, defect calculation follows forward Gauss-Seidel (GS) smoothing and both operations require matrix-vector operations. For symmetric stiffness matrices, parts of the matrix-vector operation from the last GS smoothing can be efficiently stored and reused by the defect calculation, a merging which leads to a reduction of the operation count. The AMG-procedure on the next coarser level is called with a zero- initial correction vector. This can be used, too, so that the first forward GS smoothing sweep on the coarser levels is reduced to half of the arithmetic and memory operations. If a V-cycle is chosen, i.e., only one pre-smoothing sweep is performed, the special structure of this smoother on the coarser levels furthermore leads to a reduction of the subsequent defect computation.

Since the RHSs in the lead field bases approach are computed beforehand and the stiffness matrix remains the same, we can simultaneously solve for a whole block of RHSs. The most computationally expensive operations in the AMG-CG method are the matrix-vector operations within the CG and within the AMG components smoothing, defect calculation, interpolation and prolongation. If the vector for one RHS is exchanged against a whole block of vectors for multiple RHSs and if this block is not stored as a matrix, but as a long vector (first the first entries of the RHSs, then the second entries etc., resulting in a long vector), then each matrix entry only has to be accessed once and can be multiplied to all corresponding values in the block-vector. This procedure results in much higher cache hit rates, which speeds the computations. For the simultaneous treatment of 3 RHS, the inner loops were manually unrolled, leading to a further reduction of the solver time

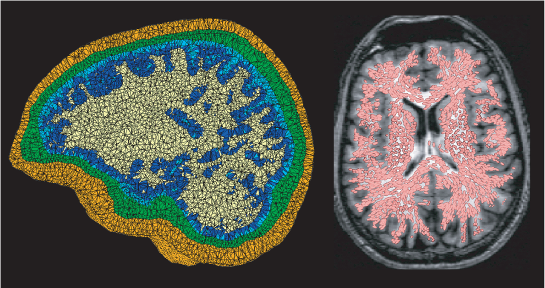

Figure 1 Left: Tetrahedral head model with 147,287 nodes and 892,115 elements. Right: Conductivity tensors in barycenters of white matter finite elements.

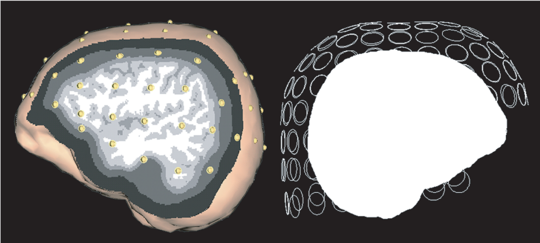

Figure 2: 71 electrode configuration (left) and 147 channel BTI MEG (right).

RESULTS

As a basis for our computations, we chose a realistic tetrahedral FE model with 147,287 nodes and 892,115 elements and anisotropic layers skull and white matter (Fig.1), a 71 electrode EEG and a 147 channel MEG configuration (Fig.2).

In a first experiment on a single processor machine, we compared the computation time for the construction of the EEG and MEG lead field bases for the Jakobi-preconditioned CG (J-CG), the symmetric Incomplete Cholesky preconditioned CG without fill-in (symIC(0)-CG) and the AMG-CG (see [3] for these approaches) with the new MultiRHS-AMG-CG while varying the number of simultaneously treated RHSs. Speedup tests were performed on three different platforms, a Mac-OSX with PowerBook G4 proc (1Ghz, 512 KB cache), a Red-Hat Linux PC with Xeon proc (2.4Ghz, 512 KB cache) and a Red-Hat Linux PC with Pentium 4 proc (3.2 Ghz, 1024 KB cache). The computation time for the EEG/MEG lead field bases with J-CG (symIC(0)-CG) on those three platforms were 83/223 min. (44/121 min.), 26/68 min. (12/35 min.) and 15/42 min. (8/22 min.). The results for AMG-CG and MultiRHS-AMG-CG for EEG and MEG are shown in Figs. 3 and 4. The computation time for a specific number of simultaneous RHSs is indicated above the curves.

The second experiment was performed on a Red Hat Linux PC-cluster with 32 nodes, each equipped with a Xeon proc (1.7GHz, 256 KB cache, 1GB memory) and a 1Gbit ethernet. We measured the performance of the parallel MultiRHS-AMG-CG forward simulations on 1, 2, 4, 6 and 8 nodes while varying the number of simultaneously treated RHSs (Fig. 5) . The number of iterations is shown over the curves (Fig.5). In Fig.6, we evaluated the speedup of the parallel MultiRHS-AMG-CG through an improved cache hit rate. The simultaneous treatment of 3 RHSs leads to a speedup of at least 1.6 for all tested numbers of nodes.

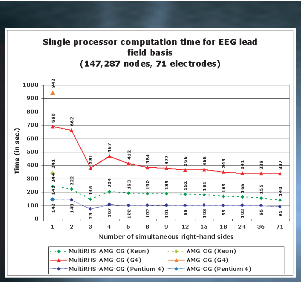

Figure 3 Time for the computation of the EEG lead field basis on a Mac-OSX PowerBook G4 and on Red-Hat Linux PC's with either Xeon or Pentium 4 architecture using the conventional AMG-CG and the new MultiRHS-AMG-CG.

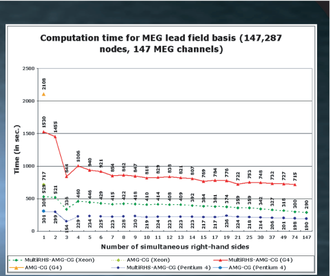

Figure 4 Time for the computation of the MEG lead field basis on a Mac-OSX PowerBook G4 and on Red-Hat Linux PC's with either Xeon or Pentium 4 architecture using the conventional AMG-CG and the new MultiRHS-AMG-CG.

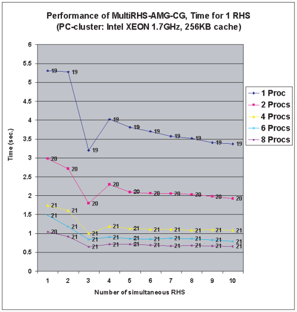

Figure 5 Time for the computation of one RHS on a Red-Hat Linux PC cluster using the new parallel MultiRHS-AMG-CG. The number of iterations is indicated above the curves.[1]

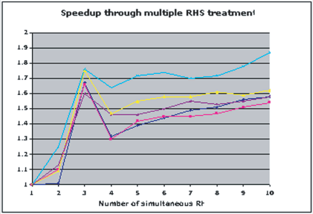

Figure 6 Speedup through an improved cache hit rate using the new parallel MultiRHS-AMG-CG.

DISCUSSION

The combination of the lead field bases concept [1,2] with the presented parallel MultiRHS-AMG-CG solver for the setup phase reduces significantly the complexity of high resolution anisotropic high resolution finite element head modeling. If FE meshes with some few hundred thousand nodes are used, the computation of the lead field bases in the setup phase for each individual head model can now be performed in a few minutes on a standard single processor PC. Furthermore, by means of the distribution of the memory on multiple computational nodes, our parallel approach allows the computation of the lead field bases for FE meshes with even some million nodes in a reasonable time, so that the problem of FE meshing can be omitted by simply transforming each voxel from a segmented MR or CT image into a finite element. The FE forward approach then reduces to the multiplication of the lead field bases to the FE source load RHS vector [1,2], which can easily be parallelized for FE computations with millions of unknowns. The presented concept can be used by all inverse methods in continuous or discrete source parameter space.

The treatment of multiple RHSs within the new MultiRHS-AMG-CG reduced the computation time for the EEG and MEG lead field bases by at least a factor of 2 (serial version) or 1.6 (parallel version). With the manual unrollment of inner loops for the simultaneous treatment of 3 RHS, this approach belongs to the fastest on all platforms so that we would recommend that choice for the RHS parameter in our software NeuroFEM [7]. On single processor platforms with a smaller cache and a slower access to the main memory, the improvement of the data-structures by means of the CSR storage for stiffness and interpolation matrices led to a further speedup factor of up to 1.38.

The source code of our software NeuroFEM [7] will be available on a fee-free basis upon request.

ACKNOWLEDGEMENTS

This work has been supported by the MPI for Mathematics in the Sciences Leipzig, by the IST-program of the European Commission, project No.10378, http://www.simbio.de, by the NIH NCRR center for Bioelectric Field Modeling, Simulation and Visualization and by the Austrian Science Fund-'Fonds zur Förderung der wissenschaftlichen Forschung'-within the SFB F013 "Numerical and Symbolic Scientific Computing".

REFERENCES

[1] Wolters, C.H., Grasedyck, L. and Hackbusch, H. Efficient Computation of Lead Field Bases and Influence Matrix for the FEM-based EEG and MEG Inverse Problem. Inverse Problems. 20(4), 2004, p.1099-1116.
[2] Wolters, C.H., Grasedyck, L., Anwander, A. and Hackbusch, H. Efficient Computation of Lead Field Bases and Influence Matrix for the FEM-based EEG and MEG Inverse Problem. in Biomag2004, Proc. 14th Int. Conf. on Biomagnetism, Boston, USA, August 2004.
[3] Wolters, C.H., Reitzinger, S., Basermann, A., Burkhardt, S., Hartmann, U., Kruggel, F. , Anwander, A. Improved Tissue Modeling and Fast Solver Methods for High Resolution FE-Modeling in EEG/MEG-Source Localization. in Biomag2000, Proc. 12th Int. Conf. on Biomagnetism, J. Nenonen, R.J. Ilmoniemi, and T. Katila, eds. (Helsinki Univ. of Tech., Espoo, Finland, 2001, http://biomag2000.hut.fi), p.655-58.
[4] Mohr, M. and Vamrunste, B. Comparing Iterative Solvers for Linear Systems associated with the Finite Difference Discretisation of the Forward Problem in Electro-Encephalographic Source Analysis. Med. Biol. Eng. Comp., 2003. p. 75-84.
[5] Wolters, C.H., Kuhn, M., Anwander, A., Reitzinger, S. A Parallel Algebraic Multigrid Solver for Finite Element Method Based Source Localization in the Human Brain. Comp.Vis.Sci. 5(3), 2002, p.165-77.
[6] Haase, G., Reitzinger, S. Cache Issues of Algebraic Multigrid Methods for Linear Systems with Multiple Right-Hand Sides SFB-Report 02-05, University Linz, SFB F013, June 2002.
[7] Anwander, A. and Wolters, C.H. NeuroFEM: A C++ class structured software package for anisotropic high resolution Finite Element Method based EEG and MEG source reconstruction, http://www.neurofem.com.

---

[1] Since the communication time could have been negatively impacted by other users jobs, we did not include it in Fig.5. To make an example: The pure computation time for 3RHS-AMG-CG on 8 procs was 1.94 sec. (thus 0.65 sec. for one RHS as shown in Fig.5). When including communication time, the computation took 2.3 sec. (thus 0.77 sec. per RHS).