Phys. Med. Biol. 50 (2005) 4695-4710

Fast multipole acceleration of the MEG/EEG boundary element method

Jan Kybic^{1,2}, Maureen Clerc², Olivier Faugeras², Renaud Keriven² and Théo Papadopoulo²

 ¹ Center for Machine Perception, Faculty of Electrical Engineering, Czech Technical University in Prague, Czech Republic
 ² Odyssée Laboratory—ENPC/ENS/INRIA, INRIA, 2004 Route des Lucioles, BP93, 06902 Sophia-Antipolis, France

E-mail: kybic@fel.cvut.cz and Maureen.Clerc@sophia.inria.fr

Received 6 March 2005, in final form 17 June 2005 Published 21 September 2005 Online at stacks.iop.org/PMB/50/4695

Abstract

The accurate solution of the forward electrostatic problem is an essential first step before solving the inverse problem of magneto- and electroencephalography (MEG/EEG). The symmetric Galerkin boundary element method is accurate but cannot be used for very large problems because of its computational complexity and memory requirements. We describe a fast multipole-based acceleration for the symmetric boundary element method (BEM). It creates a hierarchical structure of the elements and approximates far interactions using spherical harmonics expansions. The accelerated method is shown to be as accurate as the direct method, yet for large problems it is both faster and more economical in terms of memory consumption.

1. Introduction

The boundary element method (BEM) is widely used for solving the forward and inverse problems of magneto- and electroencephalography (MEG/EEG) on realistic geometries (Phillips *et al* 1997, Hämäläinen *et al* 1993). It unfortunately leads to huge and dense linear systems which can be hard to handle.

Using fine models is essential for accurately modelling the electromagnetic behaviour of the head. The geometry of the brain, especially the cortex containing the sources, is so complex and convoluted that very small elements (of the order of 1 mm) are needed to model it accurately. Fine models are also needed to accurately represent the fine details of the spatially varying fields. Finally, most BEMs have a precision that severely drops when the sources are close to an interface (Rahola and Tissari 2002, 1998, Ferguson and Stroink 1997, Clerc *et al* 2002). Therefore, as the brain sources are supposed to lie in a very thin layer of the cortex

(several mm at most) and thus very close to the surface of the brain, the layers involved must be discretized finely.

The fast multipole method (FMM) (Cheng *et al* 1999) acceleration significantly decreases the asymptotic time and memory complexity of solving the forward MEG/EEG problem. We described earlier a preliminary single-level FMM for the double-layer approach (Clerc *et al* 2002) that we extend here to a multi-level FMM and adapt for the symmetric BEM (Kybic *et al* 2005). To the best of our knowledge, this paper describes the only implementation capable of accurately solving the MEG/EEG forward problem for realistic head models described by meshes with over 30 000 points (70 000 unknowns) on a single personal computer.

There is extensive literature dealing with FMM (Cheng *et al* 1999, Beatson and Greengard 1997, Epton and Dembart 1995, Dembart and Yip 1998, Rahola 1998) for gravitational or electromagnetic scattering calculations. Some authors have considered the electrostatic Maxwell problem and the symmetric BEM approach, but have treated problems with only one interface (Of *et al* 2002).

2. Fast multipole method expansions

The fast multipole method (FMM) (Clerc *et al* 2002, Cheng *et al* 1999, Beatson and Greengard 1997, Dembart and Yip 1998) is a hierarchical approximation algorithm which significantly reduces the time and memory complexity required for the resolution of the linear system of equations $A\mathbf{u} = \mathbf{c}$, produced by the BEM (Kybic *et al* 2005). It takes advantage of the fact that interaction between surface elements decreases quickly with distance. We use an iterative method (MINRES) that accesses the matrix A only through matrix–vector multiplications $A\mathbf{u}$. With the FMM, the matrix does not need to be formed explicitly and the overall complexity of calculating the product $A\mathbf{u}$ representing the pairwise interactions between *N* elements is decreased from $O(N^2)$ to O(N).

We first briefly recall the operators appearing in the symmetric BEM formulation (section 2.1) and introduce the spherical harmonic expansion. Then the general FMM framework is presented (section 3) and applied. We refer the reader to our report in Kybic and Clerc (2004) for additional details about the symmetric BEM and our FMM algorithm, which had to be omitted here for space reasons. We compare our approach to the pre-corrected-FFT acceleration (Tissari and Rahola 2003) in section 4.1. It is also possible to hierarchically simplify the system matrix once it has been computed (Börm and Hackbusch 2003).

2.1. Symmetric boundary element method operators

To accelerate the symmetric BEM, we shall need to quickly evaluate the matrix vector products (Nx, D^*y, Dx, Sy) that appear in the symmetric BEM system (Kybic *et al* 2005, 2004, Kybic and Clerc 2003)

$$\underbrace{\begin{pmatrix} \mathsf{N} & \mathsf{D}^* \\ \mathsf{D} & \mathsf{S} \end{pmatrix}}_{\mathsf{A}} \underbrace{\begin{pmatrix} \mathsf{x} \\ \mathsf{y} \end{pmatrix}}_{\mathsf{u}} = \underbrace{\begin{pmatrix} \mathsf{w} \\ \mathsf{z} \end{pmatrix}}_{\mathsf{c}},\tag{1}$$

where the unknowns **x** and **y** represent a discretized version of the electric potential V and flow $p = \sigma \partial_{\mathbf{n}} V$, respectively, and the right-hand side terms **w** and **z** represent the known free-space field (potential and flow) corresponding to the sources. The matrices N, D, D^{*} and S are the discretized versions of the following continuous operators (all mapping a scalar function f on

an interface $\partial \Omega$ to another scalar function). For **r** on $\partial \Omega$,

$$(\mathcal{N}f)(\mathbf{r}) = \int_{\partial\Omega} \partial_{\mathbf{n},\mathbf{n}'}^2 G(\mathbf{r} - \mathbf{r}') f(\mathbf{r}') \, ds(\mathbf{r}') \qquad (\mathcal{D}f)(\mathbf{r}) = \int_{\partial\Omega} \partial_{\mathbf{n}'} G(\mathbf{r} - \mathbf{r}') f(\mathbf{r}') \, ds(\mathbf{r}')$$
$$(\mathcal{D}^*f)(\mathbf{r}) = \int_{\partial\Omega} \partial_{\mathbf{n}} G(\mathbf{r} - \mathbf{r}') f(\mathbf{r}') \, ds(\mathbf{r}') \qquad (\mathcal{S}f)(\mathbf{r}) = \int_{\partial\Omega} G(\mathbf{r} - \mathbf{r}') f(\mathbf{r}') \, ds(\mathbf{r}')$$

with a Green function $G(\mathbf{r}) = 1/(4\pi \|\mathbf{r}\|)$ and where, for example, $\partial_{\mathbf{n}}$, $\partial_{\mathbf{n}'}$ stands for a partial derivative with respect to the surface normal at \mathbf{r}, \mathbf{r}' . We obtain matrices

$$(\mathsf{N})_{ik} = \langle \mathcal{N}\varphi_k, \varphi_i \rangle, \qquad (\mathsf{S})_{jl} = \langle \mathcal{S}\psi_l, \psi_j \rangle$$
(2)

$$(\mathsf{D})_{jk} = (\mathsf{D}^*)_{kj} = \langle \mathcal{D}\varphi_k, \psi_j \rangle = \langle \mathcal{D}^*\psi_j, \varphi_k \rangle,$$
(3)

where ψ , respectively φ , are P0 (piecewise constant on each triangle), respectively P1 (piecewise linear on each triangle) boundary elements. As an example, the discretized system for a nested three-layer model is

$(\sigma_1+\sigma_2)N_{11}$	$-\sigma_2 N_{12}$	0	$-2D_{11}^{*}$	D_{12}^{*}		(\mathbf{x}_1)		(\mathbf{w}_1)	1
$-\sigma_2 N_{21}$	$(\sigma_2+\sigma_3)\mathbf{N}_{22}$	$-\sigma_3 N_{23}$	D_{21}^{*}	$-2D_{22}^{*}$		x ₂		w ₂	
0	$-\sigma_3 N_{32}$	$\sigma_3 N_{33}$	0	D_{32}^{*}	.	x ₃	=	w ₃	
$-2D_{11}$	D ₁₂	0	$(\sigma_1^{-1}+\sigma_2^{-1})S_{11}$	$-\sigma_2^{-1}S_{12}$		\mathbf{y}_1		\mathbf{z}_1	
D ₂₁	$-2D_{22}$	D ₂₃	$-\sigma_2^{-1} S_{21}$	$(\sigma_1^{-1}+\sigma_2^{-1})S_{22}$		(\mathbf{y}_2)	, .	(\mathbf{z}_2)	ļ
		Ă				u	- `	c	-

where we denote the surfaces with 1, 2, 3, the operator subscripts denote targets and sources (e.g., N₂₃ relates potential on surface 3 with free-space potential on surface 2), and σ_{α} are the conductivities of enclosed volumes, numbered from inside.

2.2. Spherical harmonics

The Green function $G(\mathbf{r} - \mathbf{r}') \sim 1/\|\mathbf{r} - \mathbf{r}'\|$ can be decomposed as (Epton and Dembart 1995)

$$\underbrace{\|\mathbf{r}' - \mathbf{C}\| > \lambda \|\mathbf{r} - \mathbf{C}\|}_{\text{well-separateness}} \Longrightarrow \frac{1}{\|\mathbf{r} - \mathbf{r}'\|} = \sum_{n=0}^{L} \sum_{m=-n}^{n} I_n^{-m} (\mathbf{C} - \mathbf{r}) O_n^m (\mathbf{r}' - \mathbf{C}) + \text{error}$$
(4)

where **C** is the centre of expansion and I_n^m , respectively O_n^m , are the inner, respectively outer, spherical harmonics (see the appendix). To obtain a practical expression, the series is truncated to order *L*. Acceptable accuracy is guaranteed for **r**, **r**' sufficiently far apart. The approximation error can be bounded by choosing a suitable parameter $\lambda > 1$.

2.2.1. Operator S. To approximate the discretized operator S, we integrate (4) with the P0 basis functions ψ_i . For each element (triangle) *i* we define the outer-field, respectively inner-field, expansion coefficients ³

$${}^{i}\mathbf{a}_{n}^{m}(\mathbf{C}) = \int I_{n}^{m}(\mathbf{C} - \mathbf{r})\psi_{i}(\mathbf{r}) \,\mathrm{d}\mathbf{r}$$
⁽⁵⁾

$${}^{i}\widetilde{\mathbf{a}}_{n}^{m}(\mathbf{C}) = \int O_{n}^{m}(\mathbf{C} - \mathbf{r})\psi_{i}(\mathbf{r}) \,\mathrm{d}\mathbf{r}.$$
(6)

³ With a deliberate, though unfortunate, conflict in notation: the outer field concerns the information propagating out, from elements **r** close to a centre of expansion **C**, and the inner field concerns the information coming into an element **r** from a remote centre of expansion **C**.

The operator S is then approximated (if elements *i* and *j* are well separated) by

$$4\pi \langle \mathcal{S}\psi_i, \psi_j \rangle \approx {}^{i}\mathbf{a}(\mathbf{C}) \odot {}^{j}\widetilde{\mathbf{a}}(\mathbf{C}) \stackrel{\text{def}}{=} \sum_{\substack{n=0\dots L\\m=-n\dots n}} (-1)^{n} {}^{i}\mathbf{a}_n^{-m}(\mathbf{C}) {}^{j}\widetilde{\mathbf{a}}_n^m(\mathbf{C})$$
(7)

where we have used the symmetry relation $O_n^m(-\mathbf{r}) = (-1)^n O_n^m(\mathbf{r})$.

2.2.2. Operator \mathcal{D} . We apply a gradient with respect to **r** to (4) and integrate the result with the basis functions ψ_i (P0) and $\varphi'_{j'}$ (partial P1). The partial P1 function $\varphi'_{j'}$ is identical to some φ_i on one triangle and zero elsewhere. We define coefficients

$${}^{i'}\mathbf{b}_{n}^{m}(\mathbf{C}) = \int \nabla I_{n}^{m}(\mathbf{C} - \mathbf{r}) \cdot \mathbf{n}_{i'} \varphi_{i'}'(\mathbf{r}) \,\mathrm{d}\mathbf{r}$$
(8)

$${}^{i'}\tilde{\mathbf{b}}_{n}^{m}(\mathbf{C}) = \int \nabla O_{n}^{m}(\mathbf{C} - \mathbf{r}) \cdot \mathbf{n}_{i'} \varphi_{i'}^{\prime}(\mathbf{r}) \,\mathrm{d}\mathbf{r}$$
⁽⁹⁾

with b, \tilde{b} again representable by $(L + 1)^2$ -dimensional complex vectors. Then we have

$$-4\pi \langle \mathcal{D}\varphi'_{i'}, \psi_j \rangle = -4\pi \langle \varphi'_{i'}, \mathcal{D}^* \psi_j \rangle \approx {}^{i'} \mathbf{b} \odot {}^{j} \mathbf{\widetilde{a}} = \sum_{n,m} (-1)^n {}^{i'} \mathbf{b}_n^{-m}(\mathbf{C}) {}^{j} \mathbf{\widetilde{a}}_n^m(\mathbf{C})$$
$$= {}^{j} \mathbf{a} \odot {}^{i'} \mathbf{\widetilde{b}} = \sum_{n,m} (-1)^n {}^{j} \mathbf{a}_n^m(\mathbf{C}) {}^{i'} \mathbf{\widetilde{b}}_n^{-m}(\mathbf{C}).$$
(10)

Coefficients ^{*i*} b corresponding to a complete P1 basis function φ_i at vertex *i* are calculated by aggregating the coefficients ^{*i*} b on all triangles $T_{i'}$ sharing vertex *i*.

2.2.3. Operator \mathcal{N} . To calculate $(N)_{i'j'} = \langle \mathcal{N}\varphi'_{i'}, \varphi'_{j'} \rangle$ we define coefficients

$${}^{i}\!\!\!\!\operatorname{c}_{n}^{m} = \left(\mathbf{q}_{i'} \times \mathbf{n}_{i'}\right)^{i} \mathbf{a}_{n}^{m} \qquad \text{and} \qquad {}^{j'}\!\!\!\!\!\operatorname{c}_{n}^{m} = \left(\mathbf{q}_{j'} \times \mathbf{n}_{j'}\right)^{j'}\!\!\!\!\!\!\!\operatorname{a}_{n}^{m} \tag{11}$$

represented by $(L + 1)^2$ 3D complex vectors. We obtain

$$4\pi \langle \mathcal{N}\varphi'_{i'},\varphi'_{j'}\rangle \approx {}^{i'}\mathsf{c}(\mathbf{C}) \odot {}^{j'}\widetilde{\mathsf{c}}(\mathbf{C}) = \sum_{\substack{n=0\dots L\\m=-n\dots n}} (-1)^{n\,i'} \mathsf{c}_n^{-m}(\mathbf{C}) \cdot {}^{j'}\widetilde{\mathsf{c}}_n^m(\mathbf{C})$$
(12)

with a complex scalar product '.'. The coefficients c for complete P1 elements are again aggregated from constituting triangles.

2.3. Translating multipolar representations

The formulae for the outer–outer, inner–inner and outer–inner operators R, S, and T (13), respectively, are generalized from Epton and Dembart (1995) and are identical for all three types of expansion coefficients (a, b, c, here represented by x) provided that elementwise complex multiplication and addition are used for coefficients c:

$$\mathbf{x}_{n'}^{-m'}(\mathbf{M}) = (\mathbf{R}_{\mathbf{N}\mathbf{M}}\mathbf{x})_{n'}^{-m'} = \sum_{n=0...n',m=-n...n} I_{n'-n}^{m-m'}(\mathbf{M}-\mathbf{N})\mathbf{x}_{n}^{-m}(\mathbf{N})$$

$$\widetilde{\mathbf{x}}_{n'}^{-m'}(\mathbf{M}) = (\mathbf{S}_{\mathbf{N}\mathbf{M}}\widetilde{\mathbf{x}})_{n'}^{-m'} = \sum_{n=n'...L,m=-n...n} I_{n-n'}^{m'-m'}(\mathbf{M}-\mathbf{N})\widetilde{\mathbf{x}}_{n}^{-m}(\mathbf{N})$$

$$\widetilde{\mathbf{x}}_{n'}^{m'}(\mathbf{M}) = (\mathbf{T}_{\mathbf{N}\mathbf{M}}\mathbf{x})_{n'}^{m'} = \sum_{n=0...L,m=-n...n} O_{n+n'}^{m+m'}(\mathbf{M}-\mathbf{N})\mathbf{x}_{n}^{-m}(\mathbf{N}).$$
(13)

It is advantageous to pre-compute the values $I_n^m(\mathbf{M} - \mathbf{N})$ (usable for both R, S) and $O_n^m(\mathbf{M} - \mathbf{N})$.

The formulae above all require $O(L^4)$ operations, which makes them a bottleneck of the FMM, especially the operator T. Several acceleration techniques were proposed, reducing the complexity to $O(L^3)$ or even $O(L^2 \log L)$: using FFT (Epton and Dembart 1995), rotation to the *z*-axis where translation is simpler (Beatson and Greengard 1997), or a plane wave representation (Cheng *et al* 1999). However, they are significantly more complex than (13) and thus faster only for high values of *L*; e.g., for $L \ge 20$ in the case of the FFT approach according to our tests.

3. Fast multipole method algorithms

Given the expansion and translation formulae for operators S, D, D^* and N, derived above, we can now use them to formulate the FMM. The particularity here is that the FMM needs to be applied to all four operators, so we first formulate it in generic terms.

3.1. FMM basics

We are to calculate the interaction between two groups of elements, A and B,

$$y_j = \sum_{i \in \mathcal{A}} f_{ij} x_i, \quad \text{for all} \quad j \in \mathcal{B}.$$
 (14)

The values f_{ij} correspond to the elements of the matrices S, D, D* and N (2), (3). The elements from \mathcal{A} and \mathcal{B} correspond to the support of the basis functions, i.e., to either triangles (P0 elements ψ), or to sets of triangles with a common vertex (P1 elements φ).

For each sufficiently *well-separated* pair of elements *i* in A, respectively *j* in B, the term f_{ij} can be approximated with an *a priori* given precision ε ,

$$\underbrace{(d(i, \mathbf{C}) > \lambda d(j, \mathbf{C}) \stackrel{\text{def}}{\Longrightarrow} D_{\mathbf{C}}(i, j))}_{\text{well-separated}} \Longrightarrow |f_{ij} - \phi_j(\mathbf{C}) \odot \widetilde{\phi}_i(\mathbf{C})| \leqslant \varepsilon, \quad (15)$$

where $\phi_i(\mathbf{C})$ and $\tilde{\phi}_i(\mathbf{C})$ are called an outer (far, or multipole) respectively inner (near, or local) expansion, **C** is the centre of expansion, and $d(i, \mathbf{C})$ is a distance from element *i* to **C**. The well-separateness condition (with $\lambda > 1$) can be identified in (4), the expansion formula corresponds to (7), (10), (12) and the expansions $\phi, \tilde{\phi}$ to coefficients a, b, c, respectively $\tilde{a}, \tilde{b}, \tilde{c}$.

The expansions are functions $\phi_j : \mathbb{R}^3 \to \mathcal{Q}, \widetilde{\phi}_i : \mathbb{R}^3 \to \widetilde{\mathcal{Q}}$, with suitable domains $\mathcal{Q}, \widetilde{\mathcal{Q}}$, in our case 2D tables of complex numbers (coefficients a, b) or complex vectors (coefficients c). The operator $\odot : \mathcal{Q} \times \widetilde{\mathcal{Q}} \to \mathbb{R}$ is bilinear and not necessarily commutative. We also define addition operators $\oplus : \mathcal{Q} \times \mathcal{Q} \to \mathcal{Q}$ and $\widetilde{\oplus} : \widetilde{\mathcal{Q}} \times \widetilde{\mathcal{Q}} \to \widetilde{\mathcal{Q}}$, distributive with respect to \odot , and shortcuts for summation Σ^{\oplus} and $\Sigma^{\widetilde{\oplus}}$. Finally, there is a multiplication (scaling) operation $\mathbb{R} \times \mathcal{Q} \to \mathcal{Q}$ with the natural semantics.

Our four interpretations of \oplus , \odot and scaling should be clear from (7), (10), (12), (13), i.e., they are standard complex vector operators, with an extra structure. The only subtlety is that while the approximation formulae (7), (10), (12) yield complex numbers in general, in our case the results are real, thanks to the properties of spherical harmonics (4).

Since the error ε in (15) can be bounded by choosing a suitable minimum relative distance λ , the error of the FMM algorithms described later can also be bounded (see also section 3.7).

We shall state that groups A_k and B_l are *well separated*, denoted by $\mathscr{D}_{\mathbf{C}}(A_k, \mathcal{B}_l)$, iff $d(A_k, \mathbf{C}) > \lambda \varrho_{\mathcal{B}_l}(\mathbf{C})$ with distance $d(A_k, \mathbf{C}) = \min_{i \in A_k} d(i, \mathbf{C})$ and radius $\varrho_{\mathcal{B}_l}(\mathbf{C}) = \max_{i \in \mathcal{B}_l} d(j, \mathbf{C})$. It implies that all elements $i \in A_k, j \in \mathcal{B}_l$ are also well separated, $\mathscr{D}_{\mathbf{C}}(i, j)$.

3.2. Single-level FMM algorithms

The simplest of the FMM algorithms is the 'grouping' or 'middle-man' algorithm. It is based on dividing the elements in A and B into spatially constrained cells, typically by partitioning the space into rectangular cells of identical size (Beatson and Greengard 1997). The interaction between far (well-separated) cells A_k , B_l is then carried out using an approximation

$$\sum_{\epsilon \in \mathcal{A}_k} x_i f_{ij} \approx \phi_j(\mathbf{C}_l) \odot \widetilde{\Phi}(\mathcal{A}_k, \mathbf{C}_l) = \phi_j(\mathbf{C}_l) \odot \underbrace{\left(\sum_{i \in \mathcal{A}_k}^{\widetilde{\Phi}} x_i \widetilde{\phi_i}(\mathbf{C}_l)\right)}_{\widetilde{\Phi}(\mathcal{A}_k, \mathbf{C}_l)}, \quad (16)$$

derived from (15), where C_l is the centre of the cell \mathcal{B}_l . This algorithm has complexity of $O(N^{3/2})$, where *N* is the number of elements $||\mathcal{A}|| \approx ||\mathcal{B}||$. We shall use a translation operator T (13) that can convert an outer expansion at point C_k into an inner one around point C_l

$$\phi_i(\mathbf{C}_l) = \mathbf{T}_{\mathbf{C}_k \mathbf{C}_l} \phi_i(\mathbf{C}_k). \tag{17}$$

In this way, we can improve the middle-man algorithm so that the inner fields $\widetilde{\Phi}(\mathcal{A}_k, \mathbf{C}_l)$ can be calculated more efficiently. Instead of computing $\widetilde{\Phi}(\mathcal{A}_k, \mathbf{C}_l)$ for each \mathbf{C}_l , the $\widetilde{\Phi}(\mathcal{A}_k, \mathbf{C}_k)$ have to be calculated only once, and are translated to the centres of all other cells \mathcal{B}_l we want to interact with. The improved algorithm is called a single-level FMM algorithm and has an asymptotic complexity of $O(N^{4/3})$.

3.3. The multi-level FMM algorithm

In order to improve the single-level FMM, we build a hierarchy of cells of different sizes, so that an optimal cell size can be chosen depending on the interaction distance. This leads to a multi-level FMM algorithm, often called simply FMM (Beatson and Greengard 1997). We create trees A, respectively B, from the input set A respectively output set B. Children of each non-leaf cell (tree node) X are themselves cells contained in X. We shall further use an outer-to-outer translation operator R and an inner-to-inner translation operator S (13):

$$\phi_j(\mathbf{C}_l) = \mathbf{R}_{\mathbf{C}_k \mathbf{C}_l} \phi_j(\mathbf{C}_k) \tag{18}$$

$$\hat{\phi}_j(\mathbf{C}_l) = \mathbf{S}_{\mathbf{C}_k \mathbf{C}_l} \hat{\phi}_j(\mathbf{C}_k). \tag{19}$$

The operator R is used to calculate the outer field for each non-leaf cell X in the tree A by summing the outer fields of all its children Y—an *up-sweep*. Similarly, during the *down-sweep*, operator S translates the inner field from non-leaf cells in tree B to their children (figure 1).

3.3.1. The interaction plan. We define a plan \mathcal{P} comprising local and far interactions \mathcal{P}_L and $\mathcal{P}_F; \mathcal{P}_L, \mathcal{P}_F \subseteq \mathcal{A} \times \mathcal{B}$. All pairs $(X, Y) \in \mathcal{P}_L$ are to be treated locally, by explicit summation (14). A pair $(X, Y) \in \mathcal{P}_F$ (for X, Y well separated) indicates that outer field $\Phi(X, \mathbf{C}_X)$, corresponding to all elements in X, must be translated to \mathbf{C}_Y and applied to calculate the contributions of X on all elements in Y.

A plan \mathcal{P} is *well formed* if each interaction between each pair of leaves of trees \mathcal{A} and \mathcal{B} is handled exactly once. In other terms, for each pair of leaf nodes $U \in \mathcal{A}$ and $V \in \mathcal{B}$, there



Figure 1. FMM interactions (for operator S). The outer fields are propagated up the tree A during the up-sweep phase using operator R. They are transferred to tree B and converted to inner fields using operator T. The inner fields are propagated down the tree B using operator S. At leaf cells, the far interactions calculated from the inner-field coefficients are summed with local interactions. (This figure is in colour only in the electronic version)

must be exactly one path going from U up the tree \mathcal{A} to a node X, then to a node $Y \in \mathcal{B}$ such that $(X, Y) \in \mathcal{P}_L \cup \mathcal{P}_F$, and finally down the tree \mathcal{B} to V.

3.3.2. The optimal interaction plan. To minimize the number of local interactions $\|\mathcal{P}_L\|$, we use local interactions exclusively on pairs of leaf cells which are not well separated. We then want to minimize the number of far interactions $\|\mathcal{P}_F\|$. Also, for implementation reasons that will be explained later (section 3.6) we must limit the number of different $T_{C_kC_l}$ operators as measured by the number of different $C_k - C_l$ vectors.

A classical approach is to descend simultaneously from the root to the leaves in both trees \mathcal{A} and \mathcal{B} (which must be identical), and, at each level, to include all valid interactions that have not been already treated at higher levels, allowing only interactions between cells at the same level (Cheng *et al* 1999, Beatson and Greengard 1997).

Our planner (algorithm 1) is more general, not requiring the trees A and B to be identical and allowing interactions between cells at different levels (different sizes). It is based on the additional constraint that for cells $(X, Y) \in \mathcal{P}_F$, the cell X is never smaller than cell Y. As a consequence, X and Y will be almost the same size for X close to Y (but well separated), while cells X further away from Y may be bigger if the separability condition allows it. If the trees share the same top-level bounding box, then 'bigger' (line 12 in algorithm 1) is equivalent to 'at higher level'.

3.3.3. Executing the interaction plan. Once the interaction plan is created, it is executed by algorithm 2 whenever the Au product is needed, i.e., at each iteration. Note that the local-interaction coefficients f_{ij} , the outer expansions ϕ_i , and the operators R, S, T can be pre-calculated. The remaining cost of executing the plan consists in applying the operators R, S and T, and the coefficients f_{ij} .

Algorithm 1: Create an interaction plan for a Multi-Level FMM algorithm **Input**: Trees \mathcal{A} and \mathcal{B} , with cells \mathcal{A}_k and \mathcal{B}_l as leaves. **Output**: An interaction plan $\mathcal{P} = (\mathcal{P}_L, \mathcal{P}_F)$ 1 $\mathcal{P}_L \leftarrow \emptyset$; $\mathcal{P}_F \leftarrow \emptyset$ 2 return CreatePlan({ root of \mathcal{A} }, root of \mathcal{B}) // CreatePlan adds to $\mathcal{P}_F, \mathcal{P}_L$ interactions between a set ${f X}$ (from \mathcal{A}) and a cell Y (from \mathcal{B}) **3** procedure CreatePlan (\mathbf{X}, Y) : 4 $\mathbf{Z} \leftarrow \emptyset$ // Cells from \mathbf{X} to be processed later while $\mathbf{X} \neq \emptyset$ do 5 Take X from **X** 6 if $d(X, \mathbf{C}_Y) > \lambda \, \varrho_Y(\mathbf{C}_Y)$ // are X and Y well-separated? 7 8 then 9 add (X, Y) into \mathcal{P}_F 10 else if $(X \text{ not a leaf}) \land (\varrho_X(\mathbf{C}_X) \ge \varrho_Y(\mathbf{C}_Y) \lor Y \text{ is leaf})$ then 12 put children of X into **X** 13 **else** add X to \mathbf{Z} $\mathbf{14}$ 15 if Y is leaf then add $\{(Z, Y); Z \in \mathbf{Z}\}$ to \mathcal{P}_L 16 17 18 else call CreatePlan(\mathbf{Z}, Y') for all Y' children of Y.

3.3.4. Asymptotic complexity. Asymptotic complexity of the described FMM is O(N) under the following hypotheses which are simple and easy to fulfil.

- \mathscr{H}_1 : there are at most K_c elements in any leaf cell;
- \mathscr{H}_2 : the number of leaf cells from \mathcal{A} near to a given leaf cell from \mathcal{B} is at most K_n ;
- \mathscr{H}_3 : in any sphere of radius *R*, there are at most $K_s(R/\varrho)^{K_d}$ cells of radius $\varrho' > \varrho$;
- \mathscr{H}_4 : for any cell with radius ϱ , the radii of all its children are at least ϱ/K_r , $K_r > 1$.

Here K_c , K_n , K_s , K_d , K_r are constants, independent of the total number of elements N. A sketch of the proof is as follows:

Local interactions can be evaluated in O(N) time, since each element can interact with at most $K_n K_c$ others. Calculating outer fields and propagating the inner fields needs O(N)operations as well, because there are O(N) elements and nodes. Finally, we show that the number of far interactions added to \mathcal{P}_F for each cell in B by the algorithm 1 is at most $K_s(\lambda K_r)^{K_d}$, hence the total number of outer-to-inner translations (operator T), proportional to the number $||\mathcal{P}_F||$, is also O(N).

Unfortunately, due to the large constants involved, the superiority of the FMM over the brute-force approach only appears for large values of *N*.

3.4. Memory complexity

In order to execute algorithm 2 efficiently, we need to pre-calculate and store the local interactions f_{ij} involved in the local plan \mathcal{P}_L and the outer expansions ϕ_i at all elements, corresponding to storing at most $K_c K_n N/2$, respectively $2(4N_v + N_t)(L+1)^2 = 4N(L+1)^2$

Algorithm 2: Execute an interaction plan for a Multi-Level FMM algorithm

Input: Interaction plan $\mathcal{P} = (\mathcal{P}_L, \mathcal{P}_F)$; trees \mathcal{A}, \mathcal{B} ; vector x_i . Precalculated values $f_{ij}, \phi_i, \mathbf{R}, \mathbf{S}, \mathbf{T}$.

Output: A vector y_j so that approximately $y_j = \sum_{i \in \mathcal{A}} x_i f_{ij}$ for all $j \in \mathcal{B}$.

1 $y_j \leftarrow 0$ for all j

// Treat local interactions

- 2 foreach $(X,Y) \in \mathcal{P}_L$ do
- **3 foreach** element $i \in X$, element $j \in Y$ **do 4** $u_i \leftarrow u_i + x_i f_{ii}$
- $\mathbf{4} \qquad \qquad \mathbf{y}_j \leftarrow y_j + x_i f_{ij}$

// Up-sweep. Calculate outer-field recursively 5 for each cell X from tree ${\cal A}$ do

$$\Phi(X, \mathbf{C}_X) = \begin{cases} \sum_{\substack{\text{element } i \in X}}^{\oplus} x_i \phi_i(\mathbf{C}_X) & \text{if } X \text{ is a leaf} \\ \\ \sum_{\substack{X' \text{ child of } X}}^{\oplus} \mathbf{R}_{\mathbf{C}_{X'}\mathbf{C}_X} \Phi(X', \mathbf{C}_{X'}) & \text{otherwise} \end{cases}$$

// Down-sweep.

6 for each cell Y from tree ${\mathcal B}$ do

$$\widetilde{\Phi}(Y, \mathbf{C}_Y) = \underbrace{\mathbf{S}_{\mathbf{C}_Z \mathbf{C}_Y} \widetilde{\Phi}(Z, \mathbf{C}_Z)}_{Z \text{ parent of } Y} \quad \widetilde{\oplus} \sum_{(X,Y)\in\mathcal{P}_F}^{\widetilde{\oplus}} \mathbf{T}_{\mathbf{C}_X \mathbf{C}_Y} \Phi(X, \mathbf{C}_X)$$
if Y is leaf **then foreach** element $j \in Y$ **do**

$$\int_{\mathbb{C}} \widetilde{\Phi}(Y, \mathbf{C}_Y) \otimes \widetilde{\Phi}(Y, \mathbf{C}_Y)$$

real numbers, where N_v and N_t are the total number of vertices and triangles of the mesh, respectively. Furthermore, we need to store the pre-calculated spherical harmonics for operators R, S, T, the biggest one being operator T needing $2(2L + 1)^2$ real values per distinct $\mathbf{M} - \mathbf{N}$ vector.

3.5. Timing of multipolar representation

The most expensive operations (table 1) are calculating the spherical harmonic expansions and applying the translation operator T. The critical problem size above which it is advantageous to apply FMM is very big. In a very simple case of two well-separated groups of elements interacting 100 times, we need more than 200 elements in each group—this observation can guide our choice of K_c . For real geometries, not all cells are well separated; we can estimate that the FMM only starts to be competitive for problems with more than 10^4 elements.

3.6. The tree structure

A *median tree* is an ideally balanced binary tree, based on splitting node elements using axisparallel planes into equal halves. However, it requires calculating (and storing) the translation operators for all pairs of interacting cells, which soon becomes prohibitive.

Table 1. Timings (on a standard 1.4 GHz Pentium PC) of some operations involved in the FMM for L = 10, sorted by the elapsed time. By a we indicate the scalar expansion coefficients, while c stands for the vector expansion coefficients needed to approximate operator \mathcal{N} .

Operation to calculate	Time (μs)
a, b, c for 1 triangle	4218.75
Apply T for c	3183.59
Apply T for a	1054.69
Apply R for c	933.10
Apply S for c	788.92
Apply R for a	310.06
Apply S for a	261.23
O_n^m for T	83.62
D _{ij'} directly	47.00
S _{ij} directly	43.34
I_n^m for R or S	22.13
⊙ for c	12.28
⊙ for a	3.66
$y \leftarrow y + x \text{ for } y \in \mathbb{R}$	< 0.10
$y \leftarrow y + x$ for a	< 0.10
$y \leftarrow y + x$ for c	< 0.10
$\alpha \mathtt{a}, \alpha \in \mathbb{R}$	< 0.10

For this reason, we have adopted a classical adaptive *octtree* structure. A bounding box of all elements becomes a cell box of the root element. At each level, a parent cell box is divided into eight identical subboxes. Each of the eight children then receives elements whose centre of gravity falls into its box. The subdivision is stopped for nodes with less than K_c elements. Empty branches are pruned. Expansion centres C_X are put into geometrical centres of each cell box. Apart from the cell boxes, each cell also has a tight bounding box, used to determine separateness.

An octtree may not be balanced. It is usually shallow, for example, for our spherical head model with $N = 71\,686$ elements only five levels are needed with $K_c = 100$. Its major advantage is that the expansion centres are guaranteed to lie on a Cartesian grid with known spacing. Therefore only a limited number of distinct translation operators need to be pre-computed. For our model with $N = 71\,686$ only 3776 operators T are needed.

3.7. Choice of parameters

The choice of λ and *L* is guided by time and accuracy considerations. The truncation error of (4) is proportional to $(\|\mathbf{r} - \mathbf{C}\| / \|\mathbf{r}' - \mathbf{C}\|)^{L+1}$ which is bounded by $\lambda^{-(L+1)}$. In figure 2 (left) we show the relative accuracy of the approximation of $1/\|\mathbf{r}' - \mathbf{r}\|$ using (4) for various values of *L* as a function of the relative distance $\|\mathbf{r}' - \mathbf{C}\| / \|\mathbf{r} - \mathbf{C}\|$ for 10^4 random points.

For octtree-type cells the minimum useful value of λ is $\lambda_{\min} = \sqrt{3} \approx 1.73$ with a localinteraction neighbourhood of $3^3 = 27$ cells. We have tested various values of λ and *L* for our three-layer sphere models. The optimal value of λ was always between λ_{\min} and 3. In order to limit the amount of memory needed we therefore decided to set $\lambda = 2$, in agreement with Dembart and Yip (1998).

There is no consensus about what accuracy is fundamentally required to calculate the MEG/EEG BEM interactions for the inverse (source identification) problem, due to the measurement and modelling errors. We have decided to require that the difference between



Figure 2. Left, from top to bottom: the relative error to approximate $1/||\mathbf{r}' - \mathbf{r}||$ using (4), as a function of the relative distance, for 10⁴ random points, and for L = 5, 10, 15. Right, from top to bottom: the relative error to approximate elements of $D_{i'j}$, for 10⁴ random triangles and L = 10, as a function of the relative distance using three quadrature rules (QR): 16 point symmetric triangle QR, 25 point product Gauss QR, 144 point product Gauss QR.

the FMM and non-FMM implementations be less than 1% of the BEM discretization error (known for our spherical models). This corresponds to calculating the BEM interactions with a relative accuracy of about 10^{-4} which in turn required us to set L = 10.

3.8. Accuracy of the operator approximation

In order to choose the appropriate numerical quadrature procedure to implement (5), (8), we evaluated the relative error of approximating S_{ij} , $D_{i'j}$, $N_{i'j'}$ as a function of relative distance between the elements (triangles) and the quadrature rule (QR) used, figure 2, right. Integration of spherical harmonics at a desired level of accuracy requires a $5 \times 5 = 25$ point tensor product

Gauss quadrature. Increasing the order further does not bring any improvement. A 16 point QR is used for the direct computation (section 2.1) (Kybic *et al* 2005, Kybic *et al* 2004).

3.9. Interaction between several trees

While the standard FMM only considers one tree, for the symmetric BEM we need to treat multiple interacting trees because each surface is handled separately and we have two types of variables: potential V(P1) at vertices and flow p(P0) at faces. For external surfaces only the P1 tree is built, since the flow p is known to be zero there (Kybic *et al* 2005, Kybic *et al* 2004). First, an up-sweep phase is performed separately for each tree and outer fields are stored. Then a down-sweep phase is performed for each pair of trees that corresponds to surfaces delimiting a common volume, i.e., for all non-zero blocks in the system matrix A (Kybic *et al* 2005, Kybic *et al* 2004). All elements (vertices and faces) have a globally unique identification number that becomes an index of the corresponding variable.

To treat the multiple tree interactions efficiently, the direct interactions S, D, N are shared across all trees, as well as the R, S and T operators—this requires sharing a common grid. Evaluating expansion coefficients a, b, c also shares many intermediate results.

4. Experiments

The superior accuracy of the symmetric BEM was already demonstrated in Kybic *et al* (2005). The purpose of this section is to demonstrate that this accuracy is not compromised by the FMM acceleration proposed, and that it allows us to treat on a single computer, in a reasonable amount of time, far larger problems then the direct (non-accelerated) implementation. We have repeated the experiments from Kybic *et al* (2005) and verified that the relative ℓ_2 difference between the non-accelerated implementation and the new FMM implementation with $\lambda = 2$, L = 10 is less than 10^{-3} , which is less than the error of the BEM method itself.

We have used spherical head models, as in Kybic *et al* (2005), since analytical solutions are available for them. They consist of three concentric spheres with radii 0.87, 0.92 and 1.0, delimiting volumes with conductivities 1.0, 0.0125, 1.0 and 0.0, from inside towards outside. Different resolution meshes were used with 6422 562 and 10 242 vertices per surface, corresponding to a total number of unknowns for the symmetric BEM equal to 4486, 17926 and 71 686, respectively. A dipolar source was placed at distance 0.425 from the centre.

The experiments were performed on a computer with a 1.6 GHz 64 bit AMD Opteron processor with 5 GB of physical memory.

4.1. Single-sphere head models

The first series of experiments (tables 2 and 3), were performed on a simplified head model containing one surface only, for a meaningful comparison with timings reported by Tissari and Rahola (2003) using pre-corrected-FFT method. The programs were run with the maximum number of elements per cell $K_c = 200$, expansion order L = 10, minimum relative distance $\lambda = 1.7$, and MINRES relative stopping threshold $\varepsilon = 10^{-6}$. The results of Tissari and Rahola are reported with parameters p = 3, 4 since this seems to correspond to the precision required. Non-accelerated direct method results are also shown.

According to the timings for the direct problem, their computer and implementation seem to be comparable to ours for the same number of unknowns, even though they use the double-layer formulation. It apparently has the advantage of involving a well-conditioned matrix with easy preconditioning, never requiring more than 6 or 7 iterations of the optimizer

Table 2. The elapsed CPU time and number of iterations needed to solve the forward problem for the single-sphere models as a function of the number of unknowns for the direct (no FMM) and FMM-accelerated symmetric BEM methods. Time needed to solve the problem using the double-layer direct and pre-corrected-FFT method reported by Tissari and Rahola (2003) for parameters p = 3/p = 4 is marked with^{*}.

	Time (s)				Iterations		
Unknowns	Direct symmetric	Direct double layer*	FMM	Pre-calculated FFT*	FMM	Pre-calculated FFT*	
362		15		32/165		≼7	
642	63		69		12		
1002		123		98/539		≼7	
2252		714		270/1336		≼7	
2562	1309		812		17		
5762		4760		724/4012		≼7	
9002		12715		1223/6706		≼7	
10 242	26 0 60		6325		26		
12962		27 685		1934/9860		≼7	

Table 3. The memory requirements (in MB) to solve the forward problem for the single-sphere head models as a function of the number of unknowns for the direct (no FMM) and FMM-accelerated symmetric BEM methods (the actual memory usage fluctuates due to the garbage collector). Memory requirements for the double-layer direct and pre-corrected-FFT method reported by Tissari and Rahola (2003) for parameters p = 3/p = 4 are marked with^{*}.

Unknowns	Direct symmetric	Direct double layer*	FMM	Pre-calculated FFT*
362		20		22/25
642	65		21	
1002		34		32/40
2252		98		66/71
2562	288		95	
5762		539		126/169
9002		1288		202/268
10 242	2546		881	
12 962		2649		296/389

(Tissari and Rahola 2003). This emphasizes the assembly time with respect to the matrixvector product evaluation time. For the largest problem, the assembly time is over 95% of the total time, each iteration takes only about 10 s. This means that once the preprocessing is done, subsequent calculations for different sources can be relatively fast (<5 min). Our FMM algorithm is always faster than the pre-corrected FMM for p = 4 and they are comparable even for p = 3. On the other hand, the FMM algorithm seems to need more memory. Note that in this case (the low number of iterations), the FMM version is faster and uses less memory than the direct version even for moderately sized problems and the subquadratic time complexity shows nicely.

The only other published FMM implementation we have found for the symmetric BEM is by Of *et al* (2002). After correcting for their use of only 7 point integration rules (the degree of their spherical harmonic expansion was not reported), their timings are comparable to ours as well.

Table 4. The elapsed CPU time and number of iterations needed to solve the forward problem for the three-sphere models as a function of the number of unknowns for the direct (no FMM) and FMM-accelerated symmetric BEM methods. Relative error with respect to the analytical solution is also reported. We put between parentheses the extrapolated value for the largest problem, which could not be solved by the direct method due to lack of memory.

	Time (s)		Iterations		Relative error (%)	
Unknowns	Direct symmetric FMM		Direct symmetric FMM		Direct symmetric FMM	
4486	2628	3030	238	238	0.989	0.989
17 926	33 928	70378	384	625	0.252	0.245
/1686	(542575)	453 600	N/A	900	N/A	0.090

Table 5. The memory requirements (in MB) to solve the forward problem for the three-sphere head models as a function of the number of unknowns for the direct (no FMM) and FMM-accelerated symmetric BEM methods in MB. We put between parentheses the extrapolated value for the largest problem, which could not be solved by the direct method due to lack of memory.

Unknowns	Direct symmetric (MB)	FMM		
4486	91	123		
17 926	1390	1106		
71 686	(22 229)	4400		

4.2. Three-sphere head models

The results shown in tables 4 and 5 correspond to $K_c = 300$, L = 10, $\lambda = 1.7$ and $\varepsilon = 10^{-6}$. We observe that no FMM acceleration takes place for the smallest mesh (N = 4486). For the middle one (N = 17926), FMM brings a memory saving but is still slower. For the largest mesh, FMM enables us to produce a valid result, which could not be calculated by the direct method for lack of memory. The assembly time is about 80% of the total time, each iteration takes about 1.5 min. Calculation for another source would therefore take about 22 h.

5. Discussion

The finite element method (FEM) implementation (Clerc *et al* 2002) is capable of solving a problem of this size but requires a well-formed and topologically correct 3D mesh which is difficult to automatically generate from the head MRI scans or the surface meshes⁴.

FMM performs much better for the single-sphere model compared to the three-sphere model, mainly because the conditioning of the system for three-sphere model system is one order of magnitude worse, due to the low conductivity of the middle layer corresponding to the skull. Moreover, because the three surfaces are close to one another, there are many more elements to be treated locally. Consequently, the subquadratic time complexity only manifests itself for very large problems. Nevertheless, the FMM enables us to solve middle-sized problems, which barely fit in the memory of current computers. Indeed, the most critical point in implementing the FMM turns out to be memory management. Most of the memory is used to store the pre-calculated local interactions.

⁴ Most 3D meshing software is commercial (not freely available) and does not support adaptive meshing, making the resulting models extremely large.

Our FMM is tailored for the symmetric BEM and uses several mutually interacting trees of two different types (P0/P1). Caching and a predetermined interaction plan are established to eliminate overhead.

From the timings (section 3.5) and by analysing the required number of most costly basic operations, we conclude that there is unfortunately little hope of significantly algorithmically accelerating the FMM for small or medium BEM problems, unless computationally more efficient expansions and translations can be established. Replacing spherical harmonics by pseudoparticles (Makino 1999) or Cartesian polynomials (Apalkov and Visscher 2003) might be an alternative to consider.

6. Conclusion

We have developed a fast multipole method to accelerate the symmetric BEM, with application to the forward MEG/EEG problem. The FMM is as accurate as the symmetric BEM with direct assembly, and with increasing problem size it gets faster and requires less memory than the direct method.

Acknowledgment

Sponsored by the Czech Ministry of Education under Project no MSM6840770012.

Appendix. Spherical harmonics

There are several definitions of spherical harmonics and Legendre polynomials, differing mainly in normalization and sign convention, and we follow the convention of Epton and Dembart (1995). The spherical harmonic $Y_n^m(\theta, \phi)$ is

$$Y_n^m(\theta, \phi) = \sqrt{\frac{(n - |m|)!}{(n + |m|)!}} (-1)^m P_n^{|m|}(\cos \theta) e^{im\phi}$$

where $P_n^m(x)$ are the associated Legendre polynomials

$$P_n^m(x) = \frac{1}{2^n n!} \frac{(n+m)!}{(n-m)!} (1-x^2)^{-m/2} \frac{\mathrm{d}^{n-m}}{\mathrm{d}x^{n-m}} (x^2-1)^n.$$

Given a vector $\mathbf{x} = (x, y, z) = (r \cos \phi \sin \theta, r \sin \phi \sin \theta, r \cos \theta)$, the inner and outer spherical harmonics are defined by

$$O_n^m(\mathbf{x}) = \frac{(-1)^n i^{|m|}}{A_n^m} \frac{Y_n^m(\theta, \phi)}{r^{n+1}}$$

$$I_n^m(\mathbf{x}) = i^{-|m|} A_n^m r^n Y_n^m(\theta, \phi) \quad \text{with} \quad A_n^m = \frac{(-1)^n}{\sqrt{(n-m)!(n+m)!}}.$$

References

Apalkov D and Visscher P 2003 Fast multipole method for micromagnetic simulation of periodic systems IEEE Trans. Magn. 39 3478–80

Beatson R K and Greengard L 1997 A short course on fast multipole methods *Wavelets, Multilevel Methods and Elliptic PDEs* ed M Ainsworth, J Levesley, W Light and M Marletta (Oxford: Oxford University Press) pp 1–37 (http://www.math.canterbury.ac.nz/mathrkb/pdfs/beatson+greengard/beatso ngreengard.pdf)

Börm S and Hackbusch W 2003 Introduction to hierarchical matrices with applications *Eng. Anal. Bound. Elem.* 27 405–22

Cheng H, Greengard L and Rokhlin V 1999 A fast adaptive multipole algorithm in three dimensions *J. Comput. Phys.* **155** 468–98 (http://amath.colorado.edu/courses/7400/2003fall/005/papers/ChengGreenga rdRokhlin.pdf)

Clerc M, Dervieux A, Faugeras O, Keriven R, Kybic J and Papadopoulo T 2002 Comparison of BEM and FEM methods for the E/MEG problem *Proc. BIOMAG*

Clerc M, Keriven R, Faugeras O, Kybic J and Papadopoulo T 2002 The fast multipole method for the direct E/MEG problem Proc. ISBI (Washington, DC: IEEE–NIH) (http://www.biomedicalimaging.org/)

Dembart B and Yip E 1998 The accuracy of fast multipole methods for Maxwell's equations *IEEE Comput. Sci. Eng.* **5** 48–56 (http://dx.doi.org/10.1109/99.714593)

Epton M A and Dembart B 1995 Multipole translation theory for the three-dimensional Laplace and Helmholtz equations *SIAM J. Sci. Comput.* **16** 865–97

Ferguson A S and Stroink G 1997 Factors affecting the accuracy of the boundary element method in the forward problem: I. Calculating surface potentials *IEEE Trans. Biomed. Eng.* **44** 1139–55

Hämäläinen M, Hari R, IImoniemi R J, Knuutila J and Lounasmaa O V 1993 Magnetoencephalography—theory, instrumentation, and applications to noninvasive studies of the working human brain *Rev. Mod. Phys.* **65** 413–97

Kybic J and Clerc M 2003 Symmetric BEM and multiscale fast multipole method for the E/MEG problem *NFSI 2003*: *Proc. 4th Int. Symp. on Noninvasive Functional Source Imaging Within the Human Heart and Brain (Berlin)* ed V Pizzella and G L Romani pp 122–4

Kybic J, Clerc M, Abboud T, Faugeras O, Keriven R and Papadopoulo T 2005 A common formalism for the integral formulations of the forward EEG problem *IEEE Trans. Med. Imaging* **24** 12–28

Kybic J, Clerc M, Faugeras O, Keriven R and Papadopoulo T 2004 Fast multipole method for the symmetric boundary element method in meg/eeg *INRIA Technical Report* 5415

Makino J 1999 Yet another fast multipole method without multipoles—pseudoparticle multipole method *J. Comput. Phys.* **151** 910–20

Of G, Steinbach O and Wendland W L 2002 A fast multipole boundary element method for the symmetric boundary integral formulation *Proc. IABEM (Austin, TX, USA)* (http://cavity.ce.utexas.edu/iabem2002/fullpapers/of.pdf)

Phillips J, Leahy R, Mosher J and Timsari B 1997 Imaging neural activity using MEG and eEG *IEEE Eng. Med. Biol.* Mag. 16 34–42

Rahola J 1998 Experiments on iterative methods and the fast multipole method in electromagnetic scattering calculations *CERFACS Technical Report* TR/PA/98/49

Rahola J and Tissari S 1998 Iterative solution of dense linear systems arising from boundary element formulations of the biomagnetic inverse problem *CERFACS Technical Report* TR/PA/98/40, toulouse, France

Rahola J and Tissari S 2002 Iterative solution of dense linear systems arising from the electrostatic integral equation *Phys. Med. Biol.* **47** 961–75

Tissari S and Rahola J 2003 A precorrected-FFT method to accelerate the solution of the forward problem in magnetoencephalography *Phys. Med. Biol.* **48** 523–41