Pierre-Henri Tournier, Marcella Bonazzoli, Victorita Dolean, Francesca Rapetti, Frédéric Hecht, Frédéric Nataf, Iannis Aliferis, Ibtissam El Kanfoud, Claire Migliaccio, Maya de Buhan, Marion Darbas, Serguei Semenov, and Christian Pichot

Numerical Modeling and High-Speed Parallel Computing

New perspectives on tomographic microwave imaging for brain stroke detection and monitoring.

This article deals with microwave tomography for brain stroke imaging using state-of-the-art numerical modeling and massively parallel computing. Iterative microwave tomographic imaging requires the solution of an inverse problem based on a minimization algorithm (e.g., gradient based) with successive solutions of a direct problem such as the accurate modeling of a whole-microwave measurement system. Moreover, a sufficiently high number of unknowns is required to accurately represent the solution. As the system will be used for detecting a brain stroke (ischemic or hemorrhagic) as well as for monitoring during the treatment, the running times for the reconstructions should be reasonable. The method used is based on high-order finite elements, parallel preconditioners from the domain decomposition method and domain-specific language with the open-source FreeFEM++ solver.

THE NEED FOR RAPID STROKE DIFFERENTIATION

A stroke, or a cerebrovascular accident, is classically characterized as a neurological deficit attributed to an acute focal injury of the central nervous system by a vascular cause, including a cerebral infarction, an intracerebral hemorrhage, or a subarachnoid hemorrhage, and it is a major cause of disability and death worldwide [1]. About 85% of strokes are ischemic due to a cerebral infarction, caused by an interruption of the blood supply to some part of the brain, and 15% are hemorrhagic (10% primary and 5% subarachnoid hemorrhage) [2]. Differentiating between these different types of strokes is an essential part of the initial workup of patients because the subsequent management and treatment of each patient is vastly different. Rapid and accurate diagnosis is crucial since the only drug currently approved by the U.S. Food and Drug Administration for treatment of acute ischemic strokes is the intravenous tissue plasminogen activator, which is administered within 3 h of the stroke onset. Neuroimaging has to play a vital role in the workup of an acute stroke by providing information essential to accurately triage patients and expedite clinical decision making with regards to treatment. Computed tomography (CT) and magnetic resonance imaging (MRI) [3] are the gold standards, but they are bulky diagnostic instruments and cannot be used in continuous brain monitoring. A noninvasive and transportable or portable device would have clear clinical applications at the bedside in a neurological intensive care unit.

Microwave tomography is a novel, early-development stage imaging modality with a large number of potential

Digital Object Identifier 10.1109/MAP.2017.2731199
Date of publication: 22 August 2017

attractive medical applications. A difference between the dielectric properties (complex permittivity) of normal and diseased brain tissues is a great potential for this imaging modality. Detecting and identifying strokes is challenging, as it corresponds to a small opposite variation of the permittivity values of brain tissues of about +/–10% of the baseline tissue values for the two types of strokes (ischemic or hemorrhagic) [4]. The rapid data acquisition time is another attractive feature of microwave tomography, but rapid tomographic reconstructions are mandatory for developing a novel imaging modality with a new paradigm: detecting, identifying, and monitoring a stroke continuously during treatments by exposing head tissues to low-level microwave incident field and capturing the scattered signal by an array of antennas. Iterative tomographic imaging requires the solution of an inverse problem based on a minimization algorithm. Reconstruction algorithms are computationally intensive with successive solutions of the forward problem needing efficient numerical modeling and high-performance parallel computing. A majority of works in the literature has made use of geometrically simple phantoms or with only a limited amount of tissue-mimicking materials. The modeling must have to accurately take account of the high heterogeneity and complexity of head tissues (skin, fat, skull, bone marrow, brain/white matter, brain/grey matter, cerebrospinal fluid, arteries, and more) for normal cases and different possible brain pathology cases (ischemic and hemorrhagic strokes, brain injuries, and more). Another major point refers to the accurate modeling of the incident field from transmitting and receiving antennas. This interaction is very complex because it must be seen as a coupling problem between the antennas and the head rather than a simple scattering problem. In addition, the electric field is measured by means of receiving antennas (sensors). Therefore, we do not have access directly to the electric field but only via antenna S parameters. The purpose of this article is to solve the inverse problem associated to a prototype developed by EMTensor GmbH, Vienna, Austria, [5] using state-of-the-art modeling and high-performance and massively parallel computing.

THE TOMOGRAPHIC SYSTEM

The model of microwave imaging is based on Brain Imaging Generation 1 (BRIMG1), a tomographic microwave system developed by EMTensor GmbH [5]. The system consists of a cylindrical metallic chamber composed of five rings of 32 transmitting/receiving antennas (Figure 1). The antennas are ceramic ($\varepsilon_r = 59$), loaded, open-ended waveguides. The diameter of the chamber is 285 mm with a height of 280 mm. The rings are 30 mm, equally spaced, with the first one located 40 mm from the top of the chamber. The chamber is filled with a matching liquid medium during measurements. The operating frequency of the system is 0.9–1.8 GHz. The data acquisition cycle of the system is fully electronically controlled, allowing for a total data acquisition of about 30 s. The imaging chamber is in a horizontal position, allowing easy positioning of a human head within an imaging domain (Figure 2). The patient's head is entered into

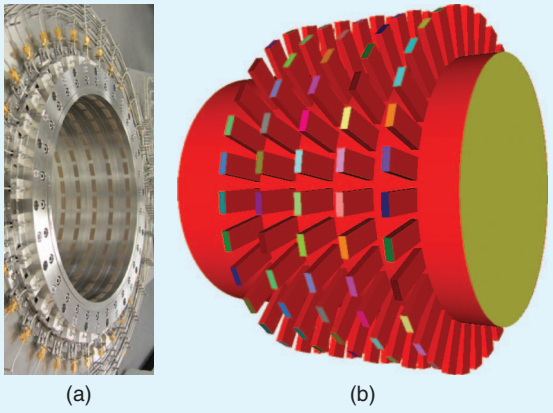

FIGURE 1. (a) A general view of BRIMG1. (Photo courtesy of EMTensor GmbH.) (b) The computational domain.

FIGURE 2. The BRIMG1: A human head measurement. (Photo courtesy of EMTensor GmbH.)

the chamber, as shown in Figure 2. A special thin membrane is used for isolating the human head from the matching liquid and keeping the liquid within the chamber. Carbon-loaded silicone rubber (CLSR) is also used to reduce reflection from boundary conditions (Figure 3).

A switching matrix connected to a network analyzer selects the transmitting and receiving antennas. The system is potentially delivering a 160 × 160 matrix of S parameters. The measured S parameters due to the scattered field of an object under investigation are obtained by complex subtraction between two measurements with an empty chamber and with the head,

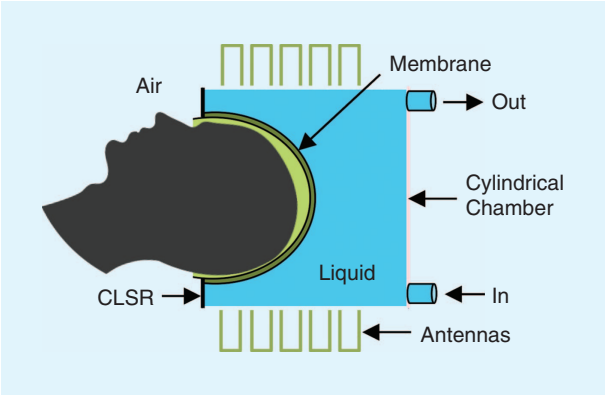

FIGURE 3. The BRIMG1: A side sketch. (Image courtesy of EMTensor GmbH.)

respectively. The raw data can be wirelessly transferred to a remote computing center.

The high-performance computing (HPC) machine will compute the tomographic images, which can be quickly transferred from the computing center to the hospital.

FORWARD MODELING

We consider the domain $\Omega \subset R^3$ for representing the whole-chamber (Figure 1) as an inhomogeneous, dissipative, nonmagnetic medium of complex permittivity, $\varepsilon(x) = \varepsilon'(x) + i\varepsilon''(x)$. For each transmitting antenna, $j = 1, 2, \ldots, N$, at radial frequency ω, the wave equation for the electric field vector, $E_j(x)$, with an $e^{-i\omega t}$ time dependence is

$$\nabla \times (\nabla \times E_j) - k^2 E_j = 0 \qquad \text{in } \Omega, \tag{1}$$

with $k^2 = k^2(x) = \omega^2 \varepsilon(x) \mu_0 = \omega^2 \varepsilon_r(x) \varepsilon_0 \mu_0 = k_0^2 \varepsilon_r(x)$, where $k(x)$ is the complex wavenumber of the inhomogeneous medium, where ε_0, μ_0, and k_0 are the permittivity, permeability, and wavenumber of free space, respectively, and where $\varepsilon_r(x)$ is the relative complex permittivity.

The boundary conditions on the perfectly conducting parts Γ_c of the walls of the chamber are

$$E_j \times n = 0 \quad \text{on } \Gamma_c, \tag{2}$$

where n is the unit outward normal to $\partial\Omega$.

The impedance boundary conditions on the aperture of the transmitting, open-ended waveguide, j, and the receiving waveguide, $i = 1, 2, \ldots, N$, $j \neq i$, are

$$(\nabla \times E_j) \times n - i\beta n \times (E_j \times n) = g_j \qquad \text{on } \Gamma_i, \tag{3}$$

$$(\nabla \times E_j) \times n - i\beta n \times (E_j \times n) = 0 \qquad \text{on } \Gamma_i,\ i \neq j, \tag{4}$$

where β is the propagation constant of the TE_{10} fundamental mode of the waveguide. In (3), we impose an incident wave corresponding to the excitation of the fundamental mode E_j^0 of the jth waveguide with

$$g_j = (\nabla \times E_j^0) \times n - i\beta n \times (E_j^0 \times n). \tag{5}$$

On the other hand, (4) corresponds to a first-order, Silver–Müller absorbing boundary condition, approximating a transparent boundary condition on the aperture of the receiving waveguide antenna, $i = 1, 2, \ldots, N$, $j \neq i$. On the bottom of the chamber, we impose a metallic boundary condition, whereas we impose an impedance boundary condition on the top of the chamber. As a result, the whole boundary value problem for each transmitting antenna, $j = 1, \ldots, N$, is to find E_j such that

$$\begin{cases} \nabla \times (\nabla \times E_j) - k^2 E_j = 0 & \text{in } \Omega \\ E_j \times n = 0 & \text{on } \Gamma_c \\ (\nabla \times E_j) \times n - i\beta n \times (E_j \times n) = g_j & \text{on } \Gamma_j \\ (\nabla \times E_j) \times n - i\beta n \times (E_j \times n) = 0 & \text{on } \Gamma_i,\ i \neq j \end{cases}. \tag{6}$$

Now, let $V = \{v \in H(\text{curl}, \Omega), v \times n = 0 \text{ on } \Gamma_c\}$, where $H(\text{curl}, \Omega) = \{v \in L_2(\Omega)^3, \nabla \times v \in L_2(\Omega)^3\}$ is the space of square integrable functions whose curl is also square integrable. For each transmitting antenna, $j = 1, \ldots, N$, the variational form of problem (6) is: find $E_j \in V$ such that

$$\int_\Omega [(\nabla \times E_j).(\nabla \times v) - k^2 E_j.v] - \int_{\bigcup_{i=1}^N \Gamma_i} i\beta (E_j \times n).(v \times n) = \int_{\Gamma_j} g_j.v \qquad \forall v \in V. \tag{7}$$

HIGH-ORDER EDGE FINITE ELEMENTS

For using a finite element discretization of the variational problem, we introduce a tetrahedral mesh T_h of the domain Ω and a finite dimensional subspace, $V_h \subset H(\text{curl}, \Omega)$. A simple conformal discretization for space $H(\text{curl}, \Omega)$ is given by low-order, Nédélec edge finite elements of polynomial degree $r = 1$ [6].

To have a higher numerical accuracy with the same total number of unknowns, we consider a high-order edge element discretization, choosing the high-order extension of Nédélec elements presented in [7].

We implemented edge elements of degrees $r = 2$ and $r = 3$ in FreeFem++, an open-source domain-specific language specialized for solving boundary value problems by using variational discretizations (finite elements, discontinuous Galerkin, hybrid methods, and more) [8]. High-order elements can be used by loading the plugin *Element Mixte3d* and declaring the finite element space *fespace* using the keywords *Edge13d* and *Edge23d*, respectively (standard edge elements of degree $r = 1$ are already present in FreeFem++ and called *Edge03d*).

DOMAIN DECOMPOSITION PRECONDITIONING

The discretization of the problem presented in the section "Forward Modeling" using the high-order edge finite elements described in the section "High-Order Edge Finite Elements" produces a linear system:

$$Au_j = b_j, \tag{8}$$

for each transmitting antenna j. Direct solvers are not suited for such large linear systems arising from complex three-dimensional (3-D) models because of their high memory cost.

However, matrices resulting from high-order discretizations are ill conditioned as shown numerically in [6] for similar problems, and preconditioning becomes necessary when using iterative solvers. Parallel solvers for electromagnetic (EM) problems can be based, for example, on parallel fast Fourier transforms [9] or nonoverlapping ([10] and references therein) or overlapping domain decomposition, which are considered in this article.

Domain decomposition preconditioners are naturally suited to parallel computing and make it possible to deal with smaller subproblems [11]. The domain decomposition preconditioner that we employ is called *Optimized Restricted Additive Schwarz* (*ORAS*) [12]:

$$M_{\mathrm{ORAS}}^{-1} = \sum_{s=1}^{N_{\mathrm{sub}}} R_s^T D_s A_s^{-1} R_s, \tag{9}$$

where N_{sub} is the number of overlapping subdomains, Ω_s, into which the domain, Ω, is decomposed (Figure 4). Here, the matrices, A_s, are the local matrices of the subproblems with impedance boundary conditions $(\nabla \times E) \times n - i\omega n \times (E \times n)$ as transmission conditions at the interfaces between subdomains. This preconditioner is an extension of the restricted additive Schwarz method proposed by Cai and Sarkis [13] but with more efficient transmission conditions between subdomains than Dirichlet conditions [14].

To describe the matrices R_s and D_s, let N be an ordered set of the unknowns of the whole domain, and let $N = \cup_{s=1}^{N_{\mathrm{sub}}} N_s$ be its decomposition into the (nondisjoint) ordered subsets corresponding to the different (overlapping) subdomains Ω_s. The matrix R_s is the restriction matrix from Ω to the subdomain Ω_s: it is a $N_s \times N$ Boolean matrix, and its (i, j) entry is equal to one if the ith unknown in N_s is the jth one in N. Notice that R_s^T is then the extension matrix from the subdomain Ω_s to Ω. The matrix D_s is an $N_s \times N_s$ diagonal matrix that gives a discrete partition of unity, i.e., $\Sigma_{s=1}^{N_{sub}} R_s^T D_s R_s = I$; in particular, the matrices D deal with the unknowns that belong to the overlap between subdomains. The preconditioner without the partition of unity matrices D_s, $M_{OAS}^{-1} = \Sigma_{s=1}^{N_{sub}} R_s^T A_s^{-1} R_s$, which is called *Optimized Additive Schwarz*, would be symmetric for symmetric problems, but, in practice, it gives a slower convergence with respect to M_{ORAS}^{-1} [15]. These domain decomposition preconditioners are implemented in the HPDDM library [16], an open-source high-performance unified framework for domain decomposition methods. HPDDM can be interfaced with various programming languages and open-source finite element libraries such as FreeFem++, which we use in the simulations.

NUMERICAL RESULTS

COMPARISON WITH EXPERIMENTAL MEASUREMENTS

The measured physical quantities are the S parameters of the scattering matrix, which are the complex reflection and transmission coefficients measured by the 160 receiving antennas when a signal is transmitted by one of the 160 transmitting antennas. A set of measurements then consists in a complex matrix of size 160 × 160. To compute the numerical counterparts of these reflection and transmission coefficients, we use the following formula, which is appropriate in the case of open-ended waveguides:

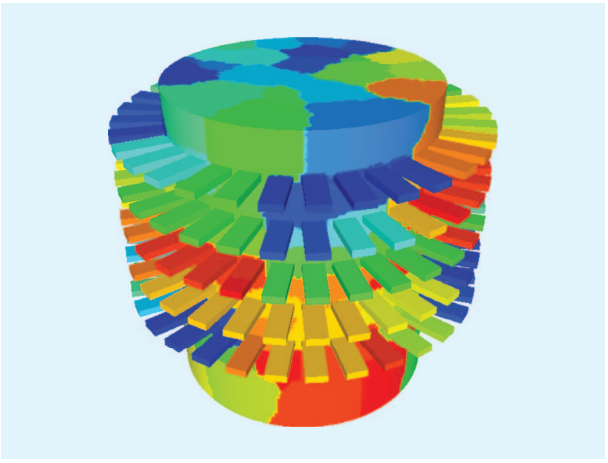

FIGURE 4. The computational domain divided into 128 subdomains.

$$S_{ij} = \frac{\int_{\Gamma_i} \tilde{E}_j . E_i^0}{\int_{\Gamma_i} |E_i^0|^2}, \; i \neq j, \tag{10}$$

where E_j is the solution of the problem (6) when the jth waveguide antenna transmits the signal, and E_i^0 is the TE_{10} fundamental mode of the ith receiving waveguide ($\tilde{E}_j$ denotes the complex conjugate of E_j). The S_{ij} with $i \neq j$ denote the transmission coefficients, and S_{ii} the reflection coefficients.

For a comparison of the computed coefficients S_{ij} with the measured ones, the imaging chamber is filled with a homogenous matching solution to reduce the return loss of the ceramic-loaded waveguide antennas and to match with the average brain tissues. The relative complex permittivity of the matching solution chosen for the experiments and numerical solution at frequency $f = 1$ GHz is $\varepsilon_r^{\mathrm{matching}} = 44 + i20$. The relative permittivity inside the ceramic-loaded waveguides is $\varepsilon_r^{\mathrm{ceramics}} = 59$, assuming a lossless ceramic material.

For this test case, the set of experimental data in S parameters consists of the 160 receiving antennas when each antenna from the second ring from the top is transmitting. Figure 4 shows the normalized magnitude (in decibels) and phase (in degrees) of the complex coefficients, S_{ij}, corresponding to a transmitting antenna in the second ring from the top and to the 31 receiving antennas in the middle ring (note that measured coefficients are available only for 17 receiving antennas). The computed coefficients are obtained by solving the direct problem with edge finite elements of polynomial degree, $r = 2$. The normalization is done by dividing every transmission coefficient by the transmission coefficient corresponding to the receiving antenna

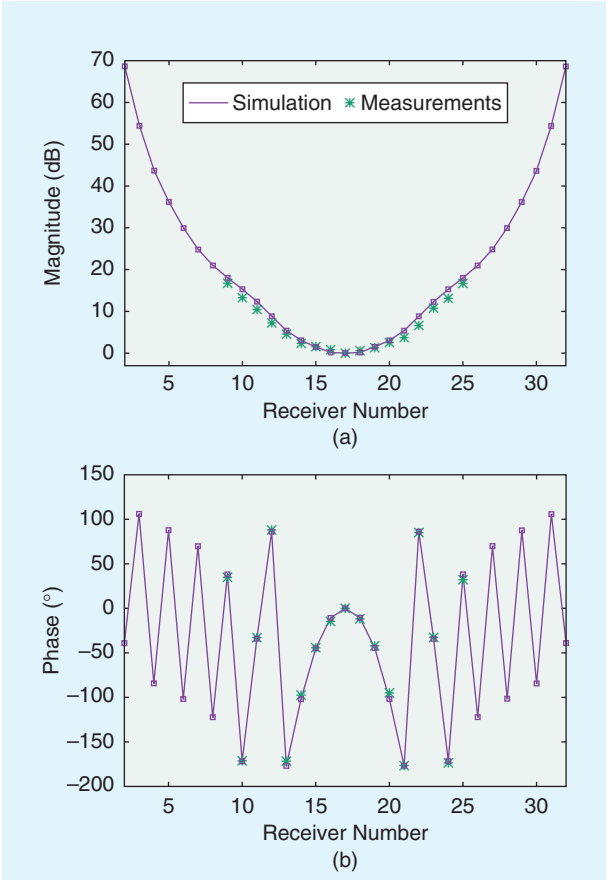

FIGURE 5. (a) The normalized amplitude and (b) the phase between the computed and measured S parameters.

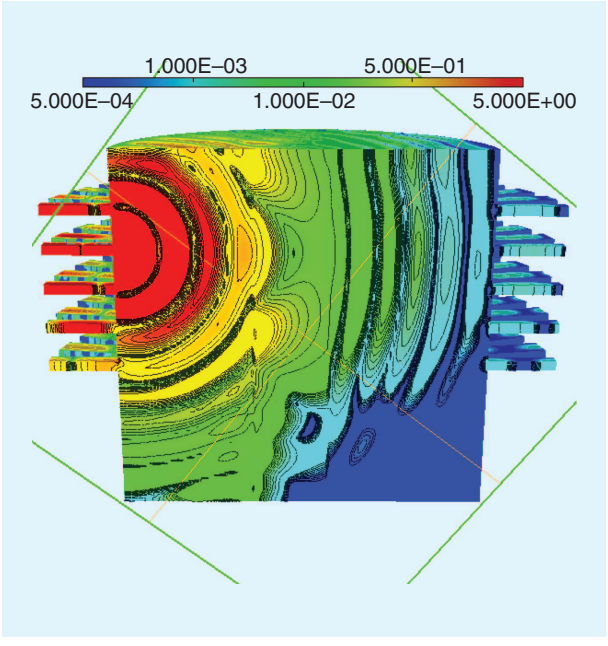

FIGURE 6. A cross section of the chamber showing the magnitude of the real part of the total field, E, in the chamber with the plastic tube.

directly opposite of the transmitting antenna, which is thus set to one. Since we normalize with respect to the coefficient having the lowest expected magnitude, the magnitude of the transmission coefficients shown in Figure 5 is larger than 0 dB. We can see that the transmission coefficients computed from the simulation are in very good agreement with the measurements.

HIGH-ORDER ELEMENT EFFICIENCY

The goal of the following numerical experiments is to assess the efficiency of the high-order finite elements described in the section "High-Order Edge Finite Elements" compared to the classical lowest-order edge elements in terms of accuracy and computing time, which are of great importance for such an application for brain imaging. For this test case, a nondissipative, plastic-filled cylinder with a 6-cm diameter and a relative permittivity of $\varepsilon_r^{cyl} = 3$ is inserted in the imaging chamber with the same background matching medium as defined in the section "Comparison with Experimental Measurements." We consider the 32 antennas of the second ring as transmitting antennas at a frequency of $f = 1$ GHz, and all of the 160 antennas are receiving.

We evaluate the error on the reflection and transmission coefficients, S_{ij}, with respect to the coefficients, S_{ij}^{ref}, computed from a reference solution. The error is calculated with the following formula:

$$\text{err} = \frac{\sqrt{\sum_{j,i} |S_{ij} - S_{ij}^{\text{ref}}|^2}}{\sqrt{\sum_{j,i} |S_{ij}^{\text{ref}}|^2}}. \tag{11}$$

The reference solution is computed on a fine mesh of approximately 18 million tetrahedra using edge finite elements of degree $r = 2$, resulting in 114 million unknowns. The section in Figure 6 shows the computational domain and the magnitude of the real part of the total field, E, over the cross section when one antenna of the second ring from the top is transmitting. We compare the computing time and the relative error (11) for different numbers of unknowns corresponding to several mesh sizes, for approximation degrees, $r = 1$ (15 points per λ) and $r = 2$ (ten points per λ) (Table 1), where λ is the wavelength in the ceramic $\lambda = (\lambda_0/\sqrt{59})$. We report the results in Figures 7 and 8. All of these simulations are carried out using 512 subdomains with one message passing interface (MPI) process and two OpenMP threads per subdomain, for a total of 1,024 cores on the Curie supercomputer.

The higher-order approximation $(r = 2)$ allows a given accuracy with much fewer unknowns and much less computation time than the lowest-order approximation $(r = 1)$. For example, at a given accuracy err of $E \approx 0.1$, the finite element discretization of degree $r = 1$ requires 21 million unknowns and a computing time of 130 s, while the high-order finite element discretization $(r = 2)$ only needs 5 million unknowns, with a corresponding computing time of 62 s. It turns out that we obtain the same accuracy with ten points per wavelength with a degree of $r = 2$ than with

TABLE 1. THE TOTAL NUMBER OF UNKNOWNS, GENERALIZED MINIMAL RESIDUAL METHOD (GMRES) ITERATIONS, COMPUTING TIME, AND RELATIVE ERROR ON COMPUTED S_{IJ}.

Degree 1				Degree 2			
Unknowns	**Iterations**	**Time (s)**	**Error**	**Unknowns**	**Iterations**	**Time (s)**	**Error**
2,373,214	35	22	0.384	1,508,916	29	39	0.242
8,513,191	46	53	0.184	5,181,678	34	62	0.099
21,146,710	60	130	0.117	12,693,924	41	122	0.057
42,538,268	70	268	0.083	26,896,130	47	236	0.036
73,889,953	86	519	0.068	45,781,986	57	396	0.019

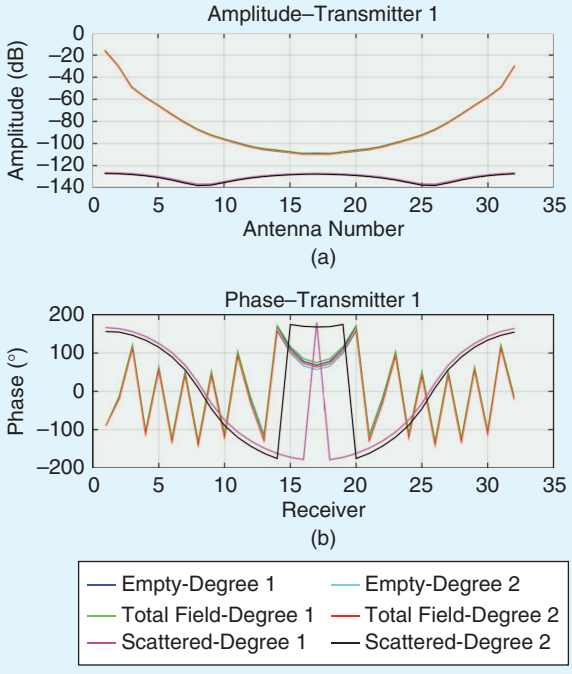

FIGURE 7. A comparison between degrees $r = 1$ and $r = 2$ on empty, total, and scattered fields (amplitude and phase).

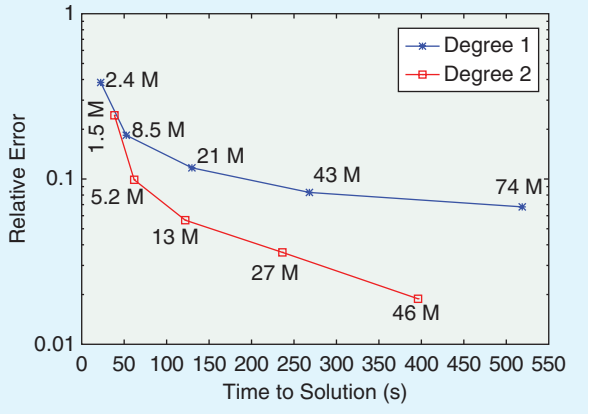

FIGURE 8. The computation time (in seconds) and relative error on computed S_{ij} using elements of degree $r = 1$ and $r = 2$ for different mesh sizes and number of unknowns in millions.

20 points per wavelength with a degree of $r = 1$. Moreover, note that greater accuracy could be achieved by using high-order geometric tetrahedral elements to reduce the geometric approximation error of the circular boundary of the chamber. However, we can deduce from Table 1 that the geometric approximation error is largely dominated by the finite element approximation error, as the error on the computed S_{ij} is significantly smaller using degree $r = 2$ edge elements compared to degree $r = 1$ edge elements, for a fixed number of degrees of freedom, even though meshes for degree $r = 2$ are coarser (leading to larger geometric approximation errors). We also report here additional information for a representative run since it can be useful to the community. For the test case corresponding to 27 million unknowns using degree $r = 2$ edge elements, the average and maximum subdomain size is 124,300 and 162,410 degrees of freedom, respectively. The total memory consumption for this run is approximately 2 TB. To give an idea of the parallel performance of the algorithm, the problem was solved using 512 cores. The computing time using 512 cores (256 subdomains with two threads per subdomain) was 366 s compared to 236 s using 1,024 cores (512 subdomains with two threads per subdomain), yielding a speedup of 1.55.

INVERSE PROBLEM

MATHEMATICAL FORMULATION

The inverse problem that we consider consists of finding the complex relative permittivity, $\varepsilon_r(x)$, in Ω, such that the solutions, $E_j(x)$, $j = 1, 2, \ldots, N$, of problem (10) lead to corresponding scattering parameters S (13) that coincide with the measured scattering parameters S_{ij}^{meas} for $i, j = 1, \ldots, N$.

Let $\kappa = k^2 = k_0^2 \varepsilon_r(x)$ be the unknown complex parameter of the inverse problem, and let us denote by $E_j(\kappa)$ the solution of the direct problem (6) with the complex dielectric permittivity ε. The corresponding scattering parameters will be denoted by $S_{ij}(\kappa)$ for $i, j = 1, 2, \ldots, N$.

The misfit of the parameter κ to the data can be defined with the following cost functional:

$$J(\kappa) = \frac{1}{2}\sum_{j=1}^{N}\sum_{i=1}^{N} |\, S_{ij}(\kappa) - S_{ij}^{\text{meas}} \,|^2 = \frac{1}{2}\sum_{j=1}^{N}\sum_{i=1}^{N} \left| \frac{\int_{\Gamma_i} \tilde{E}_j(\kappa).E_i^0}{\int_{\Gamma_i} |\, E_i^0 \,|^2} - S_{ij}^{\text{meas}} \right|^2. \quad (12)$$

Solving the inverse problem involves minimizing the functional J with respect to the parameter κ. Computing the differential of J in a given arbitrary direction $\delta\kappa$ yields

$$DJ(\kappa,\delta\kappa) = \sum_{j=1}^{N}\sum_{i=1}^{N} \text{Re}\left[\overline{(S_{ij}(\kappa) - S_{ij}^{\text{meas}})} \frac{\int_{\Gamma_i} \overline{\delta E_j(\kappa)}.E_i^0}{\int_{\Gamma_i} |\, E_i^0 \,|^2} \right], \quad (13)$$

for $\delta\kappa \in C$ and where $\delta E_j(\kappa)$ is the solution of the following linearized problem:

$$\begin{cases} \nabla\times(\nabla\times\delta E_j) - \kappa\delta E_j = \delta\kappa E_j & \text{in } \Omega \\ \delta E_j \times n = 0 & \text{on } \Gamma_c \\ (\nabla\times\delta E_j)\times n - i\beta n\times(\delta E_j \times n) = 0 & \text{on } \Gamma_i, i = 1,2,...,N. \end{cases} \quad (14)$$

We now use the adjoint approach to simplify the expression of DJ. This will allow us to compute the gradient efficiently after discretization, with a number of computations independent of the size of the parameter space.

Introducing the solution $F_j(\kappa)$ of the following adjoint problem:

$$\begin{aligned} &\nabla\times(\nabla\times F_j) - \kappa F_j = 0 \quad \text{in } \Omega \\ &F_j \times n = 0 \quad \text{on } \Gamma_c \\ &(\nabla\times F_j)\times n + i\beta n\times(F_j\times n) = \\ &\frac{(S_{ij}(\kappa) - S_{ij}^{\text{meas}})}{\int_{\Gamma_i} |\, E_i^0 \,|^2} \overline{E_i^0} \quad \text{on } \Gamma_i, i = 1,\ldots,N, \end{aligned} \quad (15)$$

we get, after some integration by parts (not detailed here),

$$\int_\Omega \delta\kappa E_j.F_j = \sum_{i=1}^{N} \left[(S_{ij}(\kappa) - S_{ij}^{\text{meas}}) \frac{\int_{\Gamma_i} E_i^0.\overline{\delta E_j}}{\int_{\Gamma_i} |\, E_i^0 \,|^2} \right]. \quad (16)$$

Finally, the differential of J can be computed as

$$DJ(\kappa,\delta\kappa) = \sum_{j=1}^{N} \text{Re}\Big[\int_\Omega \overline{\delta\kappa E_j.F_j} \Big]. \quad (17)$$

We can then compute the gradient to use in a gradient-based local optimization algorithm. The numerical results presented in the section "High-Order Element Efficiency" are obtained using a limited-memory Broyden–Fletcher–Goldfarb–Shanno (L-BFGS) algorithm [17]. Note that every evaluation of J requires the solution of the state problem (6) while the computation of the gradient requires the solution of (6) as well as the solution of the adjoint problem (15). Moreover, the state and adjoint problems use the same operator. Therefore, the computation of the gradient only needs the assembly of one matrix and its associated domain decomposition preconditioner.

The numerical results for the reconstruction of a hemorrhagic stroke from synthetic data are presented in the next section. The cost functional J considered in the numerical results is slightly different from (12) because we add a normalization term for each pair (i, j) as well as a Tikhonov regularization term [18]:

$$J(\kappa) = \frac{1}{2}\sum_{j=1}^{N}\sum_{i=1}^{N} \left| \frac{S_{ij}(\kappa)}{S_{ij}^{\text{empty}}} - \frac{S_{ij}^{\text{meas}}}{S_{ij}^{\text{meas,empty}}} \right| + \frac{\alpha}{2}\int_\Omega |\, \nabla\kappa \,|^2, \quad (18)$$

where S_{ij}^{empty} (respectively $S_{ij}^{\text{meas,empty}}$) refers to the computed coefficients (respectively measured) with empty chamber, only filled with the homogeneous matching solution, as described in the previous section, with no object. In this way, the contribution of each (i, j) pair in the cost functional is normalized and does not depend on the amplitude of the coefficient, which can vary greatly between (i, j) pairs, as shown in Figure 5. Moreover, normalizing both sides by the incident field S parameters helps eliminate multiplicative systematic errors, such as phase shifts [19]. The Tikhonov regularization term aims at reducing the effects of noise in the data. For now, the regularization parameter, α, is chosen empirically so as to obtain a visually good compromise between reducing the effects of noise and keeping the reconstructed image pertinent. All calculations carried out in this section can be accommodated in a straightforward manner to the definition (18) of the cost functional.

Finally, we can exhibit an additional level of parallelism by performing the reconstruction cross section by cross section. For the actual imaging system, one cross section corresponds to one of the five rings of 32 antennas, and we can define and solve an inverse problem independently for each of the five rings in parallel. We impose absorbing boundary conditions on the artificial boundaries of the truncated computational domain. Therefore, we are able to define five independent smaller inverse problems by reducing the computational complexity in terms of degrees of freedom and the number of relevant transmitters that must be simulated for each problem. Each of the five inverse problems takes into account 32 transmitters and (at most) 96 receivers. The reconstructed images obtained in this way are essentially identical to the images obtained by solving the full inverse problem, but they are performed faster.

NUMERICAL RESULTS

The results in this article were obtained on a Curie supercomputer (http://www-hpc.cea.fr/fr/complexe/tgcc-curie.htm), a system comprising 5,040 nodes composed of two eight-core Intel Sandy Bridge processors clocked at 2.7 GHz. The interconnect is an InfiniBand quad data rate full fat tree, and the MPI implementation used was BullxMPI version 1.2.8.4. Intel compilers and Math Kernel Library in their version 16.0.2.181 were used for all binaries and shared libraries and as the linear algebra backend for dense computations. One-level preconditioners such as (9) assembled by HPDDM require the use of a sparse direct solver.

HPDDM is interfaced with various direct solvers, such as PARDISO [20] from Intel MKL or MUMPS [21]. All numerical experiments reported in this article were performed with MUMPS. The GMRES algorithm is stopped once the unpreconditioned relative residual is lower than 10–8. The linear systems with multiple right-hand sides (one per transmitter) are solved using a pseudoblock method, where multiple operations are fused to achieve higher arithmetic intensity. In the inverse problem, during successive iterations of the minimization algorithm, the GMRES solution for each transmitter is initialized with the solution for this transmitter at the previous iteration. Moreover, block methods and Krylov subspace recycling techniques have been recently implemented in HPDDM. They are not used in the results presented in this article, but incorporating these techniques in the inversion algorithm to speed up computations is an ongoing work.

VIRTUAL HEAD MODEL

We want to assess the feasibility of the microwave imaging technique presented in this article for stroke detection and monitoring through a numerical example in a realistic configuration. Therefore, we use synthetic data corresponding to an accurate numerical model of a human head with a simulated hemorrhagic stroke as input for the inverse problem. The numerical model of the virtual head comes from CT and MRI scans and consists of a complex permittivity map of 362 × 434 × 362 data points with a spatial resolution of 500 μm. In the simulation, the head is immersed in the imaging chamber, as shown in Figure 9.

RECONSTRUCTIONS OF A HEMORRHAGIC STROKE

To simulate the evolution of a hemorrhagic stroke, we use a synthetic ellipsoid-shaped stroke whose size (principal axes) increases over time, from 3.9 cm × 2.3 cm × 2.3 cm (small stroke) to 7.7 cm × 4.6 cm × 4.6 cm (large stroke). For this test case, the relative complex permittivity of the ellipsoid is assumed to be inhomogeneous, where the relative complex permittivity at each quadrature point of the mesh is taken as the mean value between the original healthy brain permittivity values (baseline values) and the permittivity of blood ($\varepsilon_r^{\text{blood}} = 68 + i44$) at $f = 1$ GHz. The imaging chamber is filled with the matching solution $\varepsilon_r^{\text{matching}} = 44 + i20$. In a real experiment, a special membrane fitting the shape of the head is used to isolate the head from the matching medium (Figure 3). We do not take this membrane into account in this synthetic test case. The synthetic data are obtained by solving the direct problem using a mesh composed of 17.6 million tetrahedra (corresponding to approximately 20 points per wavelength) and consist in the computed transmission and reflection coefficients, S_{ij}. We subsequently add noise to the real and imaginary parts of the coefficients, S_{ij}, using a kind of multiplicative white Gaussian noise. We independently apply Gaussian noise to the real and imaginary parts of each S_{ij} coefficient as

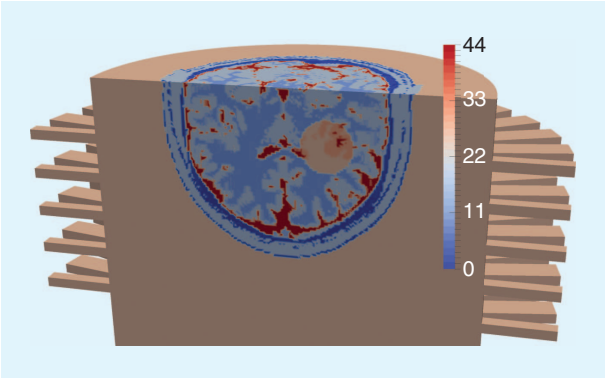

FIGURE 9. An imaginary part of the relative complex permittivity of the virtual head model immersed in the imaging chamber with a simulated ellipsoid-shaped hemorrhagic stroke.

$$S_{ij}^{\text{corrupted}} = S_{ij}(1 + p(G_{ij}^{\text{Re}} + jG_{ij}^{\text{Im}})), \tag{19}$$

where G_{ij}^{Re} and G_{ij}^{Im} are independent and identically distributed random variables drawn from the standard normal distribution $\mathcal{N}(0,1)$. In this experiment, we corrupt the data with 10% noise, i.e., $p = 0.1$.

The corrupted data, $S_{ij}^{\text{corrupted}}$, are then used as input for the inverse problem. Furthermore, we do not assume any prior knowledge on the input data, and we set the initial guess for the inverse problem as the homogeneous matching solution everywhere inside the chamber. We use a piecewise linear approximation of the unknown parameter, κ, defined on the same mesh used to solve the state and adjoint problems. For the purpose of parallel computations, the partitioning introduced by the domain decomposition method is also used to compute and store locally in each subdomain every entity involved in the inverse problem, such as the parameter κ and the gradient.

In the following, we only show reconstructed images corresponding to a cross section where the stroke is located due to space limitations; although, we are able to reconstruct the permittivity of the head in the whole chamber. Figures 10 and 11 show the real and imaginary parts of the reconstructed relative permittivity, respectively, for the three evolution steps of the hemorrhagic stroke, from a healthy brain to a brain with a small stroke and then a large stroke. Increasing the size of the ellipsoid simulates the evolution of the stroke. Each reconstruction corresponds to the solution of an inverse problem in the truncated domain containing only the first two rings of antennas from the top. The transmitting antennas are on the first ring and receiving antennas on first and second rings. Therefore, the scattering matrix contains only 64 × 32 S_{ij} coefficients. Each reconstruction starts from an initial guess consisting of the homogeneous matching solution. The solution is obtained at about 30 iteration steps when reaching the convergence criterion of 10^{-2} for the value of the cost functional using the L-BFGS algorithm. Subsequent iterations do not decrease the cost functional nor improve the reconstructed image, due to the level of noise and Tikhonov regularization.

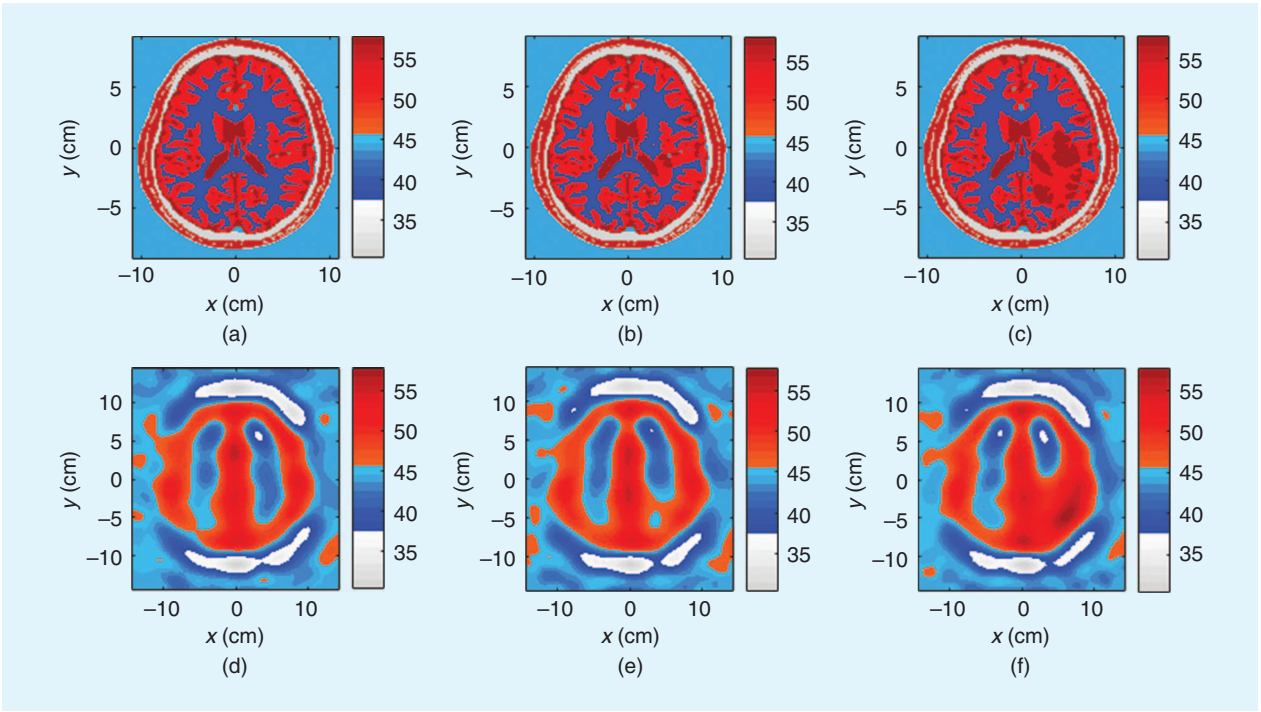

FIGURE 10. The transverse cross section of the virtual brain during the evolution of a simulated hemorrhagic stroke: the real part of the relative complex permittivity. (a) The virtual brain model of a healthy brain. (b) The virtual brain model of a brain with a small stroke. (c) The virtual brain model of a brain with a large stroke. (d) The reconstructed permittivity of a healthy brain. (e) The reconstructed permittivity of a brain with a small stroke. (f) The reconstructed permittivity of a brain with a large stroke.

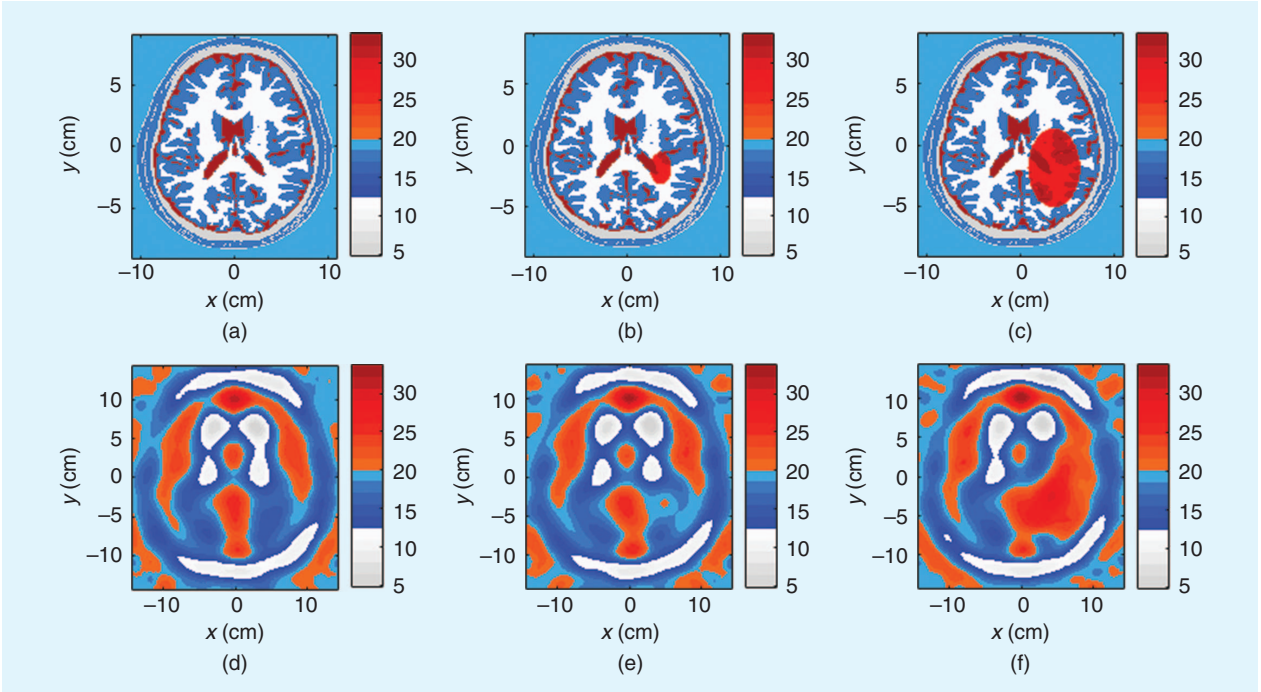

FIGURE 11. The transverse cross section of the virtual brain during the evolution of a simulated hemorrhagic stroke: the imaginary part of the relative permittivity. (a) The virtual brain model of a healthy brain. (b) The virtual brain model of a brain with a small stroke. (c) The virtual brain model of a brain with a large stroke. (d) The reconstructed permittivity of a healthy brain. (e) The reconstructed permittivity of a brain with a small stroke. (f) The reconstructed permittivity of a brain with a large stroke.

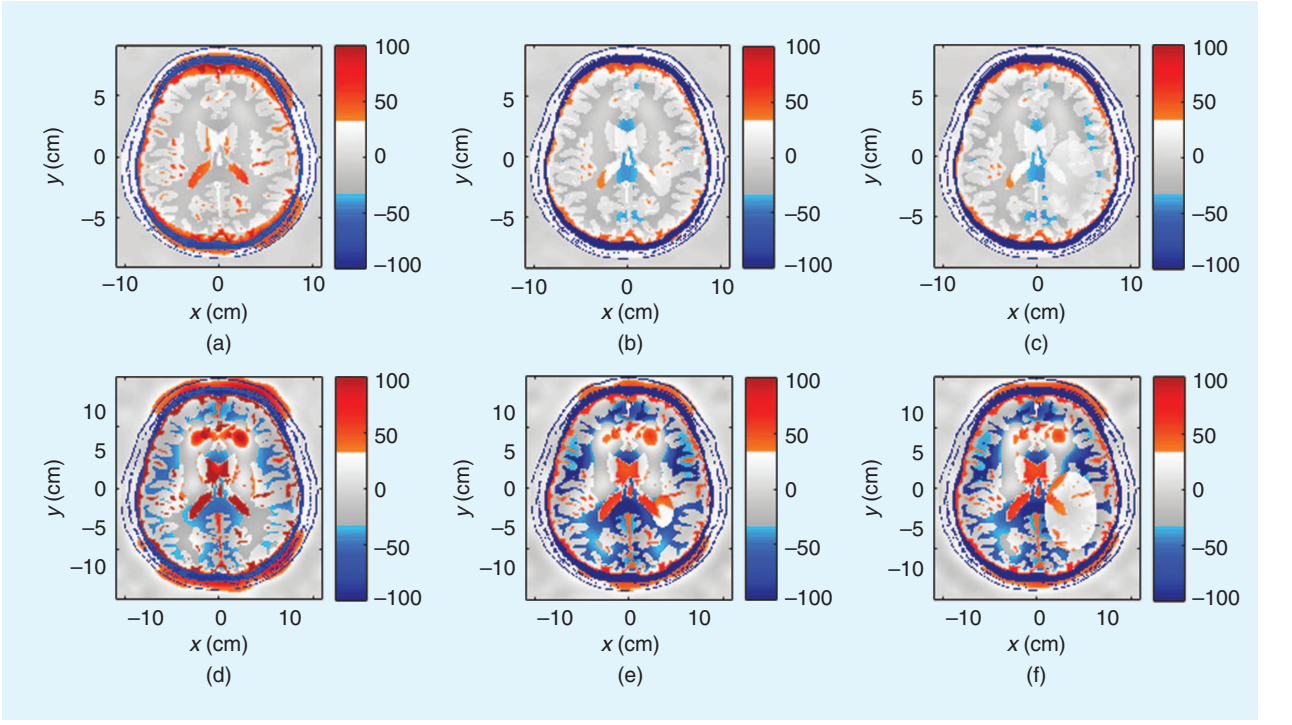

FIGURE 12. A small stroke: the distribution of the relative error on the real and imaginary parts of the reconstructed complex permittivity. (a) The real part of the reconstructed complex permittivity in a healthy brain. (b) The real part of the reconstructed complex permittivity in a brain with a small stroke. (c) The real part of the reconstructed complex permittivity in a brain with a large stroke. (d) The imaginary part of the reconstructed complex permittivity in a healthy brain. (e) The imaginary part of the reconstructed complex permittivity in a brain with a small stroke. (f) The imaginary part of the reconstructed complex permittivity in a brain with a large stroke.

The evolution of the stroke can be visually monitored from the real and imaginary parts of the reconstructed complex permittivity. Nevertheless, the threshold to firmly conclude must only be determined from clinical studies on a large number of patients. One important point is to discriminate a hemorrhagic from an ischemic stroke. For this study case, the reconstructed values show an increase of the complex permittivity allowing the assumption of a hemorrhagic stroke versus an ischemic one. The distribution of the relative error on the real and imaginary parts of the reconstructed complex permittivity for the small stroke case is shown in Figure 12. We compute the relative error using (20) for each pixel (n, m) of the reconstructed relative complex permittivity. This error can be positive or negative:

$$\mathrm{err}_{\mathrm{relative}}(m,n) = \frac{\varepsilon_r^{\mathrm{reconstructed}}(m,n) - \varepsilon_r^{\mathrm{exact}}(m,n)}{\varepsilon_r^{\mathrm{exact}}(m,n)}. \qquad (20)$$

We note that the lowest errors are located outside the brain and in the stroke. This can be expected, as the inversion algorithm performs better for homogeneous media, such as the matching liquid, but also for the stroke, as the complex permittivity value of the stroke is calculated as the mean value between the healthy tissues and the blood. This process tends to average the values, which is more favorable for the inversion algorithm. But even if the brain is highly heterogeneous, the stroke can be detected and monitored with the proposed algorithm.

TABLE 2. THE AVERAGE ERROR ON THE RECONSTRUCTED VALUES (THE REAL AND IMAGINARY PARTS OF THE COMPLEX PERMITTIVITY).

Relative Error	Real Part	Imaginary Part
Healthy brain	8.95%	20.74%
Brain with small stroke	8.92%	20.72%
Brain with large stroke	8.53%	18.92%

We now calculate the L^2 norm of the error of the reconstructed images in such manner as (11). Results are shown in Table 2.

The L^2 norm is interesting, as it gives a global quantitative criterion for estimating the performance of the reconstructed values. The L^2 norm confirms the results shown in Figure 12 for the small stroke. The error on the real part of the complex relative permittivity is lower than on the imaginary part. It is of the order of 10% for the real part whereas it is about 20% for the imaginary part.

It turns out that reconstructed images differ very slightly when using different discretization orders and mesh sizes in the

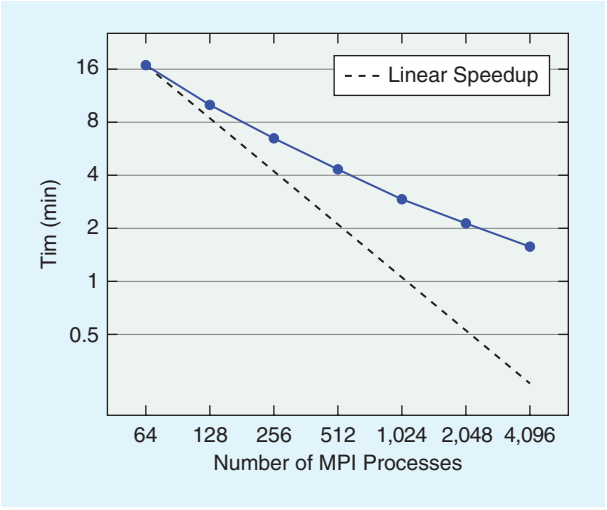

FIGURE 13. A strong scaling experiment: the total time for the large stroke case for solving the inverse problem versus the number of MPI processes.

inverse problem. In our case, elements of degree $r = 1$ with ten points per λ are sufficient for detecting the stroke. It is generally assumed that the most accurate forward model provides the best result combined with a given inverse algorithm. In this article, we have shown that it is not straightforward in this particular case with such a regularization technique. Reconstructed images for each test case shown in Figures 10 and 11 are obtained with a total computing time of less than 2 min (94 s for the large stroke case) using 4,096 cores of Curie.

Figure 13 shows the results of a strong scaling experiment for the large stroke case, which include solutions of the inverse problem corresponding to the top ring with an increasing number of MPI processes. We report the total computing time needed to obtain the reconstructed image. These preliminary results are very encouraging, as we are already able to achieve a satisfactory reconstruction time in the perspective of using such an imaging technique for monitoring. This allows clinicians to obtain almost instantaneous images 24/7 or on demand. Although the reconstructed images do not feature the complex heterogeneities of the brain, which is in accordance with what we expect from microwave imaging methods, they allow the characterization of the stroke and its monitoring.

CONCLUSIONS

The idea behind this article comes from the paradigm to develop a (portable/transportable) microwave imaging system whose raw data are wirelessly transferred to an HPC. The HPC machine will then compute the 3-D image of the patient's brain. Once reconstructed, the image is quickly transmitted from the computing center to the hospital for stroke detection (including ischemic/hemorrhagic discrimination) and monitoring during treatment. We have developed a tool that reconstructs a tomographic microwave image of the brain in 94 s on 4,096 computing cores. This computational time corresponds to clinician acceptance for rapid diagnosis or medical monitoring at the hospital. These images were obtained from corrupted synthetic data from a very accurate model of the complex permittivity of the brain. To our knowledge, this is the first time that such a realistic study (operational acquisition device, highly accurate 3-D synthetic data, 10% noise) shows the feasibility of microwave imaging. This study has been possible with the use of massively parallel computers and facilitated by the HPDDM and FreeFem++ tools that we developed. The next step will be the validation of these results on clinical data.

ACKNOWLEDGMENTS

This work is supported by the French National Research Agency (ANR) grant MEDIMAX (ANR-13-MONU-0012). This work was granted access to the HPC resources of TGCC at CEA under the allocations 2016-067519 and 2016-067730 made by Grand Équipement National de Calcul Intensif (GENCI). We would like to thank ANR for its support.

AUTHOR INFORMATION

Pierre-Henri Tournier (tournier@ljll.math.upmc.fr) is with the Jacques-Louis Lions Laboratory, Pierre and Marie Curie University, Centre Nationale de la Recherche Scientifique, France. He works on devising and implementing robust and efficient parallel numerical algorithms, in particular, using domain decomposition methods for wave propagation problems.

Marcella Bonazzoli (marcella.bonazzoli@unice.fr) is a Ph.D. student in mathematics at the Jean-Alexandre Dieudonné Laboratory, Université Côte d'Azur, Centre Nationale de la Recherche Scientifique, France. Her research activity concerns high-order finite element methods and domain decomposition methods for Maxwell's equations.

Victorita Dolean (dolean@unice.fr) is with the Jean-Alexandre Dieudonné Laboratory and the Electronics, Antennas, and Telecommunications Laboratory, Université Côte d'Azur, Centre Nationale de la Recherche Scientifique, France, and the Department of Mathematics and Statistics at the University of Strathclyde, Glasgow, United Kingdom. Her research interests are in computational methods in electromagnetism and high-performance parallel computing.

Francesca Rapetti (francesca.rapetti@unice.fr) is with the Jean-Alexandre Dieudonné Laboratory, Université Côte d'Azur, Centre Nationale de la Recherche Scientifique, France. Her main research area is in Maxwell's equations in the exterior calculus formalism, high-order Whitney finite elements for computational electromagnetism, and nonconforming mortar finite/spectral element techniques.

Frédéric Hecht (frederic.hecht@upmc.fr) is a full professor at the Jacques-Louis Lions Laboratory (LJLL), Pierre and Marie Curie University, France, and a member of the joint LJLL and the French Institute for Research in Computer Science (INRIA) and Automation team Alpines.

Frédéric Nataf (nataf@ljll.math.upmc.fr) is a Centre Nationale de la Recherche Scientifique senior scientist at the Jacques-Louis Lions Laboratory (LJLL), Pierre and Marie Curie University, France, and a member of the joint LJLL and the French Institute for Research in Computer Science (INRIA) and Automation team Alpines.

Iannis Aliferis (iannis.aliferis@unice.fr) is with the Electronics, Antennas, and Telecommunications Laboratory, Université Côte d'Azur, Centre Nationale de la Recherche Scientifique, France. His research activity concerns electromagnetic imaging, and forward and inverse scattering problems. As a teacher, he is practicing and diffusing active learning (flipped classroom and faculty training).

Ibtissam El Kanfoud (ielkanfoud@unice.fr) is a Ph.D. student with the Electronics, Antennas, and Telecommunications Laboratory, Université Côte d'Azur, Centre Nationale de la Recherche Scientifique, France. Her research focuses on electromagnetic simulation of microwave imaging systems for brain stroke detection.

Claire Migliaccio (claire.migliaccio@unice.fr) is a full professor at the Electronics, Antennas, and Telecommunications Laboratory, Université Côte d'Azur, Centre Nationale de la Recherche Scientifique, France. Her research interests include antennas, and microwave and millimeter-wave measurements dedicated to radar and imaging systems.

Maya de Buhan (maya.de-buhan@parisdescartes.fr) is an applied mathematician working for the Centre Nationale de la Recherche Scientifique at MAP5 Laboratory, Paris Descartes University, France. She is interested in theoretical and numerical aspects for inverse problems, essentially for wave propagation phenomena.

Marion Darbas (marion.darbas@u-picardie.fr) is an assistant professor in applied mathematics at Laboratoire Amiénois de Mathématique Fondamentale et Appliquée, Université de Picardie Jules Verne, Amiens, France. Her research activity concerns the numerical solution of direct and inverse wave propagation problems.

Serguei Semenov (serguei.semenov@emtensor.com) is the chief executive officer and chief technology officer of EMTensor GmbH, Vienna, Austria. His research interests are within bioelectromagnetics and biomedical applications of electromagnetic tomographic imaging.

Christian Pichot (christian.pichot@unice.fr) is a Centre Nationale de la Recherche Scientifique senior scientist at the Electronics, Antennas, and Telecommunications Laboratory, Université Côte d'Azur, France. His research activity concerns computational electromagnetics, antenna design and shape optimization, inverse scattering (microwave imaging and tomography), and development of microwave imaging systems. He is a Fellow of the IEEE.

REFERENCES

[1] R. Sacco, S. E. Kasner, J. P. Broderick, L. R. Caplan, J. J. Connors, A. Culebras, M. S. V. Elkind, M. G. George, A. D. Hamdan, R. T. Higashida, B. L. Hoh, L. S. Janis, C. S. Kase, D. O. Kleindorfer, J.-M. Lee, M. E. Moseley, E. D. Peterson, T. N. Turan, A. L. Valderrama, and H. V. Vinters, "An updated definition of stroke for the 21st century: A statement for healthcare professionals from the American Heart Association/American Stroke Association," *Stroke*, vol. 44, no. 7, pp. 2064–2069, July 2013.

[2] Intercollegiate Stroke Working Party, *National Clinical Guideline for Stroke*, 4th ed. London: Royal College of Physicians, 2012.

[3] S. Wegener, "Neuroimaging of acute ischaemic stroke: Current challenges," *Eur. Med. J. Neurology*, vol. 1, pp. 49–52, July 2014.

[4] S. Semenov, R. Svenson, V. Posukh, A. Nazarov, Y. Sizov, A. Bulyshev, A. Souvorov, W. Chen, J. Kassell, and G. Tatsis, "Dielectrical spectroscopy of canine myocardium during acute ischemia and hypoxia at frequency spectrum from 100 kHz to 6 GHz," *IEEE Trans. Med. Imag.*, vol. 21, no. 6, pp. 703–707, June 2002.

[5] S. Semenov, B. Seiser, E. Stoegmann, and E. Auff, "Electromagnetic tomography for brain imaging: From virtual to human brain," in *Proc. IEEE Conf. Antenna Measurements Applications*, Antibes Juan-les-Pins, 2014, pp. 1–4.

[6] F. Rapetti, "High order edge elements on simplicial meshes," *Math. Model. Numer. Anal.*, vol. 41, no. 6, pp. 1001–1020, 2007.

[7] J.-C. Nédélec, "Mixed finite elements in R3," *Numer. Math.*, vol. 35, no. 3, pp. 315–341, 1980.

[8] F. Hecht, "New development in FreeFem++," *J. Numer. Math.*, vol. 20, no. 3–4, pp. 251–265, 2012.

[9] F. Wei and A. E. Yilmaz, "A more scalable and efficient parallelization of the adaptive integral method—Part II: BIOEM application," *IEEE Trans. Antennas Propag.*, vol. 62, no. 2, pp. 727–738, Feb. 2014.

[10] Y.-J. Li and J.-M. Jin, "A new dual-primal domain decomposition approach for finite element simulation of 3-D large-scale electromagnetic problems," *IEEE Trans. Antennas Propag.*, vol. 55, no. 10, pp. 727–738, Oct. 2007.

[11] V. Dolean, P. Jolivet, and F. Nataf, *An Introduction to Domain Decomposition Methods: Algorithms, Theory and Parallel Implementation*. Philadelphia, PA: SIAM, 2015.

[12] A. St-Cyr, M. J. Gander, and S. J. Thomas, "Optimized restrictive additive Schwarz methods," in *Domain Decomposition Methods in Science and Engineering XVI* (Lectures Notes in Computational Sciences and Engineering, vol. 55), O. B. Widlund and D. E. Keyes, Eds. Berlin: Springer-Verlag, 2007, pp. 213–220.

[13] X.-C. Cai and M. Sarkis, "A restricted additive Schwarz preconditioner for general sparse linear systems," *SIAM J. Sci. Comput.*, vol. 21, no. 2, pp. 792–797, 1999.

[14] V. Dolean, M. Gander, and L. Gerardo-Giorda, "Optimized Schwarz methods for Maxwell's equations," *SIAM J. Sci. Comput.*, vol. 31, no. 3, pp. 2193–2213, 2009.

[15] M. Bonazzoli, V. Dolean, F. Rapetti, and P.-H. Tournier. (2017). Parallel preconditioners for high order discretizations arising from full system modeling for brain microwave imaging. *Int. J. Numer. Model.* [Online]. Available: https://doi.org/10.1002/jnm.2229

[16] P. Jolivet, F. Hecht, F. Nataf, and C. Prud'Homme, "Scalable domain decomposition preconditioners for heterogeneous elliptic problems," in *Proc. Int. Conf. High Performance Computing, Networking, Storage Analysis*, Denver, CO, 2013, pp. 1–11.

[17] J. Nocedal and S. Wright, *Numerical Optimization*. New York: Springer, 2006.

[18] M. Kern, *Numerical Methods for Inverse Problems*. Hoboken, NJ: Wiley, 2016.

[19] M. Haynes, J. Stang, and M. Moghaddam, "Real-time microwave imaging of differential temperature for thermal therapy monitoring," *IEEE Trans. Biomed. Eng.*, vol. 61, no. 6, June 2014.

[20] O. Schenk and K. Gärtner, "Solving unsymmetric sparse systems of linear equations with PARDISO," *Future Gener. Comput. Syst.*, vol. 20, pp. 3475–3487, Apr. 2004.

[21] P. Amestoy, I. Duff, J.-Y. L'Excellent, and J. Koster, "A fully asynchronous multifrontal solver using distributed dynamic scheduling," *SIAM J. Matrix Anal. Appl.*, vol. 23, no. 1, pp. 15–41, 2001.