frontiers in Physiology

METHODS
published: 02 May 2017
doi: 10.3389/fphys.2017.00265

Evaluation of a Rapid Anisotropic Model for ECG Simulation

Simone Pezzuto [1,2], Peter Kal'avský [1,3], Mark Potse [1,4,5], Frits W. Prinzen [1,6], Angelo Auricchio [1,7] and Rolf Krause [1,2]*

[1] Center for Computational Medicine in Cardiology, Lugano, Switzerland, [2] Faculty of Informatics, Institute of Computational Science, Università della Svizzera Italiana, Lugano, Switzerland, [3] Department of Biomeasurements, Institute of Measurement Science, Slovak Academy of Sciences, Bratislava, Slovakia, [4] Electrophysiology and Heart Modeling Institute IHU LIRYC, Pessac, France, [5] Inria Bordeaux Sud-Ouest, Talence, France, [6] Department of Physiology, Cardiovascular Research Institute Maastricht, Maastricht University, Maastricht, Netherlands, [7] Division of Cardiology, Fondazione Cardiocentro Ticino, Lugano, Switzerland

State-of-the-art cardiac electrophysiology models that are able to deliver physiologically motivated activation maps and electrocardiograms (ECGs) can only be solved on high-performance computing architectures. This makes it nearly impossible to adopt such models in clinical practice. ECG imaging tools typically rely on simplified models, but these neglect the anisotropic electric conductivity of the tissue in the forward problem. Moreover, their results are often confined to the heart-torso interface. We propose a forward model that fully accounts for the anisotropic tissue conductivity and produces the standard 12-lead ECG in a few seconds. The activation sequence is approximated with an eikonal model in the 3d myocardium, while the ECG is computed with the lead-field approach. Both solvers were implemented on graphics processing units and massively parallelized. We studied the numerical convergence and scalability of the approach. We also compared the method to the bidomain model in terms of ECGs and activation maps, using a simplified but physiologically motivated geometry and 6 patient-specific anatomies. The proposed methods provided a good approximation of activation maps and ECGs computed with a bidomain model, in only a few seconds. Both solvers scaled very well to high-end hardware. These methods are suitable for use in ECG imaging methods, and may soon become fast enough for use in interactive simulation tools.

Keywords: ECG, eikonal model, lead fields, bidomain modeling, patient-specific modeling, electrophysiology

OPEN ACCESS

Edited by:
Alexander Panfilov, Ghent University, Belgium

Reviewed by:
Simone Scacchi, University of Milan, Italy
Chris Patrick Bradley, University of Auckland, New Zealand

****Correspondence:***
Simone Pezzuto
simone.pezzuto@usi.ch

Specialty section:
This article was submitted to Computational Physiology and Medicine, a section of the journal Frontiers in Physiology

Received: *02 December 2016*
Accepted: *11 April 2017*
Published: *02 May 2017*

Citation:
Pezzuto S, Kal'avský P, Potse M, Prinzen FW, Auricchio A and Krause R (2017) Evaluation of a Rapid Anisotropic Model for ECG Simulation. Front. Physiol. 8:265. doi: 10.3389/fphys.2017.00265

1. INTRODUCTION

The cardiac muscle or myocardium consists of electrically active muscle cells (myocytes) connected to each other by gap junctions, embedded in extracellular fluid and other structures. The current state of the art in forward simulation of the electrocardiogram (ECG) is to use a bidomain reaction-diffusion model to simulate propagating electrical activation and repolarization as well as the potential field in the torso (Vigmond et al., 2007). The bidomain model (Henriquez, 2014) describes the electrical behavior of the cardiac tissue using two co-located domains to represent the interpenetrating networks of intracellular and extracellular space. The conductivity of each domain can be described by a conductivity tensor field $\mathbf{G}_\mathrm{i}$ for the intracellular, and $\mathbf{G}_\mathrm{e}$ for the extracellular domain, respectively. These tensors have different degrees of anisotropy (Roth, 1997). The two domains are separated by the volume-averaged cell membrane, but electrically connected to one another through the ion channels in this membrane. The bidomain model can be described by a

reaction-diffusion system

$$\begin{cases} \nabla \cdot \left(\mathbf{G}_\mathrm{i} \nabla (V_\mathrm{m} + \phi_\mathrm{e})\right) = \beta \left(C_\mathrm{m} \frac{\partial V_\mathrm{m}}{\partial t} + I_\mathrm{ion}(V_\mathrm{m}, \mathbf{w})\right) - I_\mathrm{stim}\,, & (1) \\ \nabla \cdot \left((\mathbf{G}_\mathrm{i} + \mathbf{G}_\mathrm{e}) \nabla \phi_\mathrm{e}\right) = -\nabla \cdot (\mathbf{G}_\mathrm{i} \nabla V_\mathrm{m})\,, & (2) \\ \frac{\partial \mathbf{w}}{\partial t} = \mathbf{F}(V_\mathrm{m}, \mathbf{w})\,, & (3) \end{cases}$$

where V_m is the transmembrane potential, ϕ_e the potential in the extracellular domain, β the surface-to-volume ratio of the cell membrane, C_m the membrane capacitance per unit area, I_ion a non-linear function representing the sum of all transmembrane ionic currents, $\mathbf{w}$ a vector of state variables for the ionic currents, $\mathbf{F}$ a non-linear function, and I_stim an applied stimulus current. Equation (1) describes the current flow across the membrane and can be used together with Equation (3) to integrate the V_m distribution. Equation (2) describes the conservation of charge and can be used to solve ϕ_e from a given V_m.

The functions I_ion and $\mathbf{F}$ represent the so-called membrane model. If a physiologically detailed membrane model is chosen and the torso model is sufficiently detailed, this allows highly realistic ECGs to be simulated.

Much simpler models are presently being used inside inverse models, which aim at recovering cardiac events from non-invasive signals (Ramanathan et al., 2004; van Dam et al., 2009; Tysler and Svehlikova, 2013; Wang et al., 2013; Erem et al., 2014), and in interactive ECG-simulation tools (van Oosterom et al., 2011). Typically, these models assume predefined propagation velocities and predefined action potential waveforms. They account for anisotropic propagation but neglect the effects of anisotropic myocardial conductivity on the relation between activation and the ECG. Whereas a bidomain reaction-diffusion solution for a single heart beat takes in the order of a million core-seconds and is typically run on high-performance computing (HPC) resources, simplified methods can obtain a solution in less than a second on a single CPU core. Our purpose is to propose an ECG simulation method that is intermediate between these two extremes.

The conductivities play three roles in the ECG:

1. They determine the conduction velocity, which will also be anisotropic;
2. They determine the current dipole density resulting from a voltage gradient; and
3. They determine how the resulting current flows through the heart and torso and thus what voltages will be measured by the ECG electrodes.

While virtually all models nowadays account for anisotropic propagation, rapid ECG models typically neglect the second and third role of the anisotropic conductivities. The major reason for this is that the conductivities are non-uniform: the fiber orientation rotates around the cavities of the heart and also changes transmurally. To account for it, therefore, a volumetric discretization of the heart and torso is necessary. By assuming isotropy, in contrast, existing methods simplify the problem to one on the surface of the muscle (Geselowitz, 1992). However, the anisotropic conductivity has profound effects on the ECG, especially on the standard precordial electrodes, which are located near the heart (Potse et al., 2009) and which are appreciated by cardiologists for the detailed information they provide. We hypothesized that with today's computer technology one can build a rapid, fully anisotropic ECG model that is fast enough to be employed inside an inverse method.

One possible way to reduce the computational intensity of the bidomain model is to use an eikonal model (Colli Franzone and Guerri, 1993; Pullan et al., 2002) for the evolution of the excitation wavefront. The excitation wavefront represents a thin depolarized region of cardiac cells and an eikonal equation can be used to compute the activation time at which a wavefront reaches a given point in the myocardium. The eikonal equation belongs to a broad class of Hamilton-Jacobi equations (Bornemann and Rasch, 2006) and can be formally derived by a perturbation argument applied to the bidomain equations (Colli Franzone and Guerri, 1993). Therefore, both approaches can be compared (Pullan et al., 2002).

The method we propose here uses a volumetric anisotropic eikonal propagation model, pre-computed action potentials, and an ECG simulation based on dipole sources and a set of transfer functions known historically as lead fields. The volumetric eikonal model, in contrast to the surface models used by others (van Dam et al., 2009), provides us with an activation sequence throughout the cardiac volume, which we need to compute current dipoles at different locations with different fiber orientations. The numerical scheme adopted is guaranteed to converge to the correct viscosity solution of the eikonal equation (Bornemann and Rasch, 2006), which in turn correctly defines the fastest path between points through its gradient.

Precomputed action potentials are assigned and shifted according to the computed activation times, resulting in a field of transmembrane potentials at any desired time instant. For each time instant, the current dipoles are computed with an anisotropic formula, and the ECG potentials are computed from them using a set of transfer coefficients. The transfer coefficients are computed once for each model by solving a bidomain problem in the full torso, again taking all anisotropic conductivities into account. This method of ECG simulation is mathematically equivalent with full bidomain solutions of the torso potential based on monopole or dipole sources (Potse et al., 2009; Jacquemet, 2015). The only simplification of our model with respect to a full reaction-diffusion solution is the assumption of fixed but possibly heterogeneous action potential waveforms and conduction velocities that is inherent in the eikonal formulation.

The eikonal model allows for the use of a much coarser spatial resolution of 1 mm (Colli Franzone and Guerri, 1993) than the 0.1 mm scale required to reliably reconstruct the sharp upstroke of the action potential by a reaction-diffusion model (Clayton et al., 2011; Pezzuto et al., 2016). As a result much shorter computation times can be reached. The spatial discretization at a resolution of 1 mm is high enough to resolve the local fiber orientation and allow an anisotropic computation of the dipole sources. Thus, the method is able to account for the influence of anisotropic conductivity on the ECG, as well as on the propagation velocity.

The method was entirely implemented in C, C++ and CUDA so it could run efficiently on general-purpose graphics processing units (GPGPU).

Using a reaction-diffusion and bidomain model as a reference, we evaluated the proposed model in terms of the accuracy of the simulated activation sequence and ECG.

Evaluations were performed using six patient-specific heart-torso models. Finally, we evaluated the parallel scalability of the methods on GPGPUs.

2. MATERIALS AND METHODS

2.1. Reference Model

A "monodomain" reaction-diffusion model of the heart coupled to a bidomain heart-torso model was used as reference. The monodomain model can be derived from the bidomain model by assuming that the two domains have the same anisotropy ratio, $\mathbf{G}_\mathrm{e} = \lambda \mathbf{G}_\mathrm{i}$, with λ a constant scalar. Then Equations (1) and (2) reduce to a single reaction-diffusion equation

$$\nabla \cdot (\mathbf{G}_\mathrm{m} \nabla V_\mathrm{m}) = \beta \left(C_\mathrm{m} \frac{\partial V_\mathrm{m}}{\partial t} + I_\mathrm{ion} \right) - I_\mathrm{stim} \tag{4}$$

where $\mathbf{G}_\mathrm{m}$ is an effective conductivity. Under the equal anisotropy assumption, $\mathbf{G}_\mathrm{m} = \frac{\lambda}{1+\lambda}\mathbf{G}_\mathrm{i}$. A monodomain reaction-diffusion model is a very good approximation for its bidomain counterpart even when the equal anisotropy assumption does not hold (Potse et al., 2006; Nielsen et al., 2007; Bishop and Plank, 2011; Coudière et al., 2014), provided that the effective conductivity is chosen as $\mathbf{G}_\mathrm{m} = \mathbf{G}_\mathrm{e}(\mathbf{G}_\mathrm{i} + \mathbf{G}_\mathrm{e})^{-1}\mathbf{G}_\mathrm{i}$. Spatial discretization was performed with a semi-structured finite-difference mesh with 0.2 mm resolution. The time step was 0.01 ms for both the diffusion and the reaction, respectively solved with explicit Euler and Rush-Larsen schemes. The activation time was computed as the time when V_m became positive.

To compute the ECG, we solved the equation

$$\nabla \cdot [(\mathbf{G}_\mathrm{i} + \mathbf{G}_\mathrm{e}) \nabla \phi_\mathrm{e}] = -I_\mathrm{w}, \tag{5}$$

in a 3d finite-difference full torso model at 1-mm resolution, including a downsampled heart model (Potse et al., 2009). The term I_w is a projection of the transmembrane current $\nabla \cdot (\mathbf{G}_\mathrm{i} \nabla V_\mathrm{m})$ onto the coarse mesh (Potse and Kuijpers, 2010). The surface ECGs were derived by evaluating the potential at the electrode locations.

All reference computations were performed with the propag-5 software (Potse et al., 2006; Krause et al., 2012) and ran on a Cray XC40 computer.

2.2. Eikonal Model

To compute the activation sequence of the heart we used a non-linear first-order Hamilton-Jacobi eikonal equation in the form

$$\begin{cases} \dfrac{\alpha}{\sqrt{\beta}} \sqrt{\mathbf{G}_\mathrm{m} \nabla \Psi \cdot \nabla \Psi} = 1, & x \in \Omega \setminus \{s_k\}_{k=1}^K, \\ \Psi(s_k) = \tau_k & k = 1, \ldots, K. \end{cases} \tag{6}$$

where $\Psi(x)$ is the activation time at the point $x \in \Omega$ and $\alpha = \alpha(x)$ is a scaling parameter for the conduction velocity. Propagation is initiated at $K > 0$ early activation sites s_k with initial time τ_k.

The Equation (6), which assigns a locally constant front velocity, is one of the simplest eikonal models. More sophisticated versions of the eikonal equation are also available, where the velocity of the front depends also on the curvature of the front itself and higher-order curvature terms, or where the propagation is defined by a Finsler metric rather than a Riemannian metric (Colli Franzone et al., 2014).

In general, the scaling parameter α is the velocity that the front would have if the conductivity and β were set to one. Indeed, given a propagation direction $\mathbf{p}$, the conduction velocity arising from the reaction-diffusion model (Equation 4) is proportional to

$$\frac{\sqrt{\mathbf{G}_\mathrm{m} \mathbf{p} \cdot \mathbf{p}}}{\sqrt{\beta}} \tag{7}$$

which is exactly the conduction velocity given by the eikonal model for $\alpha = 1$. α depends only on the membrane capacitance and the ionic model, in particular on the sodium-channel conductivity and gating dynamics.

The implementation of our eikonal solver is based on a local variational principle proposed by Bornemann and Rasch (2006). Let $\delta(y, x)$ be the travel time between points y and x, computed with respect to a local metric, and $\mathcal{I}_\omega$ a linear Lagrange interpolation operator restricted to a patch ω. The numerical solution $\Psi_h = \{\Psi_{h,i}\}_{i=1}^N$, with $\Psi_{h,i}$ the activation time at the node x_i of the mesh, is iteratively updated for each node by a local map Λ defined as follows:

$$\Psi_{h,i}^{j+1} = \Lambda_i(\Psi_h^j) := \begin{cases} \min\limits_{y \in \partial\omega(x_i)} \left\{ \delta(y, x_i) + \mathcal{I}_{\omega(x_i)}(\Psi_h^j)(y) \right\} & \text{if } x_i \neq s_k, \\ \tau_k & \text{if } x_i = s_k, \end{cases} \tag{8}$$

for $i = 1, \ldots, N$. The nodal value $\Psi_{h,i}$ is updated using a patch $\omega(x_i)$ of tetrahedral elements from the neighboring nodes (see **Figure 1**). The minimum in Equation (8) is taken over the polyhedral boundary $\partial\omega(x_i)$ of the patch. The new value of $\Psi_{h,i}$, if not enforced as boundary condition, is the minimum time at x_i traveling from the boundary of the patch.

The quantity $\delta(y, z)$ is defined as the travel time along the fastest path γ from y to z with respect to metric $\mathbf{C}(x) = \beta(x)\alpha^{-2}(x)\mathbf{G}_\mathrm{m}^{-1}(x)$:

$$\delta(y, z) := \inf_{\gamma \in \mathcal{C}^1([0,1], \mathbb{R}^3)} \left\{ \int_0^1 \|\dot{\gamma}(t)\|_{\mathbf{C}(\gamma(t))} \, \mathrm{d}t : \begin{array}{r} \gamma(0) = y, \\ \gamma(1) = z \end{array} \right\}. \tag{9}$$

In order to compute it, we discretize the metric $\mathbf{C}^{-1}(x)$ as piecewise-constant on the patch $\omega(x)$, i.e., its value is constant within each tetrahedron. Then the fastest path for each tetrahedron is the segment connecting y and z, and the travel time is $\|z - y\|_\mathbf{C}$.

The minimization problem in Equation (8) is solved by taking the minimum of the minima for each triangular face $\mathfrak{t}$ of $\partial\omega(x_i)$:

$$\min_{\mathfrak{t} \in \partial\omega(x_i)} \min_{y \in \mathfrak{t}} \left\{ \|x_i - y\|_\mathbf{C} + \mathcal{I}_f(\Psi_h^j)(y) \right\}. \tag{10}$$

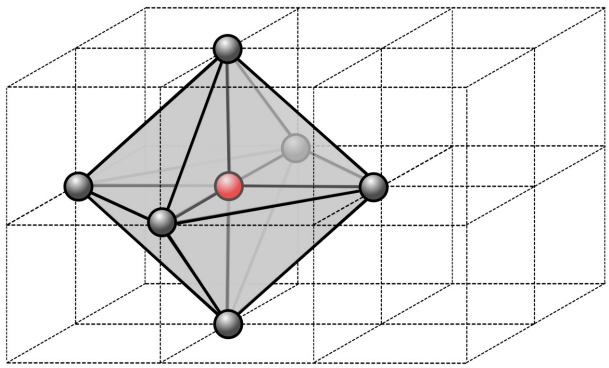

FIGURE 1 | Stencil of a single node of the computational grid. In 3D, we use 6 neighbor nodes from which we define a virtual patch of tetrahedral element. Within each tetrahedron, the parameters are assumed constant, thus the fastest path is always a segment.

The minimization over the single triangular face is a convex optimization problem, hence the minimum is unique.

It is possible to show that the map Λ has a unique fixed point (Bornemann and Rasch, 2006). Moreover, the fixed point converges to the correct viscosity solution of the eikonal equation as $h \to 0$. The convergence rate is generally $\mathcal{O}(h)$ or less, due to the low regularity of the solution, which is only Lipschitz continuous even for a regular domain and coefficients.

2.3. ECG

Each ECG lead potential $V(t)$ at time t was computed as

$$V(t) = \int_\Omega \nabla Z(x) \cdot \mathbf{G}_\mathrm{i}(x) \nabla V_\mathrm{m}(x,t)\,\mathrm{d}x \tag{11}$$

where Ω is the heart domain and $Z(x)$ is the lead field (McFee and Johnston, 1953) of the specific ECG lead. The lead field is the potential field created by a unit current applied at the electrode locations (Geselowitz, 1989)

$$\nabla \cdot (\mathbf{G}\nabla Z) = \begin{cases} -1 & \text{at the positive electrode,} \\ 1 & \text{at the negative electrode,} \\ 0 & \text{elsewhere,} \end{cases} \tag{12}$$

where $\mathbf{G} = \mathbf{G}_\mathrm{i} + \mathbf{G}_\mathrm{e}$ is the bulk conductivity ($\mathbf{G}_\mathrm{i} = 0$ outside the heart) (Potse et al., 2009). Equation (12) was solved for Z with propag-5 (Potse et al., 2006; Krause et al., 2012), using the complete heart-and-torso model at 1-mm resolution. Both the transfer function and the dipole sources ($\mathbf{G}_\mathrm{i}\nabla V_\mathrm{m}$) were computed with fully anisotropic conductivity values.

To obtain $V_\mathrm{m}(x,t)$ from $\Psi(x)$ we used a fixed action potential waveform $U(t)$:

$$V_\mathrm{m}(x,t) = U\big(t - \Psi(x)\big). \tag{13}$$

However, we may also introduce position-dependent action potentials $U\big(t - \Psi(x), x\big)$. In the case of Equation (13), thanks to the co-area formula, the ECG reads as follows

$$V(t) = \int_{-\infty}^{\infty} U'(t-\xi)\left(-\int_{\Psi^{-1}(\xi)} \nabla Z \cdot \mathbf{G}_\mathrm{i}\mathsf{n}\,\mathrm{d}\Sigma\right)\mathrm{d}\xi = (U' * w)(t), \tag{14}$$

that is, the convolution of the first derivative of $U(t)$ and the function $w(t)$ defined as:

$$w(t) := -\int_{\Psi^{-1}(\xi)} \nabla Z \cdot \mathbf{G}_\mathrm{i}\mathsf{n}\,\mathrm{d}\Sigma, \tag{15}$$

where $\Psi^{-1}(t) = \{x \in \Omega : \Psi(x) = t\}$ is the front surface within the heart and n is its normal oriented in the propagation direction. The function $w(t)$ is exactly the ECG $V(t)$ when the template action potential $U(t)$ is the Heaviside function, representing a sharp interface between the depolarized and the resting tissue.

We implemented two methods for the computation of the ECGs: one based on the general formula (11), termed "simple method" (SM), and one based on formulas (14) and (15), named "fast method" (FM). The former can account for action potential heterogeneity, which is important for the T-wave in the ECG. When focusing on the QRS complex, the main differences between the proposed methods are the way the cardiac sources are computed and resulting execution speed of the computation of the ECGs.

The SM computes the ECG by direct evaluation of Equation (11). The heart domain Ω is discretized into voxels $\mathsf{v} \in \Omega_h$. The functions Z and V_m are approximated by piecewise linear functions Z_h and $(V_\mathrm{m})_h$, respectively, thus ∇Z_h and $\nabla(V_\mathrm{m})_h(t)$ are piecewise constant vectors within each voxel v. The gradient of V_m is computed for each time step. The intracellular conductivity $\mathbf{G}_\mathrm{i}$ is assumed piecewise-constant. The time window is prescribed and the time step $\tau > 0$ is fixed (usually 1 ms). Consequently, the numerical approximation V_n of the ECG at time $t = n\tau$ reads:

$$V(n\tau) \approx V_n = \sum\nolimits_{\mathsf{v}\in\Omega_h} |\mathsf{v}|\ \nabla Z_h|_\mathsf{v} \cdot \mathbf{G}_\mathrm{i}|_\mathsf{v}\ \nabla(V_\mathrm{m})_h(n\tau)\big|_\mathsf{v}, \quad n = 0, \ldots, N_\tau. \tag{16}$$

The convergence rate with respect to h is quadratic, for smooth $Z(x)$ and $V_\mathrm{m}(x,t)$. In practice, $Z_h(x)$ is the solution of a discrete Laplace equation, which converges quadratically in the L^2 norm to the analytical solution (far from the electrodes, where the solution is not singular). The action potential depends on the activation time $\Psi_h(x)$ computed with the eikonal equation, and the convergence rate is linear or less in the L^∞ norm. Hence, the overall error of Equation (16) is dominated by the error on the activation time.

The FM, based on Equation (14), consists of two steps. First, the piecewise linear interpolants in time of $w(t)$, $w_n \approx w(n\tau)$, and of $U(t)$, $U_n \approx U(n\tau)$, are computed. Then, the surface over which the integral (Equation 15) is defined, is approximated with a triangular surface. This latter is produced from the activation map by means of the marching cubes algorithm (Lorensen and Cline, 1987; Newman and Yi, 2006), a well-known and commonly used algorithm for extracting triangulated isosurfaces from 3D

scalar fields. It has the important advantage of being intrinsically parallel.

The interpolated function $w_{n,h}$ reads as follows:

$$w(n\tau) \approx w_{n,h} = -\sum_{\mathsf{v}\in\Omega_h}\left(\nabla Z_h|_{\mathsf{v}}\cdot\mathbf{G}_{\mathrm{i}}\left(\sum_{\mathsf{t}\in\mathcal{T}_h(\mathsf{v},n)}|\mathsf{t}|\mathsf{n}_{\mathsf{t}}\right)\right), \tag{17}$$

where $\mathcal{T}_h(\mathsf{v},n)$ is the triangulation of the surface $\mathsf{v}\cap\Psi^{-1}(n\tau)$, and n_{t} is the normal to the triangle t.

It is important to note at this stage that $\Psi^{-1}(\tau n)\cap\mathsf{v}$ is empty for most of the values of n. Actually, it is non-empty for $\min_{\mathsf{v}}\Psi_h \leq \tau n \leq \max_{\mathsf{v}}\Psi_h$. This helps to reduce the computational cost significantly. For instance, with $\tau = 1$ ms, the activation front is found in the voxel on average only for 2 to 3 time instants. Moreover, the expression (Equation 17) can be further simplified since the normal n_{t} and the area $|\mathsf{t}|$ are related: let $\mathsf{v}_{1,\mathsf{t}}$, $\mathsf{v}_{2,\mathsf{t}}$ and $\mathsf{v}_{3,\mathsf{t}}$ be the vertices of the triangle t, then we have

$$|\mathsf{t}|\mathsf{n}_{\mathsf{t}} = \frac{1}{2}(\mathsf{v}_{2,\mathsf{t}}-\mathsf{v}_{1,\mathsf{t}})\wedge(\mathsf{v}_{3,\mathsf{t}}-\mathsf{v}_{1,\mathsf{t}}). \tag{18}$$

Finally, the second step is the convolution, approximated by

$$V(n\tau) \approx V_n = \frac{1}{2}\sum_{p\in\mathbb{Z}}(U_{p+1}-U_{p-1})\,w_{n-p}, \tag{19}$$

which results from a central finite difference approximation of $U'(t)$ combined with the midpoint quadrature rule.

2.4. GPGPU Implementation

The proposed computational methods for activation times and ECG were designed to run on GPGPU architectures. The power of GPGPU computation lies in a massive parallelism that can hide long-latency operations such as memory accesses. The latency hiding is possible thanks to the warp schedulers, which continuously switch among warps whose instructions are ready for execution.

The solvers were implemented in C and C++ with CUDA extension (Nickolls et al., 2008; Cheng et al., 2014), version 8.0. In all our experiments a single precision floating point representation of numbers was used, sufficient for stability and accuracy of the solvers.

The eikonal Equation (6) is solved by means of the Fast Iterative Method (FIM), proposed by Jeong and Whitaker (2008). Although a popular alternative for solving the eikonal equation is the Fast Marching Method (FMM) (Sethian, 1996), we opted for the FIM for two reasons: first, to increase the parallelization of the code and to make it suitable for GPGPU. Second, the updating of map (Equation 8) requires some care for anisotropic media in the FMM (Mirebeau, 2014).

Briefly, the FIM algorithm proceeds as follows: first, the activation map is initialized to $+\infty$ except at the boundary points; then, we iterate until a list of *active nodes* is empty. At each iteration, threads operate on single active nodes of the grid. For each node, the local minimization problem (Equation 10) is solved, and the new value is assigned to the node, according to Equation (8). When all nodes have been updated, all the neighbors of the active nodes that are not converged are appended to the active list, and activation time is re-computed. Active nodes whose activation times do not differ from their previous values are removed from the active list and marked as converged.

The implementation of ECG computation with the SM and FM is similar. Each cube is assigned to an individual thread, and each thread updates the local contribution of one cube to the ECG. For the SM, the kernel iterates over time and computes V_n according to Equation (16). The time-independent term $\mathbf{G}_{\mathrm{i}}\nabla Z$ is evaluated only once by each thread.

The FM has two kernels: one for the computation of $w(t)$ Equation from Equation (15), and another for the convolution (Equation 14). The first kernel implements the marching cubes algorithm through a look-up table that stores all the possible cases on how the front intersects the voxel. The marching cubes look-up table and the template action potential are stored in texture memory for fast indexing. The array w_n is updated concurrently by the threads. Each voxel contributes to w_n only when the front is within it (see **Figure 2**).

The computational complexity in time of the SM is $\mathcal{O}(L/h^3\cdot T_{\max}/\tau)$, where L is the characteristic length of the activation wavefront and $T_{\max}$ is the maximum activation time. In the FM, the voxel contribution is evaluated roughly $\approx \frac{h}{\tau\theta}$ times, where θ is the average conduction velocity within the voxel. Thus, the complexity in this case is $\mathcal{O}(L\,\theta^{-1}\,h^{-2}\,\tau^{-1})$. The ratio between SM and FM complexity is $\approx T_{\max}\theta/h$.

Since the computations of individual ECG leads are independent from each other, we also utilized parallel CUDA streams to support simultaneous execution of multiple kernels computing multiple ECG leads. All memory transfers between the CPU and GPU are asynchronous, in order to support also overlapping of the data transfers with the kernel executions.

2.5. Model Geometries

Patient-specific heart and torso geometries were created as described previously (Nguyên et al., 2015). Model parameters were tuned to each patient to match both the activation sequence on the endocardium and the twelve-lead ECG (Potse et al., 2014).

The model geometry was described in terms of a set of Catmull-Clark subdivision surfaces. The simulations were performed on semi-structured hexahedral meshes, i.e., while the basis of the mesh was regular, we did not store information for mesh nodes that did not play a role in the computation. These meshes were created by overlying the mesh base on the set of surfaces and assigning the properties of the mesh elements based on the surfaces in which they were included. Passive tissue characteristics were properties of elements, while active parameters and variables applied to nodes. The node properties followed from the element properties using a set of rules that ensured model consistency (Potse et al., 2006).

3. RESULTS

3.1. Numerical Assessment

In order to check the convergence of the method, we performed two numerical experiments. In both of them, we solved the

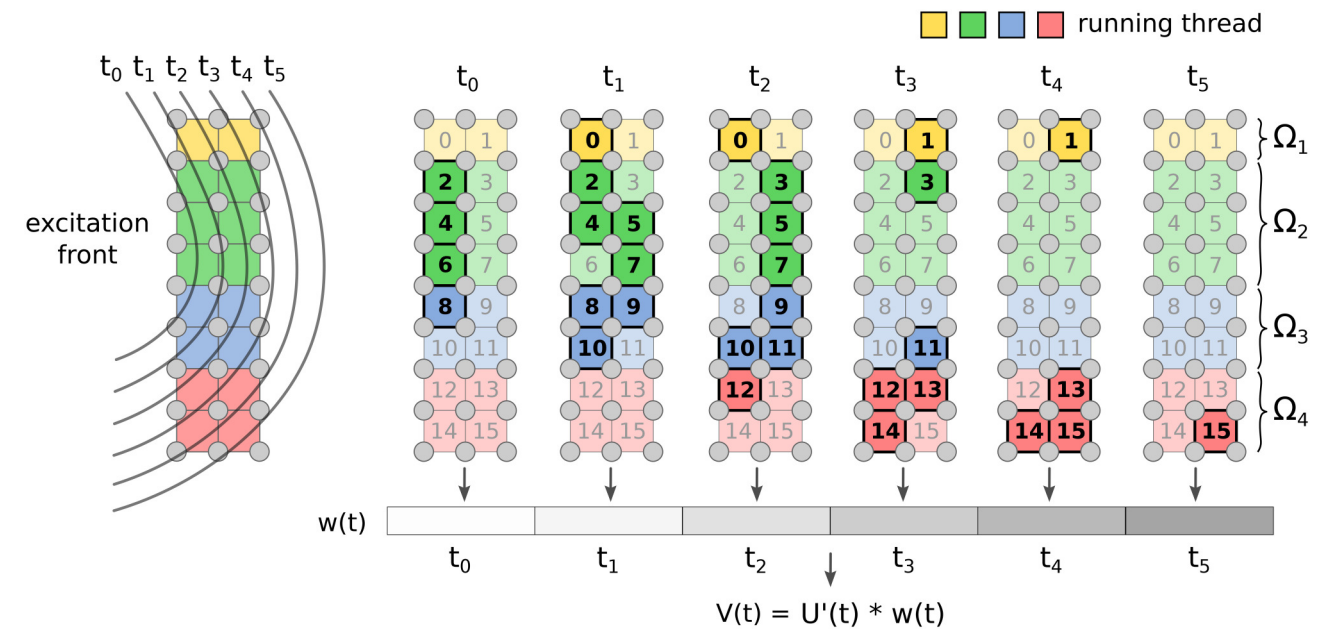

FIGURE 2 | Implementation of FM for ECG computation. Each voxel of the grid updates the value of $w(t)$ only for the time instants when the front is crossing it. Once all the voxels are done, the convolution is computed.

eikonal equation on a tissue slab of size $1.5 \times 2 \times 2$ cm, with a single pacing site at $\mathrm{x}_0 = (0, 1, 1)$ cm.

In the first experiment, parameters were selected such that an analytical solution for both the activation times and the function $w(t)$ defined in Equation (15) was available, and such that the ECG was easy to approximate with high accuracy. In particular, we considered fibers oriented along the z-axis. The ECG was computed with the lead-field function $Z(x, y, z) = -x$ and an action potential defined as follows:

$$U(t) = \mathrm{V}_{\mathrm{ref}} + \frac{\mathrm{V}_{\mathrm{dep}} - \mathrm{V}_{\mathrm{ref}}}{2}\left(\tanh\left(2\frac{t}{\varepsilon_{\mathrm{dep}}}\right) - \tanh\left(2\frac{t - \mathrm{APd}}{\varepsilon_{\mathrm{rep}}}\right)\right), \tag{20}$$

where the definitions and numerical values of the parameters are given in **Table 1**.

The exact solution for the activation times reads:

$$\Psi(\mathrm{x}) = \frac{1}{\alpha}\sqrt{\beta \mathbf{G}_{\mathrm{m}}^{-1}(\mathrm{x} - \mathrm{x}_0) \cdot (\mathrm{x} - \mathrm{x}_0)}. \tag{21}$$

The ECG solution is the convolution between the first derivative of the action potential (20) and the function $w(t)$. The latter was computed exactly, being proportional to the difference in depolarized area between the two faces of the slab orthogonal to the x-axis:

$$\begin{aligned} w(t) \sim\ & \mathrm{Area}\Big(\{(y, z) \in [0, 2]^2 \colon \Psi(0, y, z) \le t\}\Big) \\ & - \mathrm{Area}\Big(\{(y, z) \in [0, 2]^2 \colon \Psi(1.5, y, z) \le t\}\Big). \end{aligned} \tag{22}$$

The activated region on these two faces of the slab was always the intersection between an axis-aligned ellipse and a square. The corresponding area was derived from elementary geometry. See the Appendix in Supplementary Materials for further details and the explicit formula. The convolution was eventually computed numerically with high accuracy.

TABLE 1 | Values of the parameters adopted in the two convergence tests.

Parameter	Description	Value
α	Conduction velocity scaling	2.0 cm ms^{-1} mS$^{-1/2}$
β	Surface-to-volume ratio	1000 cm^{-1}
σ_{el}	Longitudinal extra-cellular conductivity	3 mS cm^{-1}
σ_{et}	Transverse extra-cellular conductivity	1.2 mS cm^{-1}
σ_{ec}	Cross extra-cellular conductivity	1.2 mS cm^{-1}
σ_{il}	Longitudinal intra-cellular conductivity	3 mS cm^{-1}
σ_{it}	Transverse intra-cellular conductivity	0.3 mS cm^{-1}
σ_{ic}	Cross intra-cellular conductivity	0.3 mS cm^{-1}
$\mathrm{V}_{\mathrm{ref}}$	Resting potential	−80 mV
$\mathrm{V}_{\mathrm{dep}}$	Depolarization potential	30 mV
$\varepsilon_{\mathrm{dep}}$	Depolarization time-scale	1 ms
$\varepsilon_{\mathrm{rep}}$	Repolarization time-scale	100 ms
APd	Action potential duration	200 ms

The convergence rate for the activation map was $\mathcal{O}(h^{0.6})$, as reported in **Figure 3**. This affected the ECG computation as well. This sub-linear convergence can be explained by the low regularity of the solution, which is singular at the source point. We also tested the ECG computation with the exact solution (Equation 21), in order to assess the convergence rate of the SM and FM alone. In this case, we observed quadratic convergence rate, as expected from the theory.

The SM and the FM perform similarly, with FM being slightly more accurate than SM at coarser grids. The high accuracy

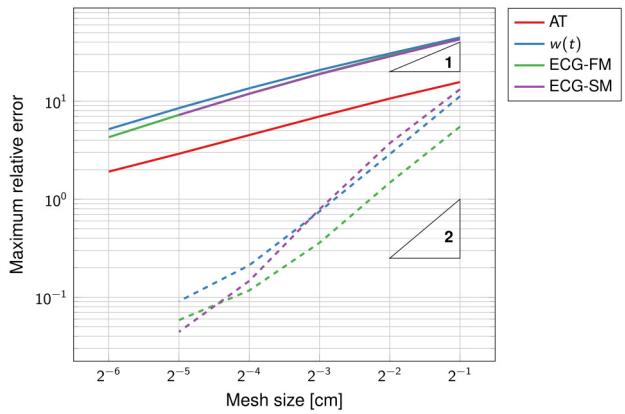

FIGURE 3 | Double-logarithmic plot of the error on the activation map, the function *w*(*t*) and the ECG. The dashed lines refer to the computation of the ECG using the exact activation map. The triangles are such that the slope of their longest side is equal to the indicated number, in logarithmic coordinates. The convergence rate of the solid lines is sub-linear.

of the SM is surprising because the template action potentials had very steep gradients in space (0.2 mm thick), clearly not captured by the coarse grid. But a closer inspection of the method showed that while the current per unit volume was not correctly approximated, the total current once integrated over the voxel was accurate.

The function $w(t)$, defined in Equation (15), was well-approximated already at the coarsest grid (**Figure 4**). As the mesh was refined, the support of the function narrowed, reflecting the fact that the total activation time was shortening.

The second numerical experiment mimicked a transmural specimen of left ventricle. Fibers were parallel to the yz-plane and formed an angle with respect to the y-axis linearly varying from $-\pi/3$ at $x = 0$ to $\pi/3$ at $x = 1.5$. Parameters were as in the previous experiment. The lead field function was the solution of the problem (Equation 12) in $\mathbb{R}^3$ with uniform conductivity $\sigma_t = 2\,\mathrm{mS\,cm^{-1}}$ and reads as follows:

$$Z(\mathsf{x}) = \frac{1}{4\pi\sigma_t}\left(\frac{1}{|\mathsf{x}-\mathsf{x}^-|} - \frac{1}{|\mathsf{x}-\mathsf{x}^+|}\right), \tag{23}$$

with negative terminal $\mathsf{x}^- = (3, 1, 1)\,\mathrm{cm}$ and positive terminal $\mathsf{x}^+ = (-1, 1, 1)\,\mathrm{cm}$.

In contrast to the first experiment, an analytical solution for the activation times, and thus for the ECG, was not available. The error was estimated using a solution obtained at a resolution of $2^{-8} \approx 0.004\,\mathrm{cm}$ as reference. The convergence rate was roughly $\mathcal{O}(h^{0.9})$; higher than in the previous experiment but still sub-linear. The function $w(t)$ and the ECG were not correctly captured at the coarser resolution, as certified by **Figure 4**. This was likely due to the too coarse representation of the fiber field.

From these experiments we concluded that a resolution of 1 or 0.5 mm for the eikonal solver provides reasonably accurate solutions, with a relative error lower than 5%.

3.2. Comparison to Bidomain with a Tissue Slab

The quantitative evaluation of the modeling error introduced by the proposed model with respect to the reference reaction-diffusion (R-D) model was conducted on a simple geometry but with physiologically motivated heterogeneities in the parameters. For this comparison we focused on the depolarization phase.

We considered a tissue slab shaped as in the previous section and embedded at the center of an electrically conductive bath of dimension $2.5 \times 3 \times 3\,\mathrm{cm}$. The fiber organization in the slab was exactly as in the second experiment above, but we also introduced a 0.5 mm thick rapidly conducting layer at the bottom face of the slab. The propagation was initiated in a $0.5 \times 1 \times 1\,\mathrm{mm}$ volume at the center of the rapidly conducting layer. The bath was composed of two homogeneous and isotropic media, one 1 mm thick placed on top of the bath, which we named "skeletal muscle," and the other, called "fluid," filling the remaining part of the volume. The activation time was measured only in the tissue slab. A bipolar ECG was also produced by taking the potential difference between two electrodes placed at the centers of the top and bottom faces of the bath. An overview of the experimental setting is depicted in **Figure 5**.

The parameters in the R-D model and the proposed model, summarized in **Table 2**, were the same when applicable. The electric conductivity in the tissue slab and the fast conducting layer was transversely isotropic, while it was assumed fully isotropic but heterogeneous in the bath. The surface-to-volume ratio was reduced in the fast layer in order to increase the conduction velocity with a factor 2. The ionic model, only present in the R-D model, was TNNP (ten Tusscher et al., 2004).

The value for α in the eikonal model was obtained from the conduction velocity θ observed in a 1d preparation of the R-D model (Equation 4) through the relation

$$\alpha = \sqrt{\frac{\beta}{\sigma}} \cdot \theta. \tag{24}$$

For instance, with $\beta = 1{,}000\,\mathrm{cm}^{-1}$, $\sigma = 1.5\,\mathrm{mS/cm}$ we observed $\theta = 0.075935\,\mathrm{cm/ms}$, the latter being 4-digit accurate, hence $\alpha = 1.961\,\mathrm{cm\,ms^{-1}\,mS^{-1/2}}$. From the same 1d preparation we also extracted the template action potential $U(t)$ by evaluating the resulting transmembrane potential.

The propagation was triggered differently in the two models. In the eikonal model, we assumed that the pacing region was activated with an initial time of 1.12 ms, the latter obtained from the monodomain simulation. This delay is associated to the time required by the current injection to start the propagation.

In order to minimize the numerical error with respect to the modeling error, we performed the comparison on a uniform grid with voxels of side 0.05 mm. The eikonal model and the monodomain equation were solved only in the tissue slab. The forward bidomain problem was eventually solved including also the bath, and injecting the transmembrane currents in the tissue slab. The lead field, depicted in the right panel of **Figure 5**, was computed from Equation (12) with a second order finite-difference scheme on a 0.1 mm spaced uniform grid and then interpolated at 0.05 mm resolution.

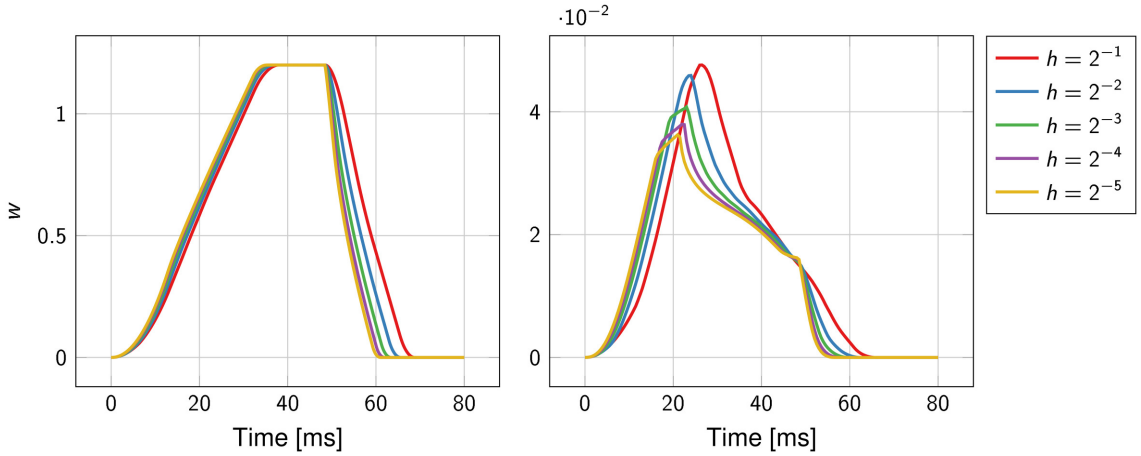

FIGURE 4 | Convergence history of the function *w*(*t*) for the first experiment (on the left) and the second experiment **(on the right)**. The mesh size *h* ranges from 2^{-1} cm to 2^{-5} cm.

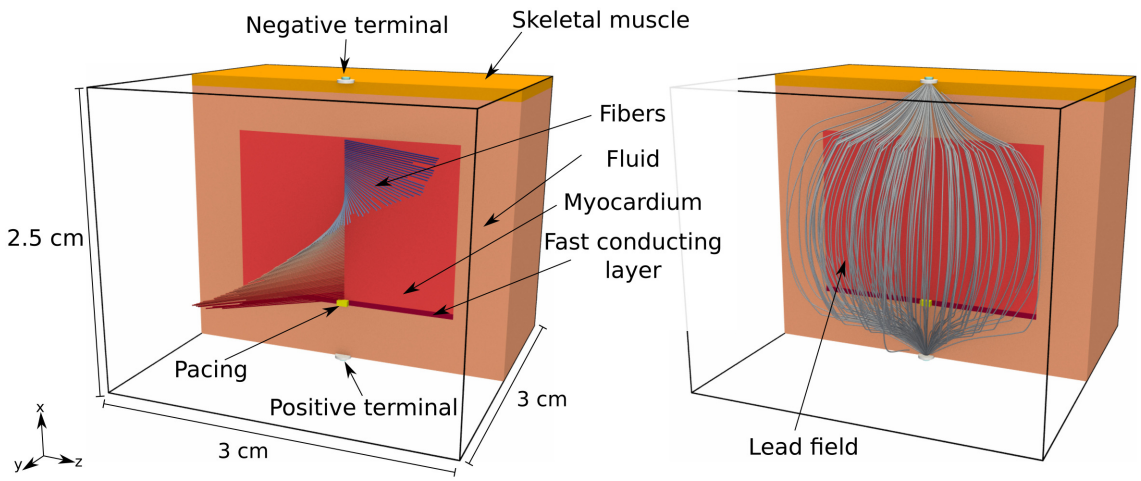

FIGURE 5 | Pictorial representation of the geometry of the tissue slab, embedded in a heterogeneous conductive bath, and the corresponding lead field. The tissue slab is 1.5 × 2 × 2 cm and positioned at the center of the bath. At the bottom of the slab there is a rapidly conducting layer, 0.5 mm thick, where the surface-to-volume ratio is reduced. At the top of the bath there is a 1 mm thick layer of skeletal muscle, with a reduced electric conductivity with respect to the underlying conductive medium (labeled "fluid" in the picture). There is a single early activation site at the bottom of the slab.

The results are reported in **Figure 6**. In the R-D model, the slab fully depolarized after 50.04 ms, while in our model the latest activation was at 49.48 ms. The maximum absolute and the L^2 errors were both around 3 ms, and for 95 % of the nodes the error was lower than 2 ms. In the R-D model, the activation occurred later than the activation from the eikonal model for all the nodes. Moreover, the largest error was localized at the boundaries and specifically in the fast layer.

We computed three ECGs in total: one with the R-D model, one with the proposed model, and one using the lead-field approach with the activation map obtained from the R-D model. The third was computed to evaluate the modeling error introduced by assuming a template action potential. The ECGs were very similar in shape and amplitude. In the terminal part of the QRS complex from the R-D model we observed a deeper inflection than in the other cases, and this seemed to be due to the assumption (Equation 13), since this was absent also in the case of the R-D model with lead field.

Finally, we performed the same comparison for a coarser grid resolution, observing a longer QRS duration in the R-D model. This is consistent with the fact that the finite-difference scheme underestimates the conduction velocity at coarser resolution (Pezzuto et al., 2016). Differences in the activation map for the eikonal model at coarser resolution were negligible.

3.3. Comparison to Bidomain with Patient-Tailored Anatomy

Similarly to the previous section, we conducted an extensive comparison between the proposed model and the reference

model on six patient-tailored anatomies, assuming the same parameters for both the models. Patients included in this study had a clinical indication for CRT; the ECG, cardiac resonance imaging, and an electrophysiological study, including 3D electro-anatomical mapping were clinically indicated. They were performed as part of the work-up for device selection, i.e., to evaluate scar location and extent, inducibility of ventricular tachycardia, and pacing site selection. The study was performed in compliance with the Declaration of Helsinki. The institutional review board approved the study protocol, and all patients gave oral and written informed consent for each investigation.

Characteristics of the heart meshes are listed in **Table 3**. In our heart models we assumed a layered structure of substances. Each ventricular model consisted of eight layers. Four transmural layers called *fast*, *endo*, *mid*, and *epi* were created in the right ventricle (RV) and another four similar transmural layers were created in the left ventricle (LV). During the experiments we set the conduction velocity to be faster in the *fast* layers of both ventricles, in order to mimic the presence of the Purkinje network, which our models did not include. Both in the eikonal model and in the reaction-diffusion model the increase in velocity was achieved by reducing the local value of β. For patient 3 only the RV was given a *fast* layer. Additional details are given in Potse et al. (2014) for patient 1 and 2 and in Nguyên et al. (2015) for patient 3 and 4.

The following electric conductivities were used for all the patients: torso 2 mS/cm; blood 6 mS/cm; lungs 0.5 mS/cm; skeletal muscle 3.55 mS/cm in the tangent plane and 0.44 mS/cm in the radial direction. The conductivities in the myocardium differed per patient as they were tailored to fit the measured ECG and activation sequence. A tuning was necessary for the activation time at the early activation sites (EASs) and the CV scaling parameter α, not present in the reaction-diffusion model. Left bundle branch block (LBBB) was assumed for all patients, and was modeled by placing the EASs only in the RV. The number of EASs for each patient is reported in **Table 3**.

The initiation of the activation in the R-D model was performed by injecting a transmembrane source current of 200 mA/cm^3 for 2 ms on 1 mm^3 of tissue centered at the corresponding EAS location. The currents can be applied at different times. In the eikonal solver, the activation time of the

TABLE 2 | Electric conductivity (mS cm^{-1}) and surface-to-volume ratio (cm^{-1}) of the media employed in the comparison test.

Medium	σ_{il}	σ_{it}	σ_{el}	σ_{et}	β
Tissue	3.0	0.3	3.0	1.2	800
Fast	3.0	0.3	3.0	1.2	356
Muscle	–	–	0.44	0.44	–
Fluid	–	–	6.0	6.0	–

TABLE 3 | Specification of the heart meshes for the eikonal model used in the experiments.

Patient	Mesh dimension	# Nodes	# Cubes	# EASs
1	165 × 91 × 119	191,734	151,959	3
2	205 × 107 × 151	292,585	235,502	1
3	186 × 123 × 126	289,252	230,968	1
4	178 × 117 × 147	350,584	290,166	2
5	156 × 91 × 108	222,951	187,534	1
6	238 × 145 × 162	493,006	413,252	4

Nodes are vertices in the mesh that represent active tissue. Cubes are non-void voxels of the geometry. The last column reports the total number of early activation sites (EASs) used in the experiments.

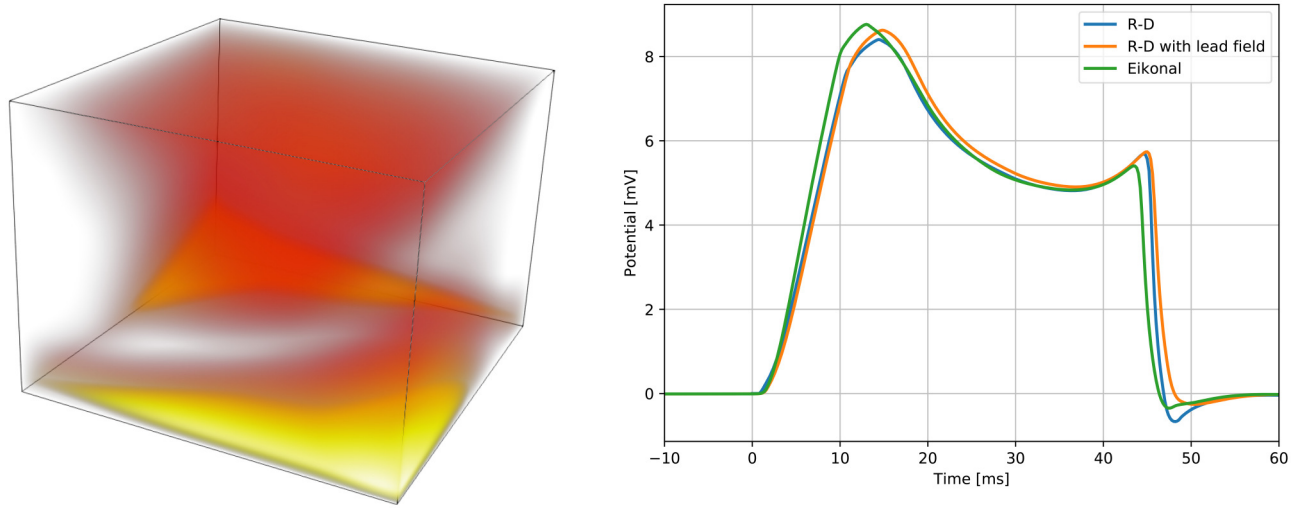

FIGURE 6 | Spatial modeling error between the proposed model and the reaction-diffusion model (left panel), and QRS complexes of the simulated ECGs **(right panel)**. On the left panel, the error is presented with a "heat map," being brighter where the error is larger. The brightest area of the map corresponds to an error of 3 ms. On the right panel, we report the QRS complex simulated by the reaction-diffusion model, by the proposed model, and by an intermediate approach that consists in using the lead field approximation with the activation map produced by the monodomain equation.

EASs was enforced as boundary condition and shifted in time in order to mimic the delay observed in the R-D model between current injection and depolarization.

The activation times were compared in terms of maximum absolute error quantiles at the nodes and specific timing markers of cardiological interest at the LV endocardium. In particular, we measured the Trans-Septal Time (TST) and the Total Activation Time (TAT), which are the time of initial and the latest time of activation of the LV endocardium, respectively, and the QRS complex duration (QRSd), defined as the last activation time. Results are reported in **Table 4**. A bull's eye plot of the LV endocardium is also provided in **Figure 7** for an average case in terms of error (Patient 1).

The overall propagation direction was captured for all the patients. The total activation time differed by 3 to 5 ms, being higher in the reaction-diffusion model. Thus, in the latter the excitation front was slower. The scaling parameter α was roughly 2 cm ms^{-1} mS$^{-1/2}$ in all cases, with a variability within 1 %.

The isochrones were very similar in shape. The absolute error was well-distributed in the anatomy, although we observed a slightly larger error at the apex and in the LV. The EASs regions also differed, with the eikonal solution more ellipsoidal than the R-D model.

On average, we found that 90% of the nodes had an absolute error lower than 10 ms. QRSd, TAT and TST were very similar in all the cases, with the exception of patient 2, where QRSd with the R-D model was very high. The difference was due to the lower resolution geometry employed in the eikonal model (1 vs. 0.2 mm of the R-D model), which may introduce artifacts in the propagation. In patient 2, the anatomy included a thin, late-activated ventricular bundle stemming from the posterior LV free-wall, observed in MRI. In the downsampled geometry, this bundle was poorly reproduced resulting in a shorter activation time.

For the sake of completeness, in **Figure 7** we also reported the electro-anatomical map performed on Patient 1 (NOGA® XP, Biologic Delivery Systems, Division of Biosense Webster a Johnson & Johnson Company, USA). Quantitatively, the mismatch between the measured time and the simulated time was on average 11 ± 8 ms for the R-D model and 14 ± 8 ms for the eikonal model. The fitting was performed on the R-D model in previous work (Potse et al., 2014; Nguyên et al., 2015).

The ECGs were computed according to the standard 12-lead ECG definition (Malmivuo and Plonsey, 1995; Macfarlane, 2012), i.e., three limb leads (I, II, III), three "augmented" limb leads (aVR, aVL, aVF) and six precordial leads (V_1 to V_6), with time resolution of 0.5 ms over a time window of 600 ms. The 12 lead fields were computed according to Equation (12) on a 3d finite-difference grid at 1 mm resolution. The total number of voxels was on average 30 million. Each computation took roughly a minute and 80 GB of memory in total using 576 MPI processes on 16 nodes of a Cray XC30 supercomputer.

The ECGs for all the patients were computed from the activation time provided by the eikonal model and then compared against the bidomain solution. We adopted the action potential template in Equation (20) with uniform parameters in space with values from **Table 1**, enabling the computation with the FM. The APD was tuned per patient to match the T-wave onset. An example for Patient 1 is provided in **Figure 8**, where we also report the measured ECG from the patient and a magnification of the QRS complex. In this case the APD was 250 ms.

The error analysis was limited to the QRS complex, since T-waves cannot be properly modeled by the eikonal approach. We compared the ECGs computed with the R-D model and the eikonal model in terms of: difference in maximum amplitude, area and positivity range, correlation, and L^2-dot product of the L^2-normalized ECGs, defined as follows:

$$\frac{\int_0^T V_1(t)\, V_2(t)\, \mathrm{d}t}{\left(\int_0^T V_1^2(t)\, \mathrm{d}t\right)^{1/2} \left(\int_0^T V_2^2(t)\, \mathrm{d}t\right)^{1/2}}, \tag{25}$$

where $V_1(t)$ and $V_2(t)$ are two signals. The closer the L^2-dot product to unity, the closer the two signals. We reported the L^2-dot product for all the leads and all the patients in **Figure 9**.

The ECGs computed with the SM and FM were very similar (not shown), with a small difference in amplitude in the QRS complex (FM was lower). The QRS complex was correctly captured in most of the leads for all the patients except for patient 6, where we observed large deviations in limb leads and moderate differences in precordial leads, whilst the activation map was fairly similar between the two approaches. This patient has a scar in the LV free wall that was modeled in the R-D model as purely passive. We suspect that in the downsampled geometry adopted by the eikonal model the scar was particularly jagged, hence affecting the ECG. In some cases our model provided a slightly larger amplitude in the signal than the bidomain model. Only in one case, for lead V5 of patient 5, we observed discordant QRS complexes. The discrepancies were particularly marked in leads with small amplitude. For patient 2, the discrepancy in the QRS duration was not present when calculated according to the ECG rather then activation map. As reported above, this was due to a ventricular bundle that did not affect the ECG. In general, the correlation between the bidomain model and the proposed model was very good in most of the leads, and the error lower than 0.5 mV.

3.4. Performance and Scaling

The runtimes of the proposed eikonal solver and ECG solvers on the GPU were tested for six patient-specific geometries with 1-mm resolution.

We tested our code on two different Nvidia GPUs: a low-end (LE) GeForce GT 650M on a laptop (384 cores, 950 MHz clock, 0.73 Tflops of theoretical peak performance), and a high-end (HE) GeForce GTX 1080 (2,560 cores, 1,607 MHz clock, 8.2 Tflops of peak performance) on a local cluster node.

In **Figure 10** runtimes of the Fast Iterative Method (FIM) for solving the eikonal model, and the Fast Method (FM) and Simple Method (SM) for the ECG computation are shown, on both the LE and HE GPU. In **Table 5** we report the speedup of HE GPU vs. LE GPU for FIM, FM and SM, and the speedup of the FM vs. SM for ECG computation. The reported variability in the computational time was due to the size of the heart and to the number of EASs. This was apparent for patient 6, whose

TABLE 4 | Activation times markers for the eikonal model compared to the reaction-diffusion equation (between parentheses), and quantiles of the absolute error between the two models, computed node-wise.

Patient	QRSd [ms]	TST [ms]	TAT [ms]	Error quantiles [ms]		
				0.50	0.75	0.90
1	155.48 (157.39)	45.93 (39.10)	109.54 (118.29)	7.11	10.77	12.91
2	184.20 (225.03)	20.15 (19.76)	136.58 (136.66)	5.02	8.95	11.99
3	137.95 (137.92)	42.04 (40.16)	87.98 (88.67)	2.58	5.06	8.06
4	143.15 (145.41)	38.47 (36.85)	104.67 (107.94)	3.76	6.15	9.15
5	144.64 (146.22)	32.01 (31.42)	105.84 (109.17)	2.64	4.52	6.83
6	165.60 (165.58)	30.04 (28.30)	128.40 (128.99)	2.88	5.02	7.72

The markers are the QRS duration (QRSd), the Trans-Septal Time (TST), which measures the time of first activation of LV endocardium, and Total Activation Time (TAT), which marks the latest activation time in the LV endocardium.

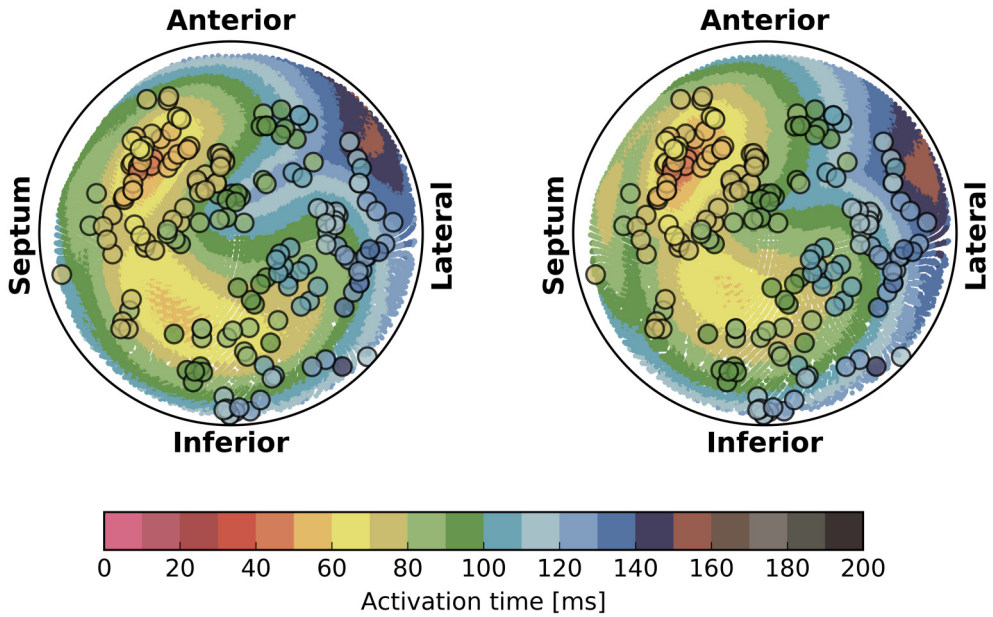

FIGURE 7 | Bull's eye plot of the activation time at the LV endocardium for the eikonal model (on the left) and the reaction-diffusion model **(on the right)**, in the case of Patient 1. The filled circles in the plot represent the measured activation time from the electro-anatomical map.

heart was extremely dilated and with 4 EASs reported by the electro-anatomical map.

4. DISCUSSION

A critical issue in computational cardiac electrophysiology is the complexity of the reaction-diffusion model. It is nearly impossible to consider this model for inverse problems, unless appropriate model reduction techniques are employed (Gerbeau et al., 2015; Quarteroni et al., 2017). Patient-tailoring of such models (Krueger et al., 2010; Potse et al., 2014) is also a tedious and time-consuming activity. The most expensive part of the procedure is the localization of early activation sites (EASs) and the determination of the local conduction velocity (CV). The proposed combination of an eikonal model for the activation sequence and an ECG computed with lead fields is a viable alternative for such a purpose. Given the EASs and the CV the activation map is found in a few seconds. Since the lead fields do not depend on these parameters, they can be computed once and repeatedly used for ECG simulation. The method can be used as part of inverse models to estimate patient-specific parameters and to make patient-specific predictions of cardiac activation patterns for complex procedures such as cardiac resynchronization therapy (CRT) (Pashaei et al., 2011). The computational power of the current parallel eikonal solver is so high, that our proposed methods are even suitable for real-time interactive simulations (Sermesant et al., 2007; van Oosterom et al., 2011) and therapy simulations based on virtual reality (Pernod et al., 2011).

The proposed method compares adequately to the bidomain model, in terms of accuracy of the delivered solution. The average modeling error in the patient-specific context was roughly 15 % for the activation map (**Table 4**), without adapting the parameters to compensate the numerical error. The ECG was less accurate in

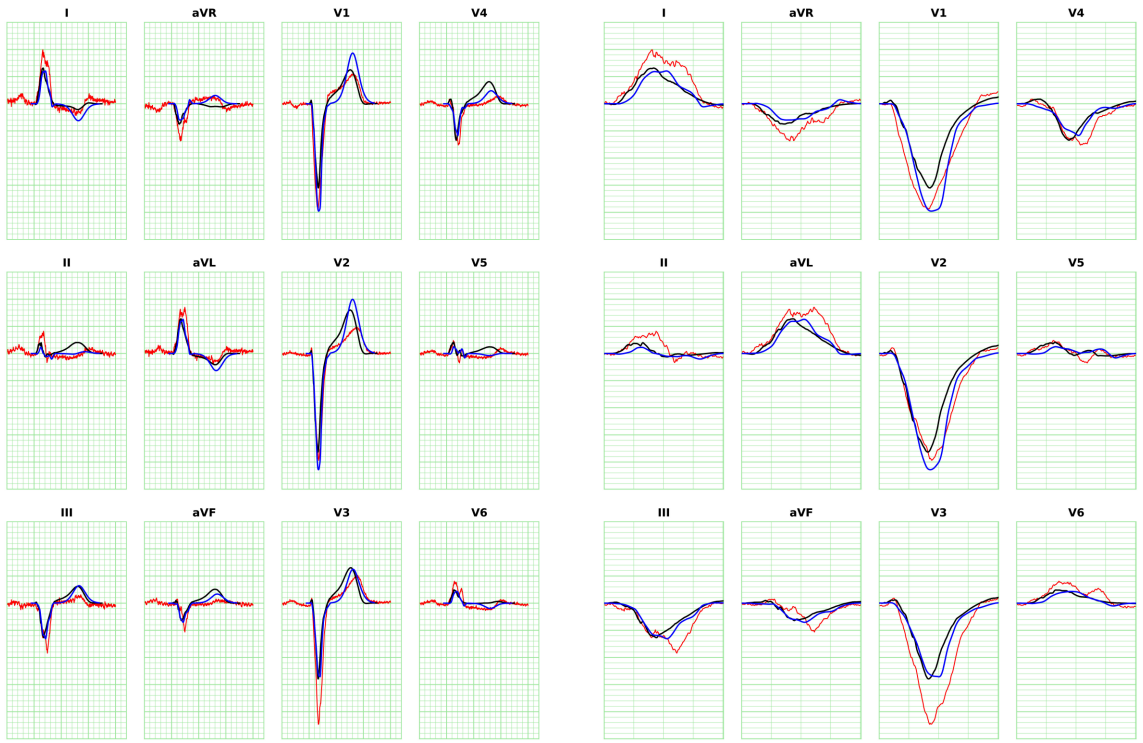

FIGURE 8 | ECG for Patient 1 (left) and magnification of QRS complex **(right)**, computed with FM. The blue curve is the solution with the eikonal solver and the lead fields formulation. The black curve is the bidomain solution. The minor grid lines are spaced 40 ms on the horizontal axis and 0.1 mV on the vertical axis. The measured ECG of the patient is shown in red.

absolute value, but the overall electrocardiographic features of the QRS complex were generally captured. It is well-known that the ECG is very sensitive even to small variations in the activation map. Thus, the mismatch between the reaction-diffusion model and the eikonal model was clearly visible, especially in leads where the potential is small.

The T-wave was almost never captured by our model. This was expected, since our model assumed that the repolarization follows exactly the depolarization, which was not the case in the reaction-diffusion model.

Finally, scalability was excellent for SM, and very good for the FIM and FM. Moreover, the SM and the FM provided very similar QRS complexes for all the patients and with different resolutions, but the FM is two orders of magnitude faster than the SM.

4.1. Appropriateness of the Eikonal Model

The eikonal model can be much faster than a reaction-diffusion equations for three reasons. First, the computational domain is reduced by one dimension, because the dependent variable is no longer a function of time. Second, the mesh can be coarser (1 vs. 0.1 mm), which results in much shorter computation times and much lower memory requirements. Last, it does not require computation of the ionic currents, which takes most of the computational time in the reaction-diffusion solution. Pullan et al. (2002) discussed that because the activation time, unlike the transmembrane potential, does not exhibit internal layers and is a smooth function in the cardiac tissue, it can be assumed that a 1-mm spatial scale is fine enough to accurately reconstruct the changes of the activation time as well as the changes in the speed of the wavefront propagation. They also discussed that sharp changes in the propagation speed can occur, for example when a wavefront collides with another wavefront or boundary, but the fine details of the wavefront shape in these small collision regions are not expected to have much influence on the overall ventricular function.

The numerical comparison on a simple geometry has shown that the eikonal equation is able to approximate the activation times of a bidomain model very accurately, and that there is no advantage in using the full bidomain formulation when modeling the excitation wavefront propagation if the conduction velocities are given.

The disadvantages of the eikonal model are obvious: the effects of ionic currents on conduction, for example partial refractoriness, cannot be simulated at all, simulation of reentrant arrhythmia is very difficult, and there is no calcium transient available for a realistic coupling to a mechanical model. Although several alternative eikonal-like models have been proposed in the literature aiming to improve the accuracy in such situations, such as fibrillation (Herlin and Jacquemet, 2011; Pernod et al., 2011) and high-order curvature effects (Dierckx et al., 2011), in our opinion reaction-diffusion models are necessary to study such phenomena.

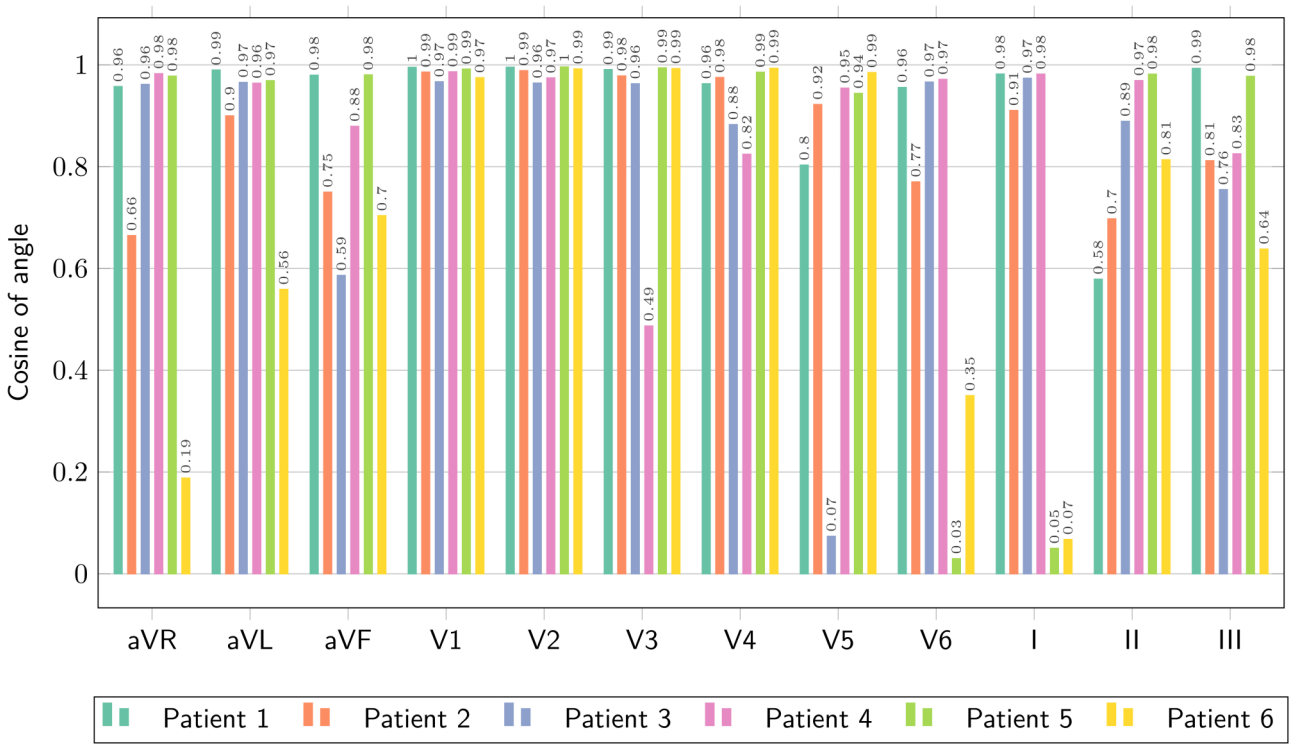

FIGURE 9 | Comparison of ECG obtained with the bidomain model and with the eikonal equation with lead fields, only limited to the QRS complex, in terms of the normalized dot-product of the two signals, as defined in Equation (25).

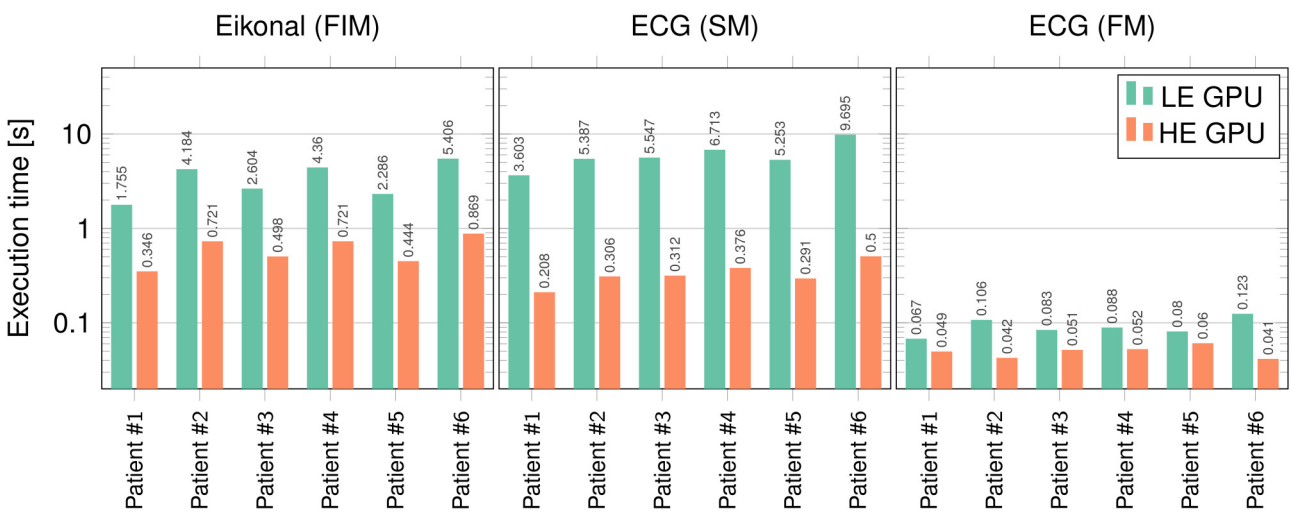

FIGURE 10 | Runtimes (in seconds) for FIM, SM, and FM on the LE and HE GPU. Logarithmic scale is used on the vertical axis.

4.2. Sources of Differences between Models

We believe that most of the discrepancy in the activation map between the R-D model and the proposed model is explained by the numerical error, and only minimally by the physiological phenomena that the eikonal model neglects.

The bidomain equation requires a mesh resolution that is comparable to the excitation front thickness, which depends on the conductivity in the propagation direction (Pezzuto et al., 2016). In practical situations it is in the order of 0.1 mm. In the cross-fiber direction, however, the conductivity is 10-fold lower, resulting in a steeper wavefront with thickness of 25 μm or lower, which the grid could not resolve. The effect is a significantly

TABLE 5 | Speedup factors of HE GPU vs. LE GPU for FIM, SM and FM, and FM vs. SM on both HE and LE GPU.

		Speedup factor						
	Patient	1	2	3	4	5	6	Average
HE vs. LE GPU	FIM	5.1	5.9	5.3	6.1	5.2	6.2	5.6
	SM	17.3	17.6	17.8	17.9	18.1	19.4	18.0
	FM	1.4	2.5	1.6	1.7	1.5	3.0	1.9
FM vs. SM	HE GPU	4.2	7.3	6.1	7.2	4.8	12.2	7.0
	LE GPU	53.8	51.3	67.7	76.3	57.1	79.5	64.3

slower CV and a higher total activation time. In this situation, we estimated a 20 % numerical error in the relative maximum norm, which is substantial.

The eikonal solver was affected by a 5 % error at 1 mm resolution on a simple geometry. Moreover, the activation time was generally overestimated by the solver, similarly to the slower propagation observed in the reaction-diffusion model when solved with a finite-difference scheme. However, it is hard to tell whether the differences between the reaction-diffusion model and the proposed model are due to numerical errors or modeling errors. Very likely it is a mix of the two.

In our opinion, the eikonal model may have been more accurate than the reaction-diffusion model in the presented experiments. A comparison of the two models in Section 3.2 reported a modeling error of about 6 %. A small modeling error is also reported by others for an idealized LV geometry with physiologically motivated fiber architecture (Colli Franzone et al., 2014).

4.3. Scalability Issues for the FIM and FM

The FIM ran on the HE-GPU on average only 5.6 times faster than on the LE-GPU. This is about half the ideal scaling. This may be explained by the fact that, at each iteration, only a fraction of the threads are active. In the FIM, only the activation times at the active nodes are evaluated, and hence only a few thousands of threads are running concurrently. A lower number of threads implies a lower number of active CUDA blocks and warps, and this may cause a paucity of warps prepared to hide the latencies of the long-lasting operations. The critical parts are the initial phase of the excitation process when the front starts to propagate from the EASs and the final phase when the last cells are activated. In these two phases only tens or hundreds of threads evaluate the solution and the GPU resources are not fully exploited.

Scalability of the FM was also not optimal. The FM ran on the HE-GPU on average only 4.5 times faster than on the LE-GPU. This means that the performance of the FM is more limited by the internal structure of the kernel than by the number of available GPU cores. The kernel generates one thread per cube in the computational domain. Warp divergence arises because individual threads compute their contribution only when the excitation front intersects the cube. With respect to the scaling, not the warp divergence itself is problematic, but rather the fact that different warps may require different numbers of clock cycles to finish their execution. The higher the number of time instants for which the warp thread evaluates the solution, the higher the warp execution time. Due to unbalanced execution times of different warps in the computational domain the linear scaling of the FM was broken.

5. CONCLUSIONS

We proposed a combination of an eikonal model for action potential propagation and an ECG simulation method based on lead fields. It can simulate an activation sequence and ECG in about 3 s on a GPGPU desktop platform (less than a second on HE-GPGPU). This method is suitable as a component of an inverse electrocardiographic model in an HPC context, but may in the near future also become practical within interactive ECG simulation tools. While we based our work largely on classical ideas, we believe that the following aspects are novel.

- We proposed a new method to compute the ECG, based on the marching cubes method, which allows for several orders of magnitude speedup.
- We compared the results of the eikonal model to those of a monodomain reaction-diffusion model, pointing out the difficulties of such a comparison.
- We compared the simulated ECGs with those of a bidomain torso model and identified the most important causes of differences.
- We discussed how the proposed methods can be implemented on a GPGPU.
- We showed that the proposed methods can simulate a highly realistic ECG in a few seconds.

AUTHOR CONTRIBUTIONS

The ideas behind this work were initially proposed by MP, who also provided the numerical experiments with propag-5 and the patient's anatomies. SP designed the numerical method, drafted the implementation and performed the simulations. PK implemented the GPGPU code and optimized it. He wrote the initial version of the manuscript, then extended and revised by SP and MP to the present state. FP contributed on several occasions to the manuscript by reviewing it. AA and RK respectively reviewed the clinical outcome and numerical aspects of the manuscript.

FUNDING

This work was supported by grants from the Swiss National Supercomputing Centre (CSCS) under project IDs s397, s598, and s669. PK was supported by the Sciex-NMS fellowship "CATION," project code 14.158. The authors gratefully acknowledge financial support by the Theo Rossi di Montelera Foundation, the Metis Foundation Sergio Mantegazza, the Fidinam Foundation, and the Horten Foundation to the Center for Computational Medicine in Cardiology. The authors also thank the Swiss PASC initiative (Platform for Advanced Scientific Computing) for their support. This work was also partially supported by a restricted grant of Biologic Delivery Systems, Division of Biosense Webster a Johnson & Johnson Company.

ACKNOWLEDGMENTS

The authors sincerely thank Hanspeter Fischer of Biosense Webster for his contribution in data acquisition and for making the data available for offline analysis.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: http://journal.frontiersin.org/article/10.3389/fphys.2017.00265/full#supplementary-material

REFERENCES

Bishop, M. J., and Plank, G. (2011). Bidomain ECG simulations using an augmented monodomain model for the cardiac source. *IEEE Trans. Biomed. Eng.* 58, 2297–2307. doi: 10.1109/TBME.2011.2148718

Bornemann, F., and Rasch, C. (2006). Finite-element discretization of static Hamilton-Jacobi equations based on a local variational principle. *Comput. Visual. Sci.* 9, 57–69. doi: 10.1007/s00791-006-0016-y

Cheng, J., Grossman, M., and McKercher, T. (2014). *Professional CUDA C Programming*. Indianapolis: Wiley.

Clayton, R., Bernus, O., Cherry, E., Dierckx, H., Fenton, F., Mirabella, L., et al. (2011). Models of cardiac tissue electrophysiology: progress, challenges and open questions. *Prog. Biophys. Mol. Biol.* 104, 22–48. doi: 10.1016/j.pbiomolbio.2010.05.008

Colli Franzone, P., and Guerri, L. (1993). Spreading of excitation in 3-D models of the anisotropic cardiac tissue. I. validation of the eikonal model. *Math. Biosci.* 113, 145–209. doi: 10.1016/0025-5564(93)90001-Q

Colli Franzone, P., Pavarino, L. F., and Scacchi, S. (2014). *Mathematical Cardiac Electrophysiology*. Switzerland: Springer International Publishing.

Coudière, Y., Bourgault, Y., and Rioux, M. (2014). Optimal monodomain approximations of the bidomain equations used in cardiac electrophysiology. *Math. Models Methods Appl. Sci.* 24, 1115–1140. doi: 10.1142/S0218202513500784

Dierckx, H., Bernus, O., and Verschelde, H. (2011). Accurate eikonal-curvature relation for wave fronts in locally anisotropic reaction-diffusion systems. *Phys. Rev. Lett.* 107:108101. doi: 10.1103/PhysRevLett.107.108101

Erem, B., van Dam, P., and Brooks, D. H. (2014). Identifying model inaccuracies and solution uncertainties in noninvasive activation-based imaging of cardiac excitation using convex relaxation. *IEEE Trans. Biomed. Eng.* 33, 902–912. doi: 10.1109/TMI.2014.2297952

Gerbeau, J.-F., Lombardi, D., and Schenone, E. (2015). Reduced order model in cardiac electrophysiology with approximated Lax pairs. *Adv. Comput. Math.* 41, 1103–1130. doi: 10.1007/s10444-014-9393-9

Geselowitz, D. B. (1989). On the theory of the electrocardiogram. *Proc. IEEE* 77, 857–876. doi: 10.1109/5.29327

Geselowitz, D. B. (1992). Description of cardiac sources in anisotropic cardiac muscle; application of the bidomain model. *J. Electrocardiol.* 25(Suppl.), 65–67. doi: 10.1016/0022-0736(92)90063-6

Henriquez, C. S. (2014). A brief history of tissue models for cardiac electrophysiology. *IEEE Trans. Biomed. Eng.* 61, 1457–1465. doi: 10.1109/TBME.2014.2310515

Herlin, A., and Jacquemet, V. (2011). Eikonal-based initiation of fibrillatory activity in thin-walled cardiac propagation models. *Chaos Interdiscipl. J. Nonlin. Sci.* 21:043136. doi: 10.1063/1.3670060

Jacquemet, V. (2015). Modeling left and right atrial contributions to the ECG: a dipole-current source approach. *Comp. Biol. Med.* 65, 192–199. doi: 10.1016/j.compbiomed.2015.06.007

Jeong, W.-K., and Whitaker, R. T. (2008). A fast iterative method for eikonal equations. *SIAM J. Sci. Comput.* 30, 2512–2534. doi: 10.1137/060670298

Krause, D., Potse, M., Dickopf, T., Krause, R., Auricchio, A., and Prinzen, F. W. (2012). "Hybrid parallelization of a large-scale heart model," in *Facing the Multicore-Challenge II*, *Vol. 7174* of *Lecture Notes in Computer Science*, eds R. Keller, D. Kramer, and J.-P. Weiss (Berlin: Springer), 120–132.

Krueger, M. W., Rhode, K., Weber, F. M., Keller, D. U. J., Caulfield, D., Seemann, G., et al. (2010). Patient-specific volumetric atrial models with electrophysiological components: a comparison of simulations and measurements. *Biomed. Tech.* 55(Suppl.1), 54–57.

Lorensen, W. E., and Cline, H. E. (1987). Marching cubes: a high resolution 3D surface construction algorithm. *ACM SIGGRAPH Comput. Graph.* 21, 163–169. doi: 10.1145/37402.37422

Macfarlane, P. W. (2012). *Basic Electrocardiology: Cardiac Electrophysiology, ECG Systems and Mathematical Modeling*. London: Springer Verlag.

Malmivuo, J., and Plonsey, R. (1995). *Bioelectromagnetism-Principles and Applications of Bioelectric and Biomagnetic Fields*. New York, NY: Oxford University Press.

McFee, R., and Johnston, F. D. (1953). Electrocardiographic leads; I. introduction. *Circulation* 8, 554–568. doi: 10.1161/01.CIR.8.4.554

Mirebeau, J.-M. (2014). Anisotropic fast-marching on cartesian grids using lattice basis reduction. *SIAM J. Numeri. Anal.* 52, 1573–1599. doi: 10.1137/120861667

Newman, T. S., and Yi, H. (2006). A survey of the marching cubes algorithm. *ACM SIGGRAPH Comput. Graph.* 30, 854–879. doi: 10.1016/j.cag.2006.07.021

Nguyên, U. C., Potse, M., Regoli, F., Caputo, M. L., Conte, G., Murzilli, R., et al. (2015). An *in-silico* analysis of the effect of heart position and orientation on the ECG morphology and vectorcardiogram parameters in patients with heart failure and intraventricular conduction defects. *J. Electrocardiol.* 48, 617–625. doi: 10.1016/j.jelectrocard.2015.05.004

Nickolls, J., Buck, I., Garland, M., and Skadron, K. (2008). Scalable parallel programming with CUDA. *Queue* 6, 40–53. doi: 10.1145/1365490.1365500

Nielsen, B. F., Ruud, T. S., Lines, G. T., and Tveito, A. (2007). Optimal monodomain approximations of the bidomain equations. *Appl. Math. Comput.* 184, 276–290. doi: 10.1016/j.amc.2006.05.158

Pashaei, A., Romero, D., Sebastin, R., Camara, O., and Frangi, A. F. (2011). Fast multiscale modeling of cardiac electrophysiology including Purkinje system. *IEEE Trans. Biomed. Eng.* 58, 2956–2960. doi: 10.1109/TBME.2011.2162841

Pernod, E., Sermesant, M., Konukoglu, E., Relan, J., Delingette, H., and Ayache, N. (2011). A multi-front eikonal model of cardiac electrophysiology for interactive simulation of radio-frequency ablation. *Comput. Graph.* 35, 431–440. doi: 10.1016/j.cag.2011.01.008

Pezzuto, S., Hake, J., and Sundnes, J. (2016). Space-discretization error analysis and stabilization schemes for conduction velocity in cardiac electrophysiology. *Int. J. Numer. Meth. Biomed. Eng.* 32: e02762-n/a. doi: 10.1002/cnm.2762

Potse, M., Dubé, B., Richer, J., Vinet, A., and Gulrajani, R. M. (2006). A comparison of monodomain and bidomain reaction-diffusion models for action potential propagation in the human heart. *IEEE Trans. Biomed. Eng.* 53, 2425–2435. doi: 10.1109/TBME.2006.880875

Potse, M., Dubé, B., and Vinet, A. (2009). Cardiac anisotropy in boundary-element models for the electrocardiogram. *Med. Biol. Eng. Comput.* 47, 719–729. doi: 10.1007/s11517-009-0472-x

Potse, M., Krause, D., Kroon, W., Murzilli, R., Muzzarelli, S., Regoli, F., et al. (2014). Patient-specific modelling of cardiac electrophysiology in heart-failure patients. *Europace* 16, iv56–iv61. doi: 10.1093/europace/euu257

Potse, M., and Kuijpers, N. H. L. (2010). Simulation of fractionated electrograms at low spatial resolution in large-scale heart models. *Comput. Cardiol.* 37, 849–852.

Pullan, A. J., Tomlinson, K. A., and Hunter, P. J. (2002). A finite element method for an eikonal equation model of myocardial excitation wavefront propagation. *SIAM J. Appl. Math.* 63, 324–350. doi: 10.1137/S0036139901389513

Quarteroni, A., Lassila, T., Rossi, S., and Ruiz-Baier, R. (2017). Integrated heart–coupling multiscale and multiphysics models for the simulation of the cardiac function. *Comput. Methods Appl. Mech. Eng.* 314, 345–407.

Ramanathan, C., Ghanem, R., Jia, P., Ryu, K., and Rudy, Y. (2004). Noninvasive electrocardiographic imaging for cardiac electrophysiology and arrhythmia. *Nat. Med.* 10, 422–428. doi: 10.1038/nm1011

Roth, B. J. (1997). Electrical conductivity values used with the bidomain model of cardiac tissue. *IEEE Trans. Biomed. Eng.* 44, 326–328. doi: 10.1109/10.563303

Sermesant, M., Konukoglu, E., Delingette, H., Coudière, Y., Chinchapatnam, P., Rhode, K. S., et al. (2007). "An anisotropic multi-front fast marching method for real-time simulation of cardiac electrophysiology," in *Functional Imaging and Modeling of the Heart*, *Vol. 4466* of *Lecture Notes in Computer Science*, eds F. Sachse and G. Seemann (Berlin; Heidelberg: Springer), 160–169.

Sethian, J. A. (1996). A fast marching level set method for monotonically advancing fronts. *Proc. Natl. Acad. Sci. U.S.A.* 93, 1591–1595. doi: 10.1073/pnas.93.4.1591

ten Tusscher, K. H. W. J., Noble, D., Noble, P. J., and Panfilov, A. V. (2004). A model for human ventricular tissue. *Am. J. Physiol. Heart Circ. Physiol.* 286, H1573–H1589. doi: 10.1152/ajpheart.00794.2003

Tysler, M., and Svehlikova, J. (2013). Noninvasive finding of local repolarization changes in the heart using dipole models and simplified torso geometry. *J. Electrocardiol.* 46, 284–288. doi: 10.1016/j.jelectrocard.2013.03.014

van Dam, P. M., Oostendorp, T. F., Linnenbank, A. C., and van Oosterom, A. (2009). Non-invasive imaging of cardiac activation and recovery. *Ann. Biomed. Eng.* 37, 1739–1756. doi: 10.1007/s10439-009-9747-5

van Oosterom, A., Oostendorp, T. F., and van Dam, P. M. (2011). Potential applications of the new ECGSIM. *J. Electrocardiol.* 44, 577–583. doi: 10.1016/j.jelectrocard.2011.05.006

Vigmond, E. J., Weber dos Santos, R., Prassl, A. J., Deo, M., Plank, G., and Bauer, S. (2007). Solvers for the cardiac bidomain equations. *Prog. Biophys. Mol. Biol.* 96, 3–18. doi: 10.1016/j.pbiomolbio.2007.07.012

Wang, L., Dawoud, F., Yeung, S.-K., Shi, P., Wong, K. C. L., Liu, H., et al. (2013). Transmural imaging of ventricular action potentials and post-infarction scars in swine hearts. *IEEE Trans. Med. Imaging* 32, 731–747. doi: 10.1109/TMI.2012.2236567

Conflict of Interest Statement: FP has received research grants from Medtronic, St. Jude Medical, Sorin, MSD, and Biotronik. AA has been a consultant to Medtronic, Boston Scientific, LivaNova and St. Jude, and has received speakers' fees from Medtronic, Boston Scientific and LivaNova.

The other authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2017 Pezzuto, Kal'avský, Potse, Prinzen, Auricchio and Krause. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.