Published in Chaos (2020), vol. 30, no. 3, pp. 033132

Time resolution for wavefront and phase singularity tracking using activation maps in cardiac propagation models

Samuel Gagné[1,2] and Vincent Jacquemet[1,2]

[1)] *Institut de Génie Biomédical, Département de Pharmacologie et Physiologie, Université de Montréal, Montréal, Canada*

[2)] *Centre de recherche, Hôpital du Sacré-Coeur, Montréal, Canada*

The dynamics of cardiac fibrillation can be described by the number, the trajectory, the stability and the lifespan of phase singularities (PS). Accurate PS tracking is straightforward in simple uniform tissues but becomes more challenging as fibrosis, structural heterogeneity and strong anisotropy are combined. In this paper, we derive a mathematical formulation for PS tracking in two-dimensional reaction-diffusion models. The method simultaneously tracks wavefronts and PS based on activation maps at full spatio-temporal resolution. PS tracking is formulated as a linear assignment problem solved by the Hungarian algorithm. The cost matrix incorporates information about distances between PS, chirality and wavefronts. A graph of PS trajectories is generated to represent the creations and annihilations of PS pairs. Structure-preserving graph transformations are applied to provide a simplified description at longer observation time scales. The approach is validated in 180 simulations of fibrillation in four different types of substrates featuring respectively wavebreaks, ionic heterogeneities, fibrosis and breakthrough patterns. The time step of PS tracking is studied in the range from 0.1 to 10 ms. The results show the benefits of improving time resolution from 1 to 0.1 ms. The tracking error rate decreases by an order of magnitude because the occurrence of simultaneous events become less likely. As observed on PS survival curves, the graph-based analysis facilitates the identification of macroscopically stable rotors despite wavefront fragmentation by fibrosis.

Nonlinear reaction–diffusion equations arising in cardiac electrophysiology can have complex non-stationary solutions taking the form of spiral-like patterns. The centers of these spirals are phase singularities (PS). In a heterogeneous medium, and even more so in the presence of discontinuities in medium properties, these PS may move arbitrarily fast and abruptly change direction. In addition, propagating waves may get fragmented, leading to the creation of a multitude of PS. The question arises of the time and space resolution needed to correctly track the trajectory of these PS. In the context of cardiac propagation modeling, this paper proposes and validates a PS tracking algorithm with an emphasis on micro-scale events and investigates how time resolution affects PS tracking accuracy.

I. INTRODUCTION

Cardiac fibrillation is observed as complex dynamics of depolarization waves that can be described by a reaction-diffusion equation.[1] Structural and functional reentries appear to play a central role in the perpetuation of arrhythmia.[2] The tip of a spiral-wave functional reentry is identified as a phase singularity (PS). The neighborhood of a PS includes cells in all phases of the cardiac cycle.[3] The number and the spatial and temporal stability of these PS provide measures of the dynamical complexity of fibrillation.[4] Their preferential locations highlight regions that are potentially critical for the maintenance of the arrhythmia and might be targeted for therapy.[5,6] The technique of phase mapping is used in the clinic to monitor atrial activity through catheters during ablation procedures.[7,8]

PS trajectories can be reconstructed by tracking the position of PS over time. The shape of meandering PS trajectories in a uniform substrate characterizes the properties of membrane kinetics,[9] and is notably associated with its restitution of action potential duration.[10] The stability of reentries can be quantified as the lifespan of its associated PS to reveal the presence of stable rotors.[11] The existence of such rotors is the cornerstone of the "mother rotor hypothesis"[12] and may indicate a reentrant circuit within fibrotic regions.[13] Failure to adequately track a PS may result in a wide underestimation of its lifespan, leading to potentially inaccurate clinical interpretation.

While many PS detection approaches have been proposed,[14–26] few methodologies for PS tracking have been published in details and carefully validated. Most PS tracking methods are based on the distance between matching PS and the comparison of directions of rotation,[15,27] on particle tracking,[26,28] or are unspecified in the article. Rogers proposed an algorithm that combines PS and wavefront tracking to increase the accuracy of PS tracking.[14] More advanced methods have been proposed to improve robustness of PS detection against signal noise in uniform continuous tissues.[21,26] As these methods rely on spatial filtering (moving average[21] or Gaussian[26]), their extension to heterogeneous, fibrotic structures with discrete uncoupling (e.g. collagen septa) is not straightforward and may blur or eliminate micro-scale propagation such as micro-reentries.

PS analysis is often used to compare computer modeling and experiments. In this paper, the focus is on simulations of reaction-diffusion models. The specificity of simulated data is that high spatial and temporal resolution is available and that full-scale analysis is desired. Since the membrane potential is accessible, activation times can be accurately determined. Deng et al. acknowledged the difficulty of tracking PS in atrial models with severe fibrosis and suggested that activation maps are more effective than phase maps for identifying stable rotors.[29] Several features of fibrillatory substrate models give rise to false positive PS detection and incorrect PS tracking despite the absence of noise. In addition to dynamical instability caused by beat-to-beat variation in refractoriness, heterogeneity in membrane properties induces wavebreaks and fast PS movement when hitting a refractory region. In contrast to wavefront velocity, the speed of PS trajectory is unbounded. Structural heterogeneity and fibrosis create a discrete substrate, while the theory of PS is continuous by essence.[30] These disconnections also fractionate wavefronts, which gives rise to many PS with close locations, making PS tracking more challenging. Finally, three-dimensional structures may lead to breakthrough activation patterns as well as curved scroll wave filaments.[31–34]

In this paper, we propose a mathematical formulation of PS tracking as an optimization problem. The main focus is placed on the tracking part rather than on PS detection. Extending Rogers' algorithm[14] to the analysis of activation maps at high spatio-temporal resolution, wavefront and PS tracking are performed simultaneously. A distance optimization criterion not present in Rogers' algorithm is integrated to handle cases where multiple events (wave breaks or collisions) occur in the same wavefront during the same time step. The time resolution needed for PS tracking is determined and tracking accuracy is studied

during fibrillation in a large range of simulated substrates. Although results are presented for planar sheets of tissue, special attention is devoted to generalizability to a wider range of microstructure models with discrete connections. Finally, after the graph of PS trajectories is constructed, graph manipulation tools are proposed to change the observation scale and coarsen the representation of the dynamics to extract the main macroscopic rotors, even in fibrotic tissue. These graph transforms are also essential for comparing PS tracking at different time resolutions.

II. METHODS

A. Discrete cell network

The cardiac tissue is described here as a discrete network of cells in two dimensions. Mathematically, it can be represented as a planar graph. Vertices are cells, edges are electrical connections between two neighboring cells and faces are polygonal domains delimited by edges. Vertices are indexed by $i \in \mathcal{V}$ and edges are indexed by $(i,j) \in \mathcal{E}$ with $i,j \in \mathcal{V}$. $\mathcal{E}$ is defined as containing unique edges so only one of (i,j) or (j,i) may be in $\mathcal{E}$. Two edges are adjacent if they share a common vertex. The adjacency relation will be used to identify connected sets of edges.

In a M-by-N rectangular grid, there are MN vertices, $M(N-1)+N(M-1)$ edges and $(M-1)(N-1)$ square faces. An edge in the bulk of the tissue has 6 adjacent edges, an edge in the border has 4, and the 8 edges in the corners have 3.

B. Activation maps

When a simulation of electrical propagation is run on a discrete tissue, the time course of the membrane potential $V_i(t)$ is computed for each $i \in \mathcal{V}$. Activation times of cell i are defined as time instants t_a such that $V_i(t_a) = V_{th}$ (threshold crossing) and $V_i'(t_a) > 0$ (positive time derivative). The parameter V_{th} is referred to as the activation threshold, and is here set to -40 mV. During the simulation, such pairs (i, t_a) are written to a file each time a cell activation is detected. This hard thresholding approach is usually sufficient for simulated data but it is not robust. In the presence of noise, signal filtering may be needed.[26]

Activation map analysis will be based on the following reasonable hypotheses:

1. Electrical excitation propagates from a cell to one of its neighbor within τ_d (maximum propagation delay); otherwise, a conduction block is considered to have occurred. This assumes that there is a minimum conduction velocity below which propagation is not possible. Typically, τ_d is of the order of 5 to 10 ms depending on mesh resolution.

2. Two consecutive activations of the same cell are separated by a time interval of at least τ_r. This minimal refractory period is assumed to be longer than the maximum propagation delay ($\tau_r > \tau_d$). Generally, τ_r is at least 30 to 50 ms.

3. The time resolution Δt of activation map analysis (wavefront tracking in particular) satisfies $\Delta t < \tau_r - \tau_d$. Otherwise, one may miss an entire wave.

The activation map A_i^t at time t is the time of the latest activation of each cell i prior to time $t + \tau_d$:

$$A_i^t = \sup\{t_a \mid t_a < t + \tau_d \text{ and } V_i(t_a) = V_{th} \text{ and } V_i'(t_a) > 0\} \tag{1}$$

with the convention that $A_i^t = -\infty$ if cell i has never been activated before time t or if no earlier data are available. Time dependency will be indicated by an exponent t. Activations up to time $t + \tau_d$ are included to verify if the wavefronts are still propagating at time t. As a result, if $A_i^t < t$ then $A_j^t \geq A_i^t$ if the excitation propagates from i to j.

As time evolves, only a part of the activation map is updated: $A_i^{t+\Delta t} = A_i^t$ as long as $t + \Delta t < A_i^t + \tau_r$. In particular, our third hypothesis $\Delta t < \tau_r - \tau_d$ implies that

$$A_i^t \in (t - \tau_d, t + \tau_d) \implies A_i^{t+\Delta t} = A_i^t \quad . \tag{2}$$

C. Wavefront detection

An edge $(i, j) \in \mathcal{E}$ is said to cross an isochrone at time t, and we write $(i, j) \in \mathcal{I}^t$, if and only if

$$\big(A_i^t < t \quad \text{xor} \quad A_j^t < t\big) \quad \text{and} \quad |A_j^t - A_i^t| < \tau_d \ , \tag{3}$$

where xor is the "exclusive or" logical operator. This means that one cell is behind the wavefront, the other one is in front of the wavefront, and the activation delay is small enough to assume causality.

A wavefront is a subset of $\mathcal{I}^t$ in which all edges cross the same continuous segment of an isochrone. The edge-to-edge relation "belongs-to-the-same-wavefront" can be determined in

each square face using a look-up table shown in Fig. 1. Up to rotational symmetry, there are 9 distinct cases (panels A to I). As with the marching square algorithms, in cases G to I, there are two possible choices of isochrone interpolation within the square. Our choice favors wavefront collisions over breakthroughs, although the global impact is very small. Also, cases G and H are relatively rare. Note that cases J and K are impossible (notations are illustrated in panel A). The first one is obvious since it would imply $t < t_2 \leq t$. In case K, we have $t_3 \leq t < t_0$ and $t_1 \leq t < t_2$. Moreover, $|t_0 - t_3| + |t_2 - t_1| = t_2 - t_3 + t_0 - t_1 < 2\tau_d$. This means that at least one of the two vertical edges satisfy (3) and we are back to one of the cases G, H or I.

Partitioning $\mathcal{I}^t$ into wavefronts comes down to determining the equivalence classes of the relation "belongs-to-the-same-wavefront". This can be efficiently performed using a dedicated data structure such as weighted Quick-Union with path compression.[35] It starts with each edge in $\mathcal{I}^t$ being its own wavefront. Then, iterating over the faces, wavefronts are merged whenever a connection is found according to Fig. 1. The end result is a decomposition of $\mathcal{I}^t$ into a union of pairwise-disjoint wavefronts

$$\mathcal{I}^t = \bigcup_{k=1}^{m^t} \mathcal{F}_k^t \ , \tag{4}$$

where m^t is the number of wavefronts at time t and $\mathcal{F}_k^t$ is the k-th wavefront. By construction, each edge of $\mathcal{F}_k^t$ is connected to one or two other edges. The method guarantees that there are either zero or two edges in $\mathcal{F}_k^t$ that are connected at exactly one edge: the extremities of the wavefront. These extremities may be a phase singularity or an anchor point on a boundary. Note the singular situation where $\mathcal{F}_k^t$ has a single element, which means that the wavefront has length zero. In that case, the two extremities are the same edge. It represents propagation through the thinnest possible isthmus.

D. Wavefront tracking

Assuming that wavefronts are detected every Δt, the objective is to identify the correspondence between the m^t wavefronts $\mathcal{F}_k^t$ at time t and the $m^{t+\Delta t}$ wavefronts $\mathcal{F}_l^{t+\Delta t}$ at time $t + \Delta t$.

A m^t-by-$m^{t+\Delta t}$ matrix G_{kl}^t is constructed to encode that correspondence: $G_{kl}^t = 1$ when $\mathcal{F}_l^{t+\Delta t}$ results from the propagation of $\mathcal{F}_k^t$ and $G_{kl}^t = 0$ if these wavefronts are independent.

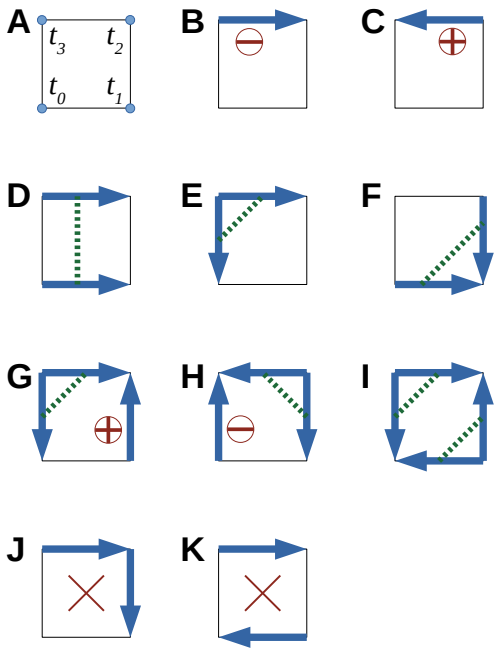

FIG. 1. Look-up table for identifying wavefronts. Arrows represent edges that cross an isochrone. Their direction indicates wavefront propagation. Dotted lines connect edges that belong to the same wavefront. The plus and minus signs show the presence of a phase singularity. (A) Vertex numbering and activation times t_0 to t_3 in a square face. (B)–(C) Wavefront endpoints. (D) Normal wavefront propagation. (E) Convex wavefront. (F) Concave wavefront. (G)–(H) Two independent wavefronts. (I) Wavefront collision or breakthrough. (J)–(K) Impossible cases (as emphasized by the big cross).

If a row of G^t is zero, i.e. $G^t_{kl} = 0\ \forall l$, the wavefront $\mathcal{F}^t_k$ disappears. If a column of G^t is zero, i.e. $G^t_{kl} = 0\ \forall k$, the wavefront $\mathcal{F}^{t+\Delta t}_l$ is a new wavefront. In the general case, several wavefronts may simultaneously merge and/or break within Δt which leads to a many-to-many relationship.

If $\mathcal{F}^t_k \cap \mathcal{F}^{t+\Delta t}_l \neq \varnothing$, then clearly $G^t_{kl} = 1$. However, if the wavefront moves too fast, the intersection may be empty so that the region between the successive fronts needs to be considered. That region is defined as the wavefronts at time t and $t + \Delta t$ plus the set of edges whose activation times fall entirely in the band from t to $t + \Delta t$:

$$\mathcal{R}^{t,\Delta t} = \mathcal{I}^t \cup \mathcal{I}^{t+\Delta t} \cup \left\{(i,j) \in \mathcal{E} \;\middle|\; A^{t+\Delta t}_i, A^{t+\Delta t}_j \in [t, t+\Delta t)\right\} \ . \tag{5}$$

As a side note, $\mathcal{I}^t$ is computed from A^t and $\mathcal{I}^{t+\Delta t}$ from $A^{t+\Delta t}$. According to (2), $\mathcal{I}^t$ would

be the same if computed from $A^{t+\Delta t}$, which ensures consistency.

The region $\mathcal{R}^{t,\Delta t}$ is decomposed into a disjoint union of connected components $\mathcal{R}_c^{t,\Delta t}$. The edge-to-edge relation "belongs-to-the-same-component" is defined as edges being adjacent (subsect. II A) or in the same wavefront (in Fig. 1D, wavefront edges are not adjacent). Connected components are identified using the same Quick-Union data structure as for wavefronts.

By definition, $\mathcal{F}_k^t \subset \mathcal{R}^{t,\Delta t}$. It is also true that $\mathcal{F}_k^t \subset \mathcal{R}_c^{t,\Delta t}$ for some c by the construction of the connected components. The same holds for $\mathcal{F}_l^{t+\Delta t}$. If $\mathsf{region}(\mathcal{F})$ is defined as the unique c for which $\mathcal{F} \subset \mathcal{R}_c^{t,\Delta t}$, a general formula for G^t is obtained:

$$G_{kl}^t = \begin{cases} 1 & \text{if } \mathsf{region}(\mathcal{F}_k^t) = \mathsf{region}(\mathcal{F}_l^{t+\Delta t}) \\ 0 & \text{otherwise} \end{cases} . \tag{6}$$

E. Phase singularity tracking

Phase singularities (PS) are identified as the edges that lie at the extremities of wavefronts based on the look-up table of Fig. 1. By abuse of language, anchor points of wavefronts on a boundary are considered to be PS. This will make it possible to continuously track PS that intermittently hit a boundary.

At time t, the number of PS is denoted by n^t. Detected PS are indexed by p and characterized by three attributes. Their location $\mathbf{x}_p^t$ is the midpoint of the edge at the extremity of a wavefront. Their chirality s_p^t expresses the direction of rotation ($+1$ or -1) and is also obtained from the look-up table of Fig. 1. Finally, they are part of a wavefront $w_p^t \in \{1, \ldots, m^t\}$. The pair (w_p^t, s_p^t) uniquely identifies a PS at time t.

PS tracking is an assignment problem between the n^t PS at time t and the $n^{t+\Delta t}$ PS at time $t + \Delta t$. Special cases $n^t = 0$ or $n^{t+\Delta t} = 0$ are straightforward but have to be handled separately. A mathematical constraint is that when a PS moves, it follows its associated wavefront and its chirality is preserved. Also, all other things being equal, shorter moves are more likely.

To specify the optimization problem, the n^t by $n^{t+\Delta t}$ distance matrix is first defined as

$$D_{pq} = g\left(\left\|\mathbf{x}_p^t - \mathbf{x}_q^{t+\Delta t}\right\|\right) \quad \text{where} \quad g(d) = \frac{d^2}{d+\epsilon} . \tag{7}$$

We used $\epsilon = 10^{-6}$ in units of grid step. Since $d - \epsilon < g(d) \leq d$, the matrix D is essentially

a distance matrix. The cost matrix incorporates a penalty for mismatched wavefronts:

$$C_{pq} = D_{pq} + B\left(1 - G^t_{w_p^t, w_q^{t+\Delta t}}\right) \ , \tag{8}$$

where the "big enough" number $B = 1+\min(n^t, n^{t+\Delta t})\cdot\max_{p,q} D_{pq}$. The cost will be summed over the matching pairs of PS.

The rationale for using the distance in the cost comes from the triangle inequality. Assume that two PS located at P_1 and P_2 at time t have to be matched with two PS located at Q_1 and Q_2 at time $t+\Delta t$. If the segments P_1Q_2 and P_2Q_1 intersect at O, then $P_1Q_2 + P_2Q_1 = P_1O + OQ_2 + P_2O + OQ_1 \geq P_1Q_1 + P_2Q_2$. In other words, two PS moving in parallel are guaranteed to give a lower cost than PS trajectories crossing each other. The inequality is strict unless the points P_1, P_2, Q_1, Q_2 are aligned (since the points are on a grid, this occurs occasionally). The function g solves this indeterminacy. For instance, if P_1, P_2, Q_1, Q_2 are on a line in that order, $P_1Q_1 + P_2Q_2 = P_1Q_2 + P_2Q_1$, but a little bit of algebra shows that $g(P_1Q_1) + g(P_2Q_2) < g(P_1Q_2) + g(P_2Q_1)$ unless $P_1P_2 = 0$ or $Q_1Q_2 = 0$ (which would mean two superimposed PS). The difference in cost is $\mathcal{O}(\epsilon^2)$. On the other hand, the function g has little effect in the general case where the difference is $\mathcal{O}(1)$.

An interesting analogy can be drawn with colloidal particle tracking.[28] If these particles undergo Brownian motion with the same diffusion coefficient, the log-likelihood of pair matching is proportional to the sum of squared distances. Cost minimization is therefore equivalent to maximum likelihood estimation (in our case with a slightly different statistical model).

For cost minimization, let us first assume that $n^t \leq n^{t+\Delta t}$ and ignore chirality. The cost function attributes a cost to each assignment function f that maps in a one-to-one way any PS at time t to a PS at time $t+\Delta t$:

$$C[f] = \sum_{p=1}^{n^t} C_{p,f(p)} \ . \tag{9}$$

The linear assignment problem consists in finding the optimal assignment $f_{opt} = \mathrm{argmin}_f\, C[f]$ that minimizes the cost. It can be efficiently solved in at most $\mathcal{O}(n^3)$ time using the Hungarian algorithm.[36] The PS of index $f_{opt}(p)$ is considered to be the evolution of the PS of index p if $C_{p,f_{opt}(p)} < B$. Otherwise, the PS of index p is said to have been annihilated. Reciprocally, a PS of index q at time $t+\Delta t$ is said to have been created if there is no p such that $f_{opt}(p) = q$ and $C_{p,q} < B$. The justification for the value of B is the following. If the

assignment f has no wavefront mismatch, i.e. $G^t_{w_p^t, w_{f(p)}^{t+\Delta t}} = 1\ \forall p$, then $C[f] < B$ whatever the configuration and distance matrix. The algorithm will naturally provide the solution that minimizes the number of wavefront mismatches.

If $n^t < n^{t+\Delta t}$, the problem is simply solved by applying the algorithm to the transposed cost matrix. To take into account chirality, the whole optimization procedure is applied separately to the subset of PS with positive chirality at time t and $t + \Delta t$ and then to the ones with negative chirality.

The outcome of PS tracking in the time interval $[t, t + \Delta t]$ is: (1) a list of pairs of PS (p, q) such that p becomes q at time $t + \Delta t$, (2) a list of PS pairs that annihilate each other (see also Subsect. II G), and (3) a list of PS pairs that are created.

To compare with previous works,[15,27] we implemented an alternative strategy for cost minimization. The smallest entry in the cost matrix, i.e. the closest pair of PS at time t and $t+\Delta t$, is first identified. Then, PS assignment is updated accordingly and the corresponding row and column are removed from the matrix. This procedure is then applied iteratively until the matrix is empty. This will be referred to as the "greedy" algorithm. It approximates the optimal solution. Its time complexity is also $\mathcal{O}(n^3)$, but may be reduced to $\mathcal{O}(n^2 \log n)$ by sorting entries of the cost matrix.

F. Ambiguous cases

Ambiguity may arise when multiple events occur within the same wavefront during one single time step. The assignment algorithm always tries to match as many PS pairs as possible, which sometimes leads to plausible but unlikely assignment. In the example illustrated in Fig. 2, two wavefronts 1 and 2 at time t move forward and become wavefronts 3 and 4 at time $t + \Delta t$ (thus $G^t_{1,3} = G^t_{2,4} = 1$). The algorithm will predict that the PS at P_- moves to Q_- and the PS at P_+ moves to Q_+ whatever the distance between the two pairs of PS. Another, equivalently possible scenario would be that P_- and P_+ annihilate and a pair is created at Q_- and Q_+.

Although a PS can move very fast when hitting a refractory or non-propagating region, it is less likely that two close PS move very fast in parallel. Therefore, the assignment $P_- \mapsto Q_-$ and $P_+ \mapsto Q_+$ is turned into an annihilation and a creation when the following conditions are met: (1) two PS of opposite chirality simultaneously jump to locations further

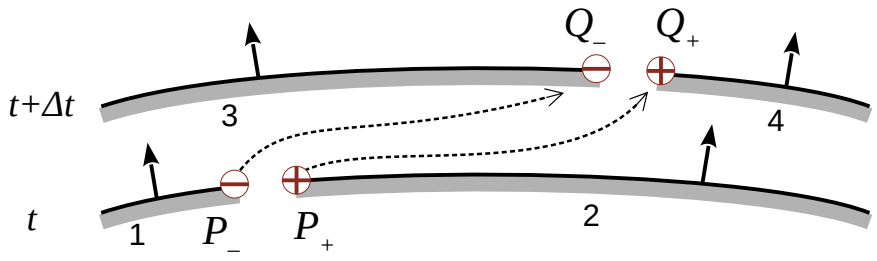

FIG. 2. Ambiguous case for the assignment problem involving a pair of phase singularities (PS). Is it a long distance jump or a wavefront merging followed by a wavebreak? The wavefronts are numbered from 1 to 4. The thick arrows show the direction of wavefront propagation. At time t, there are two PS located at P_- and P_+, and, at time $t + \Delta t$, two PS at Q_- and Q_+. The signs $+$ and $-$ indicate chirality and the dashed arrows a possible PS assignment.

than λ_{jump}, i.e. $P_-Q_- > \lambda_{\text{jump}}$ and $P_+Q_+ > \lambda_{\text{jump}}$, (2) the move would be valid according to the criteria of the assignment algorithm, (3) the two PS remain close enough along the way, i.e. $(P_-P_+)^2 + (Q_-Q_+)^2 < 2\lambda^2_{\text{neighbor}}$.

For a maximum of flexibility these operations are performed in a post-processing stage. Unless otherwise stated, the values $\lambda_{\text{jump}} = 4$ mm and $\lambda_{\text{neighbor}} = 2$ mm are used. The conditions are actually rarely met, but this additional step fixes a few transitions that would appear surprising by visual inspection.

G. Graph of phase singularity trajectories

After PS tracking has been performed every Δt during the whole simulation, PS trajectories can be reconstructed. Each trajectory starts with PS creation and ends with PS annihilation. PS are created and annihilated by pairs of opposite chirality (usually not the same pair). In order to form the pairs of annihilated PS at time t, a squared distance matrix D_{pq} similar to (7) is computed, where p runs over the list of PS with positive chirality annihilated at time t and q runs over the list of PS with negative chirality annihilated at time t. The Hungarian algorithm is applied to match the pairs. The same procedure is used to find the pairs of created PS at time $t + \Delta t$.

An undirected (multi)graph of PS trajectories is then generated.[4,14,37] Each node of the graph corresponds to a PS trajectory. Two nodes are linked if they have been created or

annihilated together, so each node is linked to exactly two nodes. In the practical implementation, each node contains the creation and annihilation times of the PS and the successive PS locations for each time instant.

Two methods for simplifying the graph while preserving its structure are proposed: the first downsamples the time resolution of the trajectories, and the second iteratively eliminates the trajectories with the shortest lifespans. The objective is to automatically provide representations of PS dynamics at different spatio-temporal scales, either showing the full details of microstructural propagation or concentrating on the macro-scale, coarse-grained rotors that sustained the arrhythmia. PS trajectories form a set of continuous curves in space-time, either closed or having their endpoints at the initial or final time of simulation. The graph transformations presented below always preserve the continuity of these curves and do not add new endpoints. It comes down to resampling these curves in space-time.

Downsampling will be used to estimate PS tracking accuracy as a function of time resolution. It is performed as illustrated in Fig. 3. Multiple time steps are reduced to one single time step. The downsampling process is applied to the PS connectivity graph and does not depend on actual PS locations. PS trajectories are followed until they reach the beginning or the end of the downsampled interval. A PS may go through ($1 \mapsto 1'$ in Fig. 3) and be considered as a simple displacement. When the path goes back to the initial time ($3 \mapsto 4$) it is interpreted as an annihilation. The trajectories also have to be followed backward to identify creations ($3' \mapsto 4'$). Along the way, the path may need to transiently go back in time ($5 \mapsto 5'$). When downsampling by a larger factor, paths may be longer and include multiple "time travels". Then the intermediate layers are removed (Fig. 3B). The precision of creation and annihilation times is reduced but the structure of the graph remains consistent.

The second graph simplification approach consists in eliminating the PS trajectory with the shortest lifespan until all lifespans are longer than a threshold τ_{ps}, as illustrated in Fig. 4. The curve segment $3 \mapsto 2$ to be eliminated has by definition a shorter lifespan than the segments $1 \mapsto 2$ and $3 \mapsto 4$. These three curve segments are merged into a single chronological trajectory $1 \mapsto 2' \mapsto 3' \mapsto 4$ that includes a jump halfway between the creation time (3) and the annihilation time (2). In case the segments form a loop ($5 \mapsto 6$), all segments are eliminated.

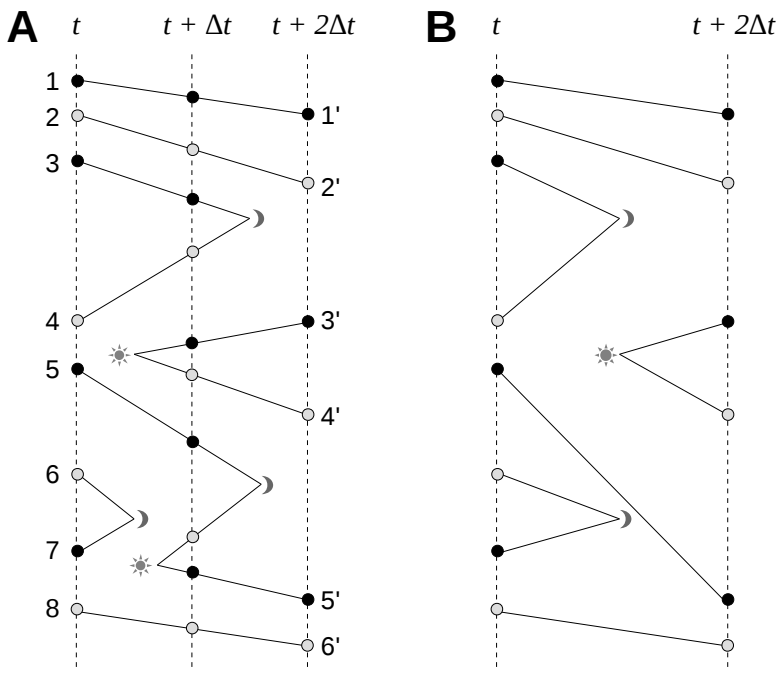

FIG. 3. Downsampling of phase singularity (PS) trajectory graph: (A) before and (B) after downsampling by a factor of two. Time is horizontal and space vertical. PS are displayed along a line for graphical purpose but trajectories are actually processed in two spatial dimensions. PS at each time instant t, $t + \Delta t$ and $t + 2\Delta t$ are numbered and represented as circles whose color indicates chirality. Connecting lines show PS tracking. The little sun symbolizes PS creation and the moon crescent PS annihilation.

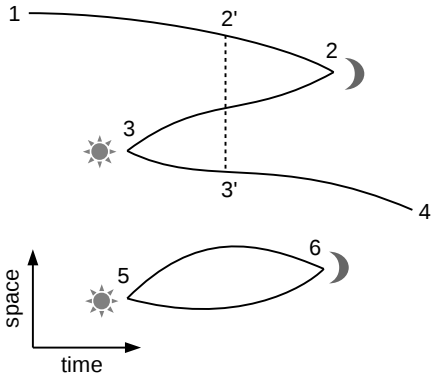

FIG. 4. Simplification of phase singularity (PS) trajectory graph. The shortest curve segment $3 \mapsto 2$ is replaced by $1 \mapsto 2' \mapsto 3' \mapsto 4$. Short loops such as $5 \mapsto 6$ are simply removed. The little sun symbolizes PS creation and the moon crescent PS annihilation.

H. Simulation protocols

PS detection and tracking algorithms are validated and evaluated in a set of complex episodes of reentrant arrhythmias simulated using the monodomain equation in a 5×10 cm rectangular sheet of atrial tissue with a spatial resolution of 0.02 cm. Four types of arrhythmogenic substrates (described below) are created. In each of them, a reentry is initiated by cross-shock stimulation. Simulations are run for 2500 ms with 40 or more different sets of parameters, providing a database of 180 episodes of arrhythmia. Activation times are computed by linear interpolation with a temporal resolution of 0.1 ms and a threshold $V_{th} = -40$ mV. The last 2000 ms are used for the analyses.

The substrates used are meant to reproduce a variety of difficulties typically encountered in PS tracking, including situations where weak coupling might create discrete propagation effects. The four substrates are associated with different hypothesized forms of atrial fibrillation and are referred to as "wavebreaks," "cholinergic," "fibrosis," and "breakthroughs."

The first substrate ("Wavebreaks"), based on a modified Luo-Rudy model, has a steep repolarization restitution that causes spontaneous wavebreaks,[38] thus providing highly complex activation patterns despite homogeneous properties. Two parameters are changed across the simulation set: the conductance of the slow inward current (G_{si}: 8 values between 0.05 and 0.075 mS/cm^2) and the anisotropy ratio (5 conductivity values from 0.5 to 4 mS/cm in the longitudinal direction and 0.5 mS/cm in the transverse direction).

The second substrate ("Cholinergic") reproduces the Kneller et al. canine model of atrial fibrillation.[39] A two-dimensional sinusoidal distribution of acethylcholine (ACh) concentration is introduced to create repolarization gradients that lead to both transient and stable rotors. Longitudinal/transverse conductivities are 5.8 and 1 mS/cm. The peak-to-peak length scale of the spatial distribution is 2.5 cm. The mean ACh concentration is either 0.3 or 3 μM and the amplitude of the sine wave is varied between 0 and 100% of mean value.

The third substrate ("Fibrosis") uses the Courtemanche model as modified by Gharaviri et al.[32] with a random distribution of non-conductive obstacles. Longitudinal/transverse conductivities are 4 and 1 mS/cm. The distribution is defined by a Markov random field with controllable spatial correlation (parameter α in the code by Bruno Sciolla that can be downloaded at `github.com/bsciolla`) and percentage of non-conductive tissue (prc). Sixty realizations are generated to represent diffuse ($\alpha = 1$, $prc = 50\%$), patchy ($\alpha = 3$,

$prc = 20\%$) or intermediate fibrosis ($\alpha = 2$, $prc = 7.5\%$).

The fourth substrate ("Breakthroughs") reproduces the Gharaviri et al. bilayer model of atrial fibrillation.[32,40] In contrast with the previous substrates, two layers (epicardium and endocardium) are simulated, but only the dynamics in the epicardial layer is analyzed to facilitate direct comparison with other substrates. Within-layer conduction is isotropic with a conductivity of 0.5 mS/cm. The two layers are coupled by 6, 12, 24, 48 or 96 discrete connections randomly distributed throughout the tissue.[32] As in the original study, a reentry is first initiated in the isolated epicardial layer. Then, at 8 different time points, the connections are reintroduced to provide pathways for breakthroughs.

I. Analysis of simulations

The tracking algorithm is applied to each of the 180 simulations with a tracking time step of $\Delta t = 0.1$, 0.2, 0.5, 1, 2, 5 and 10 ms and $\tau_d = 10$ ms. The assignment constraints include chirality conservation, wavefront conservation, or both. The assignment method is either the Hungarian or the greedy algorithm. In total, there are $180 \times 7 \times 3 \times 2 = 7560$ cases.

The reference for quantifying tracking accuracy is the Hungarian algorithm with chirality and wavefront conservation and $\Delta t = 0.1$ ms. To compare with a tracking result at larger Δt values, the reference PS trajectories are down-sampled (e.g. 0.1 vs 2 ms: down-sampling by a factor of 20) following the method of Subsect. II G. Note that PS detection is not affected at all by Δt; only PS tracking is.

III. RESULTS

A. Simulated episodes of fibrillation

The fibrillation episodes simulated in the four substrates differ by the number and the frequency of their reentries. Table I reports these values averaged over all simulations. The number of PS includes very short lived ones as well as wavefront extremities on a boundary. It is always smaller than twice the number of wavefronts. The relative discrepancy is largest in the "Breakthroughs" substrate because more wavefronts are shaped as closed curves (breakthrough patterns). The larger number of PS in the "Fibrosis" substrate stems from

wavefront fractionation caused by non-propagating obstacles. Figure 5 shows examples of membrane potential maps along with detected PS trajectories computed using our algorithm with $\Delta t = 0.1$ ms, $\tau_d = 10$ ms and $\tau_{ps} = 10$ ms.

TABLE I. Characteristics of simulated fibrillation episodes (mean ± standard deviation).

Substrate	Wavebreaks	Cholinergic	Fibrosis	Breakthroughs
Number of cases	40	40	60	40
Cycle length (ms)	96 ± 46	99 ± 17	205 ± 33	189 ± 18
1st percentile (ms)	38	72	164	160
Number of wavefronts	14.6 ± 8.1	2.9 ± 2.0	38 ± 35	5.1 ± 2.9
Number of PS	29 ± 16	5.8 ± 4.1	74 ± 68	8.9 ± 5.4

B. Parameter selection

Besides the post-processing parameters (λ_{jump} and $\lambda_{\text{neighbor}}$), the PS tracking algorithm mainly depends on two parameters: Δt and τ_d. The tracking time step Δt will be varied from 0.1 to 10 ms. The maximum propagation delay τ_d is set to 10 ms.

Histograms of inter-cellular activation delays are shown in Fig. 6. Most data points are below 2 ms which corresponds to a velocity of 10 cm/s. The value $\tau_d = 10$ ms is well above that threshold while being well below the first percentile of cycle length (Table I) which may provide an (over)estimate of τ_r. Another argument is that the average number of detected PS in the substrate with shortest cycle length is minimal near $\tau_d = 10$ ms (Table II), presumably due to fewer false positives. PS detection is however not very sensitive to τ_d over a wide range from 5 to 20 ms.

C. Consistency checks

After each tracking procedure, the consistency of the resulting PS trajectory graph is checked to validate that mathematically true assertions are also true in the output of the algorithm. Checks include conservation of chirality and wavefront, the number of PS per wavefront being 0 or 2, and the connectivity between PS trajectories through PS creation and annihilation. When chirality or wavefront conservation is not strictly enforced by the

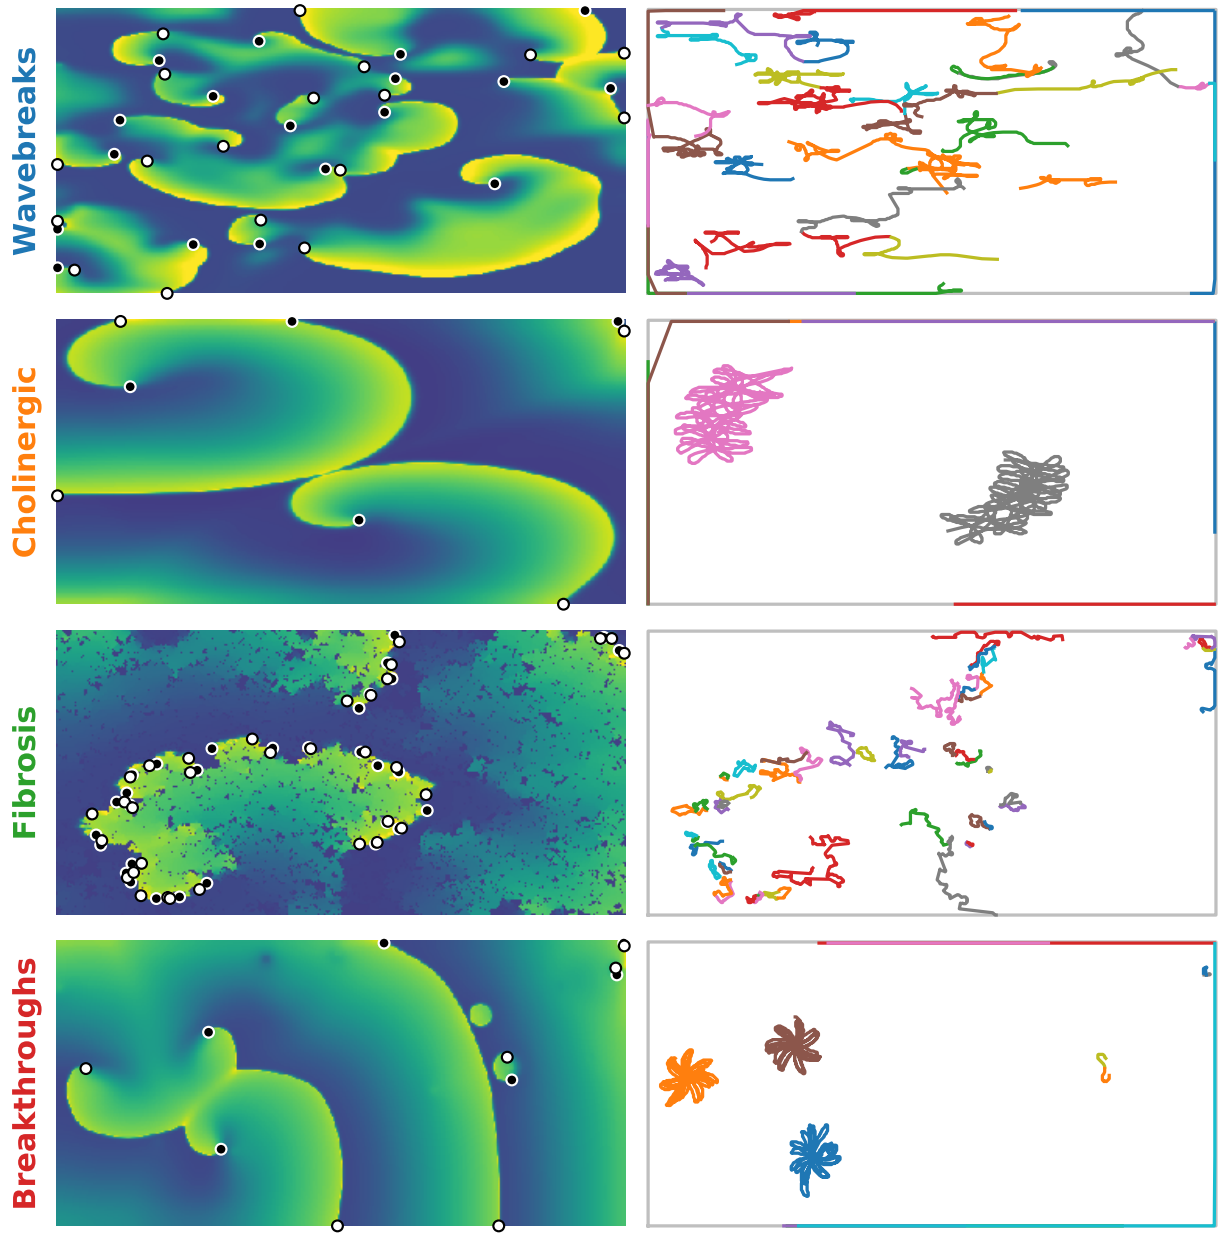

FIG. 5. *Left panels*: Examples of color-coded membrane potential maps for each substrate. Phase singularities (PS) are displayed as white or black circles according to their chirality. *Right panels*: Trajectory of the PS present at the time of the snapshot of the left panel.

algorithm, at least one inconsistency is found in 97% of the simulations. In contrast, the algorithm that accounts for chirality and wavefront tracking passes all the tests in all 180 cases.

PS trajectories ($\Delta t = 0.1$ ms) and appropriate automatic handling of ambiguous cases (Subsect. II F) were validated by visual inspection, with particular attention devoted to

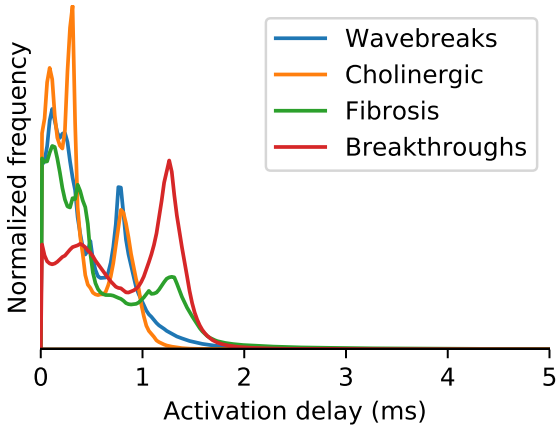

FIG. 6. Normalized histograms of delays in activation time between neighboring cell units. Each histogram is accumulated over all simulations for each substrate.

TABLE II. Effect of the maximum propagation delay τ_d on the number of phase singularities (PS) detected in the "Wavebreaks" substrate ($\Delta t = 0.1$ ms).

τ_d	Number of PS
2 ms	42 ± 25
3 ms	31 ± 17
5 ms	30 ± 17
10 ms	29 ± 16
15 ms	29 ± 16
20 ms	30 ± 16

fast-moving PS. When $\Delta t = 0.1$ ms, ambiguous cases are present in <0.06% of the time steps. This percentage increases along with Δt and reaches about 4% in the substrate "Fibrosis" when $\Delta t = 1$ ms, but remains <0.06% in the other substrates. Changing λ_{jump} from 4 mm to 6 mm only slightly reduced the percentages. The heuristic approach to ambiguous cases appears satisfactory by visual inspection, as attested by the absence of unexplained discontinuity (i.e. not the result of post-processing as in Fig. 4) in PS trajectory.

D. Computational requirements

For a model with 125,000 nodes and 249,250 edges, the complete analysis takes about 9.3 s per second of simulation for a resolution of $\Delta t = 1$ ms (single-core computations), although limited effort has been devoted to code optimization. About 30–35% of the computational time is for PS detection and tracking, 50–60% for wavefront detection, and 5% for wavefront tracking. The remaining time includes inputs/outputs and overheads. The Hungarian algorithm takes about 5–10% of the total computational time (the lower end occurs at high time resolution). Computational time is roughly proportional to Δt^{-1}. At high time resolution, it is a bit faster than this linear formula since the file of activation times is processed entirely at any resolutions (e.g. 81 s per second of simulation for $\Delta t = 0.1$ ms).

E. Wavefront tracking

A one-to-many wavefront transition is defined as a wavefront k at time t such that $\sum_{l=1}^{m^{t+\Delta t}} G_{kl}^{t} > 1$, i.e., a wavefront that breaks into two or more parts within a time step. The fraction of wavefronts that undergo a one-to-many transition (Fig. 7) reflects the difficulty of PS tracking since one-to-one wavefront transitions lead to exact PS tracking. This fraction varies from 0.9–3% with a time step of $\Delta t = 0.1$ ms to 6–18% when $\Delta t = 1$ ms, which illustrates the advantage of reducing Δt.

F. Phase singularity tracking

Figure 8 shows the error rates of PS tracking as a function of Δt. The complete method with the Hungarian algorithm at $\Delta t = 0.1$ ms is used as a reference. The error rate expresses the fraction of PS assignments that differ from reference calculations. Comparison involving different time steps uses the down-sampling technique of Subsect. II G.

Error rates at $\Delta t = 1$ ms are around 0.05% ("Cholinergic" substrate), 0.3% ("Breakthroughs"), 0.5% ("Wavebreaks") and 4.5% ("Fibrosis"). Differences in error rates between substrates reflect the complexity of the dynamics and the number of PS. Reducing Δt from 1 to 0.2 ms improves error rates by an order of magnitude, except in the "Cholinergic" substrate where the absolute number of errors is already small at 1 ms. Releasing the wavefront constraint significantly increases error rates, more so than releasing the chirality constraint.

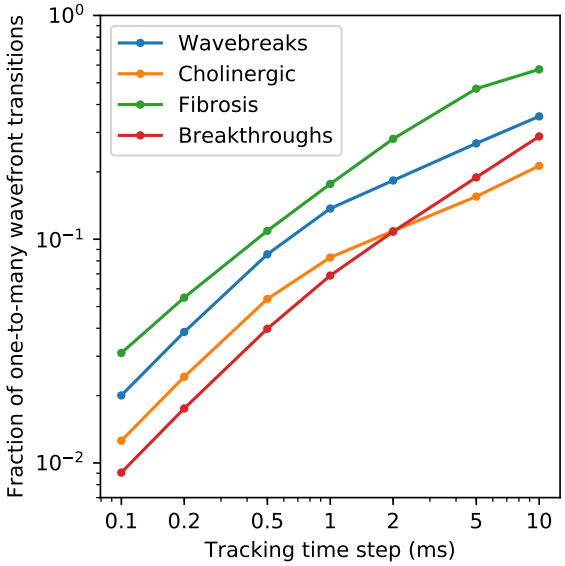

FIG. 7. Fraction of one-to-many wavefront transitions (wavefront that breaks into two or more parts within a time step) as a function of the tracking time step. The means over all wavefronts and all simulations are shown for each substrate.

Membership to wavefront sets $\mathcal{F}_k^t$ indeed encodes more bits of information than chirality.

The Hungarian algorithm used for PS pair assignment consists of at most n iterations, each of which runs in $\mathcal{O}(n^2)$ time, where n is the number of PS. If taking the closest admissible PS is an optimal solution, the loop halts after one iteration. Figure 9A shows the average number of iterations as a function of the tracking time step Δt. When $\Delta t = 0.1$ ms, only one iteration is needed 87% of the time. With larger time steps, the use of the Hungarian algorithm is necessary, particularly when the dynamics is complex ("Fibrosis" substrate). When PS tracking is compared at the same Δt between the Hungarian and the greedy assignment algorithm, the difference in terms of PS assignment is negligible or even zero at $\Delta t = 0.1$ ms. This near-perfect correspondence may be seen as an indication that Δt is small enough. The difference increases to 0.1–2% when $\Delta t = 1$ ms.

G. Lifespan of phase singularities

The lifespan of a PS is the time interval between its creation and annihilation. The time scale at which a reentry is observed influences lifespans. The PS graph simplification method of Subsect. II G enables discarding short-lived PS (lifespan $< \tau_{ps}$) and reconnect

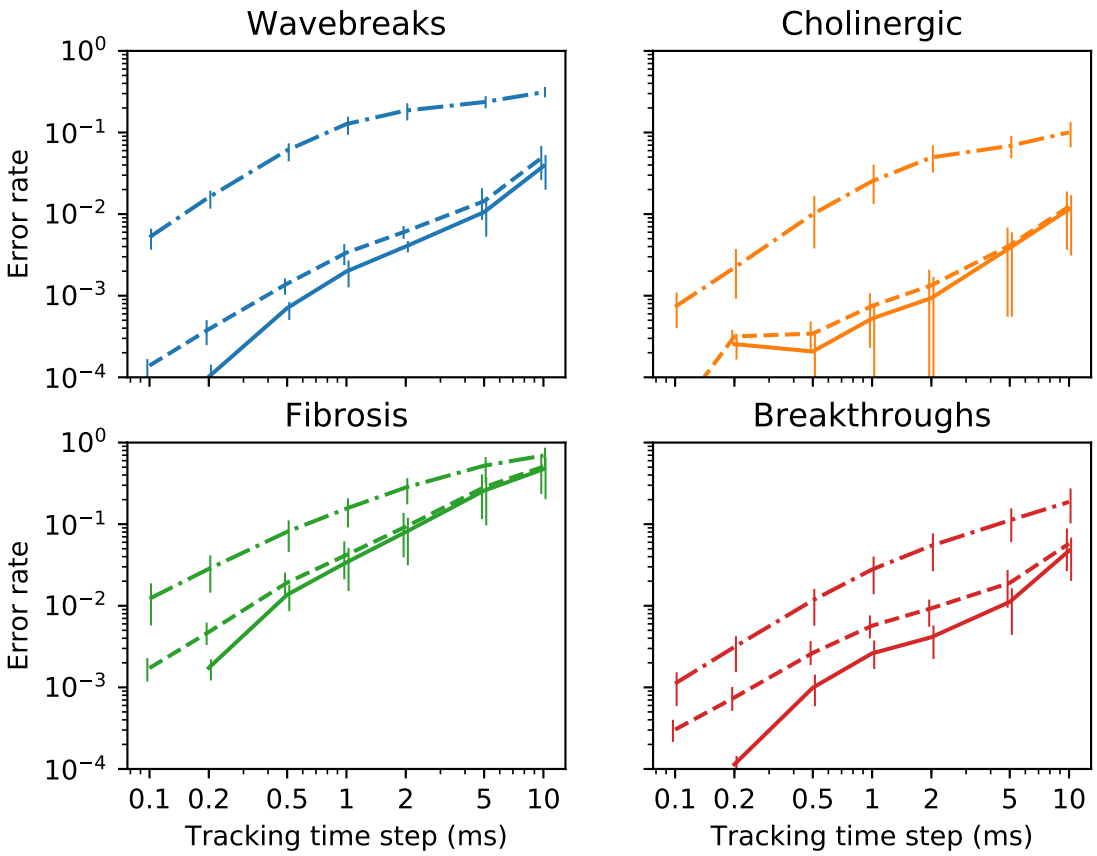

FIG. 8. Error rate of phase singularity (PS) assignment in the tracking procedure as a function of the tracking time step for the four substrates. Error bars represent inter-quartile intervals. *Solid lines*: PS tracking with the Hungarian algorithm taking into account chirality and wavefronts. Its outcome at $\Delta t = 0.1$ ms is used as a reference for computing error rates, so that data point is absent on the graph. *Dashed lines*: PS tracking without taking into account chirality. *Dash-dotted lines*: PS tracking without taking into account wavefronts.

PS trajectories, thus prolonging the lifespan of remaining PS. An example is shown in Fig. 10. Micro-conduction blocks and micro-scale non-propagating obstacles fragment the wavefronts and create a multitude of PS (panels A and B). Through an increase of the observation time scale τ_{ps} (panels C to F), PS trajectories are recombined, revealing a reentry that appears stable at the macroscopic level and that lasts the entire simulation (pink meandering trajectory), even if at the micro-scale, wave breaks occur in the vicinity of the spiral tip. This provides a reduced, macro-scale description of the dynamics comprising

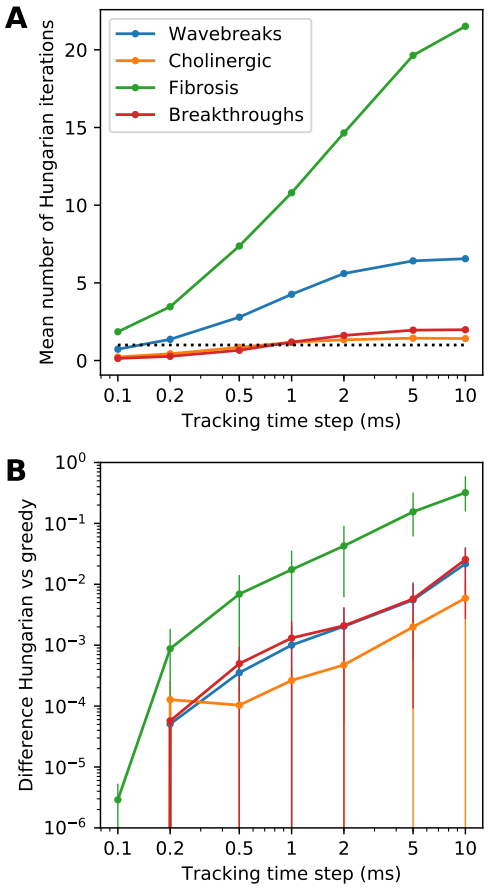

FIG. 9. (A) Mean number of iterations needed by the Hungarian algorithm for the phase singularity (PS) assignment problem. Each curve corresponds to a different substrate. The dotted horizontal line has an ordinate of 1. (B) Mean relative differences (error rates) in PS tracking using the Hungarian algorithm vs the greedy algorithm. Color code is the same as panel A. Error bars represent first and third quartiles. Some data points at $\Delta t = 0.1$ ms are not visible since the difference is zero.

much fewer PS trajectories. Note that if only PS hot spots are sought, statistics on PS location[41] or processing of local signals[25] may provide the desired information.

The distribution of lifespans can be represented by a histogram.[39] To improve resolution, PS lifespans are here accumulated over all simulations of each substrate to build survival curves showing the fraction of PS whose lifespan is longer than a given time. Without PS graph simplification, most PS are short-lived. The larger the value of τ_{ps}, the longer the lifespans. When $\tau_{ps} > 5$ ms in the "Cholinergic" and "Fibrosis" substrates, survival curve

analysis successfully identifies rotors that appear stable at the macroscopic scale (lifespan as large as the time window of analysis) whose existence is easily checked by visual inspection. In contrast, the two other substrates do not feature macroscopically stable reentries, as known from previous works.[32,38]

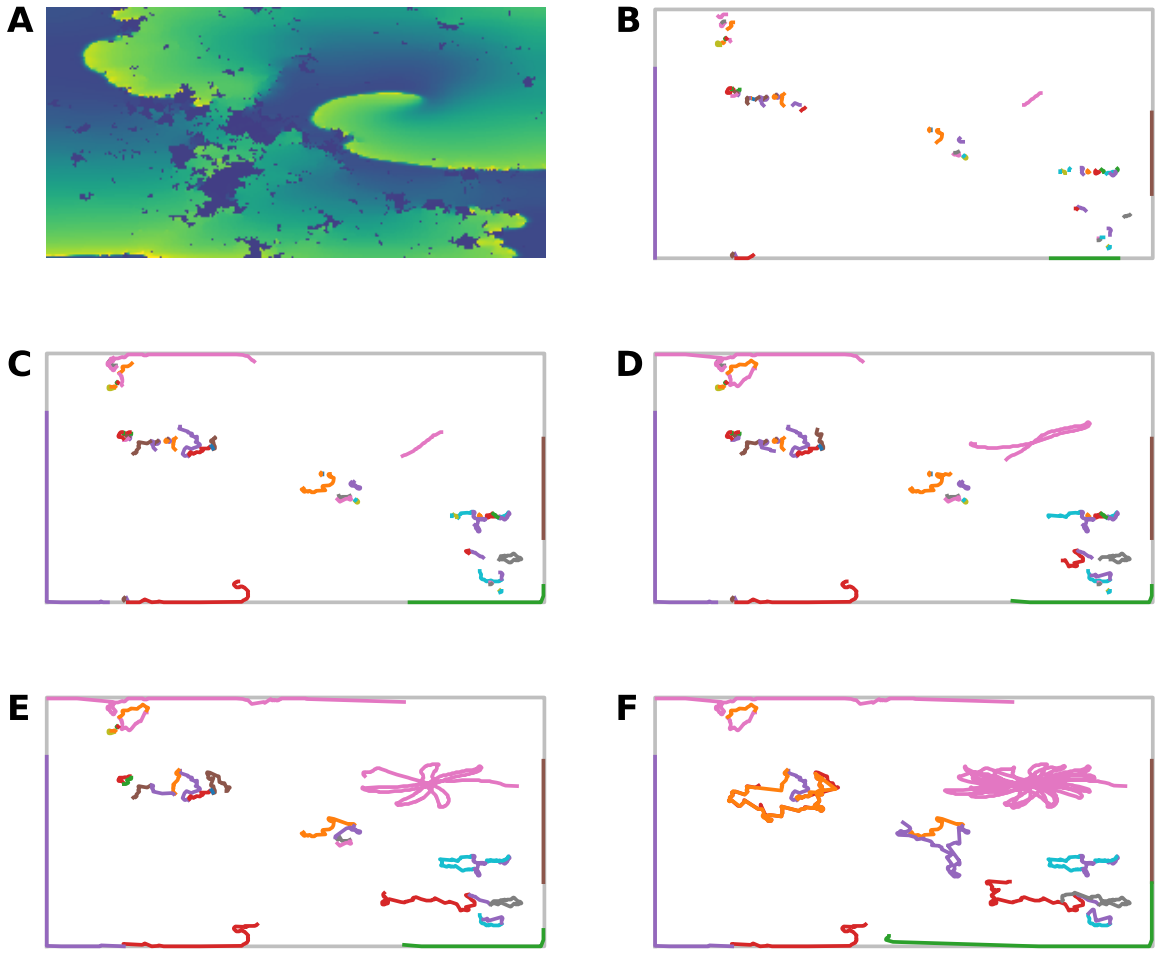

FIG. 10. Effect of phase singularity (PS) trajectory graph simplification with different values of τ_{ps} and $\Delta t = 0.1$ ms. (A) Membrane potential map in a tissue with islands of non-conducting tissue. (B) Trajectories of PS present at time $t = 2045$ ms with no graph simplification. (C) Same PS trajectories after graph simplification with $\tau_{ps} = 2$ ms, (D) $\tau_{ps} = 4$ ms, (E) $\tau_{ps} = 10$ ms, (F) $\tau_{ps} = 20$ ms.

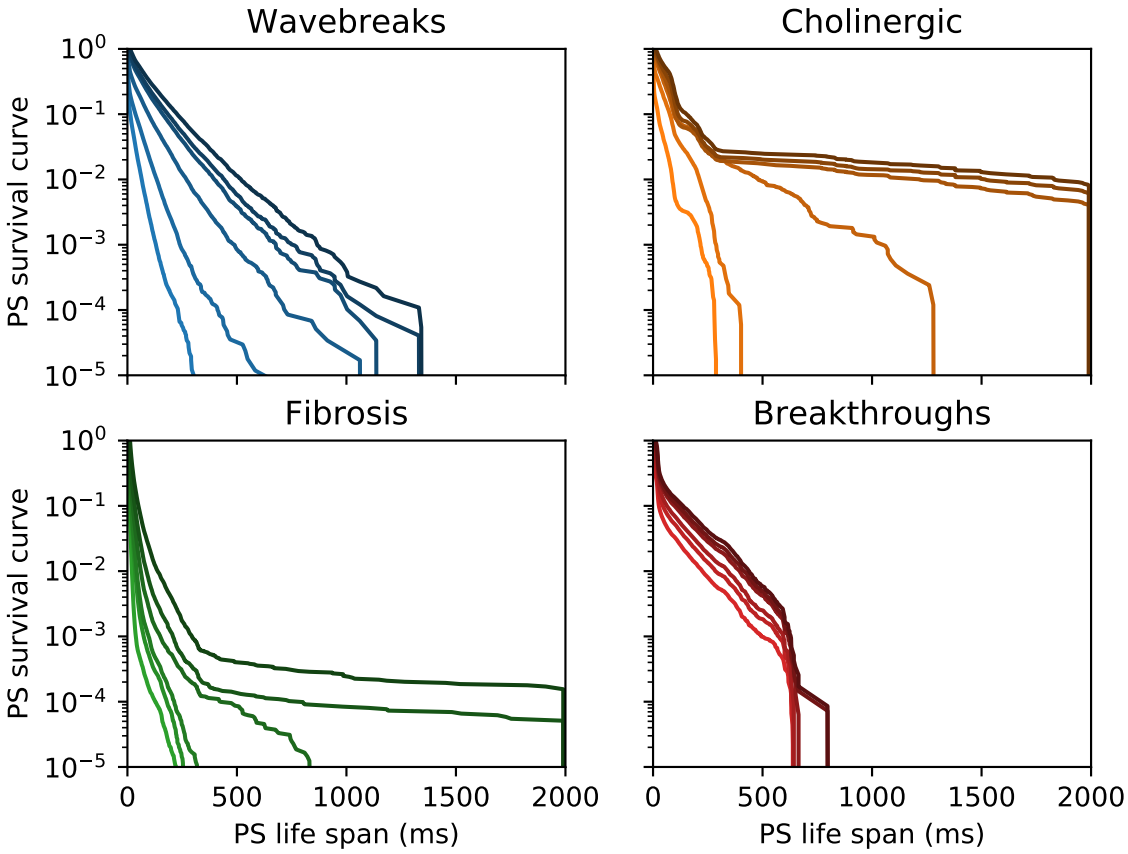

FIG. 11. Phase singularity (PS) survival probability for the four substrates analyzed with $\Delta t = 0.1$ ms. The six curves from light to dark color in each graph correspond to increasing values of the PS trajectory graph simplification parameter τ_{ps}, namely 0, 1, 2, 5, 8 and 12 ms.

H. Impact of noise

Inaccuracy in the determination of activation times may result in incorrect PS detection, particularly false positives.[42] In experimental data, noise would have a similar effect. To estimate the robustness of PS tracking with respect to a small jitter in activation times, white noise uniformly distributed between $\pm\tau_{\text{jitter}}$ is added to the activation maps of a representative simulation of each of the four substrates. For 7 values of τ_{jitter} between 0 and 3 ms and 4 realizations of the white noise process, PS tracking is performed with $\Delta t = 1$ ms. The graph of PS trajectories is then simplified to eliminate trajectories with a lifespan shorter than $\tau_{ps} = 10$ ms. Figure 12 shows the mean number of PS obtained before and after graph simplification as a function of τ_{jitter}. Additive noise on activation

times generates false positives. These PS are however short-lived. After simplification of the PS trajectory graph, the mean number of PS becomes robust to noise at least up to $\tau_{\text{jitter}} = 1$ ms. The maximum difference in mean number of PS is 5–10% for $\tau_{\text{jitter}} < 3$ ms, except in the "Fibrosis" substrate. Note that the maximum jitter used here creates an additional delay between neighboring nodes up to 6 ms (to be compared with Fig. 6), which would correspond to a conduction velocity of 3.3 cm/s. In the presence of such noise level, spatial filtering or more robust detection methods would be recommended.[21,26]

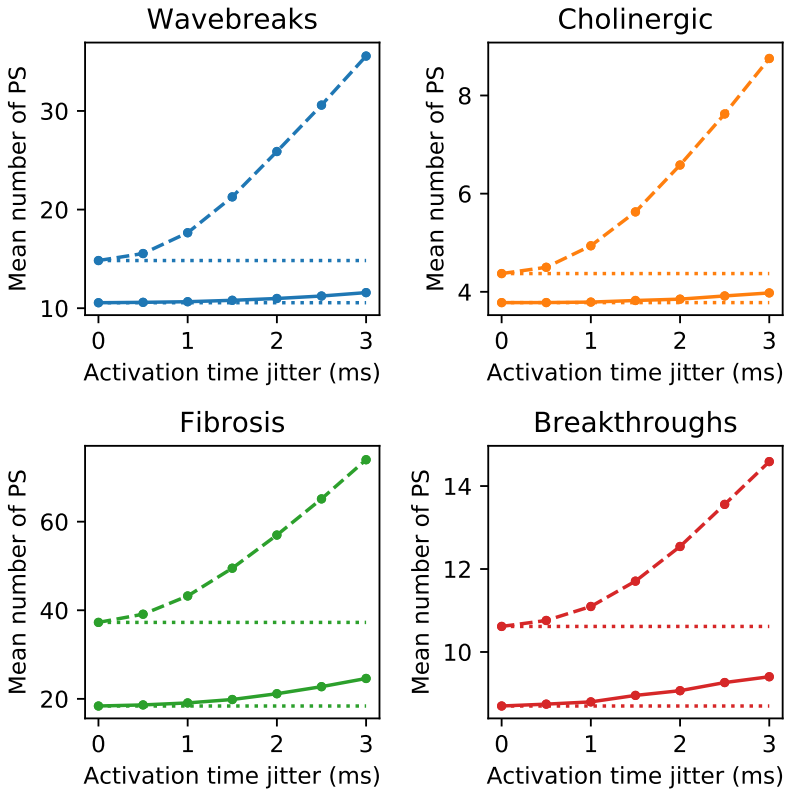

FIG. 12. Effect of noise on the mean number of detected phase singularities (PS). The dots represents the mean number of PS (averaged over time) as a function of the noise level (τ_{jitter}), computed before (dashed lines) and after (solid lines) simplification of the PS trajectory graph (elimination of PS with a lifespan < 10 ms). The results are averaged over four realizations of the noise. Standard deviations, displayed as error bars, are smaller than the size of the dots.

I. Comparison with common methods

Fenton et al.[43] proposed a PS detection technique based on the intersection between the curves $\mathrm{d}V/\mathrm{d}t = 0$ and $V = V_{th}$. The wavefronts in our approach are isopotential lines where the time derivative is positive. The endpoints of the wavefronts, i.e. the phase singularities, therefore match those of Fenton et al. A comparison between that method and the classical phase map method by Iyer et al.[16] is provided in Gong et al.[44] Further comparison of similar methods can be found in Li et al.[21]

As a basis for comparison, an implementation[42] of the Iyer-Gray method[16] is used. Phase maps are reconstructed from membrane potential maps using an embedding delay of 5 ms. It is worth noting that while the file containing the activation times is about 30 MB, the membrane potential maps for $\Delta t = 1$ ms require 1.2 GB of storage for each simulation. In the absence of wavefront tracking, PS tracking is based on the shortest distance (greedy assignment algorithm) and matching chirality. Only pairs within a distance shorter than 5 mm are matched. This PS tracking algorithm is similar to Zou et al.[45] The maximal distance is reduced to 2 mm in the "Fibrosis" and "Wavebreaks" substrates because of the larger density of PS.

Figure 13 compares the mean number of PS over time in our edge-based isopotential line method and in the Iyer-Gray approach. Unlike in Table I, wavefront endpoints on the border of the tissue are not counted as PS. When the conductive medium is uniform ("Cholinergic" and 'Wavebreaks" substrates), the two methods are consistent. When discrete coupling with an endocardial layer is included ("Breakthrough"), small differences are observed in the most complex cases with a large number of discrete connections. In the presence of discrete uncoupling and fibrosis, the mean numbers of PS are correlated but significant discrepancies are found when PS density is large. This illustrates the difficulty of PS tracking in discrete propagation models. It remains unclear what should be considered the gold standard of PS detection at the micro-scale in discrete models.

The Jacobian-determinant method has been shown to give better PS detection accuracy than most methods in a FitzHugh-Nagumo-based uniform tissue model.[21] This method requires to identify the local extrema of a function that depends on membrane potential gradients. Our simulations, in contrast, are characterized by faster upstroke velocity and conduction heterogeneity, which means that the finite difference formula for the Jacobian

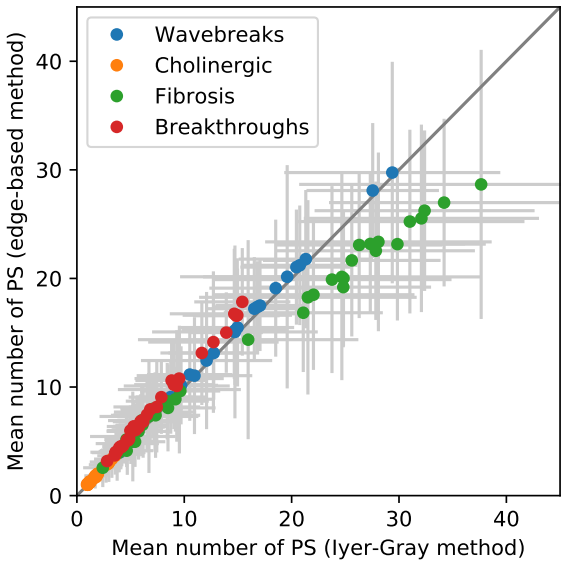

FIG. 13. Comparison between our phase singularity (PS) detection approach and that of Iyer and Gray.[16] Mean number of PS lasting more than 10 ms in simulations from the four substrates (each colored point is a simulation). Error bars represent standard deviations over time.

sometimes involves nodes that are electrically uncoupled. The structural complexity of the substrate therefore hampers reliable peak identification. A similar conclusion is reached for the Gurevich et al. method,[26] which is more robust and well-suited to tissue with smooth conductivity (e.g. "Wavebreaks" substrate). In the "Fibrosis" substrate, however, the use of a Gaussian filter of the size of micro-heterogeneities prevents the ability to track PS around these obstacles. On the other hand, such filtering might be used to remove these micro-scale propagation effects if they are deemed inconsequential.

IV. DISCUSSION

A. Tracking algorithms

We have developed a PS tracking algorithm that extends Rogers' method to address the challenges of analyzing at high resolution the fibrillation dynamics in heterogeneous mathematical models of arrhythmogenic substrates. While most PS detection methods rely on phase signals in order to apply the same methodology to both simulated and experimental signals, our approach uses activation times derived from membrane potentials, deliberately

focusing our attention on simulations. The computational advantage is that by calculating activation times during the simulation, the reduced output size (by a factor 40 or more) enables high spatial and temporal resolution (0.2 mm and 0.1 ms) which may be needed to track complex reentries in fibrotic tissue. The disadvantage is that activation times are sensitive to noise and therefore the PS detection method is not appropriate for experimental data. The PS tracking optimization (Hungarian algorithm) and graph simplification methods, on the other hand, might be incorporated in existing PS analysis techniques aimed at experimental or clinical mapping.

The mathematical formulation is edge-based rather than vertex-based. This approach is required to guarantee that each wavefront has exactly 0 or 2 extremities or PS. The edges of two distinct wavefronts $\mathcal{F}_k^t$ and $\mathcal{F}_l^t$ may indeed share a common vertex (at the exact time when two wavefronts collide), which would make it difficult to differentiate the wavefronts using a vertex-based strategy. Moreover, an isochrone (wavefront) intersects each edge at most once. Wavefronts are indeed represented as curves and not as regions covering an area in the front and in the back of the front.[46,47] This enables the analysis of micro-reentries.

The proposed wavefront tracking procedure is applicable to any discrete tissue (unstructured mesh, 3D mesh) as long as edge-to-edge neighborhood can be defined. The specificity of rectangular grid is only exploited for PS detection. The look-up table would actually be valid for any quadrilateral structured grid as the look-up table does not depend on edge lengths. Its adaptation to a triangular mesh would be straightforward since only 5 triangle cases are possible (the equivalent of cases A, B, C, E, F in Fig. 1). Generalization to the detection of filament in 3D would require additional work, notably for defining the distance between filaments.

The PS tracking algorithm is designed to use the distance between PS in successive time steps only as a last resort. In contrast to Zou et al.[15], no maximal distance constraint is imposed on PS assignment. Wavefront identification and chirality matching ensures that PS in isolated wavefronts are always correctly tracked.[14] However, multiple wavebreaks and merging occurring within a single time step necessitate additional distance-based criteria for PS tracking. This is particularly true when a wavefront is near conduction block over a significant extent, leading to fractionation of the wavefront into many short ones. Our strategy in this case is to keep all detected PS and simplify the graph of PS trajectories at the post-processing stage.

The formulation of PS assignment as an optimization problem facilitates the integration of knowledge-based terms in the cost function. The details of the cost function are motivated by the necessity to avoid ties (PS move by discrete steps on a grid). The Hungarian algorithm provides a computationally efficient solver that is not, in practical situations, significantly more complex than the greedy algorithm (which is similar to previous works[15,27]) since only a few iterations are performed. The problem is of course alleviated by reducing the time step. At $\Delta t = 0.1$ ms, the outputs of the Hungarian and greedy assignment algorithms become the same (Fig. 9), which may be used as criterion to determine whether Δt is small enough.

Despite their rarity in practical situations, ambiguous cases (Subsect. II F) are unavoidable. Heuristic criteria have to be used to differentiate between a parallel fast move of two PS and the simultaneous creation and annihilation of two pairs of PS. Our approach is to keep the continuity of PS trajectories and possibly break them afterwards in a post-processing stage if jumps are deemed not physiologically plausible.

B. Time resolution

Using high temporal resolution ($\Delta t < 1$ ms) for PS tracking has several benefits. It reduces the likelihood of seemingly simultaneous events. One-to-many relationships in wavefront tracking are less frequent (Fig. 7), so exact PS tracking (relying only on wavefront and chirality) is more likely. This is particularly helpful in the "Fibrosis" substrate, where a wavefront may collide with multiple distinct obstacles at the same time. The distance travelled by PS in a single time step is also shorter which facilitates PS tracking and decreases the number of iterations needed by the Hungarian algorithm. As a result, error rates are significantly lower at higher time resolutions in complex fibrillation dynamics (Fig. 8).

The cost of higher temporal resolution is some additional computational load, roughly inversely proportional to Δt. The main contributor to computations is wavefront tracking (notably, connected component identification) as its complexity is of the order of the number of cell units in the tissue. On the other hand, the output file size of activation times does not depend on Δt. Instead, the algorithm makes better use of available information.

C. Graph-based analysis

The analysis of PS trajectories through graphs[14,22] provides a compact synthetic description of the dynamics. An appealing feature of this approach is the ability to apply transformations that preserve the integrity of the structure, e.g., the connectivity to two other PS trajectories through the creation and annihilation events. We propose three such transformations. The first fixes the ambiguous cases. The second is the temporal down-sampling that is critical for comparing PS tracking at different Δt. The third is the simplification procedure that iteratively eliminates PS trajectories with very short lifespan, resulting in the identification of stable reentries despite wavefront fragmentation and fibrosis.

Other graph-based transformations may be defined. For example, functional and structural reentries could be differentiated by testing if PS location belongs to the list of edges located on the boundaries or on anatomical obstacles (e.g. fibrotic islands). Another application could be to coarsen the spatial resolution of the PS distribution by merging or discarding pairs of PS that are too close to each other.[42,48,49]

V. CONCLUSION

Although straightforward PS tracking methods typically work well on most simple fibrillation models, analyzing dynamics in more complex tissue substrates incorporating strong anisotropy, ionic and structural heterogeneities and fibrosis remains challenging. Instead of filtering or down-sampling the data in an attempt to reduce false positives, we used activation maps with high spatio-temporal resolution to enable tracking micro-scale wavefront propagation during simulated fibrillation. The results show that decreasing tracking time steps below 1 ms has clear benefits in terms of PS tracking accuracy.

ACKNOWLEDGMENTS

This work was supported by funding from the Natural Sciences and Engineering Research Council of Canada (RGPIN-2015-05658), and Fonds de Recherche du Québec – Santé.

REFERENCES

[1]R. A. Gray, A. M. Pertsov, and J. Jalife, "Spatial and temporal organization during cardiac fibrillation," *Nature*, vol. 392, no. 6671, pp. 75–8, 1998.

[2]F. H. Fenton, E. M. Cherry, H. M. Hastings, and S. J. Evans, "Multiple mechanisms of spiral wave breakup in a model of cardiac electrical activity," *Chaos (Woodbury, N.Y.)*, vol. 12, pp. 852–892, Sept. 2002.

[3]A. T. Winfree, *The Geometry of Biological Time*. Interdisciplinary Applied Mathematics, Springer New York, 2001.

[4]R. H. Clayton and A. V. Holden, "A method to quantify the dynamics and complexity of re-entry in computational models of ventricular fibrillation," *Physics in Medicine and Biology*, vol. 47, pp. 225–238, Jan. 2002.

[5]V. Swarup, T. Baykaner, A. Rostamian, J. P. Daubert, J. Hummel, D. E. Krummen, R. Trikha, J. M. Miller, G. F. Tomassoni, and S. M. Narayan, "Stability of rotors and focal sources for human atrial fibrillation: focal impulse and rotor mapping (firm) of af sources and fibrillatory conduction," *J Cardiovasc Electrophysiol*, vol. 25, no. 12, pp. 1284–92, 2014.

[6]S. M. Narayan, T. Baykaner, P. Clopton, A. Schricker, G. G. Lalani, D. E. Krummen, K. Shivkumar, and J. M. Miller, "Ablation of rotor and focal sources reduces late recurrence of atrial fibrillation compared with trigger ablation alone: extended follow-up of the confirm trial (conventional ablation for atrial fibrillation with or without focal impulse and rotor modulation)," *J Am Coll Cardiol*, vol. 63, no. 17, pp. 1761–8, 2014.

[7]R. H. Clayton and M. P. Nash, "Analysis of cardiac fibrillation using phase mapping," *Card Electrophysiol Clin*, vol. 7, no. 1, pp. 49–58, 2015.

[8]K. Umapathy, K. Nair, S. Masse, S. Krishnan, J. Rogers, M. P. Nash, and K. Nanthakumar, "Phase mapping of cardiac fibrillation," *Circ Arrhythm Electrophysiol*, vol. 3, no. 1, pp. 105–14, 2010.

[9]R. A. Gray, J. P. Wikswo, and N. F. Otani, "Origin choice and petal loss in the flower garden of spiral wave tip trajectories," *Chaos (Woodbury, N.Y.)*, vol. 19, p. 033118, Sept. 2009.

[10]Z. Qu, J. N. Weiss, and A. Garfinkel, "Cardiac electrical restitution properties and stability of reentrant spiral waves: a simulation study," *American Journal of Physiology-Heart and*

Circulatory Physiology, vol. 276, no. 1, pp. H269–H283, 1999.

[11]J. Kneller, J. Kalifa, R. Zou, A. V. Zaitsev, M. Warren, O. Berenfeld, E. J. Vigmond, L. J. Leon, S. Nattel, and J. Jalife, "Mechanisms of atrial fibrillation termination by pure sodium channel blockade in an ionically-realistic mathematical model," *Circ Res*, vol. 96, no. 5, pp. e35–47, 2005.

[12]S. Nattel, F. Xiong, and M. Aguilar, "Demystifying rotors and their place in clinical translation of atrial fibrillation mechanisms," *Nat Rev Cardiol*, 2017.

[13]P. M. Boyle, T. Zghaib, S. Zahid, R. L. Ali, D. Deng, W. H. Franceschi, J. B. Hakim, M. J. Murphy, A. Prakosa, S. L. Zimmerman, H. Ashikaga, J. E. Marine, A. Kolandaivelu, S. Nazarian, D. D. Spragg, H. Calkins, and N. A. Trayanova, "Computationally guided personalized targeted ablation of persistent atrial fibrillation," *Nature Biomedical Engineering*, Aug. 2019.

[14]J. M. Rogers, "Combined phase singularity and wavefront analysis for optical maps of ventricular fibrillation," *IEEE transactions on bio-medical engineering*, vol. 51, pp. 56–65, Jan. 2004.

[15]R. Zou, J. Kneller, L. J. Leon, and S. Nattel, "Development of a computer algorithm for the detection of phase singularities and initial application to analyze simulations of atrial fibrillation," *Chaos*, vol. 12, no. 3, pp. 764–778, 2002.

[16]A. N. Iyer and R. A. Gray, "An experimentalist's approach to accurate localization of phase singularities during reentry," *Ann Biomed Eng*, vol. 29, no. 1, pp. 47–59, 2001.

[17]S. Puwal, B. J. Roth, and S. Kruk, "Automating phase singularity localization in mathematical models of cardiac tissue dynamics," *Math Med Biol*, vol. 22, no. 4, pp. 335–46, 2005.

[18]M.-A. Bray and J. P. Wikswo, "Considerations in phase plane analysis for nonstationary reentrant cardiac behavior," *Phys Rev E Stat Nonlin Soft Matter Phys*, vol. 65, no. 5 Pt 1, p. 051902, 2002.

[19]M. A. Bray, S. F. Lin, R. R. Aliev, B. J. Roth, and J. P. Wikswo, "Experimental and theoretical analysis of phase singularity dynamics in cardiac tissue," *J Cardiovasc Electrophysiol*, vol. 12, no. 6, pp. 716–22, 2001.

[20]C. H. Roney, C. D. Cantwell, N. A. Qureshi, R. A. Chowdhury, E. Dupont, P. B. Lim, E. J. Vigmond, J. H. Tweedy, F. S. Ng, and N. S. Peters, "Rotor tracking using phase of electrograms recorded during atrial fibrillation," *Ann Biomed Eng*, vol. 45, no. 4, pp. 910–

923, 2017.

[21]T.-C. Li, D.-B. Pan, K. Zhou, R. Jiang, C. Jiang, B. Zheng, and H. Zhang, "Jacobian-determinant method of identifying phase singularity during reentry," *Physical Review E*, vol. 98, no. 6, p. 062405, 2018.

[22]N. Vandersickel, E. Van Nieuwenhuyse, N. Van Cleemput, J. Goedgebeur, M. E. Haddad, J. D. Neve, A. Demolder, T. Strisciuglio, M. Duytschaever, and A. V. Panfilov, "Directed networks as a novel way to describe and analyze cardiac excitation: Directed graph mapping," *bioRxiv*, 2019.

[23]Y.-S. Lee, J.-S. Song, M. Hwang, B. Lim, B. Joung, and H.-N. Pak, "A New Efficient Method for Detecting Phase Singularity in Cardiac Fibrillation," *PloS One*, vol. 11, no. 12, p. e0167567, 2016.

[24]C. D. Marcotte and R. O. Grigoriev, "Dynamical mechanism of atrial fibrillation: A topological approach," *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 27, p. 093936, Sept. 2017.

[25]E. M. Annoni, S. P. Arunachalam, S. Kapa, S. K. Mulpuru, P. A. Friedman, and E. G. Tolkacheva, "Novel Quantitative Analytical Approaches for Rotor Identification and Associated Implications for Mapping," *IEEE transactions on bio-medical engineering*, vol. 65, pp. 273–281, Feb. 2018.

[26]D. R. Gurevich and R. O. Grigoriev, "Robust approach for rotor mapping in cardiac tissue," *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 29, p. 053101, May 2019.

[27]L. Dang, *An investigation into therapies for atrial arrhythmias using a biophysical model of the human atria*. PhD thesis, Ecole Polytechnique Fédérale de Lausanne, Lausanne, Switzerland, 2005.

[28]J. C. Crocker and D. G. Grier, "Methods of digital video microscopy for colloidal studies," *Journal of colloid and interface science*, vol. 179, no. 1, pp. 298–310, 1996.

[29]D. Deng, M. J. Murphy, J. B. Hakim, W. H. Franceschi, S. Zahid, F. Pashakhanloo, N. A. Trayanova, and P. M. Boyle, "Sensitivity of reentrant driver localization to electrophysiological parameter variability in image-based computational models of persistent atrial fibrillation sustained by a fibrotic substrate," *Chaos (Woodbury, N.Y.)*, vol. 27, p. 093932, Sept. 2017.

[30]M. Saha, C. H. Roney, J. D. Bayer, M. Meo, H. Cochet, R. Dubois, and E. J. Vigmond, "Wavelength and fibrosis affect phase singularity locations during atrial fibrillation," *Frontiers in Physiology*, vol. 9, p. 1207, 2018.

[31]N. M. de Groot and M. A. Allessie, "Pathophysiology of atrial fibrillation: focal patterns of activation," *Pacing and clinical electrophysiology: PACE*, Aug. 2019.

[32]A. Gharaviri, S. Verheule, J. Eckstein, M. Potse, N. H. L. Kuijpers, and U. Schotten, "A computer model of endo-epicardial electrical dissociation and transmural conduction during atrial fibrillation," *Europace*, vol. 14 Suppl 5, pp. v10–v16, Nov. 2012.

[33]R. H. Clayton, E. A. Zhuchkova, and A. V. Panfilov, "Phase singularities and filaments: simplifying complexity in computational models of ventricular fibrillation," *Prog Biophys Mol Biol*, vol. 90, no. 1-3, pp. 378–98, 2006.

[34]M.-A. Bray and J. P. Wikswo, "Use of topological charge to determine filament location and dynamics in a numerical model of scroll wave activity," *IEEE Trans Biomed Eng*, vol. 49, no. 10, pp. 1086–93, 2002.

[35]R. Sedgewick and K. Wayne, *Algorithms (4th edn)*. Reading, MA: Addison-Wesley, 2011.

[36]J. Munkres, "Algorithms for the assignment and transportation problems," *Journal of the society for industrial and applied mathematics*, vol. 5, no. 1, pp. 32–38, 1957.

[37]R. Samtaney, D. Silver, N. J. Zabusky, and J. Cao, "Visualizing features and tracking their evolution," *Computer*, vol. 27, pp. 20–27, 1994.

[38]N. Virag, V. Jacquemet, C. S. Henriquez, S. Zozor, O. Blanc, J. M. Vesin, E. Pruvot, and L. Kappenberger, "Study of atrial arrhythmias in a computer model based on magnetic resonance images of human atria," *Chaos*, vol. 12, no. 3, pp. 754–763, 2002.

[39]J. Kneller, R. Zou, E. J. Vigmond, Z. Wang, L. J. Leon, and S. Nattel, "Cholinergic atrial fibrillation in a computer model of a two-dimensional sheet of canine atrial cells with realistic ionic properties," *Circ Res*, vol. 90, no. 9, pp. E73–87, 2002.

[40]A. Gharaviri, S. Verheule, J. Eckstein, M. Potse, P. Kuklik, N. H. Kuijpers, and U. Schotten, "How disruption of endo-epicardial electrical connections enhances endo-epicardial conduction during atrial fibrillation," *Europace*, vol. 19, no. 2, pp. 308–318, 2017.

[41]C. H. Roney, J. D. Bayer, S. Zahid, M. Meo, P. M. J. Boyle, N. A. Trayanova, M. Haissaguerre, R. Dubois, H. Cochet, and E. J. Vigmond, "Modelling methodology of atrial fibrosis affects rotor dynamics and electrograms," *Europace*, vol. 18, no. suppl 4, pp. iv146–iv155, 2016.

[42]V. Jacquemet, "A statistical model of false negative and false positive detection of phase singularities," *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 27, no. 10, p. 103124, 2017.

[43]F. Fenton and A. Karma, "Vortex dynamics in three-dimensional continuous myocardium with fiber rotation: Filament instability and fibrillation," *Chaos (Woodbury, N.Y.)*, vol. 8, pp. 20–47, Mar. 1998.

[44]Y. Gong, D. Deng, Y. Zhang, and L. Xia, "Automatic location of phase singularities in cardiac spiral wave reentry simulation," in *Computing in Cardiology*, vol. 38, pp. 473–476, 2011.

[45]R. Zou, J. Kneller, L. J. Leon, and S. Nattel, "Substrate size as a determinant of fibrillatory activity maintenance in a mathematical model of canine atrium," *Am J Physiol Heart Circ Physiol*, vol. 289, no. 3, pp. H1002–12, 2005.

[46]M. J. Bishop and G. Plank, "The role of fine-scale anatomical structure in the dynamics of reentry in computational models of the rabbit ventricles," *The Journal of Physiology*, vol. 590, pp. 4515–4535, Sept. 2012.

[47]C. H. Roney, F. S. Ng, M. T. Debney, C. Eichhorn, A. Nachiappan, R. A. Chowdhury, N. A. Qureshi, C. D. Cantwell, J. H. Tweedy, S. A. Niederer, N. S. Peters, and E. J. Vigmond, "Determinants of new wavefront locations in cholinergic atrial fibrillation," *Europace*, vol. 20, pp. iii3–iii15, Nov. 2018.

[48]W.-J. Rappel and S. M. Narayan, "Theoretical considerations for mapping activation in human cardiac fibrillation," *Chaos*, vol. 23, no. 2, p. 023113, 2013.

[49]C. H. Roney, C. D. Cantwell, J. D. Bayer, N. A. Qureshi, P. B. Lim, J. H. Tweedy, P. Kanagaratnam, N. S. Peters, E. J. Vigmond, and F. S. Ng, "Spatial resolution requirements for accurate identification of drivers of atrial fibrillation," *Circ Arrhythm Electrophysiol*, vol. 10, no. 5, p. e004899, 2017.