Title
Characterizing Excitation Patterns in Cardiac Arrhythmias

Permalink
https://escholarship.org/uc/item/61m16010

Author
Vidmar, David

Publication Date
2018

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA, SAN DIEGO

Characterizing Excitation Patterns in Cardiac Arrhythmias

A dissertation submitted in partial satisfaction of the
requirements for the degree
Doctor of Philosophy

in

Physics

by

David Michael Vidmar

Committee in charge:

Wouter-Jan Rappel, Chair
Daniel Arovas
Bo Li
Jeremie Palacci
Massimo Vergassola

2018

The dissertation of David Michael Vidmar is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

Chair

University of California, San Diego

2018

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGEMENTS

At the outset I must admit that this section, as with all acknowledgments, will remain inevitably incomplete. The respective influence that countless colleagues, friends, and acquaintances must have unwittingly had on my growth as a scientist, and by extension on this work, cannot be divvied up and distributed without the inescapable and unknowing omission. Nevertheless, I make my best effort below.

First, I would like to thank the American Heart Association for graciously funding much of this work through a predoctoral fellowship (16PRE30930015). I am grateful for the opportunity this afforded me to devote myself fully to research, and wish to thank them for taking a chance funding a project poised outside of the traditional biomedical disciplines.

Next I would like to thank my advisor, Dr. Wouter-Jan Rappel, for his guidance and mentorship. I am particularly grateful for his willingness to treat me not solely as a student, but also as a colleague. My maturation as a researcher throughout my time at UCSD has assuredly been shaped by our myriad insightful discussions. Sometimes focused, other times meandering, the ability to so freely discuss and dissect ideas and approaches has taught me a great deal about the structure of scientific discourse.

I also want to thank my family, who have been a constant source of encouragement throughout my education, from my aimless first days of college to the completion of my thesis project. Beyond material support they are owed much of the credit for the traits which got me to this point. After all, motivation and self-discipline are not exclusively innate qualities, but values imputed through observation. I am grateful to have had such pure examples of determination to observe.

Finally, I am grateful for the love and support of Daniella Bardalez Gagliuffi, who through four and half years, five submitted papers, nine apartments, and countless bottles of wine has remained my closest companion. She has been unwavering in her support throughout this time, providing stability and comfort during periods of uncertainty. And while scientific results may

exist in a vacuum, the process of exploration leading to them is undoubtedly something more personal. I count myself lucky to have had such a devoted partner by my side, with whom to share in the accomplishments which gave this journey meaning. Indeed, without someone so close to share in the moments of discovery, this entire process would surely have felt, instead, as a mere excavation of facts.

Chapter 1, in part, is currently being prepared for submission for publication of the material. Vidmar, David; Rappel, Wouter-Jan. The dissertation author was the primary investigator and author of this paper.

Chapter 2, in part, is currently being prepared for submission for publication of the material. Vidmar, David; Rappel, Wouter-Jan. The dissertation author was the primary investigator and author of this paper.

Chapter 3, in part, is a reprint of the material as it appears in (i) the American Journal of Physiology-Heart and Circulatory Physiology 2015, 309(12), H2118-H2126. Vidmar, David; Narayan, Sanjiv; Rappel, Wouter-Jan. (ii) Physical Review E 2016, 94(5), 050401. Vidmar, David; Narayan, Sanjiv; Krummen, David; Rappel, Wouter-Jan. (iii) JACC: Clinical Electrophysiology 2017, 3(12), 1437-1446. Vidmar, David; Krummen, David; Hayase, Justin; Narayan, Sanjiv; Ho, Gordon; Rappel, Wouter-Jan. The dissertation author was the primary investigator and author of these papers.

Chapter 4, in part, is currently being prepared for submission for publication of the material. Vidmar, David; Rappel, Wouter-Jan. The dissertation author was the primary investigator and author of this paper.

VITA

2012	B. Sc. in Physics, Minor in Mathematics, Pennsylvania State University
2014	M. Sc. in Physics, University of California, San Diego
2018	Ph. D. in Physics, University of California, San Diego

PUBLICATIONS

Vidmar, D., Narayan, S. M., & Rappel, W. J. (2015). Phase synchrony reveals organization in human atrial fibrillation. *American Journal of Physiology-Heart and Circulatory Physiology*, 309(12), H2118-H2126.

Vidmar, D., Narayan, S. M., Krummen, D. E., & Rappel, W. J. (2016). Determining conduction patterns on a sparse electrode grid: Implications for the analysis of clinical arrhythmias. *Physical Review E*, 94(5), 050401.

Alhusseini, M., Vidmar, D., Meckler, G. L., Kowaleski, C. A., Shenasa, F., Wang, P. J., ... & Rappel, W. J. (2017). Two Independent Mapping Techniques Identify Rotational Activity Patterns at Sites of Local Termination During Persistent Atrial Fibrillation. *Journal of Cardiovascular Electrophysiology*.

Vidmar, D., Krummen, D. E., Hayase, J., Narayan, S. M., Ho, G., & Rappel, W. J. (2017). Spatiotemporal Progression of Early Human Ventricular Fibrillation. *JACC: Clinical Electrophysiology*, 3(12), 1437-1446.

ABSTRACT OF THE DISSERTATION

Characterizing Excitation Patterns in Cardiac Arrhythmias

by

David Michael Vidmar

Doctor of Philosophy in Physics

University of California, San Diego, 2018

Wouter-Jan Rappel, Chair

Heart rhythm disorders represent a significant global health concern, affecting millions of people worldwide and serving as a risk factor for heart failure. Despite the prevalence of these arrhythmias, broadly effective treatments remain elusive. In this work, we present the cardiac conduction system as an excitable medium whose dynamics can be described with simple mathematical models. We show that the excitation of this system can become altered through well-timed stimuli which create persistent rotational activation patterns. We describe the hypothesized relationship between these patterns and the perpetuation of fibrillation, discussing the fragmenting state of multi-wavelet reentry as well as the theory of mother rotors. To distinguish between potential mechanisms, we present results using concepts of phase synchrony to characterize clinical recordings of fibrillation. Such analysis implies a level of spatiotemporal stability inconsistent with a mechanism exclusively comprised of chaotic multi-wavelet reentry. Moreover, we outline a technique to determine continuously updating vector flow fields which align with the direction of local conduction. We show that this methodology can be used to determine the source of rotational and focal excitation patterns in an automated fashion, providing a useful means to interpret otherwise complex spatiotemporal maps. Finally, we describe the spontaneous termination of fibrillation as a stochastic event, and construct a birth-death Master equation for the number of spiral tips n during the turbulent state of spiral-defect chaos. Within this framework, we can infer important statistical quantities related to fibrillation such as the mean number of spiral tips $\bar{n}$, the quasi-stationary distribution P_{qs}, and the mean episode duration τ. Our results imply that τ scales exponentially with the area A. We derive this scaling law using a WKB approach, which yields an effective Hamiltonian describing tip density $q = n/A$ and fluctuation momentum p. This stochastic approach can accurately predict τ, even for systems in which a multitude of episodes cannot be simulated due to prohibitive computational cost.

Chapter 1

Introduction

1.1 Background

Cardiac arrhythmias occur when the steady propagation of electrical activity across heart tissue is interrupted and replaced with irregular or altered excitation patterns. These heart rhythm disorders can be divided into various categories, such as tachycardia, where the heart rate is sped up, bradycardia, where the heart rate is slowed down, and fibrillation, where the heart rate becomes chaotic and uncoordinated. Moreover, arrhythmias can occur in either the atria, the upper chambers of the heart, or the ventricles, the lower chambers (Figure 1.1 A).

To better understand these altered heart rhythms, we first outline the normal process of conduction in a healthy heart. This process begins spontaneously at the sinoatrial (SA) node, known as the pace-maker of the heart, which initiates the propagation of electrical activity at a given rate. This excitation spreads out in a coherent wave across the atria, eventually reaching the atrioventricular (AV) node. This node works as a gatekeeper and collects the excitation from the atria. If the atrial excitation is disordered, due to atrial fibrillation (AF) for example, the AV node will work to block this disorderly state from propagating to the ventricles and causing ventricular fibrillation (VF). After a delay, the AV node passes this excitation wave on to propagate across

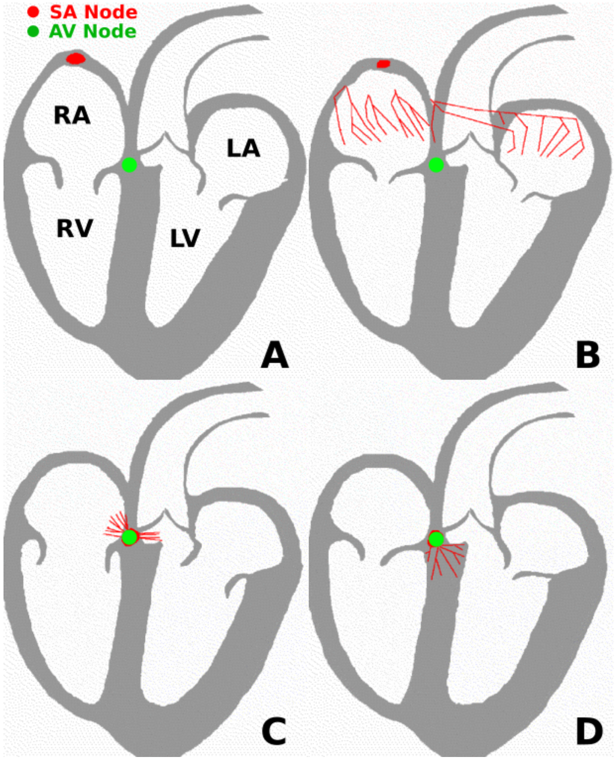

Figure 1.1: Normal conduction across a human heart. Conduction begins at the sinoatrial node (A), continues across the atria to the atrioventricular node (B and C), and proceeds across the ventricles (D). This figure has been adapted from Wikipedia Commons: ECG_Principle_fast.gif.

the ventricles. This process is illustrated in Figure 1.1 with excitation shown in red.

The ordered propagation of electrical activity described above is important to the heart's function because electrical excitations are responsible for the contraction of heart muscle and therefore for the effective circulation of blood throughout the body. During fibrillation, the characteristic mangled excitation patterns can lead to ineffective pumping of blood and can result in significant and life-threatening health concerns [FRC$^+$06, NTC$^+$08]. Not only do these heart rhythm disorders present a health risk, but they are also increasingly common worldwide. AF in particular has seen an increase in prevalence and burden over the past few decades[CRGM14]. In a 2010 study on the global epidemiology of AF, for example, it was reported that 33.5 million patients suffered from this disorder [CHN$^+$13]. As life expectancy continues to rise around the world, the burden of arrhythmias, which are often progressive disorders, is likely to become increasingly prominent.

1.2 Excitable Media

In studying abnormal electrical propagation during arrhythmias, it is instructive to view the heart as one example of a more general class of physical systems known as excitable media. This class of systems always contains individual components which remain at rest until acted upon by a stimulus of sufficient strength. Once stimulated, a component is excited for a short time before returning to its resting state, where it enters a refractory period in which it can no longer be stimulated. After an initial stimulus these excitations travel in waves across the system. Other examples of excitable media include diverse phenomenon such as Belousov-Zhabotinsky reactions, the spread of disease, and the aggregation of slime mold amoebae [TK88].

All excitable systems share important characteristics, allowing them to be described by qualitatively similar models. Because our goal is to examine electrophysiological systems, we focus our attention on the generation of action potentials and models of excitable systems derived therein. One of the first and most important of these models was proposed in 1952 by Alan Hodgkin and Andrew Huxley to describe the nerve excitations of the squid giant axon [HH52]. Their work, primarily focused on clever experimental design, implicated ionic mechanisms as responsible for the generation of action potentials. In particular, their work pointed to the importance of evolving potential-dependent ion permeabilities in the cell membrane as the culprit of rapid changes in transmembrane voltage during excitation.

To model the complex electrochemical process resulting in excitation, it is convenient to describe cell membranes with an analogous circuit [KS10]. This is because the cell membrane is selectively permeable, allowing certain ions such as Na^+, K^+, and Cl^- to pass only through specific ion channels, thereby maintaining a separation of charge and regulating the potential difference across the membrane. In this sense, the membrane can be viewed as a capacitor and a conductor in parallel, with the conductor carrying a total current I_{ion}. This gives us the basic equation for the electrical activity across the cell membrane as

$$C_m \frac{dV}{dt} + I_{ion} = 0, \tag{1.1}$$

where V is the transmembrane voltage and C_m is the membrane capacitance per unit area.

1.3 Models of Excitable Systems

The dynamics of excitation in cells can now be seen to be described completely by the ion channel dynamics responsible for I_{ion}. As Hodgkin and Huxley found, these currents depend on voltage and time and, for the squid giant axon, are show to be well described by the simple form $I = g(V - V_{rev})$ where g and V_{rev} are properties of the specific ion channel and represent conductance per unit area and the reversal potential, respectively. The Hodgkin-Huxley (HH) model comprises sodium and potassium currents as well as a group of smaller currents lumped together to represent the leakage current. This model can then be described by the following equation

$$C_m \frac{dV}{dt} = -g_{Na}(V - V_{Na}) - g_K(V - V_K) - g_L(V - V_L) + I_{ext}, \tag{1.2}$$

where I_{ext} is any external or applied current.

In order for this equation to retain the observed excitatory features we expect in our model, some of these conductances themselves must be dependent on voltage (and time). These nonlinear dependences in the HH model were ascertained through a set of voltage clamp experiments which measured the evolution of individual conductances. In particular, Hodgkin and Huxley introduced three dimensionless gating variables, n, m, and h. These variables are related to the conductances as $g_{Na} = \bar{g}_{Na} n^4$, $g_K = \bar{g}_K n^3 h$, and $g_L = \bar{g}_L$ with $\bar{g}_i$ as the maximum value of conductance for the

i^{th} ion channel. These gating variables all obey differential equations of the following form:

$$\frac{dx}{dt} = \alpha_x(V)\,(1-x) - \beta_x(V)\,x, \tag{1.3}$$

where x is the specific ion channel and α_x and β_x are rate constant functions which depend on the voltage and were chosen by Hodgkin and Huxley to match their experimental results. The HH model, then, is defined by Equation 1.2 together with the set of Equations 1.3 for each gating variable. This quantitative model can be shown to exhibit the characteristic features of an excitable system, namely the generation of an action potential in response to an external stimulus.

While the HH model can be made to fit experimental data quite well, it can be advantageous to simplify it in order to inspect and elucidate its key features. By noting that the ion channel dynamics in the HH model occurred on two distinct time scales, Richard FitzHugh proposed a simplified two-variable model of excitation in 1961 [Fit61]. This model, named the FitzHugh-Nagumo (FHN) model, can be described by the following equations

$$\frac{dv}{dt} = v - \frac{v^3}{3} - w + I_{ext} \tag{1.4}$$

$$\frac{dw}{dt} = \phi(v + a - bw), \tag{1.5}$$

where ϕ, a, and b are parameters of the model, v is the excitation variable, and w is the recovery variable.

By reducing the model to only two-variables, this generic system can be studied with simple phase-plane analysis [Koc04]. In Figure 1.2 we plot the corresponding phase plane of this system with no external current and with parameters $\phi = 0.08$, $a = 0.7$, and $b = 0.8$ [Cro87]. From this plot we see an equilibrium point at $(v,w) = (-1.20, -0.625)$, which can be shown through linearization to be an attractor of the system. If we inject an instantaneous current into this system, we see two different behaviors depending on the magnitude of the stimulus. If the

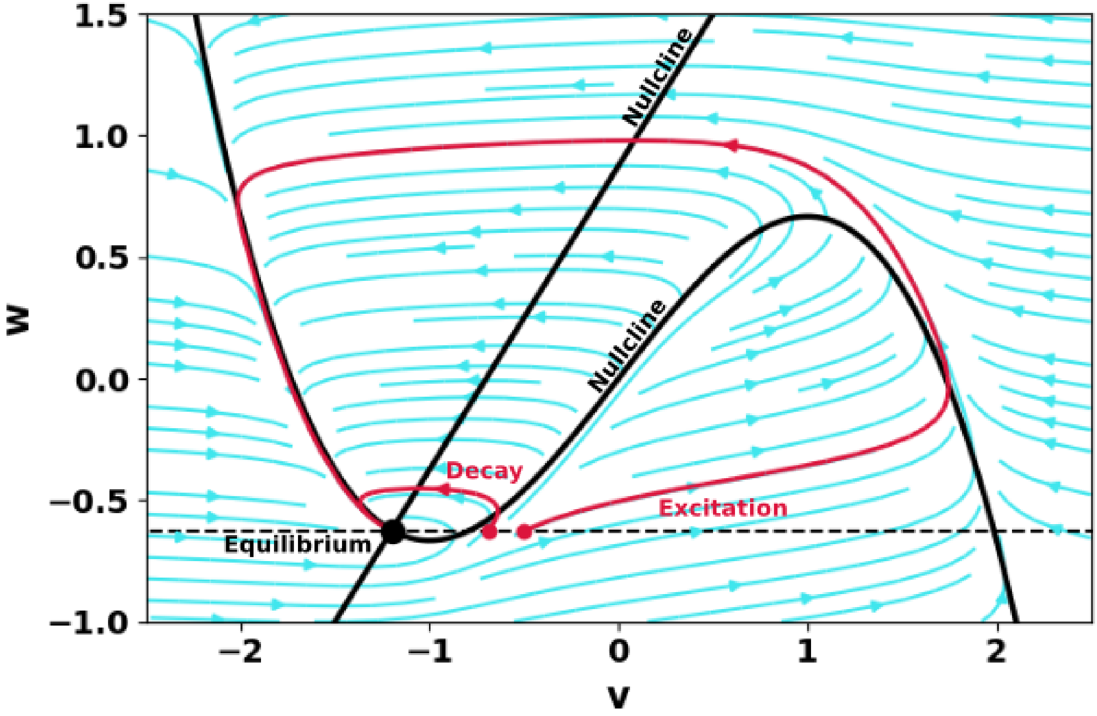

Figure 1.2: Phase plane of the FHN model. Streamlines of the phase portrait are plotted in blue with the equilibrium point as a black dot. A small stimulus initiates the decay trajectory, whereas a large stimulus initiates the excitation trajectory, both in red.

stimulus is small, the phase trajectory rapidly returns to its resting state. If, however, the stimulus is sufficiently large, the phase trajectory follows an excursion whereby v initially increases and then slowly returns to its resting state. This excitatory behavior is qualitatively similar to action potential generation, allowing the FHN model to recapitulate generic excitable phenomenon when specific ionic dynamics are not important.

Chapter 1, in part, is currently being prepared for submission for publication of the material. Vidmar, David; Rappel, Wouter-Jan. The dissertation author was the primary investigator and author of this paper.

Chapter 2

Arrhythmias in Computer Models

2.1 Fenton-Karma Model

While the FitzHugh-Nagumo model is a convenient starting point to understand excitability, it does not retain certain simple features likely to be relevant to cardiac myocyte cells. In particular, it does not allow for the different time scales known to be associated with depolarization and repolarization in cardiac action potentials [DB00]. In fact, since the pioneering work of Hodgkin and Huxley, many models have been proposed which focus on detailed descriptions of the complicated ion channel dynamics of cardiac cells. These are often fit to animal data and contain many variables, making them very computationally expensive (see [NGN12] for a comprehensive list).

Because our work is concerned with mesoscopic excitation patterns, rather than microscopic ion dynamics, simplified models which exhibit realistic pattern formation mechanisms are advantageous to complex ones. We therefore model generic cardiac excitation using the Fenton-Karma (FK) model [FK98], which includes only three variables, and yet can accurately capture the restitution properties of more complex models [FCHE02]. It achieves this by including two ionic gate variables, with one corresponding to a fast time scale and the other to a slow time scale.

The net current in the model is then described by a fast inward current, a slow inward current, and a slow outward current.

2.2 Initiation of Arrhythmias

Using the FK model of cardiac excitation we can begin to examine the genesis of interesting in silico excitation patterns on simple domains. We begin with a sheet of cardiac myocyte cells, electrically coupled via gap junctions, and with no flux boundary conditions. The electrical coupling will allow for the propagation of excitation waves in response to a stimulus, and can be described by adding a diffusion term to Equation 1.1 yielding

$$\partial_t V = D\nabla^2 V - I_{ion}/C_m, \tag{2.1}$$

where D is a diffusion tensor.

The simplest excitation patterns we can expect from this system are circular or planar waves, which propagate across tissue and eventually annihilate at a boundary. In principle this can be viewed as the type of excitation occurring during a normal heart beat, where the sinoatrial node acts as a periodic stimulus. There are, however, more interesting excitation patterns which exist as solutions to our model. In particular, excitation in the form of spiral waves can be shown to propagate indefinitely and have long been hypothesized to play an important role in heart arrhythmias [WR46, MA59, ABS73]. Once these spiral waves are initiated, they usually excite tissue at a higher frequency than the sinoatrial node and therefore become the de-facto pacemaker of surrounding tissue. In a real heart, then, these high frequency spiral waves would overtake the sinoatrial node and set the tempo for the ensuing dynamics, marking the advent of an arrhythmia.

Understanding the genesis of these spiral waves informs us of the possible mechanisms for the initiation of arrhythmias in diseased and healthy hearts. One means of initiating spiral waves, named the pinwheel experiment, was proposed by Arthur Winfree [Win87] and posits

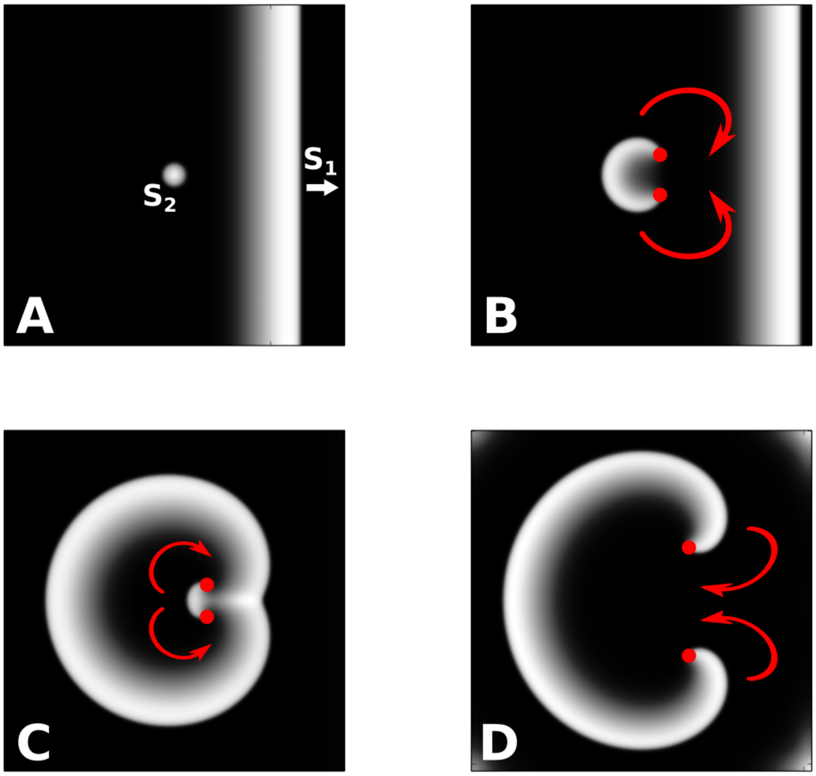

Figure 2.1: Initiation of a spiral wave in the Fenton-Karma model. (A) shows a snapshot of a plane wave S_1 propagating to the right, with a point stimulus S_2 applied during the corresponding vulnerable window. Transmembrane voltage is shown in grayscale. (B) shows the simulation at a later time where the stimulus S_2 cannot conduct to the right due to tissue still in the refractory period, causing it to rotate around two pivots shown in red. (C) and (D) show the resulting dual spiral wave pattern created due to this unidirectional conduction block.

that a well-timed pair of stimuli can create a dual spiral wave pattern. This initiation mechanism in the FK model is shown in Figure 2.1, where we have used the MBR parameter set in [FK98] with $\tau_d = 0.40$. Here, a planar wave S_1 is first stimulated at the left side of the domain after which it propagates to the right. At a later time, a point stimulus S_2 is applied at the center of the domain. This stimulus attempts to propagate radially outwards, but cannot propagate across tissue to the right, which is still in its refractory period. This unidirectional conduction block causes the stimulus to rotate around two pivots created at the interface between excitable and non-excitable tissue. As S_1 propagates rightward, and the refractoriness with it, the arms of S_2 find room to complete a full rotation and continue rotating around their respective cores indefinitely.

This initiation mechanism is insightful because it allows us to understand possible reasons for the difference in inducibility of arrhythmia between diseased and healthy hearts. As was described by Winfree, and further explored in simulation [YTW$^+$07], the second stimulus S_2 in this scenario must reside within a vulnerable window in order to successfully create a pair of long-lived spiral waves. Certainly this corresponds to the timing between the two stimuli because if S_2 is too early it will attempt to excite refractory tissue, whereas if it is too late there is no unidirectional conduction block. However, this window also depends on the intensity of the stimulus because if the depolarized region corresponding to S_2 is too small, the two arms of this stimulus will collide with each other and annihilate before completing a full rotation. It has therefore been hypothesized that diseased hearts, which more easily induce arrhythmia, may have cellular properties which make their vulnerable window larger causing them to be more susceptible to the occasional stimulus creating spiral waves.

2.3 Maintenance of Fibrillation and Pattern Formation

While the aforementioned mechanism can explain the initiation of an arrhythmia, it says nothing about the ensuing dynamical state. In particular, fibrillation is characterized by a chaotic state not well described by simple stable excitation patterns alone. Explanations of this chaotic state most commonly fall into one of two main hypothesized mechanisms. One mechanism focuses on a form of self-sustaining disorder due to the continual fragmentation of an initiated spiral wave, whereas a second mechanism involves a region of order surrounded by peripheral disorganization.

To understand the origins of these theorized mechanisms, we take a historical approach with focus on work related to understanding atrial fibrillation (AF). Early on it was largely believed that AF was a result of so-called circus movement reentry, where an excitation wave traversed a rotational circuit created around some anatomical obstacle [Min13, Lew21]. This theory explained the difference between atrial flutter, a more organized tachycardia, and AF as a difference in the size of this rotational circuit. If this circuit is large, the excitation wave extends comfortably around the circuit with wavefront and waveback well-separated, leading to an orderly tachycardia. If, however, the circuit is some critical size, the excitation wavefront will follow the waveback closely, potentially leading to disorder and AF.

2.3.1 Multi-Wavelet Reentry

The theory of circus movement reentry remained predominant until Gordon Moe, skeptical that such a theory could explain cases of chronic AF which can persist for many months at a time, proposed a radically different vision of reentry [MA59]. In his theory, AF was caused not by a single rotational circuit, but by a multitude of randomly wandering wavelets. These wavelets fragment and coalesce, continuously generating offspring and creating a state of self-sustaining turbulence accounting for the disorganization inherent in AF. In this view, such disorganization occurs homogeneously across tissue, with no regions exhibiting stable or organized excitation.

To expand on this theory of fragmenting excitation waves, termed multi-wavelet reentry (MWR), Moe proposed a computer model of such a state [MRA64]. He showed that by including heterogeneous refractory dynamics across the domain, he could induce initiated spiral waves to breakup into offspring wavelets and create self-sustaining turbulence. This state can also be initiated in the FK model by tuning the parameters to values which encourage wavebreak. In Figure 2.2 A we show the results of a simulation using parameter set 3 from [FCHE02], with periodic boundary conditions imposed to allow this turbulent state to persist for extended periods of time. In this simulation, the cores of rotational wavelets, marked with red dots, wander randomly around the domain until they are annihilated through collisions.

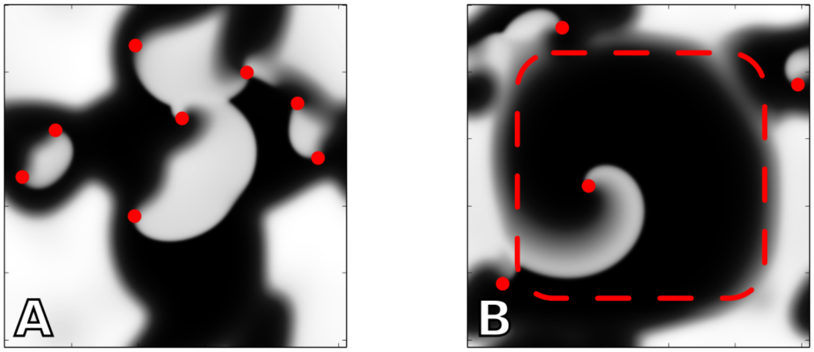

Figure 2.2: Simulation of two hypothesized mechanisms of fibrillation. (A) shows a simulation of multi-wavelet reentry, with transmembrane voltage in grayscale and rotational cores marked as red dots. Spiral waves continually fragment and traverse tissue in a stochastic fashion. (B) shows a simulation of a stable spiral wave surrounded by fibrillatory conduction in the periphery. Tissue in the central domain has parameters which allow for stable spiral patterns, whereas tissue outside this domain has parameters which generate wavebreak and fibrillatory conduction.

2.3.2 Mother Rotors and Fibrillatory Conduction

As clinical and experimental observations began to point towards underlying order during AF [DPS$^+$92, GPJ98, Jal10], as well as hierarchies of activation frequencies unexplainable by the uniform disorganization of MWR [MMB$^+$01], a second theory of AF maintenance began to emerge. In this theory, AF is perpetuated by a localized source, in the form of a stable spiral wave known as a mother rotor, which drives peripheral disorder. This disorganization occurs distant to the core of the rotor, whereby the arms of the spiral wave experience wavebreak and fragmentation as they propagate through heterogeneous tissue. Such peripheral disorder is termed fibrillatory conduction.

This theory accounts for the observed order during AF, while also explaining the disorganized electrical activity characteristic of fibrillation. We use the FK model to simulate this mechanism in Figure 2.2 B. Here we have initiated a single spiral wave inside of a central domain of tissue with parameters allowing for the propagation of stable spiral waves (the MBR parameter set in [FK98]). Outside of this central domain, which is marked with a red dotted line, we have tissue with parameters set to encourage wavebreak (parameter set 3 from [FCHE02]). In this heterogeneous substrate, then, the coherent spiral wave excitation in the central domain corresponds to the mother rotor, which drives fibrillatory conduction in tissue outside of this domain.

Chapter 2, in part, is currently being prepared for submission for publication of the material. Vidmar, David; Rappel, Wouter-Jan. The dissertation author was the primary investigator and author of this paper.

Chapter 3

Arrhythmias in Humans

3.1 Optical Mapping

Up to this point we have discussed various excitation patterns which are viable in our model of arrhythmia. To distinguish between these and elucidate the actual dynamical state underlying arrhythmia necessitates data. In this chapter we focus on data recorded during episodes of fibrillation, as it is the most difficult arrhythmia to faithfully record and poses the most significant risk to patients worldwide. There are two primary means of acquiring data to map the dynamics of fibrillation, either using potentiometric dyes or spatially-distributed electrodes.

To date, the gold standard in determining and visualizing electrical activity during fibrillation is through the use of optical mapping. In this technique, voltage-sensitive dyes are applied to the heart and recorded with optical sensors throughout induced episodes of fibrillation. The fluorescence of the dye provides a surrogate for the transmembrane voltage with surprising accuracy, matching the shape of the underlying action potential [ENS04]. The resulting time series of fluorescence is then often transformed into phase with each new cycle corresponding to an excitation event at a given location. Spatial maps of the evolution of this phase can be used to identify excitation patterns present during fibrillation. Of note, the use of optical mapping has revealed the existence of stable rotational excitation patterns during fibrillation in animals [GPJ98, ZYKB08, MSC$^+$00, PDS$^+$93] as well as in explanted human atria [HZC$^+$15].

3.2 Multipolar Contact Mapping

While optical mapping provides high resolution views of the electrical activity underlying fibrillation, it requires the use of toxic voltage-sensitive dyes and therefore is not currently viable in human patients. Clinical mapping of fibrillation must instead be conducted with recordings of electrograms from spatially-distributed electrodes. Moreover, it is important to observe fibrillation mechanisms in human patients directly, in addition to those observed in animal models, because certain electrophysiological characteristics are known to be different among different animals [BZN17].

One popular means of recording electrograms during clinical fibrillation is through the use of a 64-pole basket catheter, which is inserted into a chamber of the patient's heart. This basket catheter contains eight semi-flexible splines with eight equally spaced electrodes per spline. Each spline separates to form a spherical shape, pushing against the inner wall of the heart and creating a grid of 64 electrodes spread out across the chamber. This electrode grid can then be approximated as a simple 8x8 grid on a two-dimensional domain similar to simulations from the previous chapter. Fibrillation is then induced, and each electrode records the electrical activity of nearby tissue. These electrogram recordings can then be analyzed and used to map the ensuing episode of fibrillation.

The voltage recorded from these electrodes requires processing to be displayed as a meaningful representation of underlying excitation patterns. This is because the voltage amplitude can vary significantly between electrograms, owing to regions of scarring and recording artifacts [UNM$^+$10]. Further, electrograms represent a weighted sum of electrical activity, and therefore can be influenced by far-field activity, signal artifacts, and noise. A useful simplified representation of these complex electrograms is its set of activation times, which is a set of discrete times marking the onset of each excitation. These excitations are encoded as sharp deflections in voltage, which can be determined either manually or through the use of algorithms [NKR12, NKER12].

3.2.1 Phase Synchrony

We can use activation times marked from electrograms to explore the level of order in clinical arrhythmias as a means of distinguishing between the proposed mechanisms described in the preceding chapter. In particular, multi-wavelet reentry implies uniform disorganization whereas a mother rotor with fibrillatory conduction would have regions of tissue controlled by this rotor which would exhibit coordinated activation. We propose to use a measure of phase synchrony as a means to explore order and disorder in clinical and simulated fibrillation.

Methodology

To compute the level of phase synchrony between electrodes, we first convert the activation times from either in silico data or from clinical recordings into phase-time information. This is achieved by taking the data of an electrode or, in the case of a simulation, a virtual electrode located within the computational domain, as a periodic event, whereby the activation times mark the beginning of each new cycle. The activation time is then chosen as the time point where the phase has increased by 2π. Between activations, the phase keeps track of how far along a cycle each electrode is in time and is obtained using linear interpolation.

In particular, let the activation times for the i^{th} electrode be given by $\{t_i^k : k = 1, 2, ..., N\}$, where N is the total number of activation times. Then, the phase of electrode i between the k-th and $k+1$-th activation time is defined as:

$$\phi_i(t) = 2\pi \frac{t - t_i^k}{t_i^{k+1} - t_i^k} + 2\pi(k-1), \quad t_i^k \leq t < t_i^{k+1}. \tag{3.1}$$

Thus, the phase starts at 0 and is advanced by 2π every time the electrode records an activation event. In order to determine phase synchrony, we are not concerned with a single electrode's phase, but instead with the time-evolution of the phase difference between two electrodes. We can therefore define the time-dependent relative phase between the i^{th} and j^{th} electrode as $\psi_{ij} = n\phi_i - m\phi_j$, where n and m are integers. This gives us information about the extent to which two separate electrodes' activation dynamics tend to be proceed together in time. If the relative phase between electrodes remains constant over a given time period, those electrodes are phase locked and can be said to be synchronized. Note that we are only concerned with how constant the phase difference is with time, and therefore any global offset in either phase will not affect our results.

If the signal is purely periodic then this synchronization corresponds quantitatively to the phase locking condition $|n\phi_i - m\phi_j - \delta| < \text{const}$, where δ is an average (constant) phase shift. Because our system is inherently noisy, however, ψ_{ij} fluctuates, resulting in a statistical distribution of the cyclic relative phase

$$\Psi_{ij} = \psi_{ij} \bmod 2\pi.$$

Synchronization corresponds to a peak in this distribution while asynchronous signals result in a uniform distribution [RPK$^+$01]. For this paper we will only consider 1:1 phase locking, thus taking $n = m = 1$.

A useful way to quantify the extent of phase locking in noisy biological systems is to find the synchronization number γ of the phases in question. This is defined as the amplitude of the first Fourier mode of the cyclic relative phase distribution [RPK$^+$01]

$$\gamma_{i,j}^2 = \langle \cos \Psi(t)_{i,j} \rangle_t^2 + \langle \sin \Psi(t)_{i,j} \rangle_t^2, \tag{3.2}$$

where the brackets denote an average over time. Note that $\gamma = 1$ if the two signals are phase-locked (corresponding to a single-peaked distribution of the cyclic relative phase) and $\gamma = 0$ if the two signals are completely asynchronous (corresponding to a uniform distribution of Ψ_{ij}).

The inter-electrode synchrony between all M recording electrodes can be quantified using a network where the nodes correspond to the electrodes and where the connections represent the level of synchrony. To simplify this analysis we will assume a binary network such that the entries a_{ij} of the $M \times M$ adjacency matrix are either 1 (implying synchronization between electrodes i and j) or 0 (corresponding to asynchrony). The values of a_{ij} are determined using a threshold value, $\kappa = 0.95$, of the corresponding synchronization number $\gamma_{i,j}$, with $a_{ij} = 1$ for $\gamma_{i,j} \geq \kappa$. The synchronization index SI for electrode i is then defined as its normalized degree, namely the fraction of synchronized connections between that electrode and all other electrodes:

$$SI_i = \frac{1}{M-1} \sum_{\substack{j=1 \\ j \neq i}}^{M} a_{ij}. \tag{3.3}$$

In words, the SI of electrode i quantifies the fraction of electrodes with which i remains synchronized with over an extended period of time. Therefore it encodes information about the spatio-temporal inter-connectedness of activations during AF, and can be used to quantify regions of tissue which activate together synchronously. In particular, we can easily compute the SI for all electrodes, and make a spatially distributed SI map for a given window of time. Spatially contiguous regions with high values of SI, then, imply a region of spatially-extended synchrony.

Note that SI takes on values between 0 and 1 and depends on both number of synchronized electrodes and the total number of electrodes. For example, if there is a group of $\tilde{m}$ electrodes synchronized with each other, then each electrode within this group will have $\tilde{m} - 1$ synchronized connections and a SI value of $(\tilde{m} - 1)/(M - 1)$. Similarly if we have global synchrony over all electrodes we get $SI = 1$ for every electrode.

In Silico Data

Before applying our synchronization analysis to clinical data we examined several activation time series generated using computational modeling. In a first example, we created a single dominant spiral wave in a homogeneous domain with no-flux boundary conditions. A snapshot of this simulation, focusing on the region that contains the spiral tip, is shown in Fig. 3.1A, where the activation is shown in a gray scale with white corresponding to depolarized (active) and black corresponding to repolarized (inactive) tissue. In Fig. 3.1B we plot the corresponding *SI* map for the coarse-grained network which reveals a high value for the synchronization index everywhere except at the location of the spiral wave tip. This is due to the tip meander, shown in white in Fig. 3.1B, which has a spatial scale that is larger than the inter-electrode spacing. As a consequence, the SI value for the electrodes that include the meander domain is reduced, demonstrating how our analysis can potentially determine the location of a spiral tip.

In a second simulation we generated timing data for a spiral wave with limited meander (Fig. 3.1C). As can be seen in Fig. 3.1D, the tip trajectory remains within a domain bordered by neighboring electrodes. As a result, the SI map (Fig. 3.1D) displays uniformly high value of synchrony. Thus, these simulations show that the synchronization index is able to identify coherent domains associated with spatially conserved spiral waves and can locate the spiral tip only if its meandering path exceeds a minimum spatial scale.

In a third simulation, we simulated a spatially stable spiral wave in a 15cmx15cm tissue (Fig. 3.1E) that contains two regions with different values of one of the model parameters (see Ref. [RN13]). The parameter value within a circular coherent domain, with a radius of 6 cm and centered in the computational domain, was chosen such that a spiral wave remains spatially stable. Outside this coherent domain, the parameter value was taken such that a spiral wave inherently breaks down as can be seen in the screenshot of the simulation (Fig. 3.1E). The corresponding *SI* map is plotted in Fig. 3.1F and shows a clearly synchronized domain, surrounded by asynchronous tissue. Importantly, the size of the synchronous region corresponds to the coherent domain of the

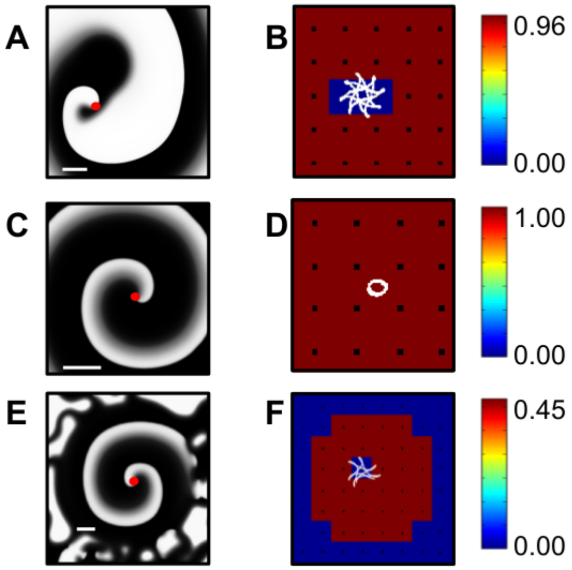

Figure 3.1: Simulated excitation patterns with their corresponding SI maps. (A) Snapshot of a simulation with a single stable spiral wave. Membrane potential is shown using a gray-scale. Parameters taken from set #1 in Ref. [FCHE02], modified to create stable spirals ($\tau_{v2}^{-} = 10$ and τ_{ω}^{+}=220). The location of the spiral wave tips is indicated by red dots. (B) The SI colormap corresponding to A, with red corresponding to highly synchronized electrodes and blue corresponding to asynchronous electrodes. The tip trajectory is indicated in white. (C) Snapshot of a simulation containing a spatially conserved spiral wave with only a small meander, obtained using the MBR parameter set from Ref. [FK98]. (D) The SI colormap corresponding to C. (E) Snapshot of a simulation in a domain consisting of a central region, harboring a single stable spiral, surrounded by a region in which this spiral breaks down into irregular, stochastic dynamics. Parameter values were taken from parameter set #3 in Ref. [FCHE02] with $\tau_r = 40$ in the inner and $\tau_r = 27$ in the outer domain. (F) The SI colormap corresponding to E, which displays a synchronized central region and an asynchronous outer region. The snapshots and corresponding SI colormaps have been scaled to emphasize tip meander, with scalebar=1cm shown in white.

spiral wave.

Clinical Data

To demonstrate the application of our synchronization analysis to clinical data, we begin by examining an episode of AF in a patient. In Fig. 3.2A we plot the *SI* map of the left atrium (LA) of a 65 year old patient for an episode corresponding to AF of duration 2.2 s. A small region in the center of the field of view displays localized synchrony and is surrounded by asynchronous

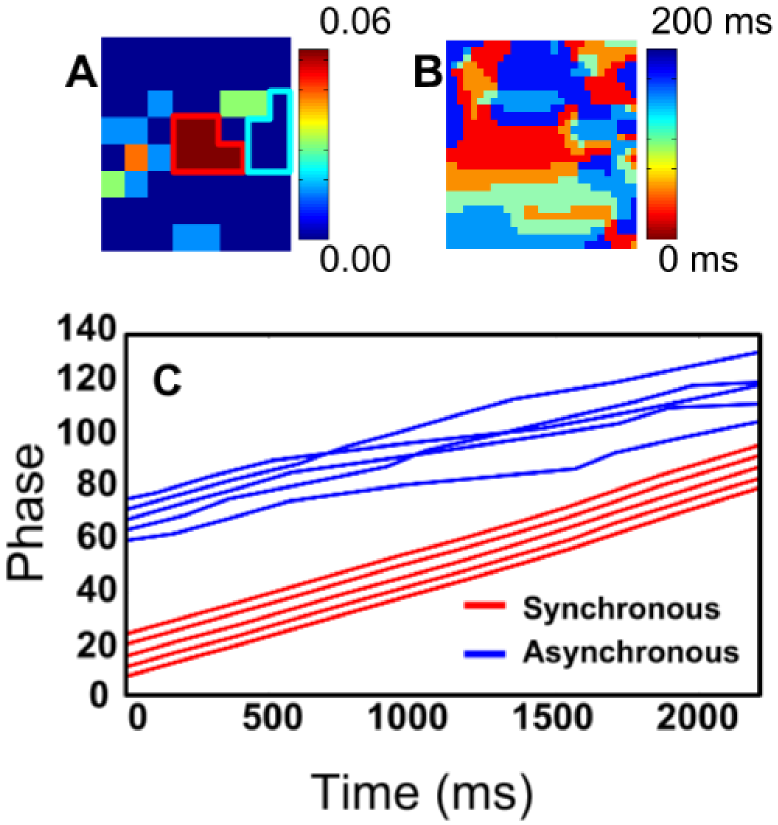

Figure 3.2: SI map of a clinical episode of atrial fibrillation. (A) The *SI* map displays a synchronized region in the posterior wall of the left atrium. (B) The corresponding isochronal map of this patient, with red (blue) corresponding to earliest (latest) activation, shows a complex activation pattern. (C) The phase as a function of time for 5 electrodes with the organized domain of A (red lines) and 5 electrodes outside this domain (blue lines). For visualization purposes, a constant offset is added to each phase.

tissue. The phase as a function of time for the 5 electrodes within this synchronized region is shown in Fig. 3.2C, along with 5 other nearby electrodes within the asynchronous domain, as indicated by the light blue border in Fig. 3.2A. For visualization purposes, each phase was phase-shifted. Clearly, the phases within the synchronous region display 1:1 phase locking such that their phase difference remains roughly constant. The other phases, in contrast, show multiple crossings, demonstrating that their phase differences vary significantly, resulting in low values of $\gamma_{i,j}$ and $a_{ij} = 0$. Interestingly, an isochronal map obtained within the same time frame shows significant variability throughout the domain, as shown in Fig. 3.2B.

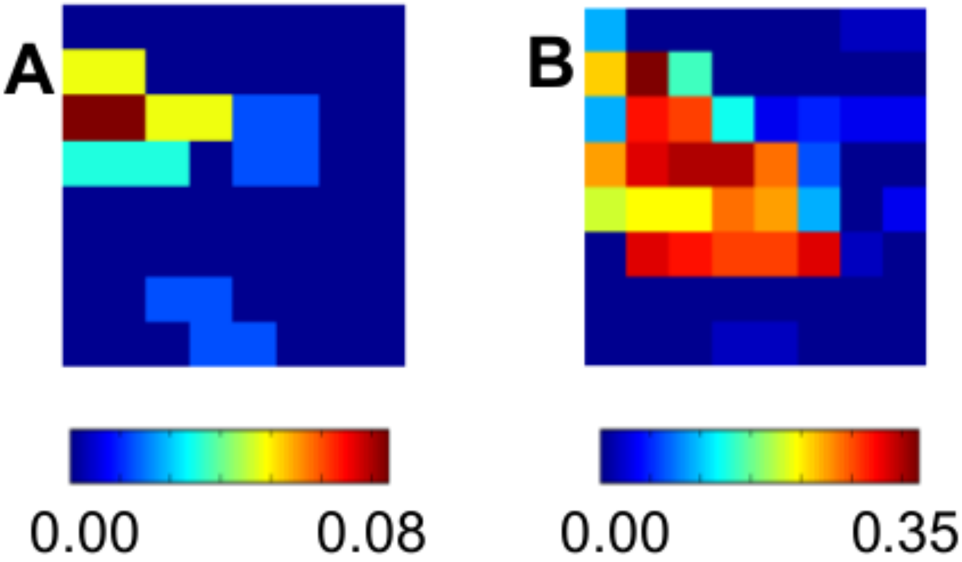

Figure 3.3: Human AF analyses, showing *SI* maps at different time intervals during an ablation procedure. (A) and (B) show this map at different stages of targeted ablation: (A) at onset and (B) after 46 min. The synchronized region in the posterior LA remains conserved, but has increased in size.

Another clinical example is provided in Fig. 3.3 where we plot the *SI* map for different 2.8 s time intervals during an ablation procedure. Fig. 3.3A corresponds to the SI map obtained at the start of the procedure and shows a synchronized region in the posterior LA. Fig. 3.3B shows a SI map from data obtained 46 min later. Between these intervals the synchronized region remains conserved, but has increased in size.

As a final application of our method, we compute the distribution of the fraction of synchronized electrodes, Ω, in patients prior to ablation. To this end, we examined the synchronization network in 24 patients for who bi-atrial electrode recordings were available. We compute these over a time period of 2.8 s, resulting in 22 RA and 15 LA recordings of sufficient length. The result of our analysis is shown in Fig. 3.4 and shows that this distribution is broad, containing patients with only a few synchronized electrodes and patients for which most electrodes are synchronized. This distribution is inconsistent with the uniform disorganization inherent to MWR, while it is compatible with the existence of rotors with peripheral disorganization.

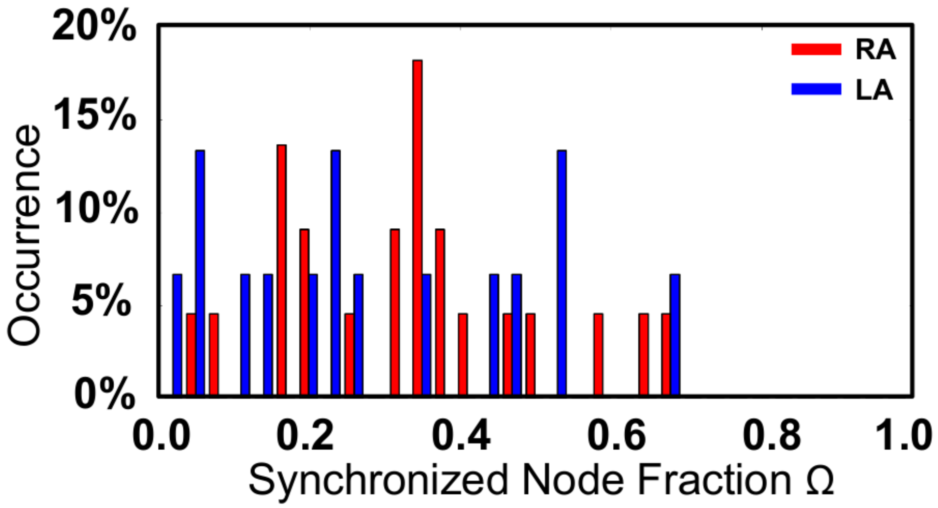

Figure 3.4: Distribution of the fraction of synchronized electrodes, Ω, for 24 patients during AF. Shown are the distributions of both the LA and RA before any ablation has occurred.

3.2.2 Asynchronous Index

The concept of phase synchrony can also be used to identify electrodes which activate in an asynchronous fashion with respect to the entire domain. This can be useful to determine spiral tips, as we see in Fig. 3.1 B where tip meander causes decreased synchrony. To this end we examine simulated activation data in our computational model as well as clinical activation data recored during ventricular fibrillation (VF). From these activation times we compute the time-dependent phase of each region of tissue as described in the previous section. This phase is then used to compute an asynchronous index (ASI) which quantifies how dynamically out-of-step a given electrode is with the surrounding tissue.

Methodology

To determine regions that are dynamically dissimilar to their surroundings, we first compute the distribution of synchronization numbers across all pairs of electrodes, γ_{global}. For each electrode i we then compute $\gamma_{local,i}$, the distribution of synchronization numbers between this electrode and all other electrodes. We can then define the ASI to quantify the difference between these two distributions and thus the extent to which electrode i records activation dynamics which are atypical from the global trend. This ASI uses the Hellinger distance $\mathcal{H}$, a standard statistical measure of dissimilarity between two distributions that ranges from 0, if two distributions are equivalent, to 1, if the two distributions are entirely distinct. It can also be written as $\sqrt{1-BC}$, where BC (the Bhattacharyya coefficient) measures the amount of overlap between two distributions. The ASI is then defined for electrode i as

$$ASI_i = \lambda_i \cdot \mathcal{H}(\gamma_{local,i}, \gamma_{global}), \tag{3.4}$$

where λ_i is a weighting factor equal to the mean synchrony between all nearest neighbors electrodes. The inclusion of this weighting factor ensures that ASI is only elevated if the asynchrony is localized in space and therefore some or all of the surrounding tissue is itself organized.

Note that we are always comparing local to global synchrony with this method and therefore it has the advantage that it is automatically normalized over different degrees of organization. Higher positive values of ASI indicate an electrode that displays vastly different synchrony compared to the global dynamics whereas values close to zero indicate an electrode that follows the global trend. Furthermore, the value of ASI_{max} encodes information about the spatio-temporal stability of a given source.

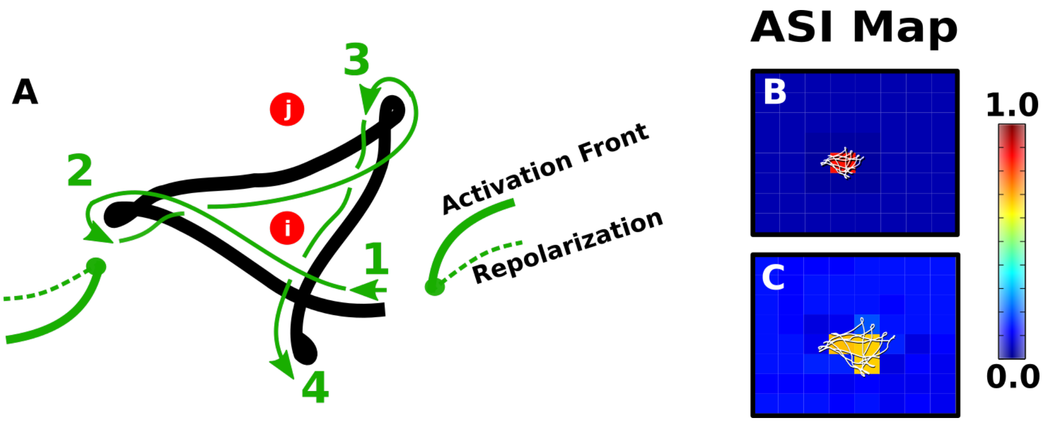

Figure 3.5: Complex tip trajectories result in asynchronous activation patterns near the region of meander. (A) Example of a complex tip trajectory (green arrow) and the resulting asynchronous activation at electrode sites located within (i) or outside (j) the trajectory. (B) and (C) demonstrate that only the electrodes that are within the tip trajectory will display an elevated ASI level. Here, the spiral tip trajectory is shown in white while the ASI value at each electrode of the 8x8 grid is shown using a color scale.

In Silico Data

Our computational models demonstrate that electrical recordings near spiral wave cores can become asynchronous with neighboring tissue due to spiral wave meander, or precession. Fig. 3.5 A shows a meandering spiral tip trajectory in black, along with electrodes placed within (i) and outside the spiral tip trajectory (j). As the spiral wave reaches point 1 and travels to point 2, both electrodes will activate as the depolarization front passes over them. As the rotor revolves around point 2 and travels to point 3, the activation will only pass over electrode i. From point 3 to point 4 the activation front will again pass over both electrodes, and the cycle restarts. Counting up the number of activations per cycle, we see that electrode positioned inside of the spiral tip trajectory will be activated 3 times. In contrast, the electrode located outside of the trajectory will only activate 2 times. Thus, this 3:2 activation pattern will cause asynchrony between electrodes inside the rotor trajectory and those outside, resulting in an elevated value of ASI.

Fig. 3.5 B shows the trajectory of a single simulated spiral wave (white line) that encompasses only a single electrode. This trajectory is plotted on top of the corresponding ASI map, computed using a 8x8 virtual electrode grid, in which the ASI value is color-coded with

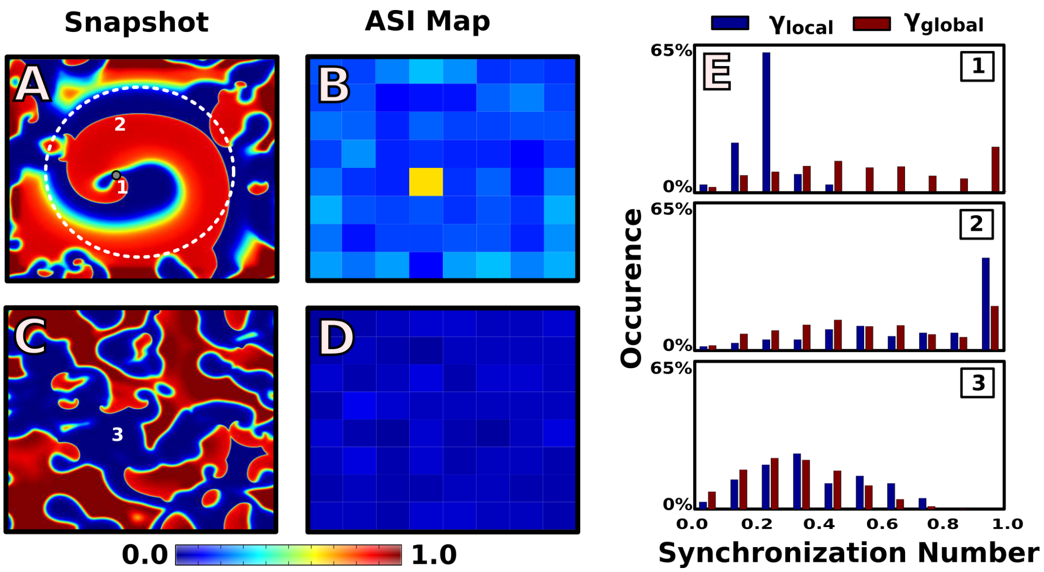

Figure 3.6: Simulations of VF dynamics and their corresponding ASI maps. (A) Snapshot of a simulation involving a stable spiral wave, inside the white dashed region, surrounded by disorganization. The colormap displays the transmembrane voltage, with warm/cold colors represent depolarized/ repolarized tissue (spiral tip: white dot). (B) Corresponding ASI map. (C) Snapshot of a simulation in which the entire domain promotes wavebreak (multi-wavelet reentry) and corresponding ASI map. Distributions of the synchronization numbers for the electrodes located at 1 and 2 (A) and 3 (C) as well as the global distributions for the simulation corresponding to (A) and (C) are also shown (weighting factors: $\lambda_1 = 0.97$,$\lambda_2 = 0.98$, and $\lambda_3 = 0.38$.

red/blue representing large/small values. As expected, the ASI value is elevated at the electrode inside the tip trajectory (0.86) but low everywhere else (0.035 ± 0.002). Changing the length scale of the spiral wave trajectory or, equivalently, decreasing the inter-electrode distance, results in more electrodes that lie within the trajectory and a larger region of elevated ASI. This is shown in Fig. 3.5 C where the tip trajectory, again shown in white, now crosses several electrodes, leading to a larger asynchronous domain.

The weighting factor λ in the definition of ASI, which computes the synchrony of surrounding tissue, increases its specificity for spiral wave cores by decreasing ASI values in areas of extended asynchrony due, instead, to spiral wave break-up. Fig. 3.6 A shows a snapshot of a localized spiral wave in the presence of spiral wave break-up. Here, the dashed line indicates the boundary of the stable domain and the location of the spiral tip is marked by a dot. We constructed local and global distributions of synchronization numbers using an 8x8 virtual electrode grid and computed the ASI map (Fig. 3.6 B) which shows that the electrode within the spiral tip trajectory has an elevated ASI value while electrodes outside that region exhibit low ASI values. This result can be understood by examining the distribution of synchronization numbers for electrode 1, $\gamma_{local,1}$ and located within the tip trajectory, as shown in the right panel, together with the global distribution of synchronization numbers, γ_{global}, also shown the right panel. While γ_{global} is almost uniform, $\gamma_{local,1}$ is sharply peaked at low synchronization numbers. Thus, these two distributions are markedly different, resulting in an elevated ASI value of 0.64. In contrast, the distribution of synchronization numbers for electrode 2, $\gamma_{local,2}$ located in the organized region is almost identical to the global distribution, leading to a small ASI value of 0.21. Electrodes located in the break-up region also record a small value of ASI (0.20 ± 0.05) because they are surrounded mostly by disorganization, leading to a small weighting factor λ.

ASI values are also low in the presence of multi-wavelet reentry. Fig. 3.6 C shows a snapshot of a simulation of multi-wavelet VF, during which VF is not driven by a source but self-perpetuates by a continuous break-up of spiral waves. Importantly, since the entire domain is driven by multi-wavelet reentry γ_{local} will be similar at each location and, thus, similar to γ_{global}. This is shown in the bottom right panel where we plot the distribution of synchronization numbers at location 3 as well as the global distribution of synchronization numbers. Since each electrode displays asynchronous behavior, the pre-factor λ is also small for each electrode, resulting in all electrodes recording low ASI values of 0.05 ± 0.02 (Fig 3.6 D).

Clinical Data

We can also compute ASI in clinical data. An example ASI map of sustained VF in a 56 year old patient, along with γ_{global} and γ_{local} for the electrode with the largest value of ASI and a corresponding isochronal map, are shown in Fig. 3.7. Raw electrograms at the site of elevated ASI, as well as a neighboring site of low ASI, are also provided for comparison. The electrodes

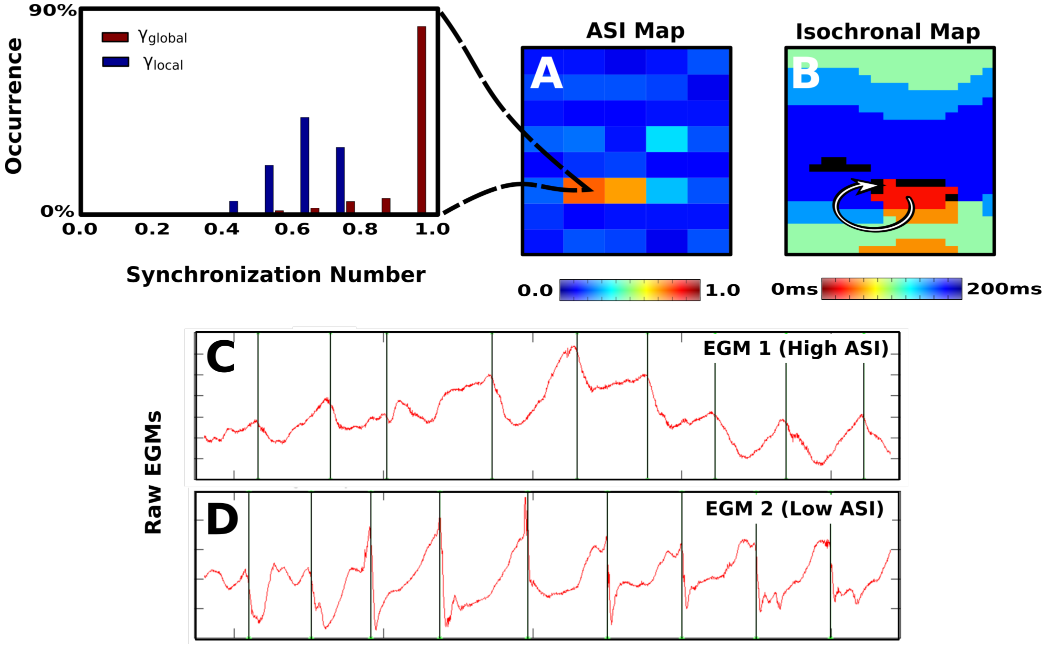

Figure 3.7: Snapshot of an ASI map during VF. (A) ASI map for the LV of a 56 year old patient, showing two regions of elevated ASI. The distributions of synchronization numbers for a single electrode in the elevated ASI region, along with the distribution of global synchronization numbers, are shown in the left panel. (B) Isochronal map, corresponding to the middle of the time interval used to compute the ASI map. (C) A sample electrogram trace at the site of elevated ASI, with activation times marked by the black lines. As expected the activation at this site is irregular, with varying cycle lengths. (D) A sample electrogram trace at a neighboring site with low ASI. The activation at this site is regular, with similar cycle lengths throughout.

with elevated ASI indicate a tissue region that is asynchronous compared to the global level of synchrony. The isochronal map reveals that this area of asynchrony corresponds to rotational activity as determined from activation times. Furthermore, this patient was ablated at a region coincident with this site of elevated ASI and VF was subsequently found to be non-inducible.

3.2.3 Computing Flow Fields

While phase synchrony affords a means of quantitative analysis of both clinical and simulated episodes of fibrillation, it cannot provide direct information related to the conduction and propagation of electrical activity across tissue. Spatio-temporal activation maps, for example, provide a visual means of exploring excitation patterns during fibrillation. These videos of electrical activity, however, yield qualitative rather than quantitative insights. It would be advantageous to determine a means of characterizing these excitation patterns in a more quantitative manner. One way of doing so is by determining the vector flow of conduction across episodes of fibrillation. This direct quantification of conduction dynamics would afford the possibility to study how arrhythmias perpetuate and evolve in time across many different patients and episodes. Furthermore, it could elucidate the relative commonality of rotors as compared to focal sources (centrifugal excitation patterns) in an objective manner.

Recordings of fibrillation from basket catheters are constrained to spatially sparse electrode grids, making the computation of the conduction propagation difficult. Moreover, standard measures of tracking propagation flow have typically been derived for image recordings in which data is graded in brightness and is available with high spatial resolution. As a result, flow is often very smooth, moving only a single grid point, or pixel, per timeframe [HS81], a qualification which is not met with many such bio-electric recordings. Instead, we present a general methodology to determine and compute a vector field, termed the wavefront flow field, describing the flow of conduction across a sparse electrode grid. We then apply our method to unique clinical data recorded during both tachycardia and fibrillation, and present results which show a strong correspondence to the underlying visual activation maps. We also present simple measures of rotor activity and focal activity derived from this vector field, and show that this automated procedure can identify regions of tissue with special clinical relevance.

Methodology

We begin by denoting the activation times for an electrode at coordinate (x,y) as the set t_{xy}. For both in-silico and clinical data these activation times are interpolated to a 29x29 grid before being used to compute the flow field. These activation times represent a simplified picture of the underlying conduction dynamics. It is assumed that some number of coherent activation fronts traverse cardiac tissue throughout these episodes, and therefore if we sufficiently coarse-grain our activation times we expect to see contiguous regions of activation. In particular, we set a time bin $\alpha = 10$ ms and compute our coarse-grained activation times as $\tau_{xy} = \lceil \frac{t_{xy}}{\alpha} \rceil$, with rescaled time referred to as t^*. The value of this time bin was chosen visually to optimize the smoothness of activation fronts, where a significantly smaller time bin leads to rough pixelated fronts and a significantly larger time bin leads to over-smoothing and the inability to distinguish individual fronts. These activation times are best described as an activation front defined by

$$\Phi(x,y,t^*) = \begin{cases} 1, & \text{if } t^* \in \tau_{xy}. \\ 0, & \text{otherwise.} \end{cases} \tag{3.5}$$

In words, the activation front is 1 if there is an activation at point (x,y) and at (rescaled) time t^*, and 0 otherwise. This activation front is then convolved with a Gaussian Kernel for spatial smoothing, and displayed as an activation map (see Figure 3.8).

To track the flow of the activation fronts in time, we utilize their spatial gradient $\vec{\nabla}_s$ as a marker in order to infer the direction of their subsequent motion. For a given time t^* we compute both $\vec{\nabla}_s\Phi(t^*)$ and $\vec{\nabla}_s\Phi(t^*+1)$ across the entire domain. We call the activation wavefront the set of points where $|\vec{\nabla}_s\Phi| > 0$ and define the instantaneous flow vector for each wavefront point i at time t^* as the normalized vector pointing from i to some point j obeying the optimization condition

$$\begin{aligned} \underset{j}{\text{minimize}} \quad & \text{mag}\Big\{[\vec{\nabla}_s\Phi(t^*)]_i - [\vec{\nabla}_s\Phi(t^*+1)]_j\Big\} \\ \text{subject to} \quad & |\vec{r}_{ij}| < R^*, \end{aligned} \tag{3.6}$$

where $\vec{r}_{ij}$ is the vector pointing from i to j and R^* is a set parameter defining the maximum possible translation per unit time.

In words, for each point i along the wavefront at a given time we search within a disk of radius R^* for the point j at the next time step which exhibits the most similar gradient vector. We assume that the point i on the wavefront has traveled towards this point j during this time step and

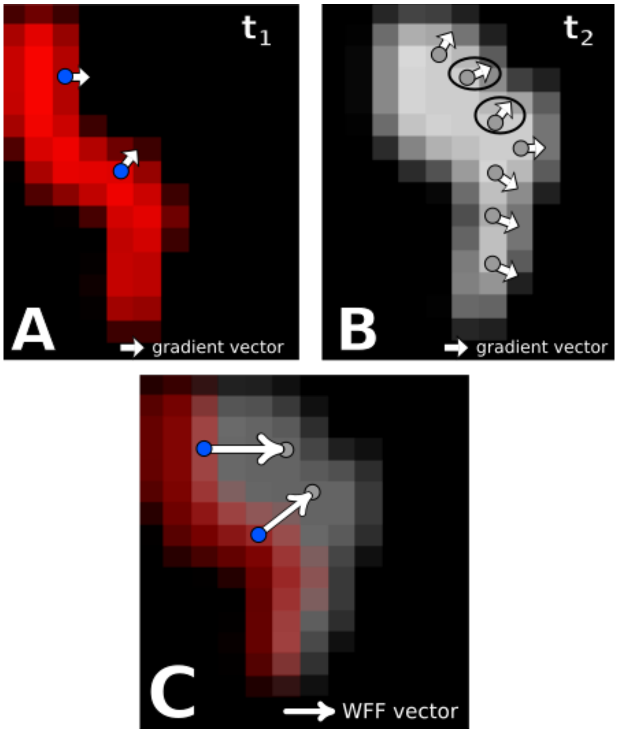

Figure 3.8: A schematic of the wavefront gradient matching procedure on clinical data. A. Snapshot at time t_1 of an activation front Φ in red. The white arrows represent the negative gradient vector at the blue points. B. Snapshot at a later time t_2 of the same activation front in gray. The instantaneous flow field at each blue point is determined by finding the point along this activation front which has the most similar negative gradient vector. These points are indicated by the circle. C. Resulting instantaneous flow vectors for the blue points in A, obtained by drawing the vector connecting these points to the circled points in B. Similar to above, the red activation front occurs at time t_1 and the gray front occurs at a later time t_2.

mark this motion with the instantaneous flow vector. This approach is outlined on two successive snapshots of a sample clinical dataset in Figure 3.8 A and B, with sample points i shown as blue dots. The activation front, smoothed by a Gaussian kernel, is shown for the two successive time points and the negative gradient vectors at the blue dots are shown as white arrows. These vectors are also displayed for the subsequent activation front in B and the ones that match best with the vectors at the blue dots in A are circled. The instantaneous flow vector is then obtained by drawing a vector connecting the blue dot to the location of the best match, as shown in C.

As a final step, we compute this instantaneous flow field across a window of time, sum together these vectors, and apply smoothing filters to get a net wavefront flow field $\vec{\Psi}(x,y)$

representing the consistent flow of conduction throughout this time period. To display a time-dependent flow field we compute this with a sliding window across our episode and interpolate through time.

We now wish to mark regions of increased rotational and focal activity from a given field $\vec{\Psi}(x,y)$ using measures of vorticity and divergence. Rather than using the standard vector curl and divergence, which compute infinitesimal rotation and focality, we want to define contiguous regions of rotational activity with a coherent spatial domain. Furthermore, rotors are expected to have some meander in space and therefore the core of rotation may cover multiple points in space. For the purposes of quantifying rotor and focal sources on our sparse grid, we define vorticity Ω and divergence Δ at a point i as

$$\Omega_i \equiv \frac{1}{\mathcal{N}_i} \sum_{j \neq i} \frac{\vec{\Psi}_j \cdot \hat{\omega}_{ij}}{|\vec{r}_{ij}|^2} \tag{3.7}$$

and

$$\Delta_i \equiv \frac{1}{\mathcal{N}_i} \sum_{j \neq i} \frac{\vec{\Psi}_j \cdot \hat{n}_{ij}}{|\vec{r}_{ij}|^2}, \tag{3.8}$$

where $\vec{r}_{ij}$ is the vector pointing from i to j, $\mathcal{N}_i$ is the normalizing factor $\sum_j \frac{|\vec{\Psi}_j|}{|\vec{r}_{ij}|^2}$, $\hat{\omega}_{ij}$ is the tangent vector to $\vec{r}_{ij}$, and $\hat{n}_{ij}$ is the normal vector to $\vec{r}_{ij}$. The factor $\frac{1}{|\vec{r}_{ij}|^2}$ acts as a weighting factor to assign decreasing importance to points further away from i and the numerator computes the alignment of the flow field and either the rotational circuit $\hat{\omega}$ or the divergent circuit $\hat{n}$.

In Silico Data

To verify that our methodology can pick up relevant activation patterns we first compute the flow field $\vec{\Psi}$ for in-silico data of a rotor and a focal source in Figure 3.9. The top four panels show the analysis on our simulated rotor, with the vorticity map clearly showing elevation at the center of the rotor core and streamlines of the flow field (green) tracking the propagation of the activation front Φ (grayscale). Similarly, the bottom four panels show the analysis on our simulated focal source, with the divergence map showing clear elevation at the site of the focal

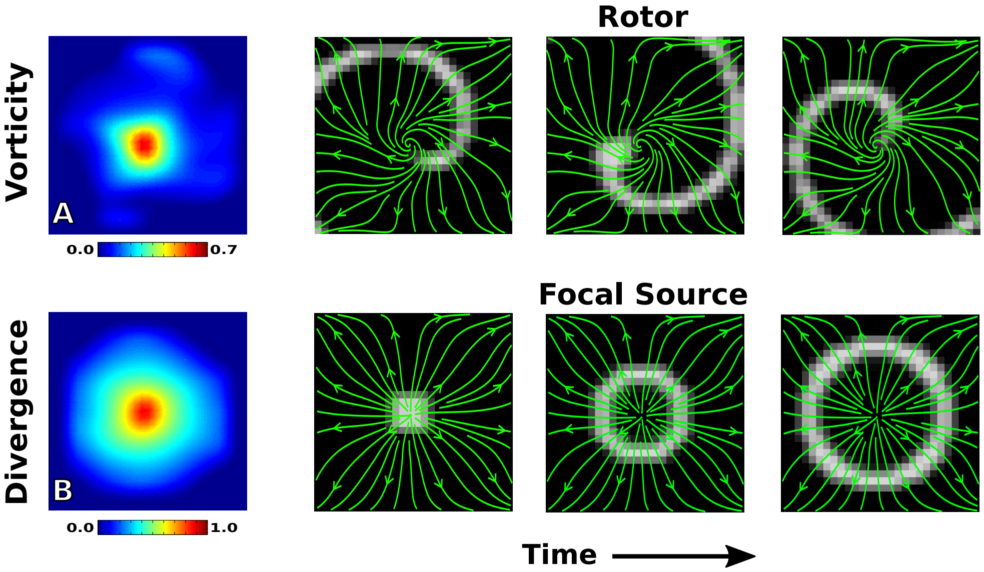

Figure 3.9: Flow field analysis of a simulated rotor (top) and focal source (bottom). A. The vorticity map records a significant elevation at the rotor core. B. The divergence map records a significant elevation at the site of the focal source. The right panels represent screenshots of the activation front at consecutive times throughout the recording, with streamlines of the flow field plotted in green.

source and the flow field tracking the motion.

It is worth noting here that while the divergence at the focal site reaches near unity, the vorticity at the rotor core is below unity. This can be attributed to the fact that rotors are expected to have some centrifugal component of motion as well as a rotational component. This is apparent when examining the simplified scenario of conduction on a homogeneous slab of tissue with the rotor's wavefront approximated as an archimedean spiral $r = \theta^{1/n}$, where n is a parameter controlling how tightly wound the spiral is. In this simple case, the conduction should flow along the normal to the wavefront given by $\hat{N} = \left[1 + \left(\frac{1}{n\theta}\right)^2\right]^{-1/2}\hat{r} + \frac{1}{n\theta}\left[1 + \left(\frac{1}{n\theta}\right)^2\right]^{-1/2}\hat{\theta}$. From this we find that the fraction of motion which is rotational is given by $\frac{1}{n\theta+1}$. The conduction flow, therefore, will become more centrifugal and less rotational as $\theta \to \infty$, whereas near the core we still have a significant rotational component as can be seen in the flow fields in Figure 3.9. The winding parameter n in human fibrillation is unknown, so we must conclude that while vorticity should be elevated at a rotor core it may not be expected to reach unity even in the ideal case of a single rotor controlling the entire domain.

Clinical Data

To show how our flow field analysis can elucidate the conduction dynamics underlying heart rhythm disorders we apply this methodology to clinical data recorded from patients experiencing two common arrhythmias. First, we examine an episode of tachycardia in the left atrium of a patient in Figure 3.10. The left panels shows the flow field Ψ plotted on top of a vorticity map for two different intervals during this episode, revealing a clear stable dual rotor activation pattern early transitioning to a single stable rotor. The panels on the right show consecutive snapshots of the activation front and streamlines of the flow field during each interval, with red dots marking the local maxima/minima of vorticity above an arbitrary threshold. Visual inspection of these frames validate the existence of this dual rotor pattern.

We also examine an episode of fibrillation occurring in the left ventricle of a 72 year old patient in Figure 3.11. This patient was enrolled for a ventricular ablation procedure, whereby fibrillation was induced, mapped, and subsequently ablated. Immediately after induction, defibrillators were charged and used to terminate fibrillation once charging was complete during which electrograms were recorded. Our flow field analysis revealed that this episode of fibrillation exhibited a focal source early after induction (Figure 3.11, top). A few seconds later, the conduction pattern changes to mostly planar activation (Figure 4, middle). Finally, near the end of this episode a dual rotor pattern appears and meanders to the top right corner (Figure 4, bottom). Similar to above, left panels show the flow field Ψ on top of either divergence or vorticity maps for each observed pattern. Right panels show snapshots of activation and flow fields, with the blue dots marking the local maxima of divergence and red dots marking the local maxima/minima of vorticity.

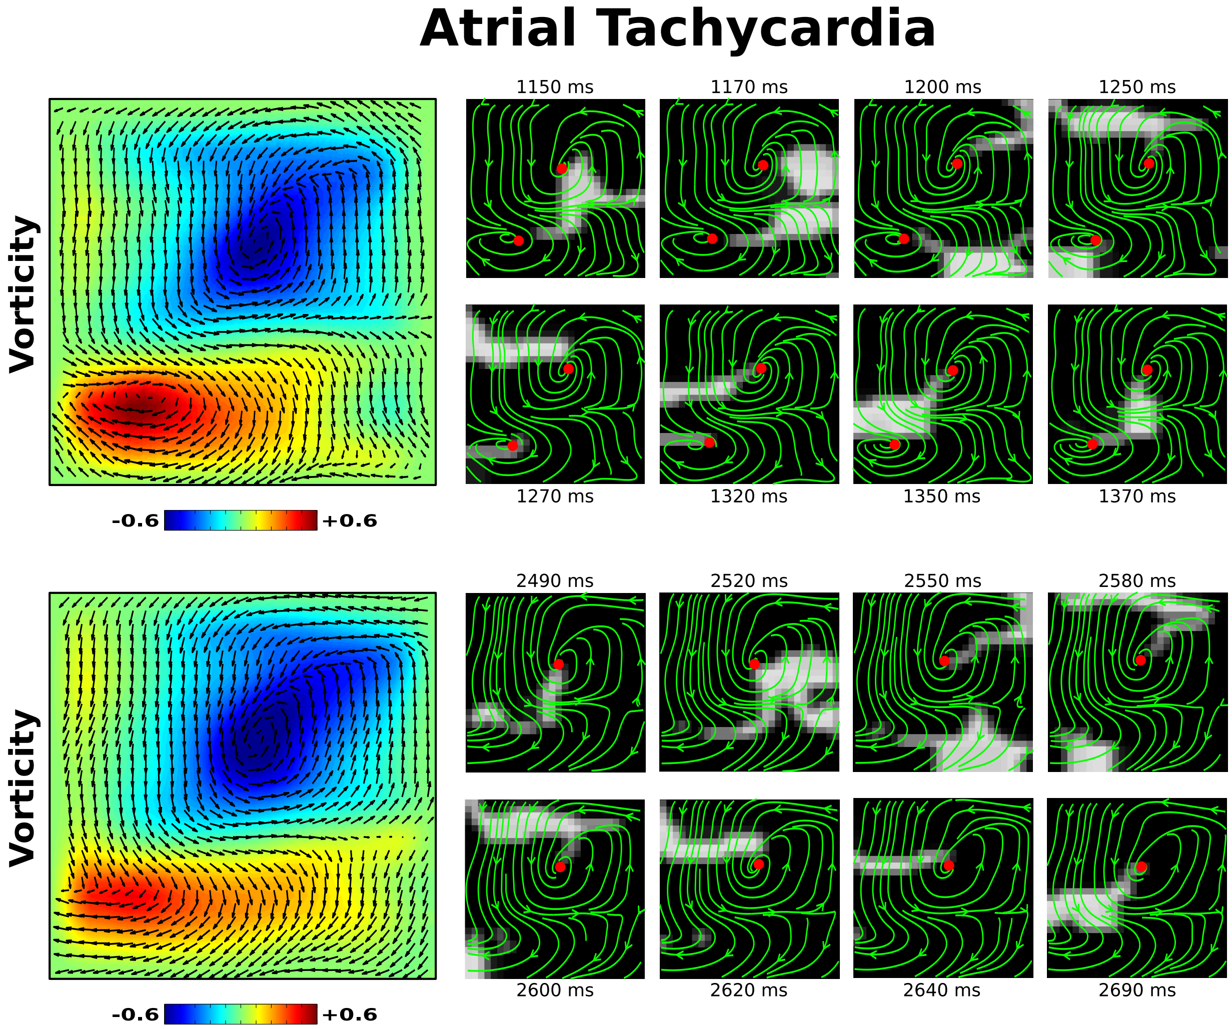

Figure 3.10: Flow field analysis during two intervals of a clinical episode of atrial tachycardia. A clear dual rotor pattern is seen in the vorticity map (top left) transitioning to a single rotor (bottom left). Corresponding activation maps (right), with streamlines of the flow field plotted in green and the local maxima/minima of vorticity above an arbitrary threshold marked as a red dot. The flow field Ψ is plotted on top of the vorticity map.

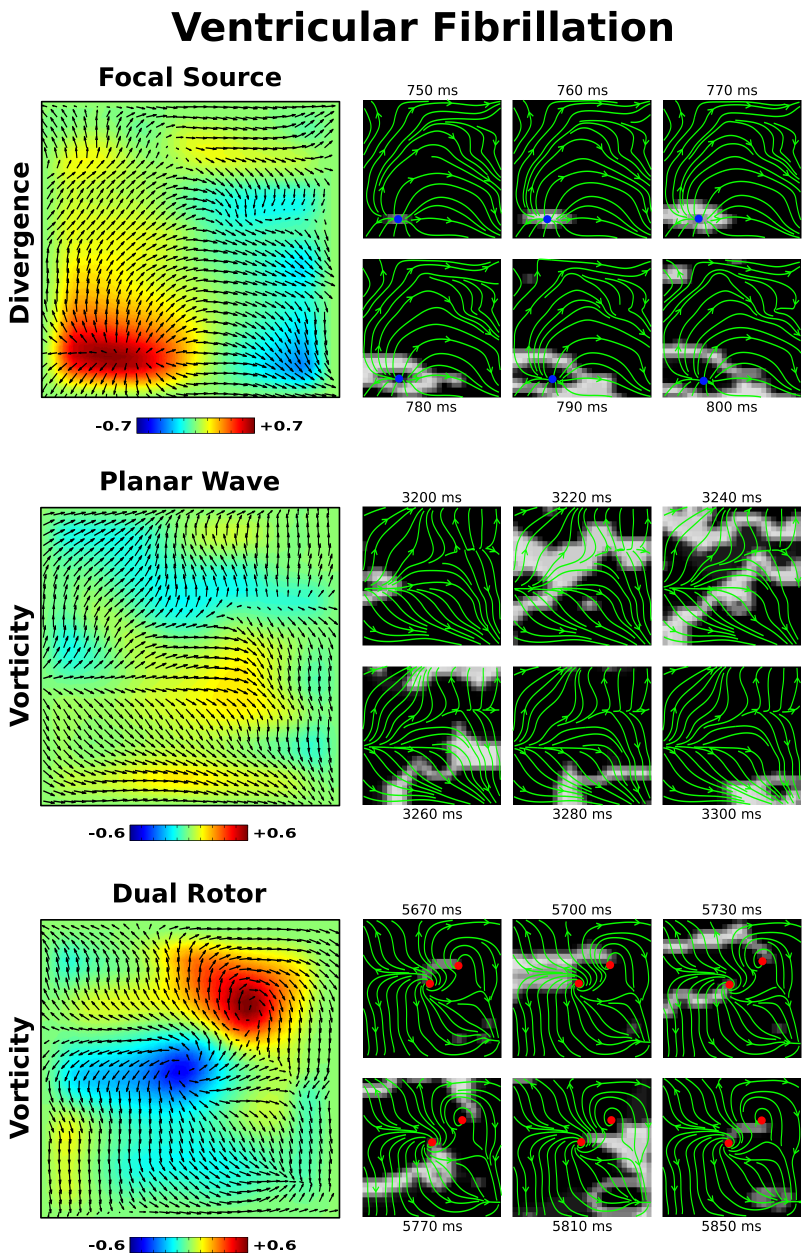

Figure 3.11: Flow field analysis during three intervals of a clinical episode of ventricular fibrillation. A focal source is seen early in the episode, with a clear region of elevated divergence shown in the divergence map (top left). Later in the episode, planar activation transitions to a dual rotor pattern, as can be seen in the vorticity map (bottom left). The flow field Ψ is plotted on top of all maps. The corresponding activation maps are also shown to the right, with streamlines of the flow field plotted in green, the local maxima of the divergence marked as a blue dot, and the local maxima/minima of the vorticity marked as a red dot.

3.2.4 Determining Phase from Electrograms

The preceding sections have described various analyses of clinical arrhythmias using discrete activation time data. It is not always trivial to mark such activation times, necessitating involved algorithms which are often proprietary. Moreover, various signal artifacts such as far-field effects, recording noise, and signal contamination from the ventricles can lead to potentially ambiguous deflections. Methods of analysis focused on analyzing the raw electrograms themselves can be desirable. While signal artifacts and variation of signal amplitude make the interpretation of raw voltage maps untenable, there exist simple ways of inferring phase from noisy and complex signals. Phase maps, instead of voltage maps, can then be useful for exploring excitation patterns in clinical arrhythmias.

Hilbert Transform and Phase Singularities

To create meaningful phase maps from electrogram data we must compute a time series of instantaneous phase $\phi(t)$ for each electrogram which should range from 0 to 2π, with each activation event resetting the phase back to zero. Assuming our electrogram signal $x(t)$ represents the periodic activation of tissue, we can compute this phase by first applying the Hilbert transform

$$H(x)(t) = \frac{1}{\pi} \int_{-\infty}^{\infty} \frac{x(\tau) d\tau}{t - \tau}. \tag{3.9}$$

This transform gives a real-valued signal which is effectively phase-shifted by $\pi/2$. This can be used to construct the so-called analytic signal, which is a complex-valued function with the real part equal to the input signal and the complex part equal to the transformed signal. The phase $\phi(t)$ is then computed as the angle of this analytic signal.

Once phase has been computed for each electrogram, spatio-temporal phase maps can be visualized and used to observe excitation patterns. These spatial maps of phase will be primarily smooth, except at points of phase singularity which correspond to the core of rotational excitation

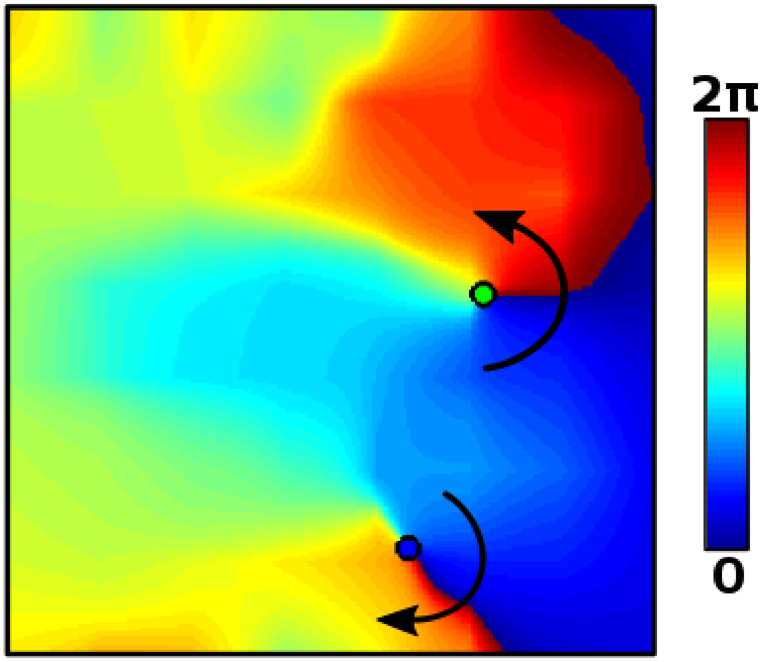

Figure 3.12: Phase map of a clinical episode of atrial fibrillation reveals a dual rotor pattern. Phase singularities are shown as dots, with blue representing clockwise rotation and green representing counter-clockwise rotation.

patterns [BLA$^+$01]. A useful means of identifying such phase singularities is by computing the topological charge as

$$n_t = \oint \vec{\nabla}\phi \cdot d\vec{l}, \tag{3.10}$$

with the contour taken around various points across the phase map. This topological charge will either be equal to zero, if it encloses no singularities, or will be $\pm 2\pi$ if encloses a phase singularity, with the sign indicating the direction of rotation. These phase singularities can then be used to identify rotor cores, though they are particularly susceptible to false positives [KZvH$^+$17].

Sinusoidal Recomposition

While using the Hilbert transform to compute phase directly from noisy electrogram data is in principle viable, it is particularly sensitive to the morphological complexities and signal artifacts often present in such data. It is therefore recommended to first simplify the electrogram signal through use of filtering and signal reconstruction before applying the Hilbert transform to compute phase. One recent methodology proposed to simplify electrogram signals uses a principle of sinusoidal reconstruction [KZM+15]. This approach is based on the idea that activation events should generally occur when the signal exhibits a large negative slope. As such, the approach consists of reconstructing the signal as a sum of single-period sinusoidal wavelets, with each wavelet having an amplitude equal to the negative slope of the signal at that point (positive slopes are set to zero). The period of these wavelets is determined through Fourier analysis to correspond with the dominant frequency of the signal. A snapshot of a phase map computed during a clinical episode of fibrillation using this method is shown in Figure 3.12, which reveals a clear dual rotor pattern. The analytic signal is interpolated to a fine grid and phase singularities are computed from the topological charge and marked as dots.

3.3 Treatment Options

To this point we have focused on analyzing episodes of arrhythmia with quantitative measures to elucidate underlying excitation patterns. Understanding these patterns is important because they point to the mechanisms perpetuating dangerous arrhythmias such as fibrillation, potentially allowing for targeted therapeutic approaches. As context, we briefly describe the current standard of care for AF and discuss therapeutic options using mapping to target AF drivers.

The first approach to treating heart rhythm disorders is often antiarrhythmic drugs. The success rate of such drugs, however, remains poor [SCC+07]. If arrhythmia persists after attempted pharmacological intervention, a common next approach is catheter ablation. This semi-invasive procedure consists of a catheter inserted into the patient's heart which uses radiofrequency energy to scar problematic regions of tissue. The current standard ablation routine is known as pulmonary vein isolation (PVI), where tissue surrounding the pulmonary veins is scarred in attempt to eliminate AF triggers that are believed to originate there [HJS^{+}98]. This routine has shown moderate, though sub-optimal, success rates [VJB^{+}15].

A recent treatment approach attempts to improve success rates by targeting each ablation procedure to patient-specific drivers of AF [NKR12], in the form of rotors or focal sources. These drivers are identified from spatio-temporal maps of induced AF, with ablation being targeted to the core of either the rotor or focal source in attempt to eliminate regions of tissue harboring AF drivers. Such procedures, often in combination with PVI, have reported high success rates [NKS^{+}12, TDM^{+}15, SKR^{+}16, MKD^{+}17, SKR^{+}17]. While many early reports were positive, some recent studies have also reported poor success rates [BST^{+}16, SSB^{+}17, BGIC^{+}17]. Randomized clinical trials are currently being conducted to clarify these discrepant results.

Chapter 3, in part, is a reprint of the material as it appears in (i) the American Journal of Physiology-Heart and Circulatory Physiology 2015, 309(12), H2118-H2126. Vidmar, David; Narayan, Sanjiv; Rappel, Wouter-Jan. (ii) Physical Review E 2016, 94(5), 050401. Vidmar, David; Narayan, Sanjiv; Krummen, David; Rappel, Wouter-Jan. (iii) JACC: Clinical Electrophysiology 2017, 3(12), 1437-1446. Vidmar, David; Krummen, David; Hayase, Justin; Narayan, Sanjiv; Ho, Gordon; Rappel, Wouter-Jan. The dissertation author was the primary investigator and author of these papers.

Chapter 4

Stochastic Dynamics of Fibrillation

While stable spiral waves have been targeted for localized ablation treatments, AF termination following targeted ablation is not always instantaneous and can take minutes to days [SKR$^+$16]. In addition, many patients experience intermittent AF episodes of variable duration, illustrating that AF can terminate spontaneously. This points towards a turbulent state of spiral defect chaos (SDC) perpetuating AF after the removal of stable sources, similar to multi-wavelet reentry described earlier.

SDC will persist until the last spiral wave is terminated, with its duration representing a stochastic event. Critical in AF management and therapy is thus the mean episode duration τ which is a statistical measure of the average time of reversal to normal sinus rhythm. The ability to infer τ from some observables of our system, particularly in the presence of different lesion sets, can be an important step towards more efficient AF therapies [RDJ$^+$07, MVB$^+$12]. Unfortunately, determining τ through direct simulations of cardiac models is challenging because a statistically significant determination of this stochastic quantity requires the time-consuming task of simulating a multitude of episodes. This becomes even more problematic for large geometry sizes since τ increases sharply as a function of the system size [Qu06].

In this section we develop a statistical description of the number of spiral tips as a function of time, treating the termination of fibrillation as a stochastic process. Specifically, we quantify the birth and death rates of spiral waves using a limited set of direct numerical simulations. Once these rates are determined, we describe the dynamics of fibrillation with a master equation, which we solve using a number of techniques, exact and approximate, from the fields of theoretical ecology and population dynamics.

4.1 Spiral Defect Chaos and the Critical Mass Hypothesis

The turbulent state of SDC is described by the continual fragmentation and collision of spiral tips, with fragmentation increasing the number of tips and collisions decreasing the number of tips. The total number of spiral tips n evolves stochastically over time due to these competing mechanisms. In the short term, this system will reach a metastable state with fluctuations around a mean number of spiral tips $\bar{n}$. Over large time scales, however, the population of spiral tips will inevitably become extinct, as an unlikely succession of annihilations will eventually cause termination.

This mechanism not only explains the diversity in observed episode durations of fibrillation, it is consistent with the so-called critical mass hypothesis [Gar14]. Based upon observations of fibrillation in animals, this hypothesis states that fibrillation can only be sustained in hearts larger than some critical size. This theory formed the rationale behind the surgical Cox-Maze procedure, during which the atria is subdivided into small compartments by a set of incisions [CSD^{+}91] and which has shown good success rates in eliminating AF [GSB^{+}05]. From a stochastic point of view, this hypothesis makes sense because with less tissue for the wavelets to traverse, collisions are increasingly likely which will cause the mean episode duration to become negligible.

4.2 Population Dynamics

The question of determining the mean time to the extinction of metastable populations has been previously studied in the context of theoretical ecology [OM10]. Similar to these ecological models of the extinction of species, termination of fibrillation can be viewed as the extinction of the population of spiral tips. As such, SDC is effectively described by the state variable n, whose evolution to zero corresponds to termination. We provide a brief overview of stochastic population dynamics here, focusing on general discrete birth-death processes (for a more detailed background see [Ass10]).

4.2.1 Discreteness Matters

It is tempting to view the dynamics of large stochastic populations solely in the mean-field limit, where we examine average densities in a continuum and can ignore fluctuations. The intuition, however, that these continuum systems will provide meaningful insights into the discrete systems of which we are concerned is not always true. It turns out that the discreteness of these populations matters when discussing dynamics, and can lead to profound departures from results in the mean field limit.

To illustrate this, we follow [KRBN10] and examine the simple system of an initial particle which generates an offspring particle at constant rate λ. Each offspring particle also generates its own offspring at the same rate. In the continuum limit, we would be concerned with the mean number of particles as a function of time which evolves as $\frac{d\bar{n}}{dt} = \lambda\bar{n}$. Solving this we get exponential growth of the mean number of particles, $\bar{n} = e^{\lambda t}$. In this formulation we have, of course, assumed fluctuations are small and can be ignored.

In order to include the discreteness and stochasticity of this system, we must examine the so-called master equation in full. The master equation describes the time evolution of the probability distribution of our state variable, denoted $P_n(t)$, in terms of transition rates between states. For this simple system, the master equation is given by

$$\frac{dP_n}{dt} = \lambda\Big[(n-1)\,P_{n-1} - n\,P_n\Big], \tag{4.1}$$

where the first term on the right-hand side represents the probability of transitioning into an n state from an $n-1$ state and the second term represents the probability of transitioning out of an n state to an $n+1$ state.

Equation 4.1 can be solved by employing an ansatz solution of the form $P_n = (1-a)a^{n-1}$, with initial condition $a(0) = 0$ corresponding to our initial condition of a single particle. Plugging this ansatz into our master equation and solving we get $P_n(t) = e^{-\lambda t}\left[1 - e^{-\lambda t}\right]^{n-1}$. Solving for the mean particle number we do indeed get $\bar{n} = e^{\lambda t}$, consistent with the simple mean-field approach. When examining the standard deviation of n, however, this concordance breaks down. In particular, now that we take into account fluctuations arising from the discreteness of particles we see that the deviation is of leading order $\sigma \simeq \bar{n}$, implying that the fluctuations in this system are in fact quite large and increase exponentially with time. Ignoring these large fluctuations is likely to lead to problems.

Many systems of stochastic populations involve not only birth events, but also death events through either pure chance or interactions such as collisions. Such systems present another problem in the mean-field limit described above because these competing birth and death processes can potentially lead to a state where the birth rate is equal to the death rate. In the continuum limit, this is a state of equilibrium and the mean number of particles will remain in this equilibrium state forever. In reality, however, the population has a nonzero chance of becoming extinct through a succession of deaths, which is described as an absorbing state. Once entering this absorbing state the population remains extinct forever, a feature not described at all in the mean field limit. Moreover in many systems this extinction is a certainty, meaning that as time goes to infinity the mean number of tips eventually decays to zero, in contradiction with the continuum result. Therefore, to treat stochastic populations with the potential for extinction one must start from the master equation in full.

For a general class of birth-death processes, where transition rates W_r change the state variable by r, we have a master equation of the form

$$\frac{dP_n}{dt} = \sum_r \Big[W_r(n-r)P_{n-r} - W_r(n)P_n \Big]. \tag{4.2}$$

Extinction corresponds to the condition $W_r(0) = 0$ and to demand the population remain positive we require $W_r(n < -r) = 0$. We restrict ourselves to exploring systems which exhibit a long lived metastable state and which have unit probability to eventually become extinct.

4.2.2 Single-Step Birth Death Processes

While a closed-form solution for Equation 4.2 does not exist in general, in the case of a single-step process certain exact solutions can be achieved. In such systems experiencing a metastable state, the birth rate exceeds the death rate at small n, whereas at large n the death rate exceeds the birth rate. In the absence of an absorbing state, this system would converge to a nontrivial stationary distribution described by $\frac{dP_s(n)}{dt} = 0$. Due to the absorbing state at $n = 0$, however, this stationary distribution will just be a delta function around zero. This system will, however, converge to a so-called quasi-stationary distribution $P_{qs}(n)$ in the interim, defined for values of $n > 0$, which will decay to zero through the $W_{-1}(1)$ rate. This quasi-stationary distribution can be determined from Equation 4.2 by setting $\frac{dP_{qs}(n)}{dt} = 0$ and ignoring $W_{-1}(1)$, resulting in

$$P_{qs}(n) = P_{qs}(1) \prod_{z=2}^{n} \frac{W_{+1}(z-1)}{W_{-1}(z)}, \tag{4.3}$$

where $P_{qs}(1)$ is found from the normalization condition $\sum_{n=1}^{\infty} P_{qs}(n) = 1$.

Another quantity of interest for this system is the mean time to extinction given n starting particles, T_n, which for our single-step process can also be determined exactly. We begin by breaking the mean time to extinction up into two parts as $T_n = T_{0\to event} + T_{event\to extinction}$, corresponding to the mean time to the first transition event and the mean time from the first transition to extinction. The mean time to the first transition will be given by $\left[W_{-1}(n) + W_{+1}(n)\right]^{-1}$, whereas the time from that transition to extinction depends on which transition event occurred. The mean time from transition to extinction, then, will be given by $p(W_{+1})T_{n+1} + p(W_{-1})T_{n-1}$ with each probability determined as $p(W_r) = W_r / \sum_r W_r$. After some algebra this yields

$$W_{+1}(n)\left[T_{n+1} - T_n\right] + W_{-1}(n)\left[T_{n-1} - T_n\right] = -1, \tag{4.4}$$

with boundary condition $T_0 = 0$. By defining $\phi(x) = \prod_{z=1}^{x} \frac{W_{-1}(z)}{W_{+1}(z)}$, this equation can be shown to give the following general solution [Gar09]

$$T_n = \sum_{x=1}^{n} \phi(x-1) \sum_{y=x}^{\infty} \frac{1}{W_{+1}(y)\phi(y)} \tag{4.5}$$

with $\phi(0) = 1$. While this solution will, in general, depend on the initial number of particles n, systems with a long-lived metastable state can be well-approximated as being independent of initial conditions due to a rapid saturation of T_n at $n > 1$ [DSS05].

4.2.3 Multi-Step Birth Death Processes

An exact solution of Equation 4.2 for multi-step processes is generally not possible because detailed balance is not satisfied. We can make progress, however, by accounting for the fluctuations present in these systems. For small fluctuations, this is traditionally performed via the van Kampen system size expansion about the mean field solution, which yields a Fokker-Planck equation. To analyze the extinction dynamics of a stochastic population, however, this small fluctuation approximation turns out to be inadequate because it does not account for fluctuations large enough to bring about such extinction [KS07]. In order to account for the large fluctuations inherent to stochastic populations, we can instead apply a WKB approximation to our master equation so long as our system size A is sufficiently large [KMK73]. We further assume that the transition rates for our population can be rewritten in van Kampen canonical form $W_r(n) = Aw_r(q)$, with density $q = n/A$.

WKB Approach

The WKB approach begins by assuming the shape function of the quasi-stationary distribution has the form $P_{qs}(q) \propto e^{-AS(q)}$, where $S(q)$ is the action. Plugging this into the master equation allows us to write, to leading order, the following effective Hamiltonian for our system:

$$H(q,p) = \sum_r w_r(q)(e^{rp} - 1), \tag{4.6}$$

with $p = S'(q)$ representing a momentum term associated with fluctuations. This Hamiltonian exhibits a single nonzero activation trajectory $p_a(q)$ satisfying $H(q, p_a) = 0$, which is responsible for the fast WKB mode. The fast-mode solution dominates the quasi-stationary distribution, which is derived in Assaf and Meerson [AM10]. After some simplification, this results in

$$P_{qs}(q) = \sqrt{\frac{|p_a'(q)|}{2\pi A}} \exp\left[\int_q^{q_0} d\bar{q}\left(A\, p_a(\bar{q}) + \frac{H_{qq}(\bar{q}, p_a)}{2H_q(\bar{q}, p_a)}\right)\right], \tag{4.7}$$

where q_0 corresponds to the attracting fixed point of the deterministic ($p = 0$) rate equation given by $\dot{q} = H_p(q_0, 0)$.

This quasi-stationary distribution can be related to the mean time to extinction τ as $\tau^{-1} = \sum_{r<0} W_{-r}(-r) P_{qs}(-r)$. Because the WKB result for P_{qs} breaks down at $O(1)$, however, we need to find the exact form of the quasi-stationary distribution at small n directly from the master equation. This can be found, up to a constant, by solving the following recursion relation

$$\sum_r \left[W_r(n-r)\,P_{qs}(n-r) - W_r(n)\,P_{qs}(n)\right] = 0. \tag{4.8}$$

By matching the asymptote of the small-n solution derived from the recursion relation above with the appropriate asymptote of the WKB solution we can then fully determine the mean time to extinction.

Random Walk Approach

The WKB approximation utilized above is a powerful tool to understand extinction dynamics of a given stochastic system. If a functional form of the transition rates W_r is not simple, however, this analysis becomes cumbersome. If we only want to determine the mean time to extinction τ, a simple approach is to represent the population dynamics as a random walk on a directed network. The nodes of this network represent the states n of our system, corresponding to the number of particles, and each edge is weighted by the transition probability between states. Note that this includes a probability that n transitions to itself as $W_0(n) = 1 - \Sigma_{r\neq 0} W_r(n)$.

The adjacency matrix for this network is given by $\omega_{ij} = W_{n_j - n_i}(n_i)$. In order to allow for computation on this matrix it should be finite, and therefore we must set a (large) cutoff N^* as a maximum allowed population. To ensure that probabilities sum to unity, any rates for $W_{n_j > n^*}$ must be set to zero. We can now represent the dynamics of our system as a random walk on this network, starting from some initial node n_0. The probability of the walker to traverse a given edge is given by its weight and extinction occurs when the walker reaches the absorbing node $n = 0$. Determining the mean time to extinction now corresponds to computing the mean first passage time of a random walk from n_0 to this absorbing node.

To compute this mean first passage time, we follow the derivation outlined in [New10] and denote the probability that a random walker is at the absorbing node at any given time as

$P_0(t)$. This also amounts to the probability that the first passage time occurred at or before time t, allowing us to write the probability that the first passage time is at t exactly as $P_0(t) - P_0(t-1)$. The mean time to extinction is then given as

$$\tau = \sum_{t=0}^{\infty} t[P_0(t) - P_0(t-1)]. \tag{4.9}$$

In order to solve this equation, we must find a relation for $P_0(t)$. We begin by noting that once a walker is at the absorbing node, it has no chance of leaving. This implies that $\omega_{0j} = 0$ for all nodes j. Further, the probability of being at any node a time t is given by

$$P_j(t) = \sum_i \omega_{ij} P_i(t-1) = \sum_{i \neq 0} \omega_{ij} P_i(t-1). \tag{4.10}$$

If we note that for all nodes $j \neq 0$ we also have no terms with ω_{i0}, we can remove the row and column corresponding to node 0 entirely. This allows us to write this as a matrix equation

$$\mathbf{P}_*(\mathbf{t}) = \mathbf{P}_*(\mathbf{t-1})\omega_*, \tag{4.11}$$

where $\mathbf{P}_*$ is a row vector of p_j with the 0^{th} element removed and ω_* is the weighted adjacency matrix with its 0^{th} column and row removed.

We can now iterate this result back in time, and we get the following relation

$$\mathbf{P}_*(\mathbf{t}) = \mathbf{P}_*(\mathbf{0})\omega_*{}^{\mathbf{t}}. \tag{4.12}$$

The only probability that this equation does not give us is that of $p_0(t)$, which can therefore infer by imposing the constraint that the walker must be at some node at any given time $\Sigma_i p_i(t) = 1$. This gives us

$$P_0(t) = 1 - \sum_{i \neq 0} P_i(t) = 1 - \mathbf{P}_*(\mathbf{t})\mathbf{e}, \tag{4.13}$$

where **e** is a column vector of ones.

Plugging the above equation into Equation 4.13 and simplifying, we get

$$\tau = \sum_{t=0}^{\infty} t[\mathbf{P}_*(\mathbf{t-1}) - \mathbf{P}_*(\mathbf{t})]\mathbf{e} = \mathbf{P}_*(\mathbf{0})[\mathbf{I} - \omega_*]^{-\mathbf{1}}\mathbf{e}, \tag{4.14}$$

where we have used the fact that

$$\sum_{t=0}^{\infty} t(\mathbf{M}^{\mathbf{t-1}} - \mathbf{M}^{\mathbf{t}}) = [\mathbf{I} - \mathbf{M}]^{-\mathbf{1}}. \tag{4.15}$$

The row vector $\mathbf{P}_*(\mathbf{0})$ here only acts to pick out a mean time to extinction for a given initial state from our full vector τ. This result can therefore be more generally represented as

$$\tau_\alpha = [I - \omega_*]^{-1}\mathbf{e}, \tag{4.16}$$

where τ_α represents a column vector signifying the mean time to extinction from all possible initial conditions.

4.3 Stochastic Model of Termination

We now apply these general techniques from population dynamics to analyze the state of turbulent SDC hypothesized to play an important role in cardiac fibrillation. In particular, we are interested in the mean episode duration τ of SDC, which is analogous to the mean time to extinction discussed above. We begin by treating the number of spiral tips n as a stochastic quantity and casting its birth-death process into a master equation. This equation describes the probability $P(n,t)$ of having n spiral tips at time t and is given by the generic master equation for birth-death processes Eq. 4.2. Here, W_r denote the transition rates for the number of spiral tips to change by r tips and can be computed directly from simulations. Since tips are created and annihilated either as pairs or as singlets, we only need to consider $r = \pm 1, \pm 2$. As a boundary condition we take $n = 0$ to be absorbing such that $W_r(0) = 0$ and $W_{-2}(1) = 0$. Once the rates are known, we can construct a transition matrix which can be used to compute τ at minimal computational cost [New10].

For large n, the death rate will exceed the birth rate since tips will have a high probability of colliding. As a result, the number of spiral tips does not grow to very large numbers. If for small n the birth rate is larger than the death rate, then a long-lived metastable state exists with a corresponding quasi-stationary distribution $P_{qs}(n)$ and mean number of tips $\bar{n}$. In this state, the number of tips fluctuates around this average value for prolonged periods of time and the mean episode duration can be computed using $1/\tau = \sum_{r<0} W_r(-r) P_{qs}(-r)$. Termination only occurs during rare escape events, corresponding to a large fluctuation away from the mean number of tips. As a consequence, standard equilibrium statistical physics approaches based on small fluctuations do not apply [AM10, DSS05]. Instead, techniques from non-equilibrium statistical physics must be invoked to determine statistical quantities corresponding to extinction, including τ.

4.3.1 Spiral Defect Chaos Simulations

To illustrate our stochastic approach to quantifying termination dynamics, we carry out simulations of SDC using electrophysiological models. The simulations use the standard reaction-diffusion equation:

$$\partial_t V = D\nabla^2 V - \Sigma I_{ion}/C_m \tag{4.17}$$

where V is the transmembrane potential, C_m (μF cm^{-2}) is the membrane capacitance and $D\nabla^2$ expresses the inter-cellular coupling via gap junctions. We perform these simulations in square two-dimensional computational domains although our approach can be equally well applied in more complex geometries. As boundary conditions, we consider both non-conducting and periodic boundaries, and we vary the area of the computational domain, which is equivalent to

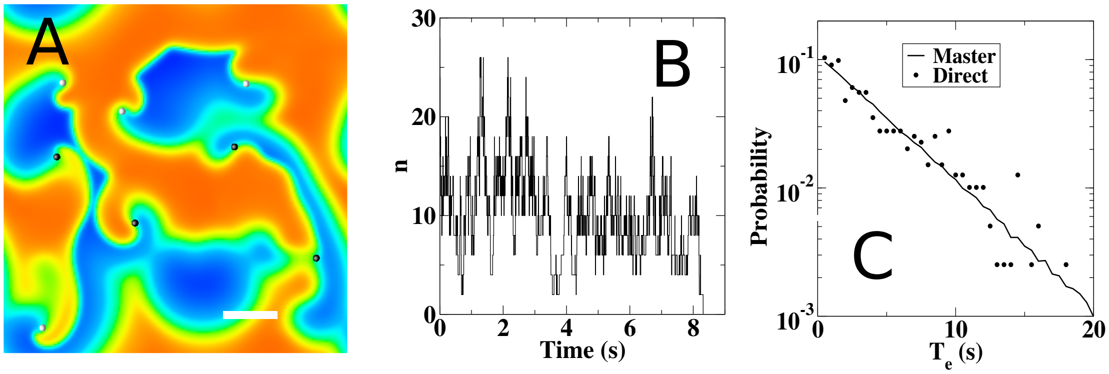

Figure 4.1: Direct numerical simulations provide statistics of spiral tip dynamics. A. Snapshot of a simulation of the LR model in a 7.5×7.5 cm computational domain with period boundaries. The voltage is represented using a color code with red (blue) corresponding to depolarized (repolarized) tissue. The location of the tips of counter- and clockwise rotation spiral waves are shown in black and white, respectively. (scale bar: 1cm) B. Typical time trace of the number of spiral tip pairs. For this particular simulation, spiral tips spontaneously extinguished after 8.3 s. C. Distribution of termination times for the direct simulations (symbols, computed using 400 termination events) and the master equation (solid line, computed using 10000 termination events).

varying the diffusion constant D while keeping the area constant. The membrane currents in the electrophysiological model are denoted by I_{ion} which are governed by nonlinear evolution equations coupled to V. For our purposes, the precise form of I_{ion} is not important and we present results using the detailed Luo-Rudy (LR) model [LR91].

Starting with a random initial condition that contains multiple spiral waves, we solve Eq. 4.17 and keep track of the number of spiral tips using a standard algorithm (Fig. 4.1A) [FK98]. The number of tips fluctuates and the simulation ends after time T_e when the number of spiral tips reaches 0 (Fig. 4.1B). We can compute the distribution of these termination times by repeating the simulations many times, starting with different and independent initial conditions. Our simulations reveal that this distribution is exponentially distributed, indicating that spiral wave termination can be well described as a Poisson process (Fig. 4.1C).

Next, we compute the birth and death rates as a function of the number of tips n using different domain sizes by quantifying the number of transitions per time interval (Fig. 4.2A-D). As a consistency check, we can use these rates to compute the distribution of termination times using the Gillespie algorithm [Gil77]. As expected, this distribution is exponentially distributed and agrees well with the one computed using direction simulations (Fig. 4.1C). Examining the computed rates, we see that W_{-1} depends linearly on the number of spiral tips for all domain sizes (Fig. 4.2A). The remaining rates, however, show a more complex dependence on the number of tips, indicating the existence of non-trivial long-range interactions between spiral tips (Fig. 4.2B-D). As a result, the rate curves are not easily fit by simple rational functions. Moreover, only rates computed based on a minimum number of transitions (here chosen to be 100) are considered. As a consequence, rates for values of n recording too few transition events are not computed. This is particularly relevant for large areas where small n rates often can not be included because they do not record enough events. Therefore, we employ a smoothing spline fit to the data to determine rates corresponding to transition events with less than the minimum number.

Importantly, we find that at large A all rates collapse onto a single curve when plotted as a function of the density $q = n/A$. Specifically, the $W_{\pm 1}$ rates scale with the perimeter L as $W_{\pm 1}(n) \sim L$ (Fig. 4.2E) while the $W_{\pm 2}$ rates are found to scale with the area as $W_{\pm 2}(n) \sim A$ (Fig. 4.2F). Note that the observed linear scaling of W_{-1} with L implies that the death rate is proportional with length of the non-conducting boundary and that creating ablation lesions will increase this rate.

4.3.2 First Passage Statistics

Once the transition rates are determined, it is straightforward to compute the quasi-stationary distribution $P_{qs}(n)$ using the transition matrix at minimal computational cost (Fig. 4.3A and B) [New10]. For small domains, this can be carried out using the rates obtained in the simulations while for larger domains, we can use the interpolated rates. As the domain size increases, the distribution shifts to larger values of n, and becomes more symmetric around the

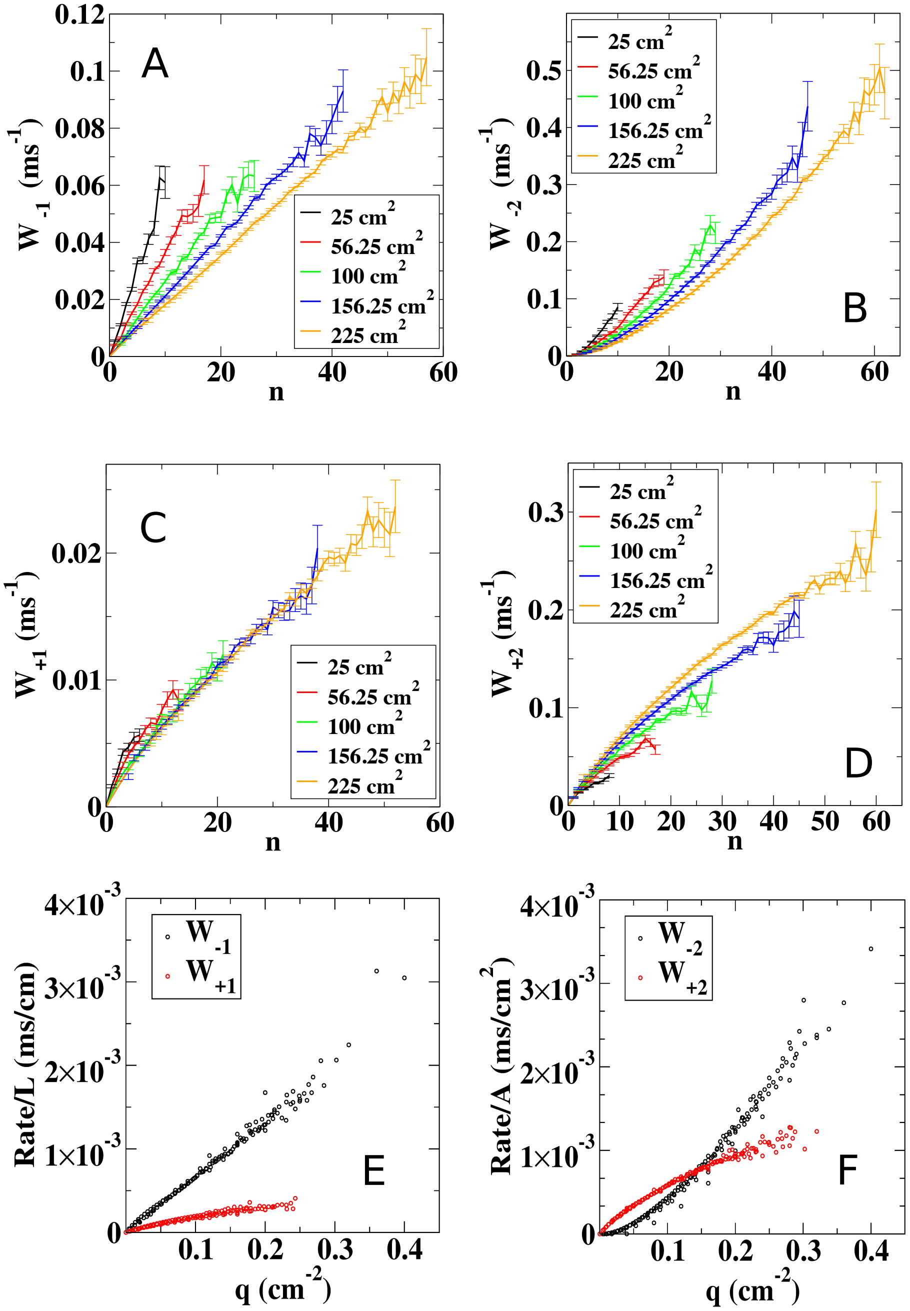

Figure 4.2: Transition rates computed using direct simulations. The birth and death rates for $n \to n-1$ (A), $n \to n-2$ (B), $n \to n+1$ (C), $n \to n+2$ (D) computed in a square geometry of various sizes with non-conducting boundaries. Error bars represent standard deviation. E: The $W_{\pm 1}$ rates, normalized by the perimeter of the domain, as a function of the density of tips, $q = n/A$. F: The $W_{\pm 2}$ rates, normalized by the area of the domain, as a function of the density of tips.

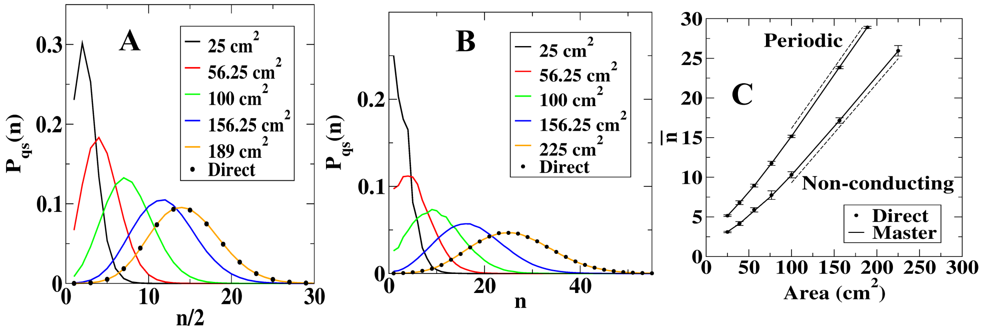

Figure 4.3: Dependence of spiral tip dynamics on the domain size. A and B. The quasi-stationary distribution for periodic (A) and non-conducting boundary conditions (B) using different domain sizes as computed using the transition matrix. The symbols show the quasi-stationary distribution as computed using the direct simulations. C. The average number of tips as a function of the area of the computational domain, computed using direct simulations (symbols) and using the quasi-stationary distribution determined from the transition matrix (line). The dashed curves are straight lines.

peak n^*. The average number of tips, $\bar{n}$, increases with system size and our simulations reveal that it depends linearly on the area of the computational domain for both boundary conditions (Fig. 4.3C).

For geometries that do not contain any non-conducting boundaries it is possible to derive closed-form solutions for the quasi-stationary distribution. In this case, n is always even and tips will be created and annihilated in pairs such that $W_{\pm1} = 0$. We can therefore compute the quasi-stationary distribution $P_{qs}(n)$ from Eq. 4.3, with the states representing pairs of tips. Furthermore, for large areas the average number of tips can be approximated by n^* which equates to $W_{+2}(n^*) = W_{-2}(n^*)$ and corresponds to the deterministic steady state [Gar09]. Using our numerically obtained scaling, this implies that the average density of tips is independent of the area and therefore $n^* \sim A$, consistent with the scaling found in the simulations (Fig. 4.3C). For domains that contain non-conducting boundaries, the ±1 rates are no longer zero and the quasi-stationary distribution cannot be written in a closed form. However, since the ±2 rates scale with A while the ±1 rates scale with $L = \sqrt{A}$, for large values of A the average number of tips will also scale linearly with the area.

To find the mean termination time τ in the direct simulations, we average the termination times T_e obtained from each independent simulation. This computation becomes more and more time consuming as A increases since termination becomes less and less likely. As a consequence, the number of determined termination events we consider vary from 400 for small domains to only 10 for the largest area still amenable to direct simulations. Our results reveal that τ displays an exponential dependence on the size of the domain, both for periodic (Fig. 4.4A) and non-conducting boundary conditions (Fig. 4.4B), consistent with earlier studies [Qu06]. Rather than using direct simulations to determine an average value for T_e, it is straightforward to use the interpolated transitions rates and the resulting transition matrix to compute τ using simple matrix operations [New10]. The resulting values for τ agree well with the direct numerical simulations (Fig. 4.4A and B). Importantly, using the transition matrix allows us to estimate the mean episode duration for system sizes that are well beyond the reach of direct simulations.

4.3.3 Scaling Behavior

For periodic boundary conditions, it is possible to obtain a closed-form expression for τ:

$$\tau(n_0) = \sum_{k=1}^{n_0/2} \phi(2k-2) \sum_{j=k}^{\infty} \frac{1}{\phi(2j)W_{+2}(2j)} \tag{4.18}$$

where n_0 is the initial number of spiral tips (must be even), $\phi(k) = \prod_{i=1}^{k/2} W_{-2}(2i)/W_{+2}(2i)$, and $\phi(0) \equiv 1$ [Gar09]. Using the numerically determined rates we find that τ quickly converges as n_0 increases. Thus, to determine the scaling with the area we can focus on the first term of this expression, $\tau(2) = \sum_{j=1}^{\infty}[\phi(2j)W_{+2}(2j)]^{-1}$. We can write $\phi(2j)$ as

$$\ln[\phi(2j)] = \sum_{z=1}^{j} \ln\left[\frac{W_{-2}(2z)}{W_{+2}(2z)}\right] \approx -\frac{A}{2}\int_{2/A}^{x} \ln\left[\frac{w_{+2}(s)}{w_{-2}(s)}\right]\, ds \tag{4.19}$$

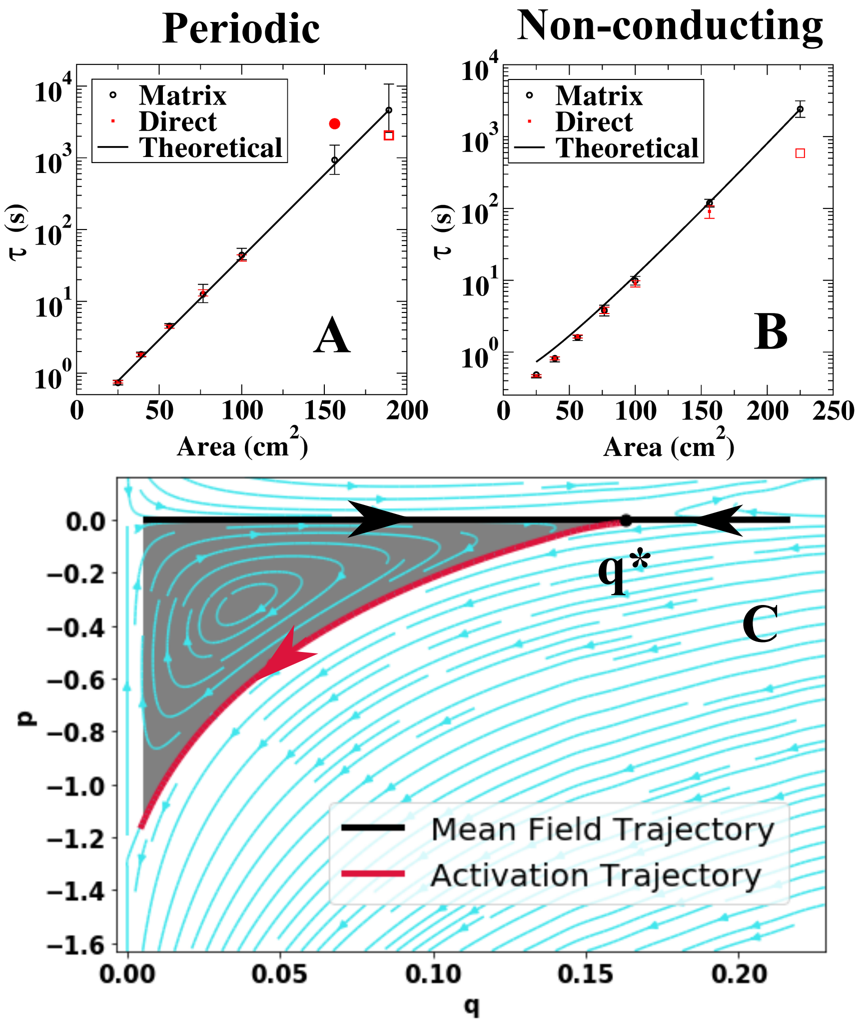

Figure 4.4: WKB approach to spiral tip dynamics. A. τ as a function of the area of the computational domain from direct simulations (red symbols) and from the master equation approach (black symbols) with error bars determined through bootstrapping. Also shown is τ computed using the analytical expression obtained using a WKB analysis (solid line). The red square represents the result of a single termination event computed using direct simulations. C: Phase portrait of the Hamiltonian dynamics in q, p space, showing the zero-energy trajectory of the WKB Hamiltonian (purple line). The shaded area represents the exponential factor $\hat{S}_0$.

where we have used the scaling of the transition rates $W_r(n) = Aw_r(n/A)$ and have defined $s = 2z/A$ and $x = 2j/A$. As a result, the mean episode duration becomes

$$\tau \approx \int_0^\infty \frac{\exp\left[A\int_{2/A}^{x} \ln\sqrt{\frac{w_{+2}(s)}{w_{-2}(s)}}ds\right]}{2w_{+2}(x)}dx. \tag{4.20}$$

The integral in the exponent is at its maximum when $x = q^*$, and for large values of A the whole integral will be sharply peaked [KRBN10]. This gives τ the following scaling behavior:

$$\tau \sim \exp\left[A\int_{2/A}^{q^*} \ln\sqrt{\frac{w_{+2}(s)}{w_{-2}(s)}}ds\right]. \tag{4.21}$$

Therefore, as an immediate consequence of the observed scaling of our transition rates we know that τ will scale exponentially with the area, consistent with our direct numerical results (Fig. 4.4A).

For domains containing non-conducting boundaries, it is no longer possible to derive an exact expression for τ. However, by viewing the number of spiral tips as a stochastic population in a metastable state we can use approximation methods to determine the scaling of the mean episode duration. As long as A, equivalent to the total population size in models of population biology, is sufficiently large, this can be achieved through a WKB approximation. In this approximation, the quasi-stationary distribution is assumed to obey $P_{qs} \sim e^{-AS}$ where S is a smooth function called the action. Substituting this into the stationary form of Eq. 4.2 yields, to leading order, the following Hamilton-Jacobi equation:

$$H(q,p) = \sum_r A^{r/2} w_r(q)(e^{rp} - 1) = 0, \tag{4.22}$$

with $p = dS/dq$ is the fluctuation momentum and where we assume the transition rates scale as $W_r(n) = A^{r/2}w_r(q)$ with $q = n/A$. The optimal path to extinction occurs along the activation trajectory $p_a(q)$ which is the non-trivial solution of $H(q, p_a) = 0$ [AM10, KMK73, DMRH94, KS07, DSS05]. This path can be determined either numerically or using approximate closed-form relations and connects the attracting fixed point $(q^*, 0)$ with the termination point along the $q = 0$ line (Fig. 4.4C). Once this path is determined, it can be used to solve for the accumulated action as $S_0 = \int_{q^*}^{2/A} p_a(q)dq$. For periodic boundary conditions, we get

$$S_0 = \int_{q^*}^{2/A} \log \gamma_0 \, dq, \tag{4.23}$$

where $\gamma_0 = \sqrt{\frac{w_{-2}}{w_{+2}}}$. For absorbing boundaries, because the scaling of the ±1 rates goes as $\sqrt{A}$ while the ±2 rates go as A, the accumulated action will be given by $S_0 + A^{-1/2}S_1$. We can then solve for S_1 perturbatively, yielding

$$S_1 = \int_{q^*}^{2/A} \frac{(\gamma_0 w_{+1} - w_{-1})(\gamma_0 - 1)}{2\gamma_0(w_{+2} - w_{-2})}. \tag{4.24}$$

The advantage of the master equation, together with the WKB approach, becomes clear when one considers the dependence of τ on area A, or equivalently the diffusion constant D. Using this formulation, we can compute the mean episode duration for any system size once the rates for a single domain $\hat{A}$ are determined. For periodic boundaries this gives

$$\tau(A) \approx \hat{\tau} e^{(A-\hat{A})\hat{S}_0}, \tag{4.25}$$

where the pre-factor $\hat{\tau}$ and the factor in the exponent $\hat{S}_0$ correspond to the termination time and action, respectivley, for the particular domain size. For absorbing boundaries, we get

$$\tau(A) \approx \hat{\tau} e^{(A-\hat{A})\hat{S}_0 + (\sqrt{A}-\sqrt{\hat{A}})\hat{S}_1}. \tag{4.26}$$

These scaling laws for τ agree well with the values of τ computed from the master equation (Fig. 4.4A, B). Moreover, the scaling can be extended to domain sizes for which computing

τ using direct simulations is not feasible. For example, directly simulating a single extinction event on a domain with area $A = 225 \mu m^2$ was found to take approximately 100 hours of CPU time. Estimating τ from this event is not useful as the error is large and generating a sufficient amount of termination events is not practical. Furthermore, for other larger domain sizes our direct simulations failed to produce a single termination event after 7 days of CPU time. Using the rates computed from this single, non-terminating event, however, we are still able to use the WKB approximation (Fig. 4.4A, B) to predict the mean episode duration, demonstrating the power of this approach. Finally, our expression for the mean episode duration reveals that it scales exponentially with the system size A, consistent with the well-established critical mass hypothesis which posits that fibrillation only occurs in hearts of a minimum size [Gar14, Qu06].

In conclusion, we present a novel approach to quantify spiral wave dynamics in spatially extended domains. This approach recast the problem into a master equation, after which statistical physics methods can be employed. Our approach is valid for any model exhibiting SDC, including electrophysiological models, and any geometry. Key in this approach are the transition rates, which can be computed numerically from a limited set of direct simulations. Alternatively, in case a priori knowledge of the dependence of the rates on the number of tips is available, it should be possible to derive expressions for these rates, resulting in analytical expression for τ. In either case, stochastic analysis of spiral wave reentry has the potential to be an important step towards determining optimal therapeutic interventions aimed at minimizing the duration of AF episodes.

Chapter 4, in part, is currently being prepared for submission for publication of the material. Vidmar, David; Rappel, Wouter-Jan. The dissertation author was the primary investigator and author of this paper.

Chapter 5

Conclusion

In this work, we have described the electrical conduction system of the heart as an excitable medium and presented results characterizing excitation patterns in both clinical and simulated arrhythmias. We began with an introduction to the theory of excitable media, with background on excitability in the Hodgkin-Huxley and FitzHugh-Nagumo models. We then discussed hypothesized mechanisms for the initiation of arrhythmias through the so-called pinwheel experiment, where a well-timed stimulus initiates spiral waves which can persist indefinitely. We moved on to discuss the hypothesized mechanisms for the maintenance of fibrillation after initiation occurs, where the dynamics of clinical fibrillation are poorly understood. In particular, we discussed the historical advent of two popular theories explaining the perpetuation of disorder during fibrillation. One of these mechanisms, multi-wavelet reentry, posits uniform disorganization due to spiral wave fragmentation whereas the other, mother rotors with fibrillatory conduction, posits the existence of a localized spiral wave with peripheral disorganization.

To better understand the dynamical state underlying clinical fibrillation, we described various sources of data which can be recorded during episodes of arrhythmia. The gold-standard measure, optical mapping, records voltage-sensitive dye with optical sensors and creates spatiotemporal maps to visualize excitation dynamics. Such studies have revealed consistent rotational patterns in animal models and explanted human hearts, but cannot be applied to human patients due to dye toxicity. In such patients, basket catheters can be inserted into the heart during therapeutic procedures from which electrograms can be recorded across an 8x8 grid. These electrograms can then be processed and used to map and visualize electrical activity occurring during recorded episodes of fibrillation.

One standard simplified representation of these electrograms, which are often complex and noisy, is their discrete set of activation times marking the onset of each excitation. From these activation times we can infer a phase, from which we proposed to use a measure of phase synchrony to quantify the ensuing dynamical state and determine regions of order and disorder across the heart. Such analysis conducted on simulated excitation patterns showed that regions controlled by stable spiral waves encoded as synchronous whereas regions of spiral wave breakup encode as asynchronous. This analysis was applied to clinical electrogram data recorded during atrial fibrillation, where it identified pockets of synchrony surrounded by asynchrony. These results are inconsistent with a mechanism experiencing uniform breakup, such as multi-wavelet reentry, but point toward a hierarchy of order consistent with the mother rotor theory. We also show in simulations that spiral wave meander can cause asynchronous activation at the spiral wave core. To that end, we proposed a measure, using phase synchrony, to quantify how dynamically out-of-step a given electrode is with respect to the entire domain. We presented results computing this measure on both simulated and clinical data, where it corresponded to sites of rotational activity.

To further characterize the complex spatiotemporal maps recorded during fibrillation, we presented a methodology to determine vector flow fields from activation time data. This technique was applied to simulated data, where it clearly identified a rotational and focal excitation pattern. When applied to clinical data, it was able to similarly identify such excitation patterns, even during complex maps with peripheral disorder. Next, the process of constructing phase maps directly from electrograms was discussed. We described the Hilbert transform and a simple method to determine phase singularities from topological charge. We also outline a recently proposed means of signal recomposition aimed at simplifying electrogram signals before application of the Hilbert transform.

Finally, we outlined a stochastic framework for analyzing the spontaneous termination of fibrillation. We discussed the state of spiral-defect chaos, similar to multi-wavelet reentry, whereby initiated spiral waves fragment and annihilate in a birth-death process. We provided a background on population dynamics of generic birth-death processes, where we outlined exact and approximate solutions to the corresponding master equation. We then used these techniques to analyze an electrophysiological model of fibrillation which exhibits spiral-defect chaos by computing transition rates from a limited number of direct simulations. We first discussed the simplified case where the boundary conditions were periodic and our system becomes a single-step process. Exact results for the quasi-stationary distribution P_{qs} and mean episode duration τ were derived for the computed transition rates. Moreover, a transition matrix was constructed from these rates and used to determine P_{qs} and mean episode duration τ, which agreed with the direct simulations. Next, we discussed the case with absorbing boundaries, where we had a multi-step process. Here, closed form solutions to the master equation do not exist, but the transition matrix can still be used to determine P_{qs} and τ. The results similarly agreed with direct simulations. We also explored the scaling behavior of τ with area A using a WKB approach. In this approach, termination occurs along the activation trajectory of an effective Hamiltonian. The accumulated action was determined from the transition rates in this model, and shown to match well with the results from the direct simulations on domains of varying size.

Future work should comprise studying the statistics of tip trajectories during spiral-defect chaos. Currently, large domains do not experience any transition events at small tip numbers, necessitating these transition rates to be determined through interpolation. Moreover, the computed transition rates exhibited a complex dependence on n, not well-fit by simple polynomials. If the rates could, instead, be fit to a known functional form, the errors could be significantly reduced. Better understanding the functional forms of these rates, as well as potentially constructing phenomenological models of the transition rates themselves, could therefore benefit this analysis.

Bibliography

[ABS73] Maurits A Allessie, Felix IM Bonke, and Francien JG Schopman. Circus movement in rabbit atrial muscle as a mechanism of tachycardia. *Circulation research*, 33(1):54–62, 1973.

[AM10] Michael Assaf and Baruch Meerson. Extinction of metastable stochastic populations. *Physical Review E*, 81(2):021116, 2010.

[Ass10] Michael Assaf. *Theory of large fluctuations in stochastic populations*. Citeseer, 2010.

[BGIC$^+$17] Muhammad Balouch, Esra Gucuk Ipek, Jonathan Chrispin, Rizma J Bajwa, Tarek Zghaib, Ronald D Berger, Hiroshi Ashikaga, Saman Nazarian, Joseph E Marine, Hugh Calkins, and David D Spragg. Impact of rotor temperospatial stability on acute and one-year atrial fibrillation ablation outcomes. *Clinical cardiology*, 40(6):383–389, 2017.

[BLA$^+$01] Mark-Anthony Bray, Shien-Fong Lin, Rubin R Aliev, Bradley J Roth, and John P Wikswo. Experimental and theoretical analysis of phase singularity dynamics in cardiac tissue. *Journal of cardiovascular electrophysiology*, 12(6):716–722, 2001.

[BST$^+$16] Eric Buch, Michael Share, Roderick Tung, Peyman Benharash, Parikshit Sharma, Jayanthi Koneru, Ravi Mandapati, Kenneth A Ellenbogen, and Kalyanam Shivkumar. Long-term clinical outcomes of focal impulse and rotor modulation for treatment of atrial fibrillation: A multicenter experience. *Heart Rhythm*, 13(3):636–641, 2016.

[BZN17] Guillaume Bassil, Manuel Zarzoso, and Sami F Noujaim. Allometric scaling of electrical excitation and propagation in the mammalian heart. *Journal of theoretical biology*, 419:238–242, 2017.

[CHN$^+$13] Sumeet S. Chugh, Rasmus Havmoeller, Kumar Narayanan, David Singh, Michiel Rienstra, Emelia J. Benjamin, Richard F. Gillum, Young-Hoon Kim, John H. McAnulty, Zhi-Jie Zheng, Mohammad H. Forouzanfar, Mohsen Naghavi, George A. Mensah, Majid Ezzati, and Christopher J. L. Murray. Worldwide epidemiology of atrial fibrillation: a global burden of disease 2010 study. *Circulation*, pages CIRCULATIONAHA–113, 2013.

[CRGM14] Sumeet S Chugh, Gregory A Roth, Richard F Gillum, and George A Mensah. Global burden of atrial fibrillation in developed and developing nations. *Global heart*, 9(1):113–119, 2014.

[Cro87] Jane Cronin. *Mathematical aspects of Hodgkin-Huxley neural theory*, volume 7. Cambridge University Press, 1987.

[CSD$^+$91] James L Cox, RB Schuessler, Jr HJ D'Agostino, CM Stone, Byung-Chul Chang, ME Cain, PB Corr, and JP Boineau. The surgical treatment of atrial fibrillation. iii. development of a definitive surgical procedure. *The Journal of thoracic and cardiovascular surgery*, 101(4):569–583, 1991.

[DB00] Graham Duckett and Dwight Barkley. Modeling the dynamics of cardiac action potentials. *Physical review letters*, 85(4):884, 2000.

[DMRH94] Mark I Dykman, Eugenia Mori, John Ross, and PM Hunt. Large fluctuations and optimal paths in chemical kinetics. *The Journal of chemical physics*, 100(8):5735–5750, 1994.

[DPS$^+$92] Jorge M Davidenko, Arcady V Pertsov, Remy Salomonsz, William Baxter, and José Jalife. Stationary and drifting spiral waves of excitation in isolated cardiac muscle. *Nature*, 355(6358):349, 1992.

[DSS05] Charles R Doering, Khachik V Sargsyan, and Leonard M Sander. Extinction times for birth-death processes: Exact results, continuum asymptotics, and the failure of the fokker–planck approximation. *Multiscale Modeling & Simulation*, 3(2):283–299, 2005.

[ENS04] Igor R Efimov, Vladimir P Nikolski, and Guy Salama. Optical imaging of the heart. *Circulation research*, 95(1):21–33, 2004.

[FCHE02] Flavio H Fenton, Elizabeth M Cherry, Harold M Hastings, and Steven J Evans. Multiple mechanisms of spiral wave breakup in a model of cardiac electrical activity. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 12(3):852–892, 2002.

[Fit61] Richard FitzHugh. Impulses and physiological states in theoretical models of nerve membrane. *Biophysical journal*, 1(6):445–466, 1961.

[FK98] Flavio Fenton and Alain Karma. Vortex dynamics in three-dimensional continuous myocardium with fiber rotation: Filament instability and fibrillation. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 8(1):20–47, 1998.

[FRC$^+$06] Valentin Fuster, Lars E Rydén, David S Cannom, Harry J Crijns, Anne B Curtis, Kenneth A Ellenbogen, Jonathan L Halperin, Jean-Yves Le Heuzey, G Neal Kay, James E Lowe, S Bertil Olsson, Eric N Prystowsky, Juan Luis Tamargo, and Samuel Wann. Acc/aha/esc 2006 guidelines for the management of patients with atrial fibrillation: full text: a report of the american college of cardiology/american heart association task force on practice guidelines and the european society of cardiology committee for practice guidelines (writing committee to revise the 2001 guidelines for the management of patients with atrial fibrillation) developed in collaboration with the european heart rhythm association and the heart rhythm society. *Europace*, 8(9):651–745, 2006.

[Gar14] Walter E Garrey. The nature of fibrillary contraction of the heart.its relation to tissue mass and form. *American Journal of Physiology–Legacy Content*, 33(3):397–414, 1914.

[Gar09] Crispin Gardiner. *Stochastic methods*, volume 4. springer Berlin, 2009.

[Gil77] Daniel T Gillespie. Exact stochastic simulation of coupled chemical reactions. *The journal of physical chemistry*, 81(25):2340–2361, 1977.

[GPJ98] Richard A Gray, Arkady M Pertsov, and José Jalife. Spatial and temporal organization during cardiac fibrillation. *Nature*, 392(6671):75–78, 1998.

[GSB$^+$05] Sydney L Gaynor, Richard B Schuessler, Marci S Bailey, Yosuke Ishii, John P Boineau, Marye J Gleva, James L Cox, and Ralph J Damiano. Surgical treatment of atrial fibrillation: predictors of late recurrence. *The Journal of thoracic and cardiovascular surgery*, 129(1):104–111, 2005.

[HH52] Alan L Hodgkin and Andrew F Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of physiology*, 117(4):500–544, 1952.

[HJS$^+$98] Michel Haissaguerre, Pierre Jaïs, Dipen C Shah, Atsushi Takahashi, Méleze Hocini, Gilles Quiniou, Stéphane Garrigue, Alain Le Mouroux, Philippe Le Métayer, and Jacques Clémenty. Spontaneous initiation of atrial fibrillation by ectopic beats originating in the pulmonary veins. *New England Journal of Medicine*, 339(10):659–666, 1998.

[HS81] Berthold KP Horn and Brian G Schunck. Determining optical flow. *Artificial intelligence*, 17(1-3):185–203, 1981.

[HZC$^+$15] Brian J Hansen, Jichao Zhao, Thomas A Csepe, Brandon T Moore, Ning Li, Laura A Jayne, Anuradha Kalyanasundaram, Praise Lim, Anna Bratasz, Kimerly A Powell, Orlando P Simonetti, Robert S D Higgins, Ahmet Kilic, Peter J Mohler, Paul M L Janssen, Raul Weiss, John D Hummel, and Vadim V Fedorov. Atrial fibrillation driven by micro-anatomic intramural re-entry revealed by simultaneous sub-epicardial and sub-endocardial optical mapping in explanted human hearts. *European heart journal*, 36(35):2390–2401, 2015.

[Jal10] José Jalife. Deja vu in the theories of atrial fibrillation dynamics. *Cardiovascular research*, 89(4):766–775, 2010.

[KMK73] Ryogo Kubo, Kazuhiro Matsuo, and Kazuo Kitahara. Fluctuation and relaxation of macrovariables. *Journal of Statistical Physics*, 9(1):51–96, 1973.

[Koc04] Christof Koch. *Biophysics of computation: information processing in single neurons*. Oxford university press, 2004.

[KRBN10] Pavel L Krapivsky, Sidney Redner, and Eli Ben-Naim. *A kinetic view of statistical physics*. Cambridge University Press, 2010.

[KS07] David A Kessler and Nadav M Shnerb. Extinction rates for fluctuation-induced metastabilities: a real-space wkb approach. *Journal of Statistical Physics*, 127(5):861–886, 2007.

[KS10] James Keener and James Sneyd. *Mathematical physiology: I: cellular physiology*. Springer Science & Business Media, 2010.

[KZM$^+$15] Pawel Kuklik, Stef Zeemering, Bart Maesen, Jos Maessen, Harry J Crijns, Sander Verheule, Anand N Ganesan, and Ulrich Schotten. Reconstruction of instantaneous phase of unipolar atrial contact electrogram using a concept of sinusoidal recomposition and hilbert transform. *IEEE transactions on biomedical engineering*, 62(1):296–302, 2015.

[KZvH$^+$17] Pawel Kuklik, Stef Zeemering, Arne van Hunnik, Bart Maesen, Laurent Pison, Dennis H Lau, Jos Maessen, Piotr Podziemski, Christian Meyer, Benjamin Schäffer, Harry Crijns, Stephan Willems, and Ulrich Schotten. Identification of rotors during human atrial fibrillation using contact mapping and phase singularity detection: technical considerations. *IEEE Transactions on Biomedical Engineering*, 64(2):310–318, 2017.

[Lew21] Thomas Lewis. Oliver-sharpey lectures on the nature of flutter and fibrillation of the auricle. *British medical journal*, 1(3146):551, 1921.

[LR91] Ching-hsing Luo and Yoram Rudy. A model of the ventricular cardiac action potential. depolarization, repolarization, and their interaction. *Circulation research*, 68(6):1501–1526, 1991.

[MA59] GK Moe and JA Abildskov. Atrial fibrillation as a self-sustaining arrhythmia independent of focal discharge. *American heart journal*, 58(1):59–70, 1959.

[Min13] George Ralph Mines. On dynamic equilibrium in the heart. *The Journal of physiology*, 46(4-5):349–383, 1913.

[MKD$^+$17] John M Miller, Vikas Kalra, Mithilesh K Das, Rahul Jain, Jason B Garlie, Jordan A Brewster, and Gopi Dandamudi. Clinical benefit of ablating localized sources for human atrial fibrillation: the indiana university firm registry. *Journal of the American College of Cardiology*, 69(10):1247–1256, 2017.

[MMB$^+$01] Moussa Mansour, Ravi Mandapati, Omer Berenfeld, Jay Chen, Faramarz H Samie, and José Jalife. Left-to-right gradient of atrial frequencies during acute atrial fibrillation in the isolated sheep heart. *Circulation*, 103(21):2631–2636, 2001.

[MRA64] Gordon K Moe, Werner C Rheinboldt, and JA Abildskov. A computer model of atrial fibrillation. *American heart journal*, 67(2):200–220, 1964.

[MSC$^+$00] Ravi Mandapati, Allan Skanes, Jay Chen, Omer Berenfeld, and José Jalife. Stable microreentrant sources as a mechanism of atrial fibrillation in the isolated sheep heart. *Circulation*, 101(2):194–199, 2000.

[MVB$^+$12] Kathleen S McDowell, Fijoy Vadakkumpadan, Robert Blake, Joshua Blauer, Gernot Plank, Rob S MacLeod, and Natalia A Trayanova. Methodology for patient-specific modeling of atrial fibrosis as a substrate for atrial fibrillation. *Journal of electrocardiology*, 45(6):640–645, 2012.

[New10] Mark Newman. *Networks: an introduction*. Oxford university press, 2010.

[NGN12] Denis Noble, Alan Garny, and Penelope J Noble. How the hodgkin–huxley equations inspired the cardiac physiome project. *The Journal of physiology*, 590(11):2613–2628, 2012.

[NKER12] Sanjiv M Narayan, David E Krummen, Michael W Enyeart, and Wouter-Jan Rappel. Computational mapping identifies localized mechanisms for ablation of atrial fibrillation. *PloS one*, 7(9):e46034, 2012.

[NKR12] Sanjiv M Narayan, David E Krummen, and Wouter-Jan Rappel. Clinical mapping approach to diagnose electrical rotors and focal impulse sources for human atrial fibrillation. *Journal of cardiovascular electrophysiology*, 23(5):447–454, 2012.

[NKS$^+$12] Sanjiv M Narayan, David E Krummen, Kalyanam Shivkumar, Paul Clopton, Wouter-Jan Rappel, and John M Miller. Treatment of atrial fibrillation by the ablation of localized sources: Confirm (conventional ablation for atrial fibrillation with or without focal impulse and rotor modulation) trial. *Journal of the American College of Cardiology*, 60(7):628–636, 2012.

[NTC$^+$08] Graham Nichol, Elizabeth Thomas, Clifton W Callaway, Jerris Hedges, Judy L Powell, Tom P Aufderheide, Tom Rea, Robert Lowe, Todd Brown, John Dreyer, Dan Davis, Ahamed Idris, and Ian Stiell. Regional variation in out-of-hospital cardiac arrest incidence and outcome. *Jama*, 300(12):1423–1431, 2008.

[OM10] Otso Ovaskainen and Baruch Meerson. Stochastic models of population extinction. *Trends in ecology & evolution*, 25(11):643–652, 2010.

[PDS$^+$93] Arkady M Pertsov, Jorge M Davidenko, Remy Salomonsz, William T Baxter, and Jose Jalife. Spiral waves of excitation underlie reentrant activity in isolated cardiac muscle. *Circulation research*, 72(3):631–650, 1993.

[Qu06] Zhilin Qu. Critical mass hypothesis revisited: role of dynamical wave stability in spontaneous termination of cardiac fibrillation. *American Journal of Physiology-Heart and Circulatory Physiology*, 290(1):H255–H263, 2006.

[RDJ+07] Martin Rotter, Lam Dang, Vincent Jacquemet, Nathalie Virag, Lukas Kappenberger, and Michel Haissaguerre. Impact of varying ablation patterns in a simulation model of persistent atrial fibrillation. *Pacing and clinical electrophysiology*, 30(3):314–321, 2007.

[RN13] Wouter-Jan Rappel and Sanjiv M Narayan. Theoretical considerations for mapping activation in human cardiac fibrillation. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 23(2):023113, 2013.

[RPK+01] Michael Rosenblum, Arkady Pikovsky, Jurgen Kurths, Carsten Schäfer, and Peter A Tass. Phase synchronization: from theory to data analysis. In *Handbook of biological physics*, volume 4, pages 279–321. Elsevier, 2001.

[SCC+07] Bramah N Singh, Stuart J Connolly, Harry JGM Crijns, Denis Roy, Peter R Kowey, Alessandro Capucci, David Radzik, Etienne M Aliot, and Stefan H Hohnloser. Dronedarone for maintenance of sinus rhythm in atrial fibrillation or flutter. *New England Journal of Medicine*, 357(10):987–999, 2007.

[SKR+16] Philipp Sommer, Simon Kircher, Sascha Rolf, Silke John, Arash Arya, Borislav Dinov, Sergio Richter, Andreas Bollmann, and Gerhard Hindricks. Successful repeat catheter ablation of recurrent longstanding persistent atrial fibrillation with rotor elimination as the procedural endpoint: a case series. *Journal of cardiovascular electrophysiology*, 27(3):274–280, 2016.

[SKR+17] Stefan Georg Spitzer, László Károlyi, Carola Rämmler, Frank Scharfe, Thomas Weinmann, Mirko Zieschank, and Anke Langbein. Treatment of recurrent non-paroxysmal atrial fibrillation using focal impulse and rotor mapping (firm)-guided rotor ablation: Early recurrence and long-term outcomes. *Journal of cardiovascular electrophysiology*, 28(1):31–38, 2017.

[SSB+17] Jonathan S Steinberg, Yash Shah, Advay Bhatt, Tina Sichrovsky, Aysha Arshad, Emily Hansinger, and Dan Musat. Focal impulse and rotor modulation: Acute procedural observations and extended clinical follow-up. *Heart Rhythm*, 14(2):192–197, 2017.

[TDM+15] Gery Tomassoni, Sandeep Duggal, Melody Muir, Lynn Hutchins, Keri Turner, AM McLoney, and Aaron Hesselson. Long-term follow-up of firm-guided ablation of atrial fibrillation: a single-center experience. *J Innov Card Rhythm Manag*, 6:2145–2151, 2015.

[TK88] John J Tyson and James P Keener. Singular perturbation theory of traveling waves in excitable media (a review). *Physica D: Nonlinear Phenomena*, 32(3):327–361, 1988.

[UNM$^+$10] Karthikeyan Umapathy, Krishnakumar Nair, Stephane Masse, Sridhar Krishnan, Jack Rogers, Martyn P Nash, and Kumaraswamy Nanthakumar. Phase mapping of cardiac fibrillation. *Circulation: Arrhythmia and Electrophysiology*, 3(1):105–114, 2010.

[VJB$^+$15] Atul Verma, Chen-yang Jiang, Timothy R Betts, Jian Chen, Isabel Deisenhofer, Roberto Mantovan, Laurent Macle, Carlos A Morillo, Wilhelm Haverkamp, Rukshen Weerasooriya, Jean-Paul Albenque, Stefano Nardi, Endrj Menardi, Paul Novak, and Prashanthan Sanders. Approaches to catheter ablation for persistent atrial fibrillation. *New England Journal of Medicine*, 372(19):1812–1822, 2015.

[Win87] Arthur T Winfree. *When time breaks down: the three-dimensional dynamics of electrochemical waves and cardiac arrhythmias*, volume 14. Princeton University Press Princeton, 1987.

[WR46] N Wiener and A Rosenblueth. The propagation of impulses in cardial muscle. *Archiv de Institut de Cardiologia di Mexico*, 16:3–4, 1946.

[YTW$^+$07] Ming-Jim Yang, Diana X Tran, James N Weiss, Alan Garfinkel, and Zhilin Qu. The pinwheel experiment revisited: effects of cellular electrophysiological properties on vulnerability to cardiac reentry. *American Journal of Physiology-Heart and Circulatory Physiology*, 293(3):H1781–H1790, 2007.

[ZYKB08] Sharon Zlochiver, Masatoshi Yamazaki, Jérôme Kalifa, and Omer Berenfeld. Rotor meandering contributes to irregularity in electrograms during atrial fibrillation. *Heart rhythm*, 5(6):846–854, 2008.