LUND UNIVERSITY
Faculty of Science

Splitting a heart

Master Thesis Equivalent to 60 ECTS

Jakob Lavröd

Project duration: 12 months
Supervised by Gunnar Ohlén

Department of Physics
Division of Mathematical Physics
January 2016

Abstract
This thesis investigates and numerically solves the Wohlfart-Arloch equations describing the electrical activity in the atrial part of the heart. This is done through operator splitting that is allowing the diffusion of ionic currents, and the local voltage-driven reactions in each cell to be de-coupled. The local reactions can be solved in closed form, and the diffusion part is solved with a spectral method by interpolating using eigenmodes. The problem of initiating the system is solved by introducing time-dependent boundary conditions and solving these parts as a series. Systematic investigations are carried out both concerning the numerical errors and also the wave speed, multiple pulse shape, and other characteristics relevant to ascertaining the validity of the numerical method used.

Abbreviations

1D: One-dimensional

2D: Two-dimensional

EKG: Electrocardiography

ODE: Ordinary differential equation

ODE45: A standard algorithm for solving ordinary differential equations in Matlab using a fourth and fifth order Runge-Kutta method.

PDE: Partial differential equation

Content

1 Introduction

The human heart is one of the most incredible machines that ever existed. It pumps without replacements or rest through the entire life of each human. Given that a car or computer without serious problems for 80 years is unheard of, it is hard not to stand in awe before it.

However, those of a more curious nature move to the next step and ask: How? How does the heart work, what enables the function, and when it ceases to work properly, what is the cause? As the average lifespan of humans increases, we enter terra incognita, humans when they first appeared on earth never lived this long, and hence we can no longer rely on nature and simply hope the heart will keep working. The increasing number of patients with cardiac arrhythmias is a testament to that.

One such effort in better understanding the heart is by building a mathematical model of the heart. Such models must have rather modest objectives to be successfully, or risk getting swallowed by complexity. In this thesis we will look at one such model, the Wohlfart-Arloch model, that attempts to model the electrical conduction in the atrium. The ultimate goal of the model is to describe a phenomenon like arrhythmias and atrial fibrillation.

Previous research has focused heavily on these phenomena, but as we shall see when delving into the world of partial differential equations needed for making computation, to a high degree completely disregarded the most subtle points of the numerical analysis and its interpretations. Our objective is, therefore, to remedy this by building a framework concerning the numerical analysis for this particular equation. In the thesis, we will try to outline the reasoning behind the choice of methods. While there are places at which one can make constructive-deductive arguments for methodological choices, there are also times at which the argument simply rests upon all other alternatives being worse, and this is an important part of understanding the core of the thesis (as well for anyone trying to replicate the results).

Hence, the goal of this thesis will be to work out a more sound theoretical foundation for the discussion of the model, and try to show how such knowledge also can help us understanding the physics of the model better.

2 Derivation of the model

The following chapter will outline the model used. While the model itself is well known and used in multiple works, there is no systematic derivation from first principles, so we seek to provide such for future reference to place the model on a solid foundation.

2.1 Introduction: Zooming in on the cell

The overall purpose of the heart is to make sure that oxygen is transported from the lungs to the rest of the cells in the body. This operation is done in two circulations: First blood is pumped to the lungs for oxygenation and then returned to the heart. Secondly, the oxygen-rich blood is pumped out to the cells of the body and finally returns to the heart. For this cycle to be carried out without mixing oxygen-rich and deoxygenated blood, the heart must be divided into a left and right section. Furthermore, to be able to perform the dual task of both collecting blood and pumping it out again, both sides must be sub-divided into a pumping section. The main one is called the ventricular, and the temporary assembly area for the returning blood is called the atria. The result is the well-known four-chamber heart seen in figure 1 [1].

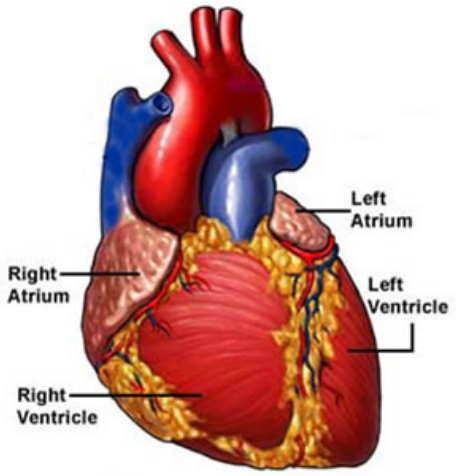

Figure 1: Overview of the human heart

From the above description, the task of the atria (that will be the focus of this work) is rather straightforward, to periodically contract so that the blood collected can be drained into the ventricular. Such process is facilitated by a muscular layer in the heart wall, the myocardium, sandwiched between the epicardium and endocardium membrane that provide lubrication to avoid friction from the contractive motion. The myocardium consists of bundles of cardiac muscle fibers, held in place by a fibrous skeleton of collagen and elastin. Each fiber, in turn, consists of cardiac muscle cells, which are rich in sarcomeres that allow for the contraction. This motion happens when the calcium ion concentration inside the cell rises since the ions binds to troponin on the muscle cell filaments that enable them so slide past each other resulting in contraction. The ions needed can either flow into the cell from the surroundings or be released the sarcoplasmic reticulum, which stores calcium ions. Electric fields will activate the later one, so the electric properties of the inside of the cell become the decisive factor in determining if contraction will take place or not [1].

The above description shows that for the purpose of understanding the motion of the atrium, the cardiac cell is the smallest unit, serving as the "atom" of the cardiac tissue. Any attempt at mathematical modeling must start from fundamental relations for each cell, which is then connected to form a model of the myocardium. This motivates the study of an isolated cardiac cell, in the same way, that solid state physics start from the properties of isolated atoms and combine them to study materials.

2.2 Modeling philosophy

The task of modeling the atrial cells is not a new one, and the area has been a very active field since the end of the 90's, spawning a wide collection of different models. Among these, one of the most extensive is the Nygren-Fiset-Firek-Clark-Lindblad-Clark-Giles model [2]. This explicitly models both the sarcoplasmic reticulum, the cleft space that handles the activation of calcium release and the surroundings. With 26 system variables, it is an impressive construction based on many years of both experimental and theoretical work. With such advanced model, it is even possible to make the explicit predictions behavior of the atria that can be investigated experimentally.

Such complex model is however not without problems. To model all the different effects, as many as 47 different parameters are needed. While the authors have tried to make admirable attempts at separately determining as many of parameter values as possible through many different experiments, this does not change the fundamental problem that with large sets of parameters, basically any behavior can be accommodated by the model. The modeling process is also complicated by the fact that atrial cells, even within the same heart, often vary rather drastically in properties, so a realistic model should use parameter fields instead of scalar values. As the complexity is increased, the transparency of the model is decreased, making it harder to identify what effect is doing what, or what factors are truly casually responsible for certain behaviors. Nor can such problem be solved by independently varying a single factor and keeping the rest constant, since there will be complicated interaction effects (e.g., if the parameters were varied in certain simultaneous combinations, the result could be virtually unchanged). This issue is known as the problem of identifiability that according to [3] serve as a major limitation of these large models.

Another risk in building such elaborate model is not to make it complex enough. This might sound contradictory, but if accuracy to a certain level aspires, neglecting any factor of that magnitude will ruin such attempts, and make it completely unnecessary to include finer effects. At the level of such large models, these sources of "missed biology" can come from virtually any direction. Since the space of functions is infinite dimensional, the experimental procedure alone can never fully provide a proper model. Instead, some theoretical assumptions must always be introduced to interpolate and extrapolate the obtained results. At such high demands for accuracy, such process soon becomes incredibly complex and increases the risk of underfitting by not having sampled enough of the experimental conditions.

The objections raised in previous paragraphs are, of course, not new, and there is an alternative direction of modeling that focuses on very structurally simplified models, so-called

toy theories [4]. Since models with a single bounded variable can only describe convergence toward a steady state, at least, two variables must be used. A prime example is the FitzHugh-Nagumo model [4]. The physiological focus is solely on the voltage of the cell, and the second so-called "recovery variable" is a phenomenological construction that does not directly correspond to some physical entity in the cell. Even so, solutions of the equations share some qualitative properties with experimentally observed behaviors. The simplicity of the equations has made it very popular with the modeling community, and it is widely used to get qualitative understanding and test methods before they are applied to more complicated models.

While the neatness of the FitzHugh-Nagumo model is admirable, such construction comes at the price of obscuring the underlying biology. Especially the recovery variable has a very unclear ontological status since it is formally derived by clustering a lot of different effects together (a method known as asymptotic reduction), assuming their timescales to be equal. Such construction complicates communication outside of the modeling community since it makes it significantly harder to explain the inner workings of the model for domain experts such as biologists whose base of knowledge is firmly grounded in the local processes of each cell. Furthermore, it makes it harder to assess which behaviors are physiologically reasonable if a system variable cannot be interpreted regarding actual system dynamics.

The natural question if, therefore, if one can achieve a balance between the two extremes presented above. The idea would be to construct a model that is structurally simple and with enough transparency for relating cause and effect, but still accurate enough for the results to be physiologically interpreted. There are two routes for constructing such model, either starting from the basic biology and physiology and build upwards, the so-called bottom-up modeling (the Nygren-Fiset-Firek-Clark-Lindblad-Clark-Giles model is example of such), or instead start from what behaviours that are expected and try to construct reasonable relations that reconstruct the observations, the so-called top-down modeling (the FitzHugh-Nagumo model). The traditional way to go about building toy theories is the latter approach since it allow for constraining the output, and the first route is in a way much riskier since it is easy to include too many effects, and hence lose the simplicity. On the other hand, following the bottom-up approach opens up the possibility of using first-order principles for constraining the form of the model, and allow us to ask what biological constituents that are needed for replicating observed behaviors.

One candidate for achieving this modeling middle position is the Wohlfart-Arlock model. It was proposed in 1992 by the eponyms to simulate EKG by studying voltage difference between different myocardial cells [5], an application of lesser interest if the goal as in our case is the dynamics of the atria, but the heart of the model is a simplified, yet physiologically interesting description, of the single cell that will be used in this study. The model has been used in several previous works on atrial simulation [6], [7], [8], [9], [10] and a variation of the model also allows for the description of the sino-atrial node that controls the rhythm of the heart [10]. A derivation of the model can be found in [11], but the model used in that study is slightly different on several accounts from the original Wohlfart-Arlock model, and the work does not provide a fair discussion on several key issues that are essential to the structure of the

model. This motivates the need for providing a review of the derivation of this work as outlined below.

2.3 The cardiac cell

Given that it is electrical effects that put the contraction of the cell in motion, the problem will be treated as an electrical problem. Furthermore, the feedback from the contractive motion (such as stretch-activated ion channels) will be neglected to simplify the description since this effects have the nature of corrections (and are not even included in most advanced models [4]). The key quantity in any electrical system is the net charge, and the lack of free electrons (due to their high reactivity with other molecules as free radicals) means that the ions fulfill the role of charge carried.

The cell contains excess charge; such will quickly migrate to the cell membrane due to the high conductivity of the intracellular fluid similar to how free charge end up on the surface of metal. The electric field that is formed by these ions will, in turn, attract oppositely charged ions from the extracellular fluid. Because the cell membrane is not permeable to charged particles due to its non-polar insides, the two sets of charges cannot neutralize directly. While the field either inside or outside the membrane will be canceled, the field in the membrane itself will not due to the charges on either side, meaning that the proteins, that contain charged amino acids, inside the membrane will be affected by the electric field. This could for example cause them to change configuration by moving the charged part, and the electrical potential difference V between the inner and outer part of the cell is therefore an appropriate variable to use. The voltage formed must be a function of the charge Q on each surface. If the charges themselves do not alter the shape and behaviour of the membrane (hence the elasticity of the membrane is neglected for simplicity, even though a small such effect exists [11]), the voltage must by superposition be proportional to the charge, $V = CQ$, where C is the membrane capacitance, since the membrane essentially act as a capacitor with the cell membrane as dielectric. Hence this can be measured experimentally for cardiac cells.

For the cell to exhibit a dynamical behavior, the excess charge on the inside must vary with time, but still slowly enough for the remaining charge to have time to redistribute and give rise to a uniform electric field. The main source of such currents is the ion channels. These are pores in the membrane consisting of proteins, with the twist that they are highly selective and only let through a particular ion [4]. Therefore, only the specific ions that can flow in and out in greater amounts needs to be considered. This shortens the list to $\mathrm{Na}^{+}, \mathrm{Ca}^{2+}, \mathrm{K}^{+}$, with Cl^{-} as the runner-up, but neglected due to it considerably weaker influence than the three preceding ones [12]. As will be seen below, this is the minimum ions that must be included to be able to replicate acceptable behaviors of the cell.

If such channel is opened in the membrane, charged ions will flow though due to the electric field through the membrane, so the current will be a function of the voltage difference, $i(V)$. However even in the absence of electric field there can be a flow due the fact that the concentration of the ions might be different on the inside and outside of the cell and ions therefore will move in though diffusion. At some voltage these to effects must cancel, and the net flow though the channel is zero even though the channel is open. This can be analyzed

using thermodynamic equilibrium, since the two states, inside and outside, that for a specific ion have different probabilities to occur (proportional to the concentrations inside c_{int} and outside c_{ext}) must have different energy, hence a voltage difference, since they follow a Boltzmann distribution. This voltage is known as the Nernst voltage E and given by:

$$E = \frac{k_B T}{ze} \ln\left[\frac{c_{ext}}{c_{int}}\right] \tag{1}$$

Here k_B is the Boltzmann constant, T is the absolute temperature, e is the fundamental charge and z the valence number of the ion. A problem with E is that as charge is transported, the concentrations will change, and so will therefore E. However a key difference is that the transport moves excess charge while the concentration includes all charged ions of that type, which is a massive difference. For both sodium and potassium, the effect of charge transport on E is very minute [4]. For calcium the matter is rather different since the concentration inside the cell initially is very low, so the argument that the influx is too small to change the concentration is not applicable. However, as pointed out by Wohlfart and Arloch [5], in the cell membrane there is also a sodium-calcium pump that by consuming energy eject calcium at a rather high rate, so the approximation of keeping the concentrations constants is not as poor as one might first be lead to assume. Furthermore, if one actually wishes to incorporate this effect into the model, such model must also include the dynamics of the pump which is highly non-linear in the voltage. For reasons to be discussed below, such term would significantly complicate the model. This means that the voltage V is the central variable in the problem, that summaries the electrical state of the cell.

The next step is to compute $i(V)$. To do so, both the diffusion and drift from the electric field must be accounted for, and furthermore the electric potential inside the channel, which depend both on the concentration of charged ion, and also on the charged amino acids inside the channel wall, meaning that both the concentration and electric potential must be solved simultaneously to compute the flux. The resulting equation is known as the Planck-Nernst-Poisson equation, and solution to this model under varying conditions and geometries is a very active research problem [4]. Since only a simplified model is sought, an asymptotic solution for the channel can be derived under fairly realistic conditions, discussed in depth in [4], resulting in a linear relationship, $i(V) = \sigma(V - E)$, were σ is the conductance of the channels of that type. This is the most common form for the voltage dependence in the literature. However it is often motivated as either a Taylor approximation or with the analogy to a resistor, neither which are correct as the range of validity of the relation is too large to argue based on linearization and the argument by analogy really does not provide any physical basis [4]. The solution to the Planck-Nernst-Poisson equation however show that the ansatz is ground in a, while simplified, still reasonable theory, and also provide some motivation to the striking success of the linear model when compared with experimental data.

If all the channels were open, the equation for the voltage is simply obtained by differentiating the voltage/charge relationship (since $\dot{Q}$ is the total charge), $C\frac{dV}{dt} = -\sum_{\text{ion}} \sigma_{\text{ion}}(V - E_{\text{ion}})$, were summation are carried out over all the channels and the minus

sign appears due to the definition of positive current as going out from the cell. It is clear that if the size of the cell membrane doubled, both the capacitance and the conductance would also double (the first from the increase in area that can store charge, the second from the doubled number of channels). This means that the equation can only depend on the ratio $g_{\text{ion}} = \sigma_{\text{ion}}/C$, called the ionic rates (as they describe at what rate a certain ion will move through the channel) . Such linear equation is simply going to converge to an equilibrium value, and hence cannot exhibit any interesting dynamics, which is expected.

To obtain a realistic behavior of the cell, it must be possible for the channels to open and close, which is achieved by a change in the structure of the channel proteins. There exists rather a broad consensus that the channel proteins have a finite number of discrete configurations that they reside in most of the time. Furthermore, there exists rather strong experimental evidence for that too excellent approximation the protein lack memory of past configurations and move between these in a stochastic manner (even though the existence of hysteresis effects is an active debate from a fundamental standpoint, it has been shown that memoryless models reproduce experimental data very well, hence the point is mute from our simplistic model build perspective). The lack of memory means that the probability distribution over the states of the protein can be modeled as a continuous Markov chain. Careful measurement of the ion channels also shows that they are in a state that either conduct or not, there is nothing in-between, meaning that the current can be obtained by multiplying the $\sigma_{\text{ion}}(V - E_{\text{ion}})$ term with the probability that the channel is conducting.

The goal of the analysis is, therefore, to first elucidate what states the channel can be in, and secondly find the topology of the Markov chain by finding what states can transition into which. This is a formidable task, and there exists a very broad literature proposing different structures with as many as 30 different states [14] for each channel. In line with the conservative attempt to model the dynamics, the absolute minimum number of states to replicate realistic behavior will be used, since anymore will cause an explosion in number of parameters as will be seen further below, for a n state chain there are $(n-1)^2$ parameters that must be prescribed from the dynamical behavior, which soon gets unmanageable. One way to reduce the number of parameters is the so-called Huxley-Hodgin ansatz, where the protein is studied as separate parts with separate Markov chains and lessen the number of parameters (even though such process, strictly speaking, is an approximation, as long as the protein sections are rather separated, it is often a surprisingly good one).

Following this recipe to its extreme, each protein section can be given the minimal number of two states, describing an open state that allows ions to pass through, and a closed state that blocks the path. Given the independence assumption between the different parts of the channel, the opening probability is simply the product of the opening probabilities of all subparts. The two-state approximation has the advantage that each Markov chain will contain a single time scale that therefore can be interested as the time until thermal equilibrium settle for that process, giving the parameter a clear physical significance (in contrast to the much more muddy affair for a larger number of states). This facilitates communication with domain experts, as biologists and chemists often have a rather good understanding of the relevant time scales involved in the change of the proteins from experiments while overall transition rates

between all states are much harder to estimate without careful experimentation. From a computational perspective, we will also see below that including more states in each chain makes computations considerably slower, more complicated and less transparent. Another advantage is that only the opening probabilities need to be considered as dynamical variables since they determine the closing probabilities from the fact that the probabilities must sum to 1.

A natural question is if the probabilistic nature of the channels must be considered, or if there are enough channels to apply the law of great numbers and simply work with the mean value of the number of open channels (hence neglecting stochastic fluctuations). A conservative estimate is that there is at least 10^5 channels of each type, and using the central limit theorem for the summation of the channel current then shows that the stochastic fluctuations are less than 1 % of the mean value, which mean that given the approximation already swallowed by the model, there is no need to include stochastic effects.

For both the sodium and calcium channel, there exist two such protein sections. The first regulate the shape of the total protein and is governed by two positively charged rods. Normally these are attracted toward the inside of the membrane, but if the potential difference is reduced, they can instead move toward the outer side of the membrane, and thereby opening the channel by their outgoing motion (with probability p_{Na} and p_{Ca}). The second protein has the form of a ball that is attached by a thin chain of amino acids in the lower end of the channel. Normally the charged ball is repelled from the end of the channel (with probability q_{Na} and q_{Ca}), but if the voltage differences change, it can come closer and effectively block the opening, closing the channel for further ion transport, known as an inactivation of the channel. These conformational changes are rather fast (less than 0.1 ms) to the time scale of the dynamics, and the Markov approximation can therefore be used. It is also clear that both processes have a clear biological origin (and not fictitious processes introduced to better fit experimental data) and it can also be argued that they are the minimal number of processes, as there must be some way for the channel to be activated, but also some way for it to deactivate and close, even without the voltage lowering (otherwise the channels would remain open and keep the voltage off the equilibrium level).

The potassium channel is slightly more complicated as some of the gates are open all the time, so-called leakage channels, which primary function is to maintain the equilibrium voltage. Some of these channels can only transmit current in one direction (they are said to be inwardly rectifying). If a voltage difference in the opposite direction is applied, an attached magnesium ion is getting sucked into the channel and block it, so there is only a probability q_K that such channel remain open. There is also the opposite case, with channels that are opened with a probability p_K, following the same mechanism as the sodium and calcium channels. Finally there is the combined kind, that has both an opening probability, and inward rectification, so its probability of being open is $p_K q_K$.

If all these results are collected, this gives the complete voltage equation:

$$\frac{dV}{dt} = g_{\mathrm{Na}} p_{\mathrm{Na}} q_{\mathrm{Na}} (E_{Na} - V) + g_{\mathrm{Ca}} p_{\mathrm{Ca}} q_{\mathrm{Ca}} (E_{Ca} - V) + \left(g_{\mathrm{K,L}} + g_{\mathrm{K,i}} q_{\mathrm{K}} + g_{\mathrm{K,a}} p_{\mathrm{K}} + g_{\mathrm{K,ai}} p_{\mathrm{K}} q_{\mathrm{K}}\right)(E_K - V) \quad (2)$$

One way to interpret this equation is as including all channels containing up to two different sub-processes. A truncation beyond that is motivated by simplicity since such terms will tend to be small far more often since they will only contribute if all probabilities are fairly close to 1, and such possibility is reduced the longer the product of factors becomes. Furthermore, for each combination, only a single term is included. Investigation of the cardiac cell has shown different types of sodium-, calcium- and potassium gates, but these are also neglected for simplicity. Fundamentally, all dynamically interestingly parts have at least some representation of the equation, although in simplified form.

The next step in the derivation process is to write down equations for each of the probabilities using the two-state Markov chains. As an example, study the opening probability for sodium channel, p_{Na}. If the probability to transition from open to closed per unit time is α_{Na}, and similar the probability per unit time to go from closed to open is β_{Na}, then the corresponding equation for the Markov chain can be written:

$$\frac{dp_{Na}}{dt} = -\alpha_{Na} p_{Na} + \beta_{Na}(1 - p_{Na}) \quad (3)$$

The equation can be rewritten by observing that as $t \to \infty$ and the system attain thermodynamic equilibrium, $\dot{p}_{Na} \to 0$, so solving the resulting equation gives the equilibrium probability $\tilde{p}_{Na} = \frac{\beta_{Na}}{\alpha_{Na} + \beta_{Na}}$, which offer easier interpretability than the rather abstract transition probabilities:

$$\frac{dp_{Na}}{dt} = k_{p_{Na}}(\tilde{p}_{Na} - p_{Na}) \quad (4)$$

Here $k_{p_{Na}} = \alpha_{Na} + \beta_{Na}$ is the rate at which the thermal equilibrium is obtained. Both $\tilde{p}_{Na}$ and $k_{p_{Na}}$ are functions of the voltage (since this change the likelihood of transition). Therefore it is natural to study the system in thermal equilibrium, since statistical mechanics than can be applied on the two-level system consisting of the open and closed state to compute $\tilde{p}_{Na}$. According to the Boltzmann distribution, the probabilities only depend on the difference in energy ΔE. Since different gates and sub-process are assumed independent, the energy difference can only depend on the voltage, and can hence be written as the series:

$$\Delta E = \Delta E_0 + \gamma_1 V + \gamma_2 V^2 + \cdots \quad (5)$$

Here ΔE_0 signify the difference in energy between the states in the absence of electrical field, and further terms are Taylor expanded around this point (with coefficients γ_n for the n:th order term). It make sense to expand around $V = 0$, as this signify were there are no electrical effects, and the larger the potential difference becomes, the more important electrical effects will be.

The linear term can be interpreted as the work done by the charges of the protein when moving between the two configurations, and will be the dominant contribution to low field

strength. The linearity means that the charges are not interacting, so higher order terms can be described as a change in the electrical potential caused by the change of other charges. The quadratic term, for example, describes the polarization of the membrane resulting from the charges. For the potential differences appearing in the cell, estimates of these effects show that they can be neglected [15], leading to the so-called low field limit where the energy difference is a linear relationship of the voltage. Inserted into the Boltzmann distribution for a two-state system, this gives:

$$\tilde{p}_{Na}(V) = \frac{1}{1 + \exp\left[\frac{V_{p_{Na}} - V}{V_{Na}}\right]} \tag{6}$$

The constant $V_{p_{Na}}$ is the voltage at which there is equal probability to be in either configuration. If $V \ll V_{p_{Na}}$ it is very unlikely that the protein is in the open state, and for $V_{p_{Na}} \ll V$, it is most likely open. The width of the transition region is characterized by the parameter V_{Na}, the change from 27 % to 73 % take place within $\pm V_{Na}$. From the Boltzmann distribution follows that $V_{Na} \propto T$, showing that the origin of V_{Na} is essential thermal, even though the potential difference $V_{p_{Na}}$ is required to change the configuration, the deficit can be provided by the surroundings. Reversely, the surroundings can absorb enough energy to preserve the system even though it should change. Hence the sigmoid curve $\tilde{p}_{Na}(V)$ for the probability can be motivated from first principles. For the inactivation process, the corresponding probability can be written:

$$\tilde{q}_{Na}(V) = \frac{1}{1 + \exp\left[\frac{V - V_{q_{Na}}}{V_{Na}}\right]} \tag{7}$$

This sign in the exponential argument has been chosen in reverse since the inactivation should come into play as the voltage increase, and the ball can be attracted. From physical reasons, it is expected that $V_{q_{Na}} < V_{p_{Na}}$. This follows from the fact that a very slow increase in the voltage (so that the probabilities can be assumed to obtain their equilibrium values at all times) should not result in any excitation. Therefore, for such perturbation, well before the channel start opening, it should be inactivated to prevent any conduction. Another way to express this is that $\tilde{p}_{Na} \ll 1 - \tilde{q}_{Na}$ since the change in $\tilde{q}_{Na}$ should take place well before the change in $\tilde{p}_{Na}$. This property also serves as a motivation for why the voltage width V_{Na} should be chosen as equal in both cases (in addition to reducing the number of parameters), since if this was not the case, the inequality could not be satisfied.

Another way to interpret the Boltzmann distribution is that it constrain the ratio β_{Na}/α_{Na}. This is however not enough to compute the transition rates independently, as some physical description of the actual time scale the transition takes place over must be introduced. A simple such model that is commonly used in the literature is to assume that there exists an energy barrier between the two states, and use transition state theory, which gives $\alpha_{Na} = k_0 \exp\left[-\vartheta \frac{V - V_{q_{Na}}}{V_{Na}}\right]$, $\beta_{Na} = k_0 \exp\left[(1 - \vartheta) \frac{V - V_{q_{Na}}}{V_{Na}}\right]$, which is known as the Kramer ansatz (with k_0 being the rate when both processes are equally likely, and ϑ that describe how quickly the different rates change with respect to each other, $\vartheta = 0$ corresponds to α_{Na} being

constant, while $\vartheta = 1$ models β_{Na} constant). This correctly model that if V is far from the value $V_{q_{Na}}$, one of the processes will exponentially disappear since it basically never process the barrier, which is physically very reasonable. The flaw with this approach however is that for $0 < \vartheta < 1$, both rates become unbounded in either of the limits $V \to \pm\infty$, which means that the transition happens instantly, which is very unphysical and does not correspond to biological data.

A more realistic model for β_{Na} should go to 0 as $V \to -\infty$ since this transition becomes unlikely, and go to some maximal value as $V \to \infty$ and vice versa for α_{Na} . If one wants to preserve the exponential tail, but bounded the opposite side, the exponential function could be replaced by a sigmoidal one. The constraint on the ration means that the parameters for the sigmodials cannot be chosen independently, instead if they are added together $k_{p_{Na}}$ must then take the form:

$$k_{p_{Na}} \propto \frac{1}{\Gamma + \exp\left[\frac{V_{p_{Na}} - V}{V_{Na}}\right]/\Gamma} + \frac{1}{1/\Gamma + \Gamma \exp\left[\frac{V - V_{p_{Na}}}{V_{Na}}\right]} \tag{8}$$

Here Γ is a parameter that describes if the total rate toward equilibrium tends to increase or decrease as the voltage is increased. For $\Gamma < 1$ $k_{p_{Na}}$ is an increasing function and for $\Gamma > 1$ $k_{p_{Na}}$ is an decreasing function. For simplicity the case $\Gamma = 1$ will be chosen, since the total rate then becomes constant. As will be seen further down, this speeds up the computation significantly and reduce the complexity of the model, so such approximation is worthwhile given the gains and the already simplified structure of the equations.

3 Single cell equation

In this chapter, we solve the single cell equations, an original contribution not reported by earlier works. In doing so, the concept of operator splitting emerges, and we apply the theory of numerical analysis to investigate the simulation error.

3.1 Studying the single cell equations

Collecting all results, the Wohlfart-Arlock equation for an atrial cell becomes (see appendix for parameter values):

$$\begin{cases} \dfrac{dV}{dt} = g_{\mathrm{Na}} p_{\mathrm{Na}} q_{\mathrm{Na}} (E_{Na} - V) + g_{\mathrm{Ca}} p_{\mathrm{Ca}} q_{\mathrm{Ca}} (E_{Ca} - V) + \\ \quad \left(g_{\mathrm{K,L}} + g_{\mathrm{K,i}} q_{\mathrm{K}} + g_{\mathrm{K,a}} p_{\mathrm{K}} + g_{\mathrm{K,ai}} p_{\mathrm{K}} q_{\mathrm{K}}\right)(E_K - V) \\ \dfrac{dp_{Na}}{dt} = k_{p_{Na}} (\tilde{p}_{Na} - p_{Na}) \\ \dfrac{dq_{Na}}{dt} = k_{q_{Na}} (\tilde{q}_{Na} - q_{Na}) \\ \dfrac{dp_{Ca}}{dt} = k_{p_{Ca}} (\tilde{p}_{Ca} - p_{Ca}) \\ \dfrac{dq_{Ca}}{dt} = k_{q_{Ca}} (\tilde{q}_{Ca} - q_{Ca}) \\ \dfrac{dp_K}{dt} = k_{p_K} (\tilde{p}_K - p_K) \\ \dfrac{dq_K}{dt} = k_{q_K} (\tilde{q}_K - q_K) \end{cases} \tag{9}$$

The addition of the so-called gating probabilities $p_{\mathrm{Na}}, q_{\mathrm{Na}}, \ldots$ have made the voltage equation non-linear, which is necessary if it is to be a realistic model of excitable cell since a linear model either must be unbounded (which is physically unreasonable) or asymptotically stable (meaning that no excitation can take place)

Before the equations can be solved to extract information about the system dynamics, it is necessary to know that such solution exists. As pointed out by [17], the argument that the model comes from a real system with a guaranteed solution is insufficient due to the approximations introduced in the modeling process. For the model to be sensible, the solution must also be unique, in other cases not enough information have been included from the underlying system. In fact, uniqueness is insufficient. If the solution changes drastically under an infinitesimal variation in the initial conditions, there is no hope of ever using the model to describe a real system that always lack that sort of extreme mathematical precision. These three conditions (existence, uniqueness, and continuity under initial conditions) have been proposed by Hadamard [16] for a set of equations to be well-posed. Therefore, they must be checked for the Wohlfart-Arlock model for a single cell.

While proving such property might seem non-trivial, in the case of initial value problems for ODE:s there is a very powerful tool: Picard-Lidelöfs theorem [16], which state that it is sufficient for all partial derivatives with respect to the system variables to be bounded for there to be a unique solution continuous in the initial condition, to exists at all times. Since all

system variables can be bounded (the voltage between the minimal and maximal Nernst voltage and the probabilities between 0 and 1), straightforward calculations show that so must also be the case for the derivatives. However, the theorem provides practically no information about what such solution might look like.

The first attack on the problem would be to try to write down a closed form solution. Using Duhamel's principle [16] a scalar nonlinear integrodifferential equation can be constructed from the equations, but the nonlinearity prevents the use of basically the entire toolbox for this type of equations. Hence to the author's knowledge, in general, such method provides no result of interest. To be able to solve the system one must either restrict the class of solution to a subset for which one can simplify the equations or instead look for approximate solutions using numerical methods.

3.2 Equilibrium and perturbations around it

Following the first route, the simplest form of the solution is the equilibrium state of the cell since this corresponds to removing all time derivatives and solving a purely algebraic problem. From physical consideration, the myocardial cell must have such state, and it must be unique since it is known from an experimental investigation that the cell relaxes back to the original configuration after a heartbeat. Since the ion gates must be in thermodynamics equilibrium, they are given by $\tilde{p}_{ion}(V)$ and $\tilde{q}_{ion}(V)$ (V). Therefore only the equation for the voltage remains. Setting $\frac{dV}{dt} = 0$ in the Wohlfart-Arloch equation (9) yields:

$$g_{\mathrm{Na}}\tilde{p}_{\mathrm{Na}}\tilde{q}_{\mathrm{Na}}(E_{Na} - V) + g_{\mathrm{Ca}}\tilde{p}_{\mathrm{Ca}}\tilde{q}_{\mathrm{Ca}}(E_{Ca} - V) + \left(g_{\mathrm{K,L}} + g_{\mathrm{K,i}}\tilde{q}_{\mathrm{K}} + g_{\mathrm{K,a}}\tilde{p}_{\mathrm{K}} + g_{\mathrm{K,ai}}\tilde{p}_{\mathrm{K}}\tilde{q}_{\mathrm{K}}\right)(E_K - V) = 0 \tag{10}$$

This can be interpreted as a balance between the current influxes created by the three gates. Hence, the equilibrium voltage must be placed in-between the lowest Nernst potential, $E_K = -90$ mV, and the highest , $E_{Ca} = 80$ mV. From physiological knowledge, we know that the equilibrium should be rather close to the -90 mV bound, since the entire purpose of the potassium gate is to create a resting state when the other gates are almost closed. Rewriting the equation to better reflect this fact, we get:

$$V = \frac{g_{\mathrm{Na}}\tilde{p}_{\mathrm{Na}}\tilde{q}_{\mathrm{Na}}E_{Na} + g_{\mathrm{Ca}}\tilde{p}_{\mathrm{Ca}}\tilde{q}_{\mathrm{Ca}}E_{Ca} + \left(g_{\mathrm{K,L}} + g_{\mathrm{K,i}}\tilde{q}_{\mathrm{K}} + g_{\mathrm{K,a}}\tilde{p}_{\mathrm{K}} + g_{\mathrm{K,ai}}\tilde{p}_{\mathrm{K}}\tilde{q}_{\mathrm{K}}\right)E_K}{g_{\mathrm{Na}}\tilde{p}_{\mathrm{Na}}\tilde{q}_{\mathrm{Na}} + g_{\mathrm{Ca}}\tilde{p}_{\mathrm{Ca}}\tilde{q}_{\mathrm{Ca}} + \left(g_{\mathrm{K,L}} + g_{\mathrm{K,i}}\tilde{q}_{\mathrm{K}} + g_{\mathrm{K,a}}\tilde{p}_{\mathrm{K}} + g_{\mathrm{K,ai}}\tilde{p}_{\mathrm{K}}\tilde{q}_{\mathrm{K}}\right)} \tag{11}$$

From here we see that V is a weighted combination of E_{Na}, E_{Ca}, E_K, with weights being determined by how open the gates are. Though we seemingly have solved the equation, the gating probabilities themselves depend on voltage. Since the equation is transcendental, it cannot be solved and hence cannot be solved exactly, but the second best is to recursively hover in on the answer. This can be done starting with $= E_K$, calculating the right-hand side in equation (11), and the repeating the process. The process rapidly converge to the equilibrium voltage $\tilde{V} = -88.7$ mV.

Having computed the equilibrium, the next question is what happens if we slightly deviate from it. The stumbling block for the attempt at solving the equations directly before was the

non-linearity. If the solution is restricted to a small neighborhood around the equilibrium, the non-linear terms are going to have no importance. The resulting linear system can be represented regarding a constant matrix whose eigenvalues is going to describe the asymptotic behavior of the system (decay to equilibrium if all eigenvalues have a negative real part, otherwise a rapid departure from the equilibrium). The Hartman-Grobmans theorem [18] guarantee that unless the real part of some eigenvalue is zero (which isn't the case), all conclusions about asymptotic behavior hold true even if the non-linear terms were included.

If the model is to be applicable at all, the equilibrium must be stable. Hence, all eigenvalues must have negative real part (called negative definite). To see this, one first note that the lack of interaction between the channels means that the matrix must contain only a diagonal (with all elements negative due to the convergence toward equilibrium) a first row and a first column, and the mirror pair of elements can therefore always be changed arbitrarily as long as the product is conserved by rescaling the variables. Secondly, the sum of a negative definite matrix and its transposed is still negative definite. If the mirror elements had different sign, this means that we could first rescale them to the same magnitude and then cancel them out by adding the transpose. This decouples that part of the matrix, and the eigenvalue is then found on the diagonal, with a negative value.

On the remaining block, one applies an extra negative sign in front and shows that the remaining matrix is positive definite. To do this, one applies Routh-Hurwitz criteria that require all sub-determinants to be positive. If one calculates these backward, saving the voltage for last, the computation is trivial since they are simple diagonal matrices. The only complicated one is the entire determinant. If one compute the Laplace expansion for the matrix with elements s_{ij} this gives:

$$\left(\prod_{i=1}^{N} s_{ii}\right)\left(1-\sum_{i=1}^{N}\frac{s_{1i}s_{i1}}{s_{ii}s_{11}}\right) \tag{12}$$

This means that the question of the positive definiteness can be reduced to the condition:

$$\sum_{i=1}^{N}\frac{s_{1i}s_{i1}}{s_{ii}} < s_{11} \tag{13}$$

This criterion can easily be checked by explicit computation of the ratios $\frac{s_{1i}s_{i1}}{s_{ii}}$ directly from the system matrix, and shows that all eigenvalues are indeed negative.

Notice that the proof of convergence of the equilibrium applies only to infinitesimal perturbations from the equilibrium. If the result were truly independent of the size of the perturbations, the cell could not be excited. Within the medical literature, the concept of a threshold voltage is well-known, a critical voltage beyond which the cell get excited. Strictly speaking, Picard-Lidelöfs theorem forbids such behavior due to continuity in initial conditions, but there can be a rapid, yet continuous, transition between the two behaviors. In other words, the maximal size of disturbance created by the perturbation will increase very rapidly as we move away from the equilibrium.

The eigenvalues themselves not only provide a guarantee of stability for the equilibrium but also contain information about the relaxation of the system. The smallest of the negative eigenvalues can be used to estimate the typical time for the system to relax, with our choice of parameters $1/\lambda_{\min} \approx 225$ ms, showing the relaxation is indeed very slow, as anticipated. The eigenvector provides us with information about which variables it is that is responsible for the slow converge. Normalizing the eigenvector so the voltage component is 1, we get:

Variable	Eigenvector component
V	1
p_{Na}	0.0018
q_{Na}	-1.7450
p_{Ca}	0.0005
q_{Ca}	-0.1283
p_K	31.9836
q_K	0.0023

We notice that among the gating variables, the largest one is p_K, which corresponds to it having the largest difference from the equilibrium in the slowest eigenmode to die out, hence it can be seen as the cause of the slow relaxation. This is reasonable and corresponds to what is known physiologically, being that the potassium gate is the slowest, and therefore take the longest time to return to equilibrium.

3.3 Numerical algorithm

With no exact solution available to the system of differential equations, we must turn to numerical analysis for solving the single cell equations. There are multiple general purpose methods for solving such systems of ordinary differential equations, such as the Runge-Kutta methods. However, before using such generic tool, it might be worth first to study if the equations have some special structure that can be taken advantage of in the construction of the numerical method.

If we study the equations, we notice three fascinating features. The first one is that even though each equation is non-linear, if all other variables are kept constant, each equation would instead be linear. The second realization is the while all the equations are coupled, the coupling is rather limited since each of the gating variables only depend on the voltage, but not on each other (since all the gates are independent of each others). The third observation is that each variable has negative feedback to itself, meaning that if all other variables are constant, it will converge to a fix-point. This is a consequence of the thermodynamic foundation the derivation of the equations rests upon.

This means that if the coupling was ignored, the fact that each equation then would be a linear differential equation means that they all would have closed form solutions. Furthermore, each of these sub-solutions would be unconditionally stable. Ignoring the coupling between the variables will, of course, generate an error, but for a short step in time Δt, the approximation will be acceptable. The entire process can then be repeated for the next step in time and so on. The method outlined here is known as operator splitting [19], and is a common technique for solving differential equations where each of the subparts is simpler to solve alone.

As pointed out above, the reason for the error is that when the voltage equation is solved, constant values are used for each of the gating variables, and vice versa. If the two processes was simply alternating, first solving the voltage equation a step Δt, then solving the gating equations a step Δt and then back to the voltage equation, the constant values used in the equation would always be either drawn from the beginning or the end of the time step, which is not very representative. As a result, such splitting procedure, known as Lie Splitting, will have a first order error in time [19].

A better solution would be to have the constant value drawn from the middle of the time step. This can be achieved by first taking a step $\Delta t/2$ with the voltage equation, and using the resulting value to solve the gating equations for a full step Δt. Finally, the new values for the gating variables are used to solve the voltage equation for a step $\Delta t/2$. In this way, the gating variables are solved using an approximation for the voltage at the middle of the time step, and the voltage is solved half of the time with the initial value and half of the time with the final values for the gating variables. The resulting scheme, known as Strang splitting, can be shown to have an error that is second order in time [19].

The next natural question is if an extended amount of trickery with the order the two sets of equations are solved can buy us even higher orders. If one always wants to take steps forward in time, the answer is, however, a resounding no, as has been proven multiple times in the literature [20]. If one accepts steps backward in time, the order can be increased indefinitely, but that means sacrificing the guarantee for unconditional stability. Another option is to take complex steps in time, but this has the disadvantage of slowing done the computation as complex numbers must now be used throughout the calculations that require twice the number of calculations and also makes the method very hard to interpret [21]. Hence, if we restrict ourselves to real steps forward in time, the second order barrier cannot be overcome directly.

At this point, it is natural to stop and reflect on what has been discussed. The proposed idea has been to turn down the generic ordinary differential equation solvers that exist and instead try to tailor a particular method using the unique features of the equation. This has the disadvantage that all the work that has gone into building algorithms like ODE45 in Matlab cannot be used. In fact, so advanced is such algorithm, that even with all the special features we have drawn upon, it will not be enough to outperform ODE45. Hence, if this thesis were solely about solving the single cell equation, we would have reached the end of the road. However, there are two reasons while it is worthwhile to follow the path we have outlined. First, the method chosen allow us to understand better the processes involved in solving each of them separately, gaining value understanding of the system behavior. Secondly, foreshadowing the coming chapter, the single cell problem is of very limited value for understanding heart tissue. In such model, we must deal with a continuum of cells, and the equations become non-linear partial differential equations instead. It is in this setting that the actual strength of operator splitting will be revealed, so the single cell case instead serves as a toy problem for which the procedure can be introduced.

3.4 Detailed derivation

Having finished with the overview of the method, the next step in using Strang splitting is to work out each operator. As we noted above, the two different processes are the conduction and the gating, and we can solve both of these exactly, as outlined below.

3.4.1 Conduction

The first sub-problem that will be dealt with is that of the voltage, which changes through channel conduction. The equation is derived by considering all gating variables set to constant values. By regrouping the terms in the Wohlfart-Arlock equation, the ODE for the voltage can be written:

$$\dot{V} = k(E - V) \tag{14}$$

Where k and E are constants. To provide a physiological interpretation for these constants, the exact expressions derived from the equation can be studied. In the case of k it is given by:

$$k = k_{Na} + k_{Ca} + k_K = \frac{\sigma_{Na}}{C} p_{Na} q_{Na} + \frac{\sigma_{Ca}}{C} p_{Ca} q_{Ca} + (1 - 0.75 q_K)\left(\frac{\sigma_K}{C} p_K + \frac{\sigma_{Kb}}{C}\right) \tag{15}$$

Hence, the rate with which the voltage changes is the sum of all the single channels rates that drive the cell to equilibrium. The relation means that the more ways the charge imbalance can be transported, the quicker the system will reach equilibrium. The expression for E is:

$$E = \frac{k_{Na} E_{Na} + k_{Ca} E_{Ca} + k_K E_K}{k_{Na} + k_{Ca} + k_K} \tag{16}$$

The equilibrium voltage is a weighted mean of all the Nernst voltages and is known as the Donnan potential [22]. If the rate of transport through the membrane for each ion would remain fixed, this is the voltage the system would settle in. A way to understand the weighted mean property is to study the case where the membrane is only permeable to a single type of ion. By definition, the system would then settle in that ion's Nernst voltage. With multiple types of ion, a compromise must, therefore, be met. Hence, the Donnan potential provides a very concise way of discussing the reasons for the change in the voltage. For example, the initial spike of the action potential can be explained by a shift in the Donnan potential due to the increased permeability of sodium ions through the membrane.

The solution of this well-known first-order equation when starting from $t = 0$ with a known $V(0)$ is:

$$V(\Delta t) = e^{-k\Delta t} \cdot V(0) + (1 - e^{-k\Delta t}) \cdot E \tag{17}$$

Since the factors in front of $V(0)$ and E sum to 1, $V(\Delta t)$ can be seen as an interpolation between the initial state and the Donnan state, with the weighting dependent on the length of time step.

3.4.2 Gating

The second sub-problem is that of the gating variables. As a proof of concept, the solution procedure for the sodium channel opening probability p_{Na} is demonstrated. The corresponding equation is $\dot{p}_{Na} = k_{p_{Na}}(\tilde{p}_{Na}(V) - p_{Na})$. By the Duhamel's principle the general solution a step $\Delta t/2$ forward in time is:

$$p_{Na}(t) = p_{Na}(0)e^{-k_{p_{Na}}\Delta t/2} + \int_0^{\Delta t/2} \tilde{p}_{Na}\left(V(t')\right) k_{p_{Na}} e^{-k_{p_{Na}}(\Delta t/2 - t)} dt' \tag{18}$$

The interpretation of this expression is very similar to that of the conduction equation, but the single equilibrium level is replaced by an exponential moving average of equilibriums $\tilde{p}_{Na}\big(V(t)\big)$. Since $V(t)$ is assumed to be frozen at the value $V(\Delta t/2)$, the integral can be computed yielding the result:

$$p_{Na}(t) = p_{Na}(0)e^{-k_{p_{Na}}\Delta t/2} + \tilde{p}_{Na}\big(V(\Delta t/2)\big)\big(1 - e^{-k\Delta t/2}\big) \tag{19}$$

As in the conducting case, the result is an interpolation between the initial gating value $p_{Na}(0)$ and the equilibrium value $\tilde{p}_{Na}\big(V(\Delta t/2)\big)$.).

3.4 Simulation of single cell

The described algorithm was coded in Matlab and consisted of a precomputation of all the interpolation coefficients, and a for loop that for each step computed the change in voltage and gating probabilities.

3.4.1 Initial simulation

Hence, with all the theoretical groundwork laid, it is finally time to study actual simulations of the single cell. To stimulate the cell, current must be injected. One could envision using different kinds of time-dependent current pulses $i(t)$ to achieve this, but experimentation showed that the essentially only important property of the pulse was if it was fast enough to trigger the opening of the sodium gate before the closing of the gate was activated. Once the sodium gate was active, the large flux of current made any further contribution from the pulse $i(t)$ almost neglectable, which also makes sense physiologically, the reason the electrical single is sustained is not that ions are transported between the cells, but rather that in each new cell getting activated there is a large in-rush of new ions. Therefore the simplest choice of current pulse is $i(t) = \big(V_0 - \tilde{V}\big)\delta(t)$, having the consequence of simply increasing the initial voltage from $\tilde{V}$ to V_0. With $V_0 = -60$ mV:

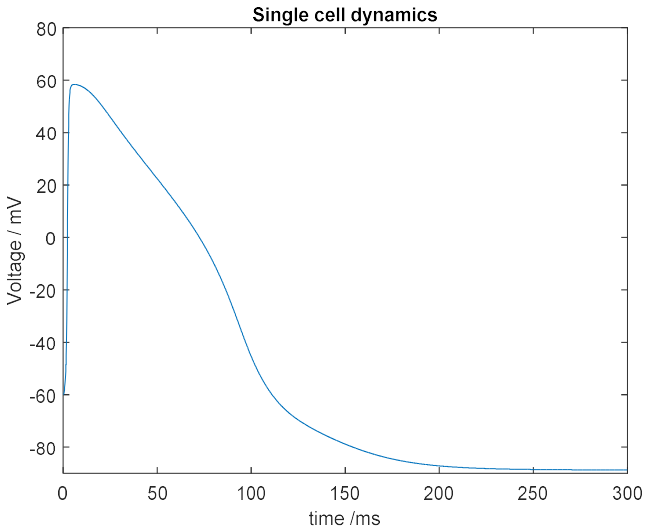

Figure 2: Simulation of single cell dynamics with initial voltage = -60 mV

The curve seen in figure 2 is the characteristic voltage curve seen by all previous work [5], [6], [7], [8], [9] on the Wohlfart-Arlock model for a single cell. It qualitatively corresponds well with what is physiologic known about cardiac cells. Figure 2 starts with a sharp increase in voltage caused by the influx of sodium ions. As the sodium gates starts closing up, the increase stops. Instead, potassium ions are flowing out of the cell, moving the voltage back to the resting value. This process is slowed down by the calcium ions flowing into the cell, creating the middle area where the voltage is rather high, but the downward slope is rather small. We can also see that the order of magnitude estimate for the relaxation of the cell from the eigenvalues of the linearized problem of 250 ms fits rather well. To verify this we can also study the gating probabilities as a function of time:

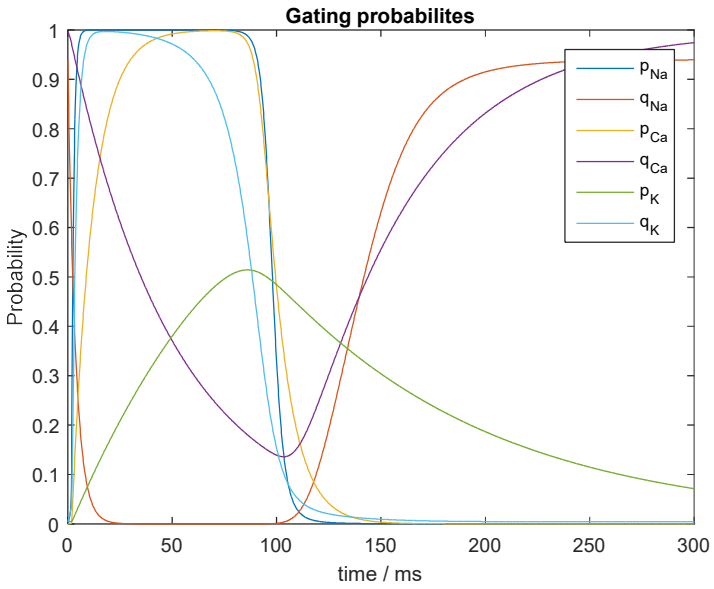

Figure 3: Simulation of single cell gating probabilities

As stated above, the red curve in figure 3 corresponding to the probability of the sodium gate to not be closed almost immediately drops, taking that gate out of the equation. The yellow curve shows how the calcium gate opens up, and the dropping of the purple curve shows it the closing of, and the product of the two being the probability to find the gate open. This shows that there is a rather precise time a short duration after the simulation that the calcium is playing a role in the problem. The potassium gate is the slowest to open, and once it has pushed the voltage down again, starts to close back. Even after 300 ms, it clearly has not reached its equilibrium, being the slowest process in the problem.

3.4.2 Comparing the qualitative properties – threshold voltage

In the derivation of the model, we mentioned that the cardiac cell has two important qualitative properties. The first one is that a certain threshold voltage must be exceeded for the cell to activate. To investigate if the Wohlfart-Arlock model exhibits such features V_0 was varied as can be seen in figure 4,5,6,7 below.

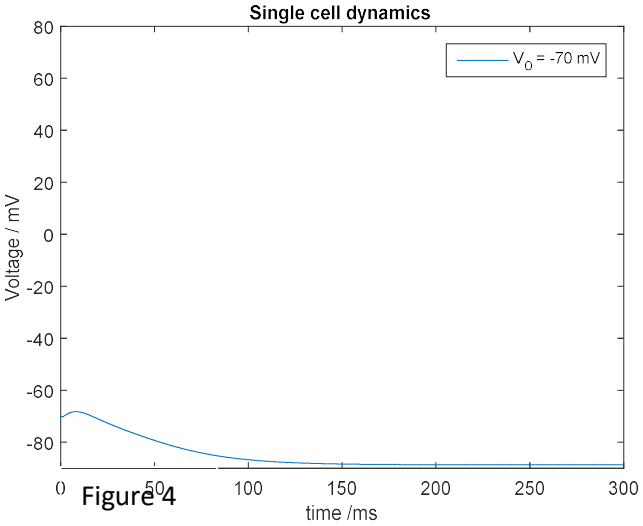

Figure 4

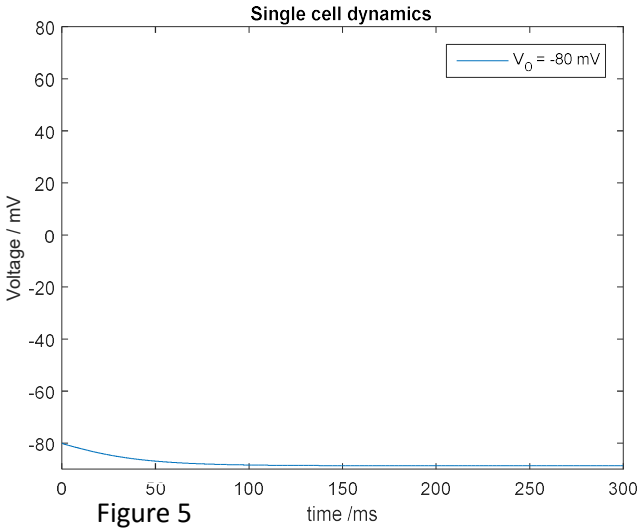

Figure 5

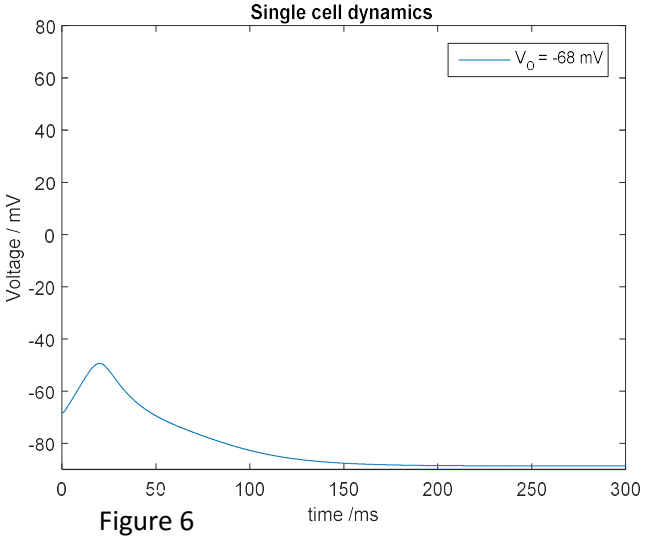

Figure 6

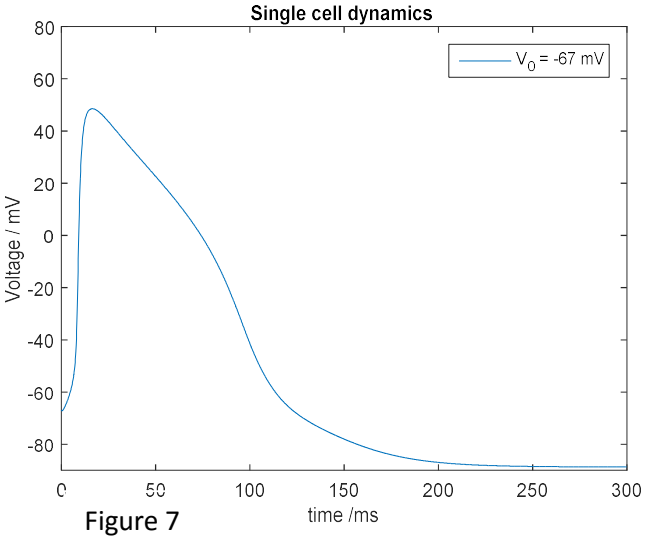

Figure 7

As can be seen in the figures 4,5,6,7, although there is no discontinuous threshold, the transition from a minuscule voltage pulse to a full excitation is very tiny. As the voltage increase, even more, the height of the pulse rises slightly, evening out at 60 mV, the Nernest

voltage of sodium. With a sodium channel, a cell can never raise higher than this, and the calcium channel is too slow to have any reasonable kind of chance to bring the voltage higher.

3.4.3 Comparing the qualitative properties – refractory time

The second property the cardiac cells are known to have is a refractory time. This means that if the current is injected a second time with a delay, a second voltage pulse will only be created if a certain time has passed. To investigate this, we use two delta pulse of height 40 mV with a variable delay t_d as can be seen in figure 8,9,10,11 below.

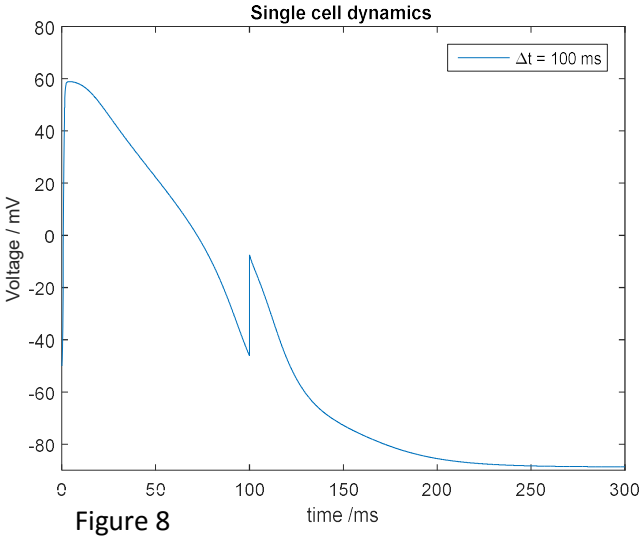

Figure 8

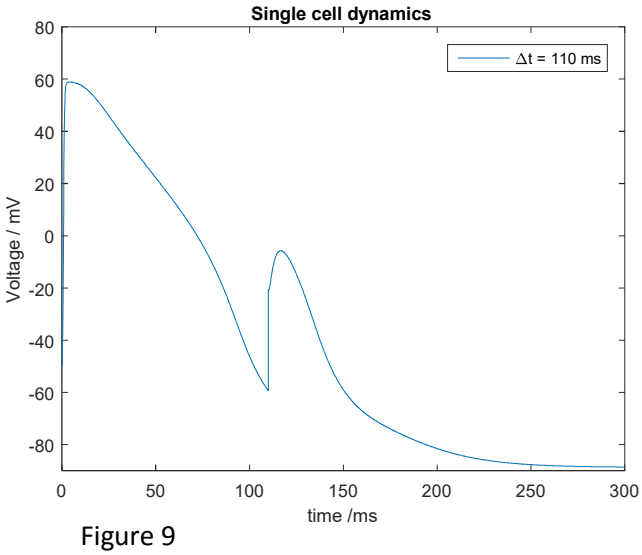

Figure 9

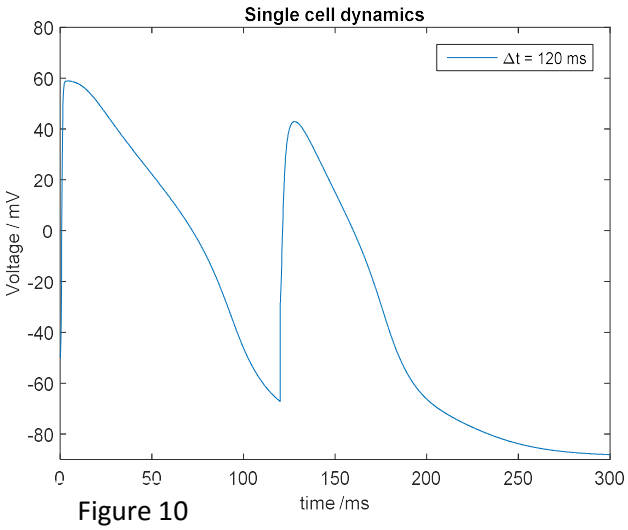

Figure 10

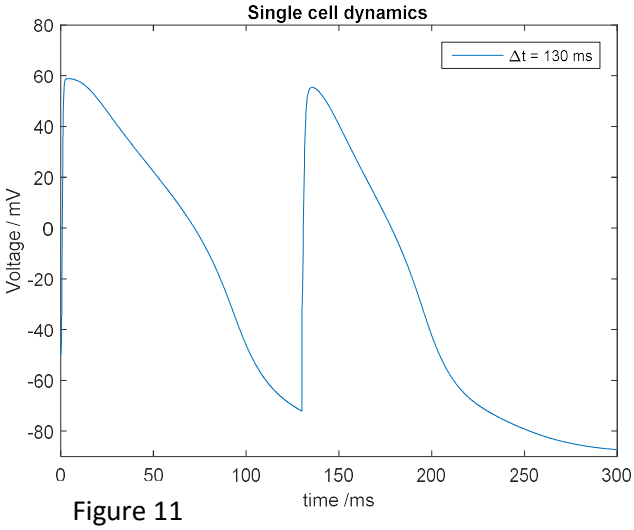

Figure 11

As we can see in figure 8,9,10,11, the refractory time is in the order of 110 ms. If we compare with the gating variables, we see that this is roughly the time for the sodium gate to open up again, which makes sense since this is the critical channel needed to start the process.

3.4.4 Analysis of numerical method

Having investigated the qualitative properties of the solution, we would like to verify that our numerical scheme performs adequately. It is worth noting that there is a glaring lack of such measures among earlier works, with consequences, we will further explore in later sections. As we can only obtain an approximation to the solution of the problem by using a numerical

method, it is critical to know how close the approximation is to the actual solution. From the theoretical analysis above we expect the error to scale as $O((\Delta t)^2)$. Before we make use of this to estimate the error, it is inded important to actually verify that this assumtion is correct.

From solving the problem with step size Δt we obtain a solution $V_1 = V + \epsilon(\Delta t)^2$, were V is the theoretical solution to the exact equations, an entity that is completely theoretical since we have no way of obtantaning it, and an error $\epsilon(\Delta t)^2$. Next step is to solve the problem again, but cutting each step up in two. This modified solution V_2 should then have an error four times as small, so $V_2 = V + \frac{1}{4}\epsilon(\Delta t)^2$. Note that there are two unknown in the problem, V and ϵ. This means that with only V_1 and V_2 we can only calculate the error $\epsilon(\Delta t)^2$, but cannot gurantee the scaling since we have two equations and two unknowns. To do this we need a third solution, created by cutting each step in four parts, yielding a solution $V_4 = V + \frac{1}{16}\epsilon(\Delta t)^2$. To check the scaling, the standard method, as described by [23] is to compute $X = \frac{V_2 - V_1}{V_4 - V_2}$. From the knowledge of the error scaling we obtain $X = 4$, so by comparing the computed values of X (as shown in figure 12) we see if the error scales as predicted.

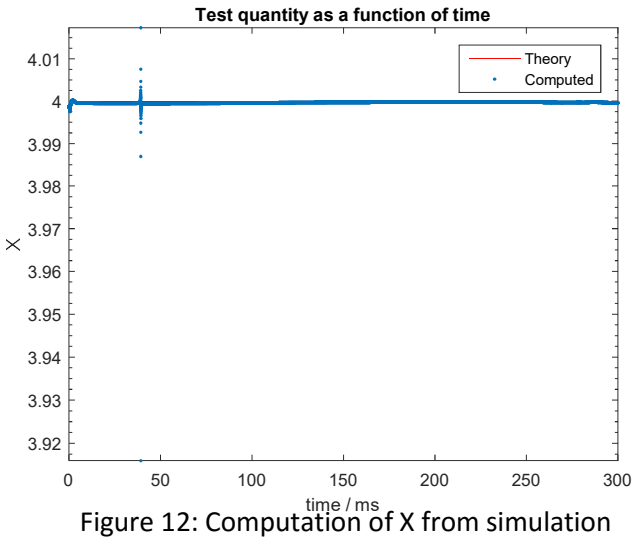

Figure 12: Computation of X from simulation

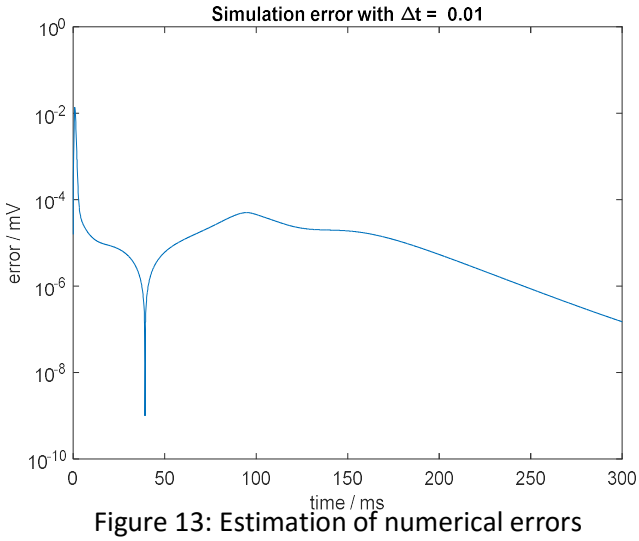

Figure 13: Estimation of numerical errors

Indeed, as seen in figure 12, the computed values stay very close to 4. To explain the deviations, one must keep in mind that apart for the quadratic error term, there are higher order terms that being left of the analysis, based on the argument that Δt is so small that no other terms are needed. It is clear that this assumption becomes less valid when the problem goes though very rapid phenomena, which actually require better resolution. This also means that any kind of improvement of the convergence by assuming the error scale as the largest term (such as Richard extrapolation) become more doubtful if we doesn't also reduce the step size.

4 Derivation of Wohlfart-Arlock equation for tissue

In this chapter, we derive the Wohlfart-Arlock for a section of cardiac tissue. While the equations themselves are well known from the previous works, all past derivation hinges on the fact that the cells are placed in a rectangular lattice, an assumption we show to be superfluous. Hence, the derivation is an original contribution of this work.

4.1 Internal equation

Having demonstrated that the Wohlfart-Arlock model for the single cardiac cell gives reasonable results, the next step is to study the interaction between the cells. As the gating variables refer to local attributes of each cell, they cannot affect each other. Instead, the only interaction between the cells is the exchange of ions. The electrical connections between the cells are known as gap junctions, and the ions go through a so-called connexin protein that forms a tunnel through which all ions can pass without selectivity. Since the concentrations inside the cells are approximated to be constant, only the excess charge particles will move, meaning that charge will distribute itself more equally among the cells. To a first approximation, these channels lack gating ability, so they remain open all the time. Furthermore utilizing the simplified linear model for the flux, the contribution to the time derivative of the voltage of cell i from cell j can be written $g_{ij}(V_j - V_i)$, with g_{ij} describing the rate at which the voltage equalize, hence the strength of the coupling between the cells.

Hence at the outset, the generalization from one cell to several might seem very simple, write up the equations for all cells and then add the correct coupling terms. This program suffers from three fatal problems. First and foremost must the distribution of couplings g_{ij} be modeled. This is a formidable task since the atrial cells are not simply arranged in any regular lattice, experiments have shown that each cell has an average of 12 neighbors, primarily concentrated to the front and back of the cell, and therefore it is very complicated to write down an explicit representation of the lattice. From a modeling purpose we, are uninterested about what is going on at this lattice scale a since it is the motion of the electrical waves that interest us. The second problem is of somewhat the opposite nature, over rather large length scales, the properties of the atrial tissue are not homogenous. Instead it has a fibrous nature were the direction of the fibers vary with position. Furthermore there a special high conductive pathway known as the Bachmann bundle that runs from the sinoatrial node to the left atrium and modeling these would require special measures which complicate the simple structure of the equations. To avoid these issues, one could restrict the modeling to a small section of the atrium were the fibers can be assumed to be oriented in the same direction, does allowing on to assume homogeneity, which greatly simply the resulting equations. The third problem is the sheer number of cells involved in the heart. Even if the area of interest is reduced, the number of cells will still be will be at least in the order 10^7, which is computationally very demanding, in fact, no simulation have yet been performed at this scale.

However, it has been said that a problem is just a misunderstood opportunity. Continuum modeling tells us that as long as the numbers of cells are large, the microscopic details will smooth out and not be crucial for understanding the macroscopic behavior. Hence, the voltage and gating variables of each cell can be replaced by a voltage field and gating fields with a value at every point. This can be seen as the limit of the equations when $a \to 0$. The key

question is then how to derive the corresponding field equations that describe the evolution of the fields. This is trivial for the gating variables since they are local, so one simply replaces each variable with the corresponding field. In the discrete case, the voltage equations was changed by adding the terms $g_{ij}(V_j - V_i)$. In the present case the cell instead interacts with a continuous surrounding, so there must be a corresponding term expressed in differential operators.

Focusing on $i = 0$ without loss of generality (due to the homogeneity assumption), the total interaction term can be written $\sum_j g_{0j}(V_j - V_0)$. Apart from homogeneity, we also assume the tissue has parity symmetry, meaning that $(x, y) \to -(x, y)$ leaves the tissue unchanged, which is experimentally supported [24]. This can be understood as staying on the fiber, but moving in the opposite direction. The symmetry means that the differential operator cannot contain any odd derivatives. This means that the existence of a cell in the forward direction is followed by a cell in the backward direction, so that the interaction can be written $\sum_{j>0} g_{0j}(V_j - 2V_0 + V_j)$. This can be expressed in the voltage field by Taylor approximation, and since the distance between the cells are proportional to a, the dimensionless factor can be absorbed into the coefficients c_n. If the direction vector describing the direction of gap junction is $\hat{e}_j$, this means that the differention operator along that direction is $(\hat{e}_j \cdot \nabla)$ so the series becomes:

$$\left[c_2 a^2 \sum_j (\hat{e}_j \cdot \nabla)^2 + c_4 a^4 \sum_j (\hat{e}_j \cdot \nabla)^4 + \cdots \right] V \tag{20}$$

Therefore, as $a \to 0$, the lowest order differential term that will describe the interaction is $\sum_j (\hat{e}_j \cdot \nabla)^2$, and therefore all higher order terms are neglected. If the coordinates of the vectors are $\hat{e}_j^x$ and $\hat{e}_j^y$, the operator can be explicitly written as $\left[\sum_j (\hat{e}_j^x)^2\right] \partial_{xx} + 2\left[\sum_j \hat{e}_j^x \hat{e}_j^y\right] \partial_{xy} + \left[\sum_j (\hat{e}_j^x)^2\right] \partial_{yy}$. By the spectral theorem, there exists a coordinate transformation which removes the cross term. Physically this corresponds to rotating the coordinates so that the faster fiber direction is taken as x-coordinate and the slower orthogonal path is taken as the y coordinate. Furthermore, by scaling each of the other coordinates independently, the coefficients can be set to unity, so that the interaction term is $\partial_{xx} + \partial_{yy} = \nabla^2$. The scaling means that if a square block of tissue is simulated, this actually corresponds to a rectangular one. This transformation means that although the operator ∇^2 is rotationally symmetric, this does not apply for the underlying tissue, so care must be taken when studying solutions with motion in both x and y. Setting the coefficients equal to 1 means that distance is measured in units of time. Although this might seem strange, this means that computed speeds of waves will be dimensionless. These can then be compared with experimentally measured speeds to determine the proper transformation back to normal units.

The result is a very simple interaction term. In some sense, this couples the cell to the continuum in the simplest way possible, as any lower order differential operators cannot be

used due to the asymmetry they induce. This is as far as one can come with a homogenous model; a more exact interaction must account for the variation in structure. However, the fact that experimentally speeds can be used to scale the results guarantees that the term gets approximately the right size. If only the interaction contribution is studied, the equation becomes $\partial_t V = \nabla^2 V$. This is the well-known diffusion equation, which describes that the voltage spread out and equalize over the cells.

4.2 Boundary conditions

The Wohlfart-Arlock model for cardiac tissue is a partial differential equation. Hence, both initial and boundary conditions are needed for the problem to have a unique solution. While the initial conditions are simply the initial voltage distribution, the boundary conditions are much more unspecified. On an infinite domain, there are no boundaries, so for simplicity, we can use the same conditions in x as in y. Focusing on x, the natural choice is that $\lim_{x\to\pm\infty} V(x,t) = 0$, since the cells far ahead have not yet been excited, and those far behind have all but return to equilibrium. However, infinite domains make very limited sense in the numerical world. Instead, x must be limited to a finite interval $[0, L]$. The goal of the boundary conditions is to try to mimic an infinite domain.

These boundary conditions must be homogenous as nothing is going on in the outskirts. Since the equation is second order in space, such conditions can be written as a linear combination of $V(0,t), \partial_x V(0,t), V(L,t), \partial_x V(L,t)$. For a second order equation, two boundary conditions are needed. Therefore both boundaries must be involved, as the only homogenous condition with only one boundary is $V(0,t) = 0$, $\partial_x V(0,t) = 0$ that yields no excitation.

To resolve these problems, there is a need for local boundary conditions, the first being a function of only $V(0,t)$ and $\partial_x V(0,t)$, and the second of only $V(L,t)$ and $\partial_x V(L,t)$. The first option is the Dirichlet boundary condition $V(0,t) = V(L,t) = 0$, making the boundary an endless current sink. This will be very unphysical, especially when the cells close to the boundary reach their peak of excitation, since there will then be a massive current gradient though the boundary. Rather than mimicking an infinite area, such choice of conditions reinforces an artificial boundary.

A second alternative is the so-called Robin boundary condition defined by $\partial_x V(0,t) = \gamma V(0,t)$. From the equilibrium analysis of the traveling wave ansatz we know that such relation is satisfied for both boundaries asymptotically with different values of γ for the left and right boundary. There are however a number of problems with such choice of boundary condition. Since the condition is only valid asymptotically, it means that it will not be valid when the wave is close to either end, which is exactly when the boundary conditions actually does play a role in determining the shape of the wave. Furthermore the fact that γ must be select differently means that the isotropic properties of the medium will be violated. Finally, it is very unclear how to properly extend such idea to a 2D setting.

This leaves only the Neumann boundary condition, $\partial_x V(0,t) = \partial_x V(L,t) = 0$. Such condition corresponds to an insulated boundary, no current can flow in or out of it, meaning that as the excitation wave come close to the boundary, the current that would have flown

though the boundary must instead flow back. While this sounds highly artificial, it is important to keep in mind that the current that flows between the cells serves the purpose of an igniting the excitation, but is relatively small compared to the inflow of currents from the ion gates. Another way of understanding this is by noting that Neumann boundary condition are sometimes called "reflective boundary conditions" in that a wave sent against such boundary will be reflect. For the cardiac cells however,, the reflected wave will not be able to propagate, as the cells have already been excited and is recovering, meaning that the wave will disappear over the boundary, like it would have done had the boundary not existed, exactly the behaviour we have been searching for all along. Finally, these boundary conditions are very pleasant to work with in a 2D setting, as a plane waves traveling parallel to the x-axis satisfy $\partial_y V(x,y,t) = 0$, since there is no dependence on y. If we apply the same conditions in both x and y, we hence gain $\partial_x V(0,y,t) = \partial_x V(L,y,t) = 0$ and $\partial_y V(x,0,t) = \partial_y V(x,L,t) = 0$, which was exactly what we wanted.

4.3 Collecting results

Summarizing all the results, the single cell equation and the Laplacian term from the charge diffusion can be added together to yield the non-linear partial differential equation that is known as the Wohlfart-Arlock model for cardiac tissue:

$$\begin{cases} \dfrac{\partial V}{\partial t} = \nabla^2 \mathrm{V} + g_{\mathrm{Na}} p_{\mathrm{Na}} q_{\mathrm{Na}} (E_{Na} - V) + g_{\mathrm{Ca}} p_{\mathrm{Ca}} q_{\mathrm{Ca}} (E_{Ca} - V) + \\ \qquad \left(g_{\mathrm{K,L}} + g_{\mathrm{K,i}} q_{\mathrm{K}} + g_{\mathrm{K,a}} p_{\mathrm{K}} + g_{\mathrm{K,ai}} p_{\mathrm{K}} q_{\mathrm{K}}\right)(E_K - V) \\ \dfrac{\partial p_{Na}}{\partial t} = k_{p_{Na}} (\tilde{p}_{Na} - p_{Na}) \\ \dfrac{\partial q_{Na}}{\partial t} = k_{q_{Na}} (\tilde{q}_{Na} - q_{Na}) \\ \dfrac{\partial p_{Ca}}{\partial t} = k_{p_{Ca}} (\tilde{p}_{Ca} - p_{Ca}) \\ \dfrac{\partial q_{Ca}}{\partial t} = k_{q_{Ca}} (\tilde{q}_{Ca} - q_{Ca}) \\ \dfrac{\partial p_K}{\partial t} = k_{p_K} (\tilde{p}_K - p_K) \\ \dfrac{\partial q_K}{\partial t} = k_{q_K} (\tilde{q}_K - q_K) \end{cases} \tag{21}$$

5 Solution of the Wohlfart-Arlock equation for tissue

In this chapter, the partial differential equation derived in the last section (equation 21) for the Wohlfart-Arlock model in tissue is solved. While previous works have done this, the novelty lies in the numerical method developed specifically to solve this problem, as all previous works have utilized the Euler-method.

5.1 Introduction

The last section extended the single cell equation to a partial differential equation (PDE) for the entire tissue. Precisely as in the case of the single cell equation, it is necessary to guarantee that the problem is well-posed before any attempt at finding a solution can be made. The Picard-Lindelöf's theorem can be generalized [19], as long as the solution *a priori* was known to be bounded, which (as discussed above) is the case. Therefore, a unique solution continuous in the initial conditions can be guaranteed to exist at any point time.

5.2 Ideas behind the numerical method

From previous work such as [8], we know that the standard "beating" solution of the equation is a traveling front in an isotropic space, so we can without a loss of generality make the front's propagation direction parallel with the x-axis. Hence, the voltage will only depend on x and not on y, and we only have a 1D PDE to solve. It is important to notice that the motivation for searching for such theory is to account for fronts, there is no intrinsic reason for studying 1D cardiac tissue, since unlike nerve signals, the cardiac tissue always have the local form of a sheet.

As noted in the single cell case, no closed form solution exist that is of practical use, so we must turn to numerical analysis. Previous works on the Wohlfart-Arlock model have heavily utilized the *Method Of Lines*, an idea based on first discretizing the space into a lattice of N^2 square pieces of side length a (the choice of squares is motivated by the symmetry of the Laplace operator induced by the scaling in the previous section). By replacing the Laplace operator by central finite differences, a system of $7N^2$ non-linear coupled ODE's is constructed. This operation allows for the application of the ODE theory discussed in the previous section. The standard approach used to solve these equations has been using the explicit Euler method, where the right hand side in all equations have been approximated to a constant value for each time step, and then trivially integrated to yield the new values.

This deceptively simple scheme suffers from the problem that while a PDE formally appears to be very similar to an ODE, it is actually of a very different nature. More precisely, there is a very intimate relation between of lattice spacing a and the length of the time-step Δt. For the non-linear equation such a relationship is rather complicated to work out, but studying only the diffusion part of the equation, the bound $\Delta t < a^2/2$ can be derived for the method to be stable. Keeping in mind that after the rescaling of space is preformed in the derivation, it is expected that in terms of the actual dynamics $\Delta t \sim a$ so the use of the Euler scheme forces unphysical short steps in time. Any explicit method that computes future values from the current ones is going to suffer from similar problems [25].

An alternative would be to switch to an implicit method (such as implicit Euler or Crank-Nicolson). There is a large variety of such schemes which are unconditionally stable, hence from the point of stability (but, of course, not regarding accuracy) there are no restrictions on Δt in relation to a. Alas there is no free lunch, since in such a scheme, one is faced with the problem of solving at least N^2 nonlinear equations in each timestep (the equations for the gating variables can be represented in terms of the voltage, leaving only one equation per lattice point), a computationally rather complicated problem. As if this result was not depressing enough, it can be shown that if solutions obtained by these finite difference methods (explicit or implicit) should contain no unphysical "undershoots" it can be of most second order, and for all methods beyond first order $\Delta t \propto a^2$ apply (so while stability can be guaranteed in the implicit case, the physicality of the solution cannot).

The root of all these problems comes from introducing the effect of charge diffusion into the equations. In contrast, to the very local single cell effects (that only act on the same grid point) diffusion very rapidly spreads out among different grid points. It would, therefore, be very advantageous to use different methods to attack each of these terms since they stem from very different physics. This is exactly the kind of problems the method of operator splitting was built to handle. To include this third process, two things are needed. First, a solver for the diffusion equation is needed, and secondly, the new process must be included together with voltage conduction and gating so that the second order accuracy is preserved.

5.2.1 Diffusion

We begin with the first task, that of solving $\partial_t V = \partial_{xx} V$. Together with the boundary conditions, solving this problem is a classical exercise in partial differential equations. Using separation of variables, the eigenfunctions of the operator ∂_{xx} will be $\cos(n\pi x/\ )$ with eigenvalues $-(n\pi/L)^2$, hence the solution can be written as an infinite series of the form:

$$V(x,t) = \sum_{n=0}^{\infty} c_n \exp\left(-\frac{n^2\pi^2 t}{L^2}\right)\cos(n\pi x/L) \tag{22}$$

Here c_n are coefficients that describe how much of the initial state is in each of the eigenmodes. These can be determined by setting $t = 0$, which gives the equation:

$$V(x,0) = \sum_{n=0}^{\infty} c_n \cos(n\pi x/L) \tag{23}$$

If the left-hand side is a known function that satisfies the boundary conditions (and also has the regularity necessary actually to be a solution to the Wohlfart-Arlock equations), then we know that the function can be expanded in a cosine series, and hence all the coefficients can be determined. However, as there is an infinite number of modes, so must we also know the initial state is an infinite number of points on the interval. In the numerical world, there is no such thing. Instead, we only know the voltage on N lattice points x_i on the interval. This means that we have an underdetermined set of equations:

$$V(x_i, 0) = \sum_{n=0}^{\infty} c_n \cos(n\pi x_i / L) \tag{24}$$

If we interpret this problem in matrix form, with N known values of V, it can be treated as a $N \times 1$ vector, and c as a $\infty \times 1$ vector. Any kind of selection relation can then in general written $c = \tilde{S}V$ for some matrix $\tilde{S}$ of size $\infty \times N$ that determines how to transform voltage data into coefficients. In similar fashion, a diagonal matrix $D(\Delta t)$ can be introduced with $\exp\left(-\frac{n^2\pi^2 t}{L^2}\right)$ on the diagonal and finally a matrix S of size $N \times \infty$ with matrix elements $S_{in} = \cos(n\pi x_i / L)$, so that the new voltage can be expressed in terms of the old as $SD(\Delta t)\tilde{S}V$. To determine $\tilde{S}$, some kind of further criteria is required.

For the purpose of operator splitting, we want to minimize the amount of error of each part of the solution. For the ODE part, we could solve the problem exactly, which expressed in another way means that if only that solution had been used, taking two steps in time of half the size should produce the same result as a single step. Such a condition is known as the flow condition. Applying this to the diffusion solution, means that

$$SD\left(\frac{\Delta t}{2}\right)\tilde{S}SD\left(\frac{\Delta t}{2}\right)\tilde{S} = SD(\Delta t)\tilde{S} \tag{25}$$

Since $D(\Delta t)$ is a diagonal matrix, $D\left(\frac{\Delta t}{2}\right)D\left(\frac{\Delta t}{2}\right) = D(\Delta t)$. Hence at first sight, we would like $\tilde{S}S = I$, so that $\tilde{S} = S^{-1}$, but that is not possible since S is not invertible (being a non square matrix). However, if $\tilde{S}$ was the inverse to the first $N \times N$ block in S, and zero for all other elements, the relation would be true, since $\tilde{S}S = I$ would be valid for the $N \times N$ block, and for all else the relation would be 0 on both sides since we premultiply with $\tilde{S}$.

Hence, what we have shown is that if we assume $c_n = 0$ for $n \geq N$, that is cutting of all higher modes, we can satisfy the requirement of independence of the step size. The reason for picking the N first modes and no else is twofold: Lower modes decay slower, so will stay important longer, and higher modes will in general carry information about details of the function while the lower modes describe the essential shape. This idea of using the first eigenmodes of the system is known as a spectral method [19]. Hence in the end the problem simplifies into the square system of equations:

$$V(x_i, 0) = \sum_{n=0}^{N-1} c_n \cos(n\pi x_i / L) \tag{26}$$

The next question is how to select x_i. While we have assumed $c_n = 0$ for $n \geq N$ to be able to perform calculations, this will not be true in practice. Hence the most significant error is most likely coming from the c_N term. If x_i is chosen so that $\cos(N\pi x_i / L) = 0$, at least that mode will not interfere. Solving the equation gives $x_i = \frac{i-1/2}{N}L$ for $1 \leq i \leq N$, so the lattice points should be chosen equidistantly. Notice that N must be selected at the beginning of the

computation, meaning that we no longer have the freedom to vary the step size as in the case of the single cell equation.

This might seem like an obvious result, and the reasoning unnecessarily complicated, but it will become essential in further on when we move onto higher different geometries were other eigenfunctions apply.

If we restrict all matrices to their $N \times N$ block, the new voltage can then be written $SD(\Delta t)S^{-1}V$. An extra nice property is that the columns in S is orthogonal, so if the columns are also normalized, the matrix becomes an ortonormal matrix $S^{-1} = S^T$. The transformation S^{-1} and S are in the literature known as the discrete cosine transform and the inverse discrete cosine transform. These can be found as standard routines in Matlab. Since all matrices are all independent of time as long as Δt remains fixed, they can all be precomputed into the matrix $K(\Delta t) = SD(\Delta t)S^{-1}$, called the propagator, so that the entire diffusion part of the problem is reduced to multiplying with a matrix. Since these matrices must be quite large for reasonable accuracy, this will be the most computer intensive step of the calculation, but the good news is that matrix multiplication is precisely the task that Matlab was built to perform best.

Although we have eliminated errors related to the step in time by making the result independent of the steps were taken in time, errors will occur since not all modes were included. Therefore, the size of the sum $\sum_{n=N}^{\infty} c_n \cos(n\pi \ _i/L)$ of excluded modes must be estimated. Using the triangle inequality, these can be bounded by $\sum_{n=N}^{\infty}|c_n|$. To determine how fast the coefficients decay, we need some knowledge of the regularity of the function. From the general theory of the equation we can assume the function to be at least analytical, meaning that the coefficients will be bounded geometrically $|c_n| \leq C\rho^n$ for some C and $0 < \rho < 1$ [26]. Inserting the bound gives:

$$\sum_{n=N}^{\infty}|c_n| \leq C\sum_{n=N}^{\infty}\rho^n = \frac{C}{1-\rho}\rho^N \tag{27}$$

Since the distance between the lattice points are $a = \frac{1}{NL}$, the error is asymptotically on the form $O\left(\rho^{\frac{1}{aL}}\right)$, hence going to 0 faster than any power of a, the trademark of spectral methods. This means that as a and Δt shrinks down together; the error from the splitting will dominate the spectral error. Note also that the proposed method is unconditionally stable, there are no longer any bounds on Δt and a. This opens up the possibility to construct a second order method, without the constraint $\Delta t \propto a^2$.

5.2.2 Reconciliation

With the diffusion subproblem solved, this process should be incorporated into the splitting procedure already including the voltage conduction and gating variables. The first option would be to mix the order in which all the three processes take part, but there is an interesting shortcut that saves us from such increase in complexity. It stems from that the diffusion equation part, $\partial_t V = \nabla^2 V$, has no coupling with the gating variables, so these two can be solved together. The gating variables however does depend on the voltage, meaning that that

part of the solution must be modified. Returning back to the basics, the formal gating solution is given by:

$$p_{Na}(t) = p_{Na}(0)e^{-k_{p_{Na}}\Delta t} + \int_0^{\Delta t} \tilde{p}_{Na}\left(V(t')\right) k_{p_{Na}} e^{-k_{p_{Na}}(\Delta t - t)} dt \tag{28}$$

This time $V(t')$ is no longer a constant, but will be the change in voltage due to diffusion. The problem is we only know $V(0)$, before the application of the diffusion, and $V(\Delta t)$, after the diffusion. However since the splitting error is of second order, $p_{Na}(t)$ does not have to be solved exactly, but only to second order accuracy, there is simply no need to strain out a gnat if a camel has already been swallowed. This approximation can be achieved by linearly interpolating $\tilde{p}_{Na}\big(V(s)\big)$ between the values $\tilde{p}_{Na}\big(V(0)\big)$ and $\tilde{p}_{Na}\big(V(\Delta t)\big)$ with the ansatz $\tilde{p}_{Na}\big(V(t)\big) = \left(1 - \frac{t}{\Delta t}\right) \cdot \tilde{p}_{Na}\big(V(0)\big) + \frac{t}{\Delta t} \cdot \tilde{p}_{Na}\big(V(\Delta t)\big)$, and by inserting this function into the integral, the gating variable at t can be computed as:

$$\begin{aligned} p_{Na}(t) = e^{-k_{p_{Na}}\Delta t} p_{Na}(0) + \left[\frac{1 - \left(1 + k_{p_{Na}}\Delta t\right)e^{-k_{p_{Na}}\Delta t}}{k_{p_{Na}}\Delta t}\right] \tilde{p}_{Na}\big(V(0)\big) \\ + \left[1 - \frac{1 - e^{-k_{p_{Na}}\Delta t}}{k_{p_{Na}}\Delta t}\right] \tilde{p}_{Na}\big(V(\Delta t)\big) \end{aligned} \tag{29}$$

As in the previous case, the result is an interpolation between the initial gating value $p_{Na}(0)$ and the equilibrium, but this time there are two equilibrium levels, the initial equilibrium $\tilde{p}_{Na}\big(V(0)\big)$ and the final equilibrium $\tilde{p}_{Na}\big(V(t)\big)$. Since $k_{p_{Na}}\Delta t$ is constant, the interpolation coefficients are the same everywhere on the lattice and can be pre-computed before the calculations to save time. Furthermore, if p_{Na} is represented as a matrix, Matlab allows the computation to be performed in a vectorized way without the need for any loops, resulting in both clarity and speed.

5.3 Computing the traveling wave solution

With all the theory and numerical analysis on our side, there is only one problem, how to select initial conditions? One might think that this is a very delicate process of finding a proper starting state. However, it turns out to be very simple. Experimentation showed that more or less any pulse-shaped form with sufficient amplitude and width will work; in the example below in figure 14 we used a Gaussian pulse, placing the maximum at $x = 0$. The maximum is $V = 60$ mV (sodium channel equilibrium voltage) and the width is 10 ms (length is measured is time as mentioned earlier).

As soon as the first simulation is done, the problem of initial conditions becomes rather mute, since we now have access to the traveling wave profile directly as can be seen in figure 15. By saving the data from the previous simulation and importing it into a new one, we can start the system directly in the right conditions.

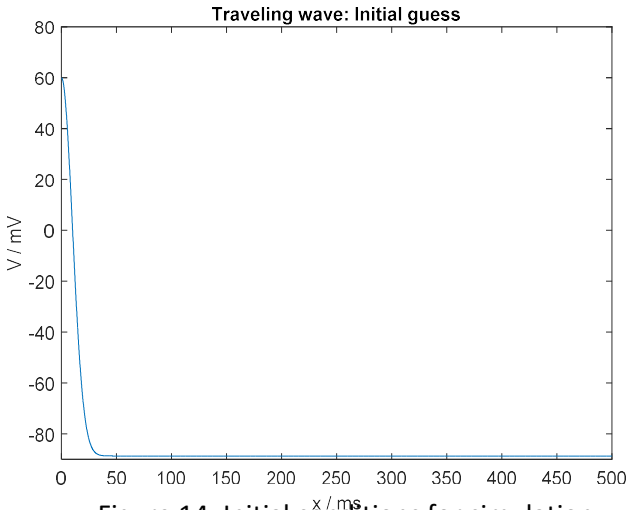

Figure 14: Initial conditions for simulation

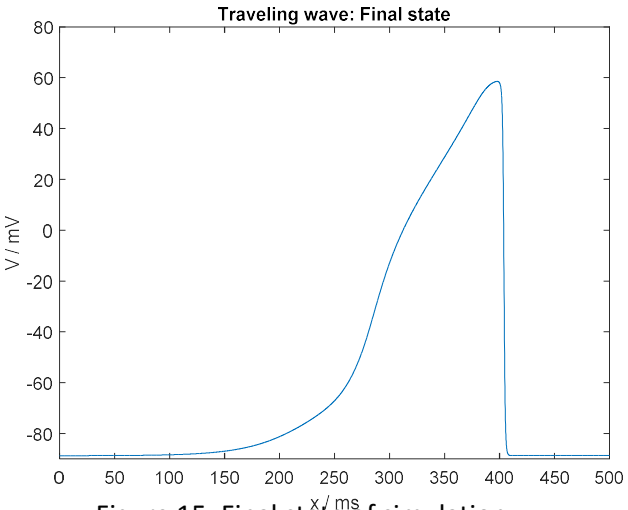

Figure 15: Final state of simulation

Looking at the final state of the travelling wave in figure 15 we see a shape very similar to the action potential of the single cell, which is a clear indication that the internal dynamics cause the majority of the cells behaviour (open of the gates est.) and diffusive flow of ions between the cells primarily serve just to trigger the cells. The reason the action potential is mirrored is because moving backward in distance means moving forward in time since a particular cell was triggered.

Previous works such as [8] have assumed the wave speed c to be constant. To verify this, we initiated the system with a traveling wave solution and measured the position of the maximum as a function of time as shown below in figure 16. The agreement with a straight line was so good that the limiting factor was the resolution of the position of the peak since it can only take discrete values. A more sophisticated version introduced to solve this was to find the maximum and use the two adjacent lattice points to fit a parabola, but the difference was barely noticeable.

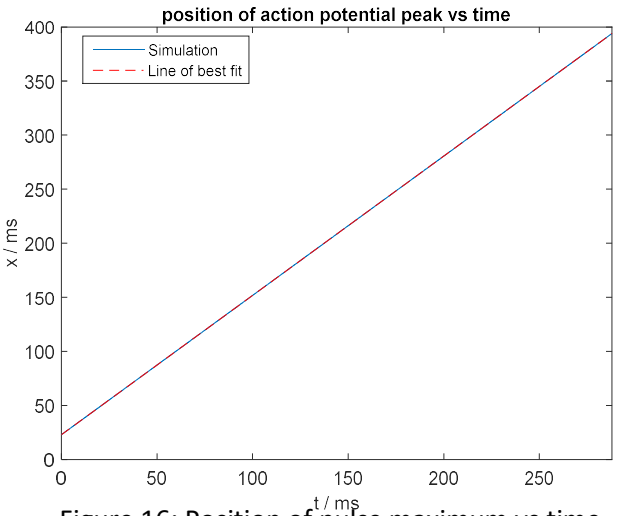

Figure 16: Position of pulse maximum vs time

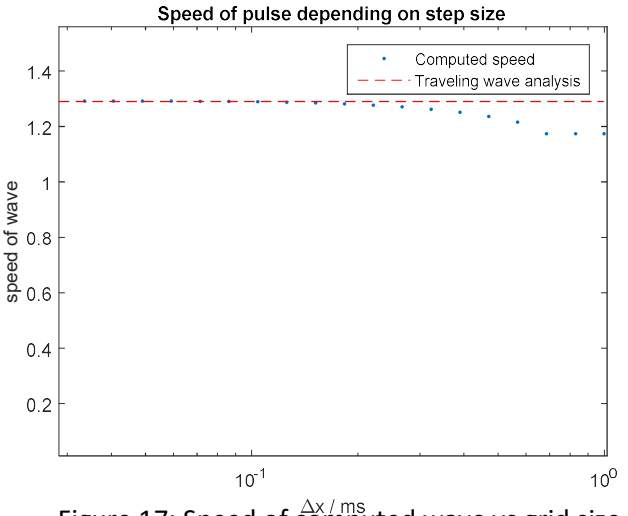

Figure 17: Speed of computed wave vs grid size

A much more interesting question is how the computed speed depend on the step sizes Δt and Δx. Since we model an object moving with constant velocity, it is reasonable to couple these steps by $\Delta x = c\Delta t$ (using the c from the highest resolution computation), but the results we get was also checked with other constants of proportionality between the two step sizes. From the relationship between step size and wave speed seen in figure 17 above we see that $\Delta x \lesssim 0.1$ ms to yield correct results. This is a rather restrictive criterion, but not very surprising given that almost the entire dynamic is described by the sodium gate, a very fleeing event that must be captured accurately. The guideline concerning step-length is especially important considering that there exists no prior analysis of this in the previous works. Just because the solution is stable, doesn't need to mean that it corresponds to the solution of the differential equation contrary to the conclusion of [10]. If one use step size $\Delta x = 1$, one will still obtain a rather reasonable action potential that will both look and feel right, but careful analysis shows that it will be moving with the wrong speed, and there is no guarantee it will behave correctly. This cast a very problematic shadow over several of the past results such [8] and [10], we will return to some of these when studying multi excitation states.

As discussed in the case of the single cell equation, a proper way to perform error analysis is to vary Δt by cutting up steps and checking that the ratio off V_1, V_2, V_4 is 4. This can be seen to be true in figure 18 below, and allow us to estimate the error of the simulation as $\frac{4}{3}(V_2 - V_1)$ as plotted in figure 19. We see that the error is extremely small except for the activation and sodium gate opening.

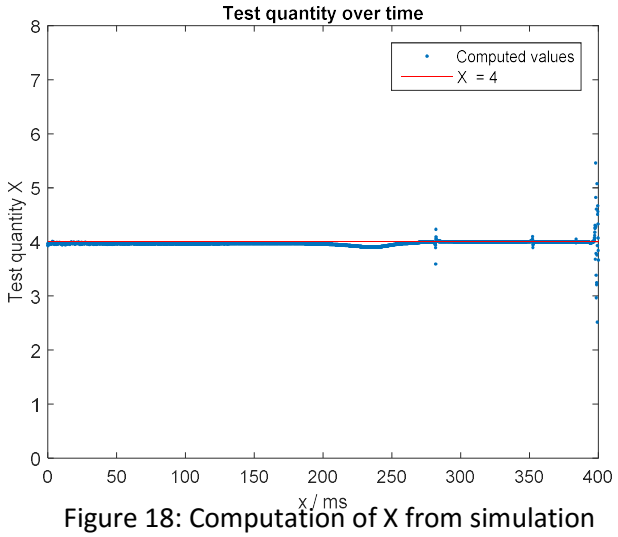

Figure 18: Computation of X from simulation

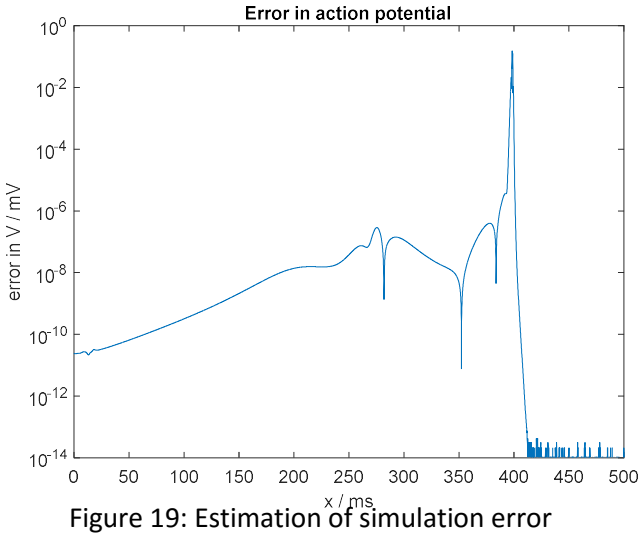

Figure 19: Estimation of simulation error

A natural question is if one could somehow adjust the step sizes to obtain better resolution of this critical area. While such idea might seem very tempting, there are indeed multiple problems with it. First and foremost, in such idea a lot of the symmetries and pre-computation properties must be dropped since the grid itself must be moving and changing. Secondly, while it is possible to imagine how to perform this in 1D with a single wave, if we like to study some arbitrary voltage configuration this soon becomes very complicated. Attempts have been made in the literature to account for this [27], but they have been rather unsuccessful for more general waveforms.

5.4 Gibbs phenomena and different interpolation schemes

One of the central ideas in the numerical method was to use cosine to interpolate the voltage when computing the diffusion. This choice was only which satisfied the flow condition that made the solution operator into a semigroup, but it is not without problems. The opening of the sodium channels gives a very sharp gradient, and it is well known that such is rather hard to treat using spectral methods. One artificial phenomenon appearing is the so-called Gibbs phenomena were the interpolating function is oscillating in an unphysical way after a sharp gradient. To study this phenomenon, we used $\Delta x = 0.1$ and then made an interpolation using only every third value. This allow us to plot the "actual" solution vs. the interpolated one shown in figure 20. While they agree rather well for the most part, just before the gradient front, we get this behaviour:

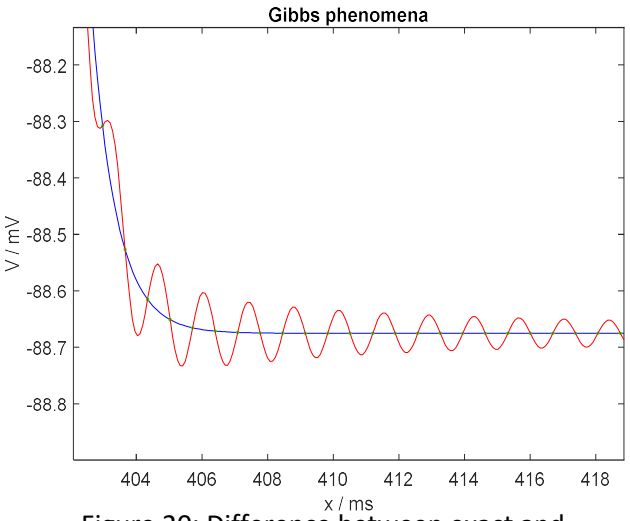

Figure 20: Difference between exact and interpolated solutions

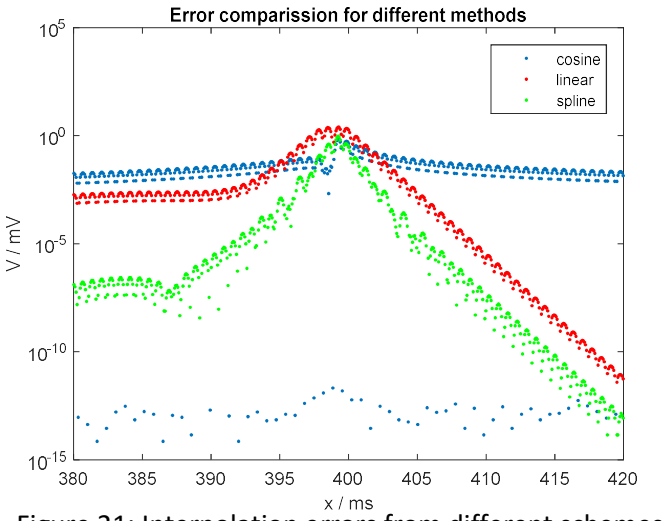

Figure 21: Interpolation errors from different schemes

To investigate better options, we also tried linear and spline interpolation. As we can see from figure 21, these clear outperforms the cosine interpolation far away from the step gradient. However, we know that at these areas the calculations are already really precise. At the most critical region, that of the sharp gradient, the cosine method is slightly better. Also, notice that the large errors are not as bad in practice, the interpolations functions are only used for computing the diffusion current, and the current total contribution from oscillations are minuscule.

6 Modeling the influx of currents

In this section, we demonstrate how the novel solution procedure from the last chapter can be used to solve the so-called "influx problem" that has not received any satisfactory treatment in the previous literature.

6.1 Numerical method

As seen in earlier chapters, any excitation wave that resembles the traveling wave solution will quickly converge toward the later. Hence, in practice, it is very easy to obtain this solution. There are however other solutions that we would also like to generate. Previous works [7] have tried to simulate an excitation quickly followed by a second one as a method of studying the properties of the equations. While we do not know the final state, the question is how to generate a sufficiently close approximation to it. The method previously used has been to inject voltage directly into one of the lattice points, with a time delay between the injections. Direct reference to the grid itself makes the scheme very dependent on the grid spacing, a very undesirable feature. The procedure is also very artificial since the two pulses would not be created this way.

If we imagine the scenario of an extra pulse being sent out after the first one, we would expect the two pulses to have traveled for quite a while until reaching the right part of the atria that we are studying. Hence, we would like the pulse to "enter" the area from the left of the interval. The way a pulse is transported is through the current transport between the cells that opens the ion channels. Hence, if a current influx similar though that of the traveling way could be simulated, it would be possible to inject a pulse in such a way that it would already be in a converged state. If two such current pulses were injected with a time delay, it would hopefully bring us close to the two wave state.

The current i is given by $-\partial_x V$ in the non-dimensionalized units. Hence the current of the excitation front can be estimated by a finite difference. If the current is measured between the two grid points n and $n+1$, $\partial_x V = \frac{V_n - V_{n+1}}{a} + O(a^2)$. Since $a \propto \Delta t$, the estimate is still second order accurate. A higher order stencil could be used, but the next two grid points would be at $\pm\frac{3}{2}a$, and given how the front has a very steep slope (which is the region primarily contributing to the current), the lower order scheme was preferred. This way, the current function $i(t)$ can be computed.

The next step is to inject the current. The left boundary condition is given by $\partial_x V(0,t) = 0$, but with current injection, it instead becomes $\partial_x V(0,t) = i(t)$. Instead of changing the entire solution of the diffusion problem, we can use the linearity of the diffusion equation to solve the part with the inhomogeneous boundary condition separately. Hence we have the partial differential equation problem:

$$\begin{cases} \partial_t V = \partial_{xx} V \\ \partial_x V(0,t) = i(t) \\ \partial_x V(L,t) = 0 \\ V(x,0) = 0 \end{cases} \tag{30}$$

We seek to compute $V(x, \Delta t)$. Since the system is linear, it must by Schwartz Kernel theorem be possible to write the answer in the form $V(x, \Delta t) = \int_0^{\Delta t} h(x, \Delta t - t) i(t) dt$ with the kernel function $h(x, t)$. This function describes how much excess voltage exists at position x a time t after a pulse of current, hence it describes the response to an influx of current. Inserting this ansatz into the problem and then Laplace transforming the equation, collecting all terms and finally applying the inverse transform using calculus of residues then yields:

$$h(x,t) = \frac{1}{L}\left[1 + \frac{1}{2}\sum_{n=1}^{\infty} \exp\left(-\frac{n^2\pi^2 t}{L^2}\right)\cos\left(\frac{n\pi x}{L}\right)\right] \tag{31}$$

To preserve the flow condition, the sum must be truncated at $n = N - 1$. If $i(t)$ was known for all times, the integral could have been computed, but only one current value per time step is known, so we only know the current at $i(0)$ and $i(\Delta t)$. Using the recipe from before, and linearly interpolating $i(t)$ between the two end values, gives us:

$$V(x, \Delta t) = \alpha_1 i(0) + \alpha_2 i(\Delta t) \tag{32}$$

$$\begin{cases} \alpha_1 = \dfrac{L^3}{2\pi^4 \Delta t}\displaystyle\sum_{n=1}^{N-1} \dfrac{1 - \left(1 + \dfrac{n^2\pi^2\Delta t}{L^2}\right)\exp\left(-\dfrac{n^2\pi^2\Delta t}{L^2}\right)}{n^4}\cos\left(\dfrac{n\pi x}{L}\right) \\ \alpha_2 = \dfrac{L^3}{2\pi^4 \Delta t}\displaystyle\sum_{n=1}^{N-1} \dfrac{\exp\left(-\dfrac{n^2\pi^2\Delta t}{L^2}\right) - 1 + \dfrac{n^2\pi^2\Delta t}{L^2}}{n^4}\cos\left(\dfrac{n\pi x}{L}\right) \end{cases} \tag{33}$$

While the coefficients in front of the current values are rather cumbersome, they can be computed once and for all, and the extra voltage contribution is then simply added to the KV term during the diffusion step.

6.2 Estimating the current

The first step of the program is to estimate the current i from a travelling action potential using the grid size $\Delta x = 0.1$ (according with the recommendation from the last section). The result can be seen in figure 22. The sharp peak corresponds to the sodium influx, and the undershoot at $t = 170$ ms corresponds the outflow of potassium ions brining the voltage back. The small current in-between the peak and the undershoot comes from the balance between the calcium and potassium currents. Hence, a correct current profile clearly requires all three currents to be reproduced.

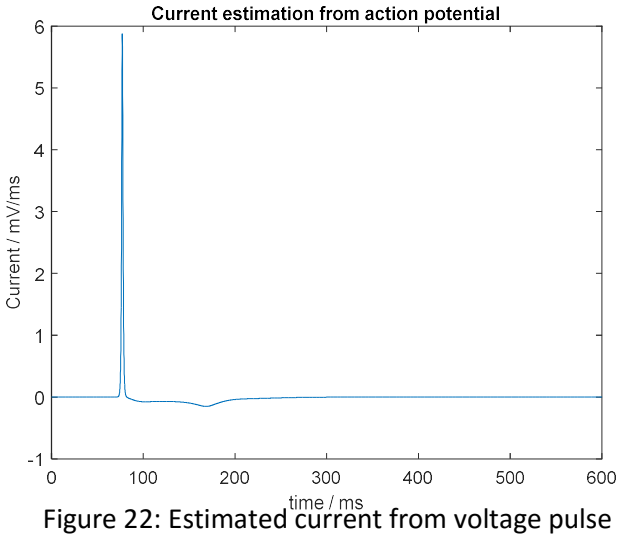

Figure 22: Estimated current from voltage pulse

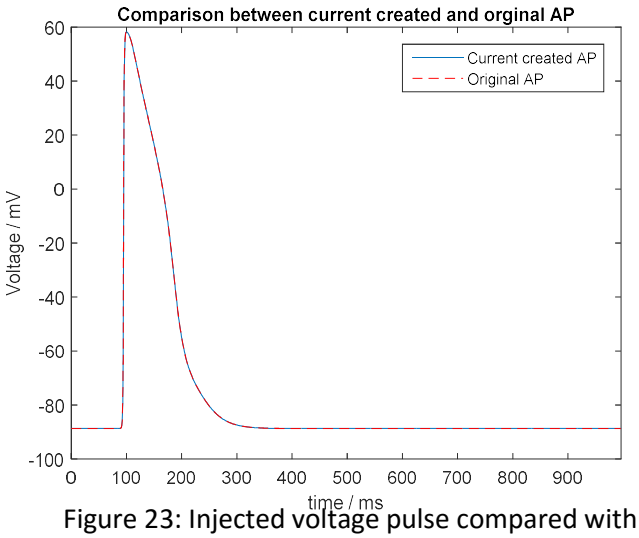

Figure 23: Injected voltage pulse compared with original

In figure 23, we see the result of the injected current by comparing the calculated waveform with the one from the 1D simulation. The agreement is excellent, as expected from the errors in both simulations are of the same magnitude.

6.3 Two pulse state

Having built a tool for injecting action potentials, the next question is where to apply it. As already hinted in the introduction to this section, the main use is multiple excitation states, in which the cells are excited again soon for the first time. We did a similar thing for the single cell, revealing that there is a refractory period, as expected. By doing the same kind of investigation here as seen in figure 24,25,26,27, we could change the time between two current pulses to see if the injected state converges to a stable state or not.

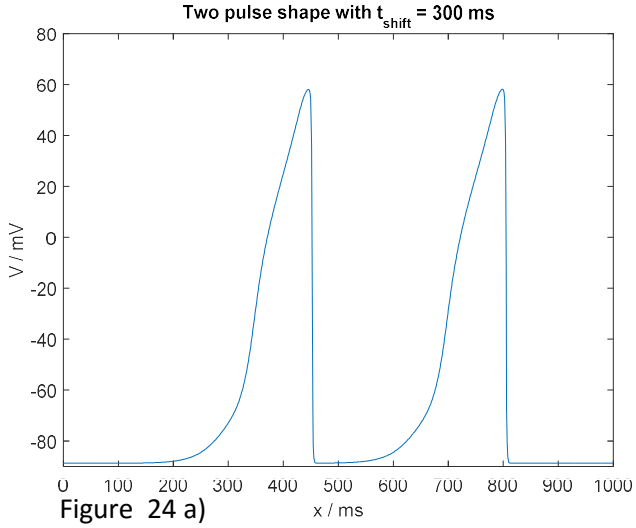

Figure 24 a)

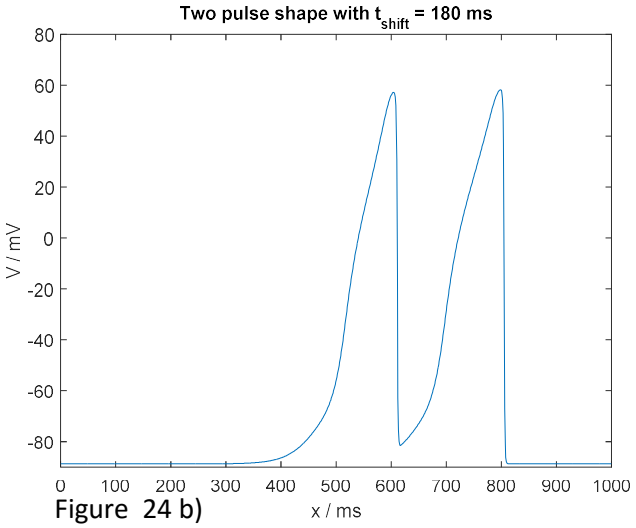

Figure 24 b)

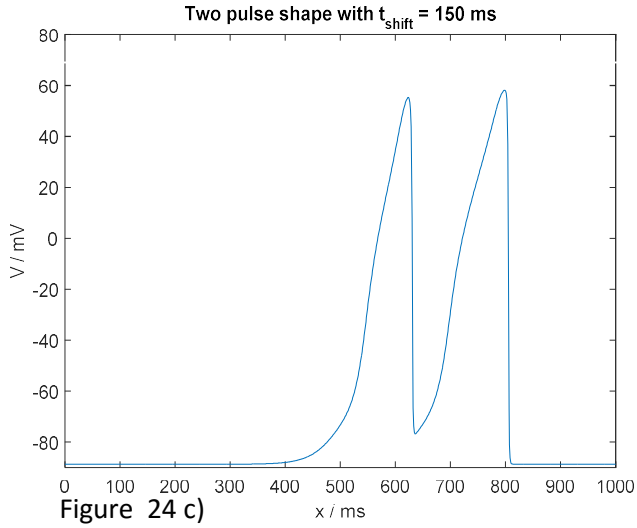

Figure 24 c)

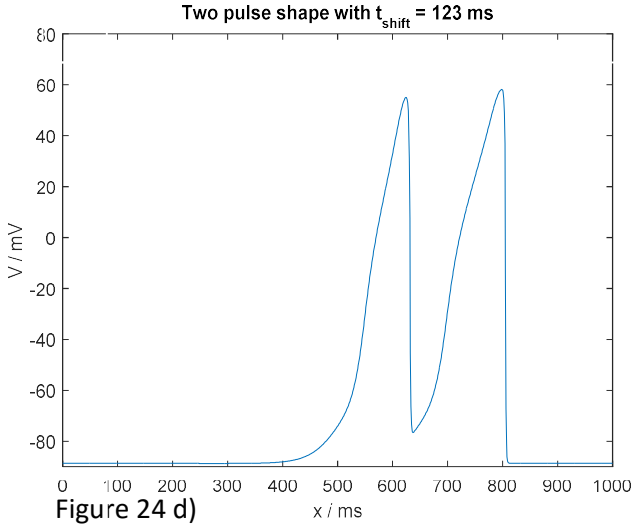

Figure 24 d)

From the different attempts in figure 24 a) – d), outlined as examples above, we see that as the pulses get closer, the second one contracts slightly and becomes slightly lower. The lowering comes from the sodium gate being partially closed combined with an increase in potassium influx, and the narrowing comes from the calcium gate no longer being able to provide sufficient outflux of calcium ions to slow down the decay of the voltage. All of these effects become stronger the closer the two pulses are. For times differences shorter than 123 ms, it was not possible to initiate the system. The natural question is if this sharp transition contradicts the statement that all transition must be continuous. To answer this we must realize that in the ODE case we could tune the time difference directly, but in the PDE case, we must instead create an approximate starting state by injecting currents. The fact that the current function is not identical to the true current is enough to break the continuity.

Another interesting use of the method is to compare with previous works on the equation. It is worth noting that while the amplitude of the second pulse is lower, this is a very slight decrease. In [7] we instead see a rather significant drop. The most likely reason is the combined use of an unphysical way to initiate the pulses and the low order low-resolution Euler method used. This provides further incentive to scrutinize earlier results and try to replicate them using the new numerical method.

7 Outlook

By this point, it is natural to ask the academic question, what have we learned? A reader who expected an in-depth analysis of the arrhythmic patterns and spirals of the equation, or a fair treatment of chaotic behavior or similar, might by this point be rather disappointed. Does this mean that the thesis has missed the target, and failed to accomplish the given task? To answer this, we should rather focus on the current state of the field of research surrounding the Wohlfart-Arloch equations. Several of the previous works has focused much attention to arrhythmias, given that describing these was one of the primary goals of the model. The main outcome from these collective works is to demonstrate the wide spectrum of behaviors exhibited.

However, all of them suffer from a chronic lack of systematic numerical analysis. Even in the case when the authors have considered things like step size and accuracy, the methods and standards of evaluation has been questionable at best. The numerical analysis has in almost all cases started and ended with the Euler method, which as we have discussed rather at depth, is not particularly suitable for this kind of problem. The partial blame falls on the continuity effect; newer works simply follow in the footstep of older ones, using the same methods and ideas. This highlights the critical need to break this chain of poor numerical analysis and outline a plausible alternative.

With this in mind, we now start to approach the red thread in this thesis, to provide an algorithm for solving the equations. We argue throughout the thesis for the use of operator splitting methods. These have the advantages of unlocking a lot of the theory for the diffusion equation by decoupling it from the non-linear, but local, terms. The fact that many different kinds of scenarios have closed form solutions for the diffusion equation unlocks a whole new world of possibilities for creating realistic, accurate and simple solutions.

Another natural question at this point is if the entire business of finding numerical methods and solving the equations is relevant from the physics point of view. Should we not simply treat numerics as a black box, and simply worry about performance? For many different kinds of equations this is the case, but as argued and showed multiple times, by interpreting the numerical method in terms of the physics, we receive a better understanding of the underlying physics. The concept of Donner level, for example, arise very naturally though the splitting, and so does the injection current diffusion and propagators. These concepts are hard to impossible to distinguish in the Euler scheme, but using operator splitting they appear as natural parts of the analysis.

Hence, we arrive back to the starting question, in what way does this work serve the field of research? The goal has been to form a stepping stone, a springboard for future work within this area to build upon. Any future researchers seeking to contribute could then hopefully save the time and effort to go beyond the Euler scheme, and directly be introduced to the next level of tools. Hopefully, this could both potentially refine the algorithm further, and perform new investigations. At the end of the day, what we have achieved is open up a little window from which the world of arrhythmias can be seen and marveled upon.

Acknowledgment
I would like to thank my thesis advisor Gunnar Ohlén for proposing this project and providing access to previous research. I would also like to thank both advisor and examiner for their patience with this work. Finally I would like to thank my parents for their continual and unconditional support.

Appendix: Parameter values

The model parameters, with values from [8]:

Parameter	Value	Description
$V_{p_{Na}}$	-40 mV	Voltage when $p_{Na} = 1/2$
$V_{q_{Na}}$	-75 mV	Voltage when $q_{Na} = 1/2$
V_{Na}	10 mV	Thermal voltage for sodium gate
E_{Na}	60 mV	Nernst potential for sodium
g_{Na}	8 ms^{-1}	Sodium ionic rate
$k_{p_{Na}}$	1 ms	Time constant for p_{Na}
$k_{q_{Na}}$	0.27 ms	Time constant for q_{Na}
$V_{p_{Ca}}$	-30 mV	Voltage when $p_{Ca} = 1/2$
$V_{q_{Ca}}$	-60 mV	Voltage when $q_{Ca} = 1/2$
V_{Ca}	10 mV	Thermal voltage for calcium gate
E_{Ca}	80 mV	Nernst potential for calcium
g_{Ca}	0.04 ms^{-1}	Calcium ionic rate
$k_{p_{Ca}}$	0.1 ms	Time constant for p_{Ca}
$k_{q_{Na}}$	0.02 ms	Time constant for q_{Ca}
V_{p_K}	-20 mV	Voltage when $p_K = 1/2$
V_{q_K}	-20 mV	Voltage when $q_K = 1/2$
V_K	12.5 mV	Thermal voltage for potassium gate
E_K	-90 mV	Nernst potential potassium
$g_{\mathrm{K,L}}$	0.05 ms^{-1}	Constant potassium ionic rate
$g_{\mathrm{K,i}}$	0.038 ms^{-1}	Potassium ionic rate for inactivation
$g_{\mathrm{K,a}}$	0.02 ms^{-1}	Potassium ionic rate for activation
$g_{\mathrm{K,ai}}$	0.015 ms^{-1}	Potassium ionic rate for activation/inactivation interaction
k_{p_K}	0.01 ms	Time constant for p_K
k_{q_K}	0.5 ms	Time constant for q_K

Appendix A: Solving the equations in two dimensions

A.1 Numerical analysis

So far all the domains the equations have been solved in have in one way or another been one-dimensional. While the cardiac tissue always truly has been two- dimensional, and there being no reason to consider otherwise, we have been able to cut down on the dimension by considering some simplification of the full 2D motion. The normal 1D corresponded to a plane wave along the x-axis. The radial case corresponded to a wave that was created by radial symmetry. Finally, it is time to break these restrictions and consider full 2D dynamics.

There are two primary differences between 1D and 2D. The first is that the different state variables are no longer vectors but must instead be stored in matrices. For the gating and conduction part of the problem, the only difference is that when we update, we instead add matrices. In Matlab, which treat vectors as matrices, there isn´t even any difference in the code. The second difference is that the Laplacian $\square_{\square\square}$ change into $\square_{\square\square} + \square_{\square\square}$. Just as we could derive a propagator that relates the voltage at $\square_{\square}$ with the voltage at $\square_{\square}$ a time $\Delta\square$ later in the 1D case, the same could be done in 2D, relating the coordinate $(\square_{\square}, \square_{\square})$ with $(\square_{\square}, \square_{\square})$. However, if we like to keep the spacing from the 1D case, with $\square$ points over the length $\square$, there are $\square^2$pairs that can be related. Meaning that the transition matrix would have $\square^4$ entries, all too many to be effectively stored (in fact anything beyond $\square^2$ is non-trivial with our grid size).

Faced with this problem, we can yet again result in operator splitting, splitting the Laplacian into $\square_{\square\square}$ and $\square_{\square\square}$. If the problem is solved on a rectangular grid, $\square_{\square\square}$ and $\square_{\square\square}$ is commuting. This means that operator splitting is exact, so one doesn´t even need to use the Strang splitting, but can simply use sequential splitting, first updating with respect to x and thereafter with respect to y. Going one step further, on a square grid, the 1D propagator matrix $\square$ becomes identical for both x and y since $\square$ is a symmetric matrix, so instead of computing diffusion though $\square\square$ as in 1D, we instead use $\square\square\square$ in 2D, resulting in minimal change between the 1D and 2D case.

It is worth noting however that although the similarity, the time consumption rises from $\square^2$ for matrix – vector multiplication to $\square^3$ for matrix – matrix multiplication. In general, none of the matrices doesn't really have some kind of further structure that allows us to cut time, $\square$ for example might be relatively spare, but still uses roughly 200 elements of the 1000 per row.

A.2 Simulating the 2D system

With the numeric's set up, we can finally solve the system in two dimensions. The first step is to use the 1D solution and verify that behaves correctly also in 2D. Such a simulation can be seen below. Notice the plane front and the sharp borders due to the boundary conditions.

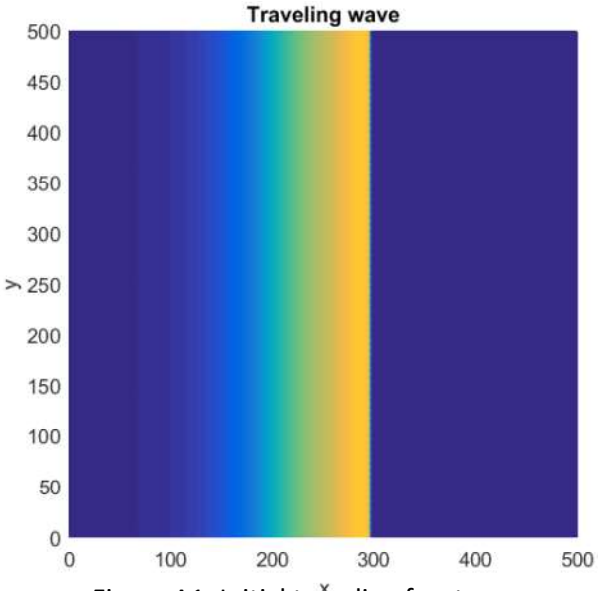

Figure A1: Initial traveling front

Next step is to replicate the spiral solutions seen in previous works [7], [9]. A good way of creating a spiral is to reset the upper half plane on the plane way. Simulations of these were run for several hours without any sign that the spirals were not stable. As can be seen in the series of images below, the resulting dynamics will create a spiral solution.

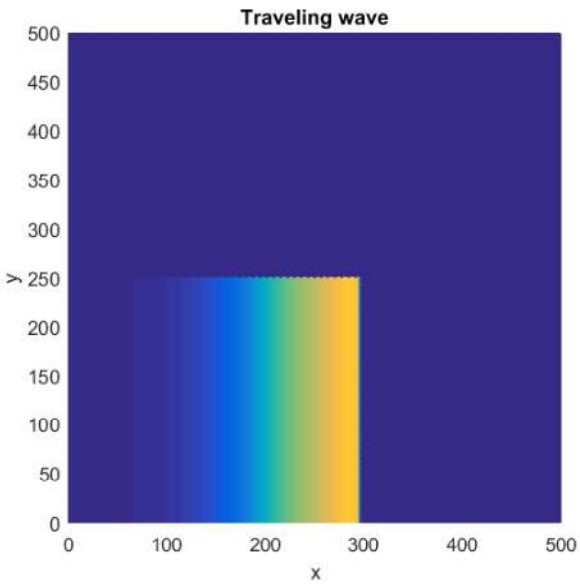

Figure A2: Resetting the upper half-plane

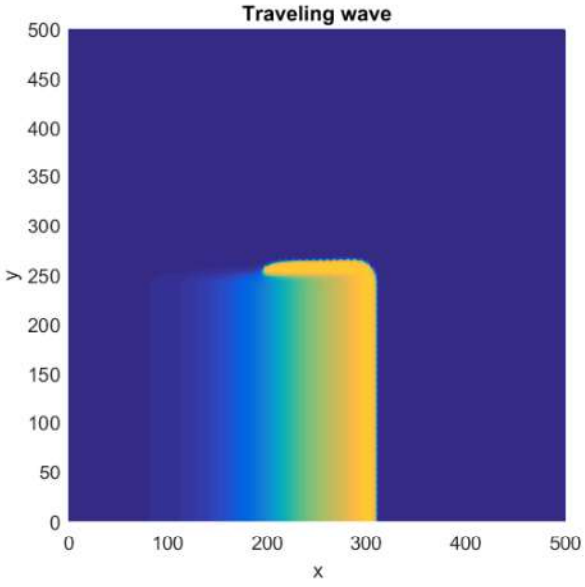

Figure A3: Creation of spiral, step 1

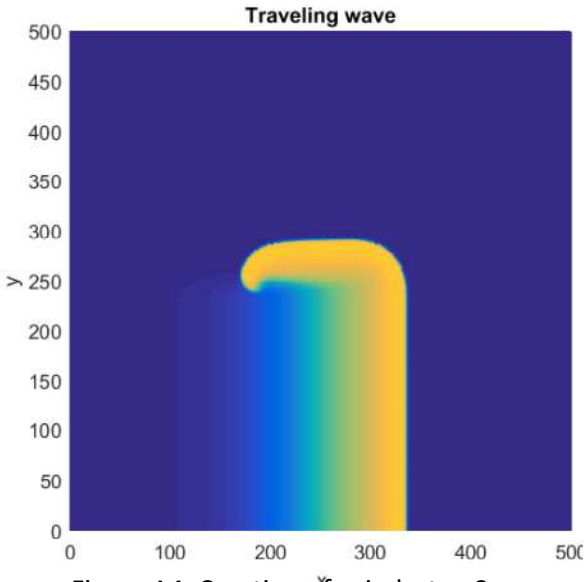

Figure A4: Creation of spiral, step 2

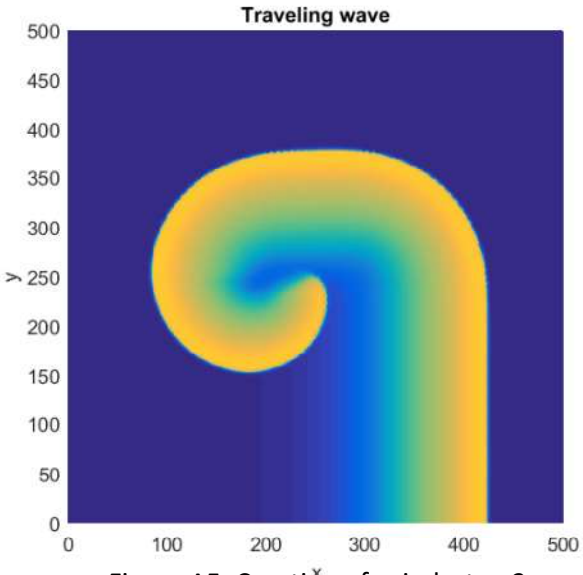

Figure A5: Creation of spiral, step 3

A double type of spiral can be constructed by cutting off both ends of the plane wavefront since the resulting solution then will exhibit reflection symmetry over the axis parallel to the front's propagation direction. The result is shown in the pictures below:

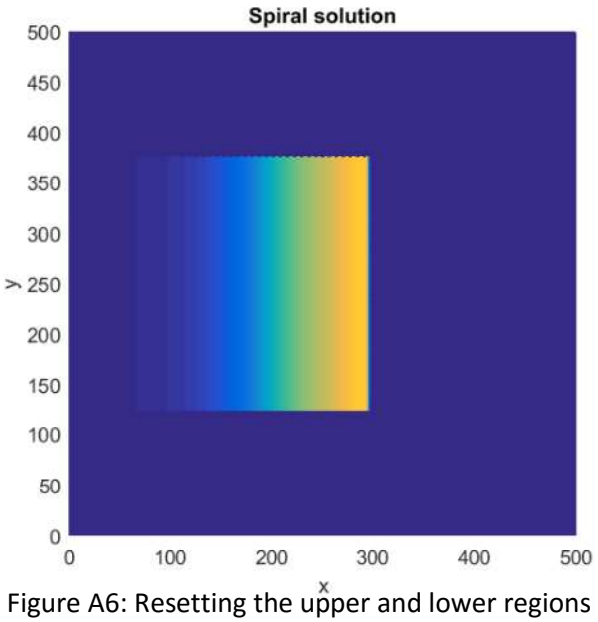

Figure A6: Resetting the upper and lower regions

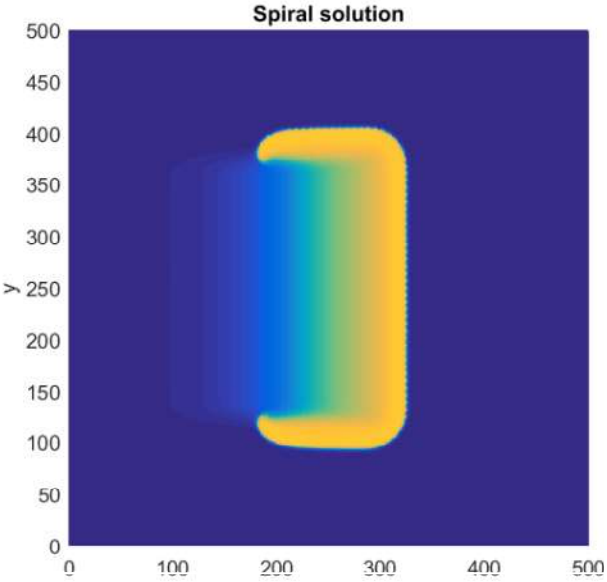

Figure A7: Formation of double spiral, step 1

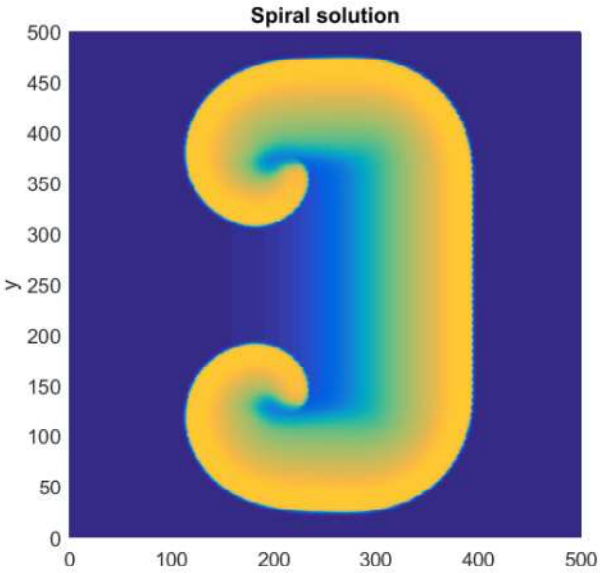

Figure A8: Formation of double spiral, step 2

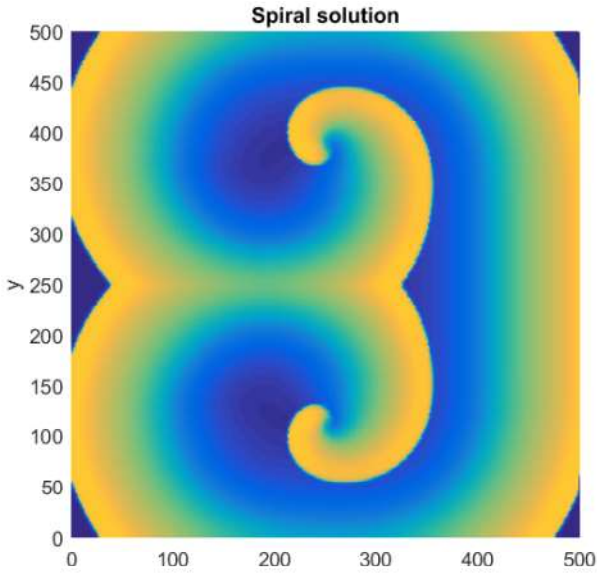

Figure A9 Formation of double spiral, step 3

It is clear that the 2D dynamics of the system is incredibly more rich then the 1D dynamics. The reason is simple: Two fronts will always annihilate when colliding in the 1D case, meaning that the two fronts must always travel in the same direction. In a plane, the two fronts can avoid each other (as the spiral motion shows) hence give rise to persistent spirals.

References

[1] R. L. Drake, A. Wayne, M. Mitchell, Gray's Anatomy for Students, Churchill Livingstone, 2014.

[2] A Nygren, C. Fiset, L. Firek, J.W. Clark, D.S. Lindblad, R.B. Clark and W.R. Giles, ”Mathematical Model of an Adult Human Atrial Cell: The Role of K+ Currents in Repolarization,” *Circulation Research,* pp. 63-81, 1982.

[3] S.L. Peck, ”Simulation as experiment: a philosophical reassessment for biological modeling,” *TRENDS in Ecology and Evolution,* pp. 530-534, 2004.

[4] J. Keener, J. Sneyd, Mathemtical physiology, Salt Lake City: Springer, 2009.

[5] M. Christensson, ”Aspects of Spiral wave Dynamics,” Lunds Universitet, Lund, 1996.

[6] M. Oskarsson, ”Creation and Propagation of Spiral waves in Myocardial Tissue,” Lunds Universitet, Lund, 1997.

[7] Å. Bergkvist, ”Modelling of Cardiac Excitation,” Lunds Universitet, Lund, 2000.

[8] L. Rydqvist, ”Complex Cardiac Excitation in inhomoheneous Tissue,” Lunds Universitet, Lund, 2002.

[9] J. Johnsson, ”Three-Dimensional Spiral Waves in Cardiac tissue,” Lunds Universitet, Lund, 2006.

[10] P. Östborn, ”Simulation of sinus nodal behaviour,” Lunds Universitet, Lund, 1998.

[11] P. Östborn, Phase transitions in Large Oscillator lattices, Lund: Lund Universitet, 2003.

[12] A. Reed, P. Kohl, R. Peyronnet, ”Molecular candidates for cardiac stretch-activated ion channels,” *Global Cardiology Science & Practice,* pp. 19-46, 2014.

[13] M. Hiraokaa, S. Kawanoa, Y. Hiranoa, T. Furukawab, ”Role of cardiac chloride currents in changes in action potential,” *Cardiovascular Research,* pp. 23-33, 1998.

[14] Y.K. The, ”Analysis of Ion Channels with Hidden Markov Models,” Albert-Ludwigs-Universität, Freiburg, 2005.

[15] C. F. Stevens, ”Interactions between intrinsic membrane protein and electric field,” *Biophysical Society,* vol. 78, pp. 295-306, 1977.

[16] M. Powell, Approximation theory and methods, Cambridge: Cambrdige university press, 2001.

[17] K.G Andersson, L-C. Böiers, Ordinära differentialekvationer, Lund: Studenlitteratur, 1992.

[18] J. D. Logan, An Introduction to nonlinear partial differential equations, New Jersey: Wiley interscience, 2008.

[19] R. J. LeVeque, Finite Difference methods for ordinary and partial differential equations, Philadelphia: Siam, 2007.

[20] F. C. Sergio Blanes, ”On the necessity of negative coefficients for operator splitting schemes of order higher than two,” *Applied Numerical Mathematics,* vol. 54, pp. 23-37, 2005.

[21] S. Blanes, F. Casas, A. Murua, ”Splitting methods with complex coefficients,” *Bol. Soc. Esp. Mat. Apl,* vol. 50, pp. 47-61, 2010.

[22] K. J. Laidler, J. H. Meiser, B. C. Sanctuary, Physical Chemistry, Cengage Learning, 2002.

[23] P. Pohl, Grundkrus i numeriska metoder, Fallköping: Liber, 2005.

[24] R. Weber, U. Steinhoff, E. Hofer, D. Sanchez-Quintana, H. Koch, ”Modelling the electrical propagation in cardiac tissue using detailed histological data,” *Biomedizinische Technik,* vol. 48, pp. 476-479, 2003.

[25] W. Hundsdorfer, ”Numerical Solution,” Thomas Stieltjes Institute, Amsterdam, 2000.

[26] G. Strang, Computational Science and Engineering, Massachusett: Wellesley Cambrdige Press, 2007.

[27] R. C. Dean, ”Implicit-Explicit Numerical Methods in Models of Cardiac Electrical Activity,” Univ. of Saskatchewan, Saskatoon, 2008.