INVITED-TALK

Tuning into my heart through wearables: Towards a formal cardiac digital twin

PARTHA S ROOP, University of Auckland, Auckland, AUK, New Zealand

NATHAN ALLEN, Auckland University of Technology, Auckland, AUK, New Zealand

SHAHAB KAZEMI, University of Auckland, Auckland, AUK, New Zealand

Open Access Support provided by:

University of Auckland

Auckland University of Technology

Published: 28 September 2025

Citation in BibTeX format

MEMOCODE '25: International Symposium on Formal Methods and Models for System Design
September 28 - October 3, 2025
Taipei, Taiwan

Conference Sponsors:
SIGDA
SIGBED

https://doi.org/10.1145/3742875.3754694
ISBN: 9798400719943

Tuning into my heart through wearables: towards a formal cardiac digital twin

Partha Roop
University of Auckland
Auckland, New Zealand
p.roop@auckland.ac.nz

Nathan Allen
Auckland University of Technology
Auckland, New Zealand
nathan.allen@aut.ac.nz

Shahab Kazemi
University of Auckland
Auckland, New Zealand
shahab.kazemi@auckland.ac.nz

ABSTRACT

Digital Twins (DTs) mimic a physical system using a digital version of the real system. While these have been explored in many domains, digital twins of human organs are yet to be created, especially those that are inspired by formal methods. To this end, we propose the first Cardiac Digital Twins (CDTs) by leveraging two key innovations from our research group.

The first is a real-time model of the heart, that is based on a network of hybrid automata to represent the cardiac conduction system that mimics the rhythmic electrical activity of a normal heart. The model can be parametrised to exhibit disease states in real-time and this approach is being used by MathWorks for closed-loop validation of pacemakers in real-time. This work has raised the interest of both device manufacturers and certification agencies, especially in the USA.

Our group has expertise in digital biomarkers obtained from wearables, such as Electrocardiograms (ECGs) and Photoplethysmograms (PPGs). These provide a window into the cardiac cycle and we have already shown that the two signals are strongly correlated. Hence, a second innovation is related to using wearables to personalise the real-time heart model, so that the model generates ECGs matching that of an individual in different states. Our approach paves the way for developing personalised therapies, real-time monitoring, and accurate estimation of heart rate variability.

CCS CONCEPTS

• **Applied computing** → **Computational biology**; • **Computer systems organization** → **Embedded and cyber-physical systems**; • **Theory of computation** → *Formal languages and automata theory*.

KEYWORDS

digital twins, cardiac modelling, wearables, electrocardiograms

ACM Reference Format:
Partha Roop, Nathan Allen, and Shahab Kazemi. 2025. Tuning into my heart through wearables: towards a formal cardiac digital twin. In *International Symposium on Formal Methods and Models for System Design (MEMOCODE '25), September 28-October 3, 2025, Taipei, Taiwan.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3742875.3754694

MEMOCODE '25, September 28-October 3, 2025, Taipei, Taiwan

ACM ISBN 979-8-4007-1994-3/2025/09
https://doi.org/10.1145/3742875.3754694

1 INTRODUCTION

DTs [13] of human organs have the potential to significantly transform healthcare through personalised therapy, empowering individuals to take full control of their health and wellbeing. Roughly 32 % of global deaths in 2019 according to the World Health Organisation (WHO) is related to cardiac diseases. Hence, our work has focussed on cardiac health, which is the leading cause of death globally. A CDT has the potential to significantly reduce the impact and progression of cardiac diseases. A CDT can also be used as a predictive model, helping to manage conditions that may otherwise lead to disability or death. However, DTs of human organs seem more like science fiction than reality [26].

The core challenges are three-fold. Firstly, high-fidelity sensors are needed that can provide accurate cardiac status through body-surface signals such as ECGs and PPGs, the latter of which provides a signature of the cardiac cycle through an optical sensor [23]. Current sensors are highly inaccurate for continuous monitoring, especially since they are susceptible to motion artefacts. Secondly, and more importantly, a cardiac model capable of personalised response-matching of a human is lacking. Finally, existing methods are reliant on solving a challenging problem, called the *inverse problem*, which is ill-posed [17] that relates body-surface signals to cardiac surface signals.

Wearable devices [21], with many sensors, are able to provide a window into human biology, giving a snapshot of many biophysical signals, unthinkable just over a decade ago. Hence, the real-time status of many physical signals are now available for continuous monitoring. For example, we can access temperature, sleep quality, Heart Rate (HR), peripheral capillary oxygen saturation (SpO2), Heart Rate Variability (HRV) and many more parameters routinely, especially to enhance healthcare [18]. Wearables can also be used for non-invasive blood pressure [22] and blood glucose [24] monitoring, which have the potential to save many lives. Moreover, smart watches are capable of real-time arrhythmia detection, which is saving lives.

Hence, in this paper, we propose to leverage our prior work on biomedical sensing and real-time cardiac modelling, especially since both are grounded in formal methods. We hypothesise that the parametrisation of a real-time heart model, using timing parameters obtained from ECG and PPG signals using sensor fusion [23], could be used for CDT development. We will investigate the link between noisy body-surface signals, obtained through wearable sensors [5], and an underlying real-time abstract mathematical model of the human heart, bypassing the need for solving the inverse problem.

1.1 Solution Overview

Our approach is summarised in Figure 1. First, in ①, we will obtain recordings from participants through an ethics approved protocol in the near future. These recordings will include ECG and PPG data during several *states* of the individual, such as during sleep, resting, walking, and so on. A pre-existing abstract cardiac model (②) is used, which is based on the University of Auckland (UoA) model [3, 6, 32]. Using the wearable data in ① we will develop novel algorithms for sensor fusion to remove the effect of motion artefacts. These will be used to compute the average timing values of key timing intervals of the cardiac cycle, which will be used for initial parametrisation of the abstract cardiac model model (③) to create the DT for the person (④).

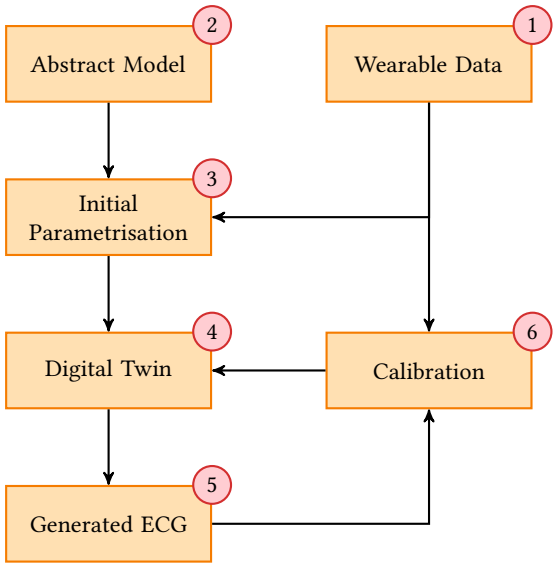

Figure 1: Proposed method of the Cardiac Digital Twin

To ensure the correctness of the generated DT, and to fine-tune its behaviour, the cardiac model is extended with the addition of ECG generation (⑤). This ECG will be used in a closed-loop calibration process (⑥) to fine-tune the parameters of the DT, such that the generated ECGs more closely resemble those from the wearable data. This process will be completed for each state of the user, to create a set of representative DTs for the person. This process will make use of multiple techniques, including genetic algorithms, such as particle swarm, or evenArtificial Intelligence (AI) techniques, such as transformer models [28].

The rest of this paper is organised as follows. In Section 2, we will provide a background on the cardiac conduction system, some formal models that have been used to model it, and elaborate on three key cardiac activity signals — PPGs, ECGs, and Electrograms (EGMs). Then, we will present a set of improvements that have been made over existing models for the purpose of EGM generation in Section 3. Following this, we will apply these same techniques to the problem of ECG generation in Section 4, along with some preliminary results and comparisons to experimentally obtained signals. The final section is devoted to concluding remarks.

2 BACKGROUND AND RELATED WORK

Formal methods [30] have been used in many safety-critical settings, especially space, automotive, and aviation. These Cyber Physical Systems (CPSs) represent highly distributed and concurrent subsystems, and formal methods have matured to levels to tackle the complexity of such systems [19]. This is an inspiration for the current work, which started taking shape around 2014 [32], especially examining the role of formal methods for the design of real-time models for the cardiac conduction system.

The inspiration for us is based on the observation that, like a man-made CPS which comprises of many distributed controllers [7], the human body has a network of distributed organs, which are not working in isolation, but are highly coordinated. For example, the heart and the gut are connected to the nervous system through the *vagus nerve*. Moreover, the heart and the lung also operate in tandem, where during inhalation the activity of the heart increases, and during expiration the opposite occurs. This is known as Respiratory Sinus Arrhythmia (RSA) [20], a desirable arrhythmia which induces HRV [12] that has been shown to be an excellent indicator of overall health and wellbeing [25]. HRV is typically reduced due to stress and may be gained back through deep sleep.

Given such complexity at both at the system (complete human) and module (organ) levels, formal methods may be explored for the modelling and verification of biological systems. Hence, there is a body of work on using formal modelling for biological cells, especially in the heart [8] and gut [29]. We will start with a review of such modelling.

2.1 Cardiac Overview

The purpose of the human heart is to pump blood through the lungs and body to provide sustenance. Internally, it is made up of pairs of atria and ventricles, each of which having a left and right variant, as shown in Figure 2. The atria, which are comparatively smaller, serve as places for the blood to accumulate before being passed through to the ventricles. The ventricles, on the other hand, are tasked with pumping the blood contained within them around the entire body which requires much stronger contractions and higher pressures. This process periodically repeats at around 60 bpm to 100 bpm to ensure nutrients are continually provided. This rhythmic activity leads to the *cardiac cycle*, comprising of the *systole*, when the heart muscles contract and expel the blood contained within them, and *diastole*, when they relax and fill with blood.

Also shown in Figure 2 is the cardiac conduction system, here simplified to just 33 nodes [16]. Each node is a circle with its name inside of it, and connections between nodes are represented by the solid lines. Our work with this model started with a Matlab Simulink model, that was kindly shared with us by Prof. Marta Kwiatkowska [10] for modelling forward conduction only. Subsequently, we refined this model to capture both forward and backward conduction [32], which is key to capturing different more complex arrhythmia. The two key components for each of these models are the *node model*, which captures the electrical activity of a group of cells, and the *path model*, which captures the conduction delay between nodes.

The foundational node of the cardiac conduction network is known as the Sinoatrial (SA) node, which determines the rhythmic

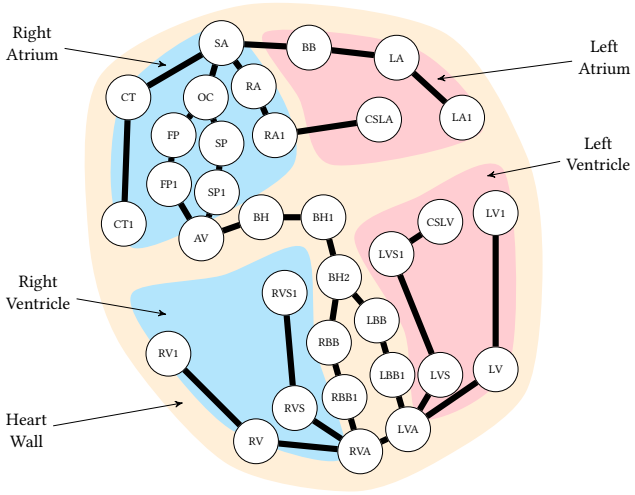

Figure 2: Diagram of the human heart and its abstracted conduction network

rate of the heart and therefore is commonly known as its *natural pacemaker*. This is one of several in the cardiac system which self-stimulate on a regular basis without any required external input, which are known as *auto-rhythmic* cells. For the SA node, this is generally between 60 bpm to 100 bpm for an average person.

Other auto-rhythmic nodes in the heart include the Atrioventicular (AV) node, which naturally beats at a rate of 40 bpm to 60 bpm, and the Bundle of His (BH), which naturally beats at around 30 bpm to 40 bpm. Under normal circumstances, the SA node's rate will be the fastest and will therefore *dominate* the pacing of the heart. However, in cases of cardiac failure these other regions exist as a sort of backup to ensure some level of cardiac function is maintained. For example, if the SA node's auto-rhythmic behaviour fails then the AV node will take over.

As mentioned earlier, synchronisation between atrial and ventricular activity is important to ensure that there is enough time for blood to flow from one to the other. This synchronisation is known as the *A-V delay*. In Figure 2, the section of the conduction network that controls this delay is the area that runs from the SA node through to the Right Ventricular Apex (RVA) and Left Ventricular Apex (LVA), including the AV node and the BH. After reaching the RVA and LVA the conduction spreads out across both the right and left ventricles, through a region known as the Purkinje fibres.

This region is often important when it comes to the production, and prevention, of various cardiac diseases. For example, the branching and joining segment which consists of the Fast Pathway (FP) and Short Pathway (SP) is another important aspect of this, and is often the cause of AV Nodal Re-entrant Tachycardia (AVNRT). Alternatively, the AV node in particular plays an important role in preventing the impacts of Atrial Fibrilation (AF), where the atria are beating too quickly, from flowing through to the ventricles so that the core function of the heart, pumping blood around the body, can be maintained. In this sense, it acts as a low-pass filter.

2.2 Formal Models

Traditionally, bioengineers model the cardiac electrical activity by capturing each individual ionic channel that makes up each cell [9]. As a result, they have developed different types of computational modelling, mainly enabled by finite element and finite volume methods [27]. These provide a high degree of fidelity albeit that such models are computationally very expensive.

Interest in *abstract modelling* started due to the need for using such models to ensure the validation of closed-loop pacing devices that treat arrhythmia, such as pacemakers. This necessitated the need for models that can elicit real-time response to such devices, while being able to mimic physical phenomena accurately from the device perspective. These models covered a range of complexities at various levels of abstraction, such as simple Timed Automaton (TA) models [15], linear Hybrid Automaton (HA) models [1], and detailed HA models [31]. In the following we will elaborate on such an abstract cell model, developed in our group.

On a physical level, cardiac cells operate through charged ions passing in and out of the cell walls through *ionic channels* and associated gates. As this happens, a difference in charge is created between the inside and outside of the cell, i.e. an electric potential, which is able to cause the surrounding muscle tissue to contract (when the cell is depolarised), and relax (when the cell is re-polarised).

The typical morphology of this electric potential is shown in Figure 3, a shape known as its Action Potential (AP) [11]. To start with, a cell is at its initial Resting Period (RP) (Phase 0). Whenever the cell is stimulated, for example due to a neighbouring cell, it begins to depolarise and its potential increases (Phase 1). Once this reaches a threshold voltage (V_T), the cell goes into the Upstroke phase where the cell rapidly continues to depolarise up to some maximum V_O. At this point, the muscle tissue also rapidly contracts.

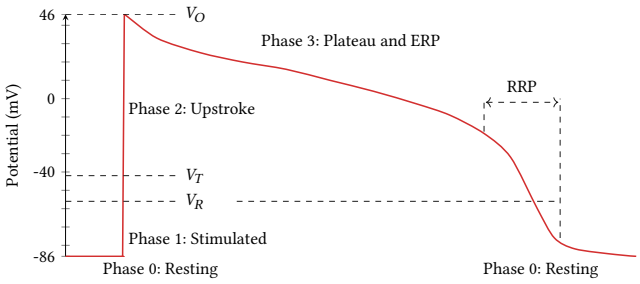

Figure 3: An example Action Potential of a human atrial cardiomyocyte cell

After this, the cell begins its recovery (re-polarisation) process, beginning with its Plateau and Effective Refractory Period (ERP) (Phase 3). Any external stimulus received by the cell during this time will be ignored. Eventually, the cell then enters its Relative Refractory Period (RRP), where re-polarisation happens more rapidly. Here, an external stimulus is able to re-stimulate the cell, but not at 100 % efficiency. Once the cell reaches some resting voltage V_R it will return to its RP (Phase 0) and the process can repeat.

As mentioned, the external stimulus that initiates this AP often comes from a neighbouring cell as the activity propagates through the network. Once a cell undergoes its own depolarisation, then this

will also continue along the network and pass to its own neighbours, causing a chain reaction throughout the network. Typically, this conduction chain will be initiated by one of the auto-rhythmic cells in the network, such as the SA node, however in cases where the person has a pacing device installed, such as a pacemaker, this initiation could also arise due to an artificial pacing event.

The phases of an AP can be represented as an HA, as shown in Figure 4 [32]. This model is an extension of earlier abstract models [1, 31] which captures complex physical phenomena such as the *restitution behaviour* of cells more accurately. As an HA, this model is divided into four locations, q_0, q_1, q_2, and q_3, which match the phases of the AP, namely resting, stimulated, upstroke and ERP respectively. The function g is used to aggregate the voltages from the neighbours. The electric potential of the cell is captured using a set of three Ordinary Differential Equations (ODEs), v_x, v_y, and v_z, with the resulting value being a linear combination of these three values, namely $v_x - v_y + v_z$.

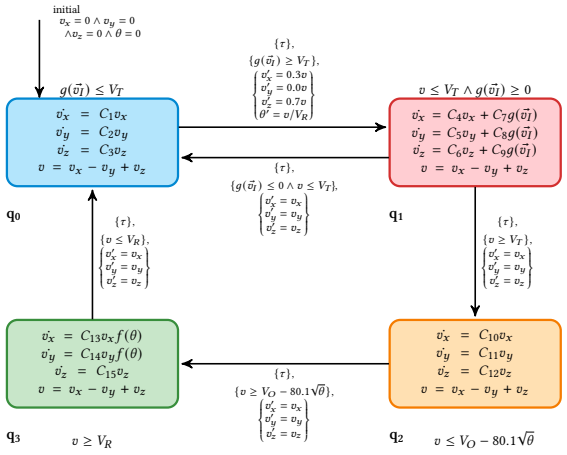

Figure 4: The Hybrid Automaton of the UoA Cell Model, reproduced from [33]

Multiple instances of these cell models are created in order to capture each node in the cardiac conduction network, which are then connected through an equivalent path model (also defined as an HA). This path model captures the conduction delay between nodes, along with the process of *annihilation*, where two wavefronts travelling in opposite directions will cancel each other out.

2.3 Signals that Expose Cardiac Activity

In almost every case, it is impossible to observe the individual cardiac activity within a person's heart. Instead, we use a range of signals that can be captured through various means, such as through wearables.

2.3.1 Photoplethysmograms. Wearable technologies, for health monitoring applications, are ubiquitous [18, 21]. In these, heart rate sensing is mainly performed by a light-based sensor. This device uses optical sensing methods to illuminate the translucent part of the skin to measure the changes in blood oxygenation relative to the cardiac activity. This change is recorded as the PPG signal, an example of which is depicted in Figure 5.

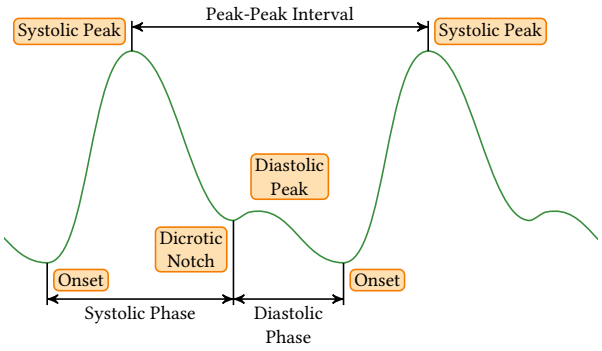

Figure 5: A typical PPG signal

These signals can be used to detect several physical properties of the heart. For example, the time taken from the *onset* to the *dicrotic notch*, captures the duration of the systole. Likewise, the time taken from the dicrotic notch to the subsequent onset determines the duration of the *diastole*. The time between successive *systolic peaks*, known as the *peak-to-peak interval*, corresponds to the heart rate.

2.3.2 Electrocardiograms. To observe more detailed cardiac electrical activity, electrodes are placed on the body surface, which produces a signal known as an ECG. Typically, there will be 10 electrodes placed on the body with 12 channels of recording, known as *leads*. Six of these leads (V1 through V6) are known as the chest or precordial leads and use electrodes that are placed in a near-horizontal pattern around where the heart is located. The remaining leads come from three of the electrodes placed on the limbs of the person, with them combining to produce three standard limb leads (I, II, and III) and three augmented limb leads (avF, avL, avR). Figure 6 shows how these six leads are created from three electrodes placed on the limbs of the person. The fourth electrode, placed on the right leg, then acts as a common ground for all others.

The three standard limb leads (top row of Figure 6) are created through calculating the difference between each pair of the three electrodes. The augmented limb leads (bottom row) are created from each limb electrode relative to a common point named *Wilson's Central Terminal*, V_W, according to eq. (1). The vectors created by each of these leads (red arrows) together allow for the direction of cardiac activity to be distinguished at roughly a 30° resolution.

$$V_W = \frac{1}{3}(RA + LA + LL) \tag{1}$$

A typical ECG signal for lead II is shown in Figure 7, and is annotated to show the key waves (P, Q, R, S, and T) along with some of the key timing intervals between these events. Here, the P wave corresponds to the atrial depolarisation, the QRS complex captures the ventricular depolarisation, and the T wave represents the ventricular re-polarisation. The intervals between these then represent the physical properties of the cardiac conduction. For example, the R-R interval represents the firing rate of the SA node, the P-R interval corresponds to the A-V delay, and the R-T interval is the Action Potential Duration (APD) of the ventricles.

2.3.3 Relationship between Photoplethysmograms and Electrocardiograms. It has been shown that the PPG and ECG signals are

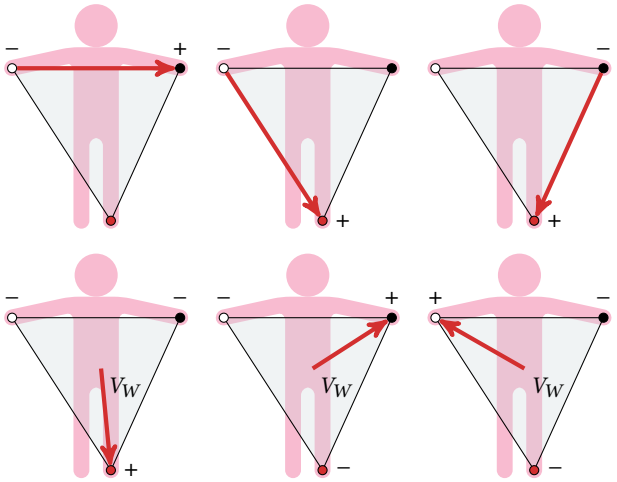

Figure 6: Diagram showing the location of limb electrodes (right arm in white, left arm in black, and left leg in red), and limb leads for a 12-lead Electrocardiogram. The first row shows standard limb leads I, II, and III, while the second row shows augmented limb leads aVF, aVL, and aVR. Wilson's Central Terminal is labelled as V_W.

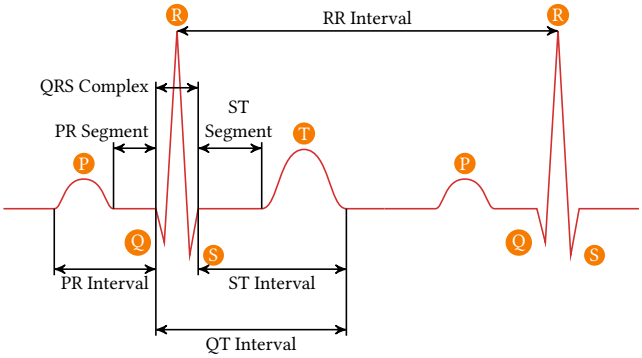

Figure 7: A typical Electrocardiogram signal

strongly correlated [23, 24], as shown in Figure 8. Here, there is a time-delayed relationship between the two signals, with this offset being known as the Pulse Arrival Time (PAT).

Nowadays, it is quite common for wearable devices to approximate an ECG, despite only having a single sensing location. As wearable sensors are noisy, modern wearables incorporate *sensor fusion* algorithms using multiple sensors to improve their accuracy. A similar approach will be taken in the development of our DTs.

2.3.4 Electrograms. In addition to the above non-invasive sensing methods, cardiac surface electrical activity can also be exposed through leads of pacing devices. For example, when a pacemaker is connected to the heart, it is done through a set of leads which attach to the interior walls of the heart. The signal seen through these leads is known as an EGM. This is the result of a movement of electrical charge (as the AP wavefront propagates through the heart) relative to the position of the lead electrodes [2, 14].

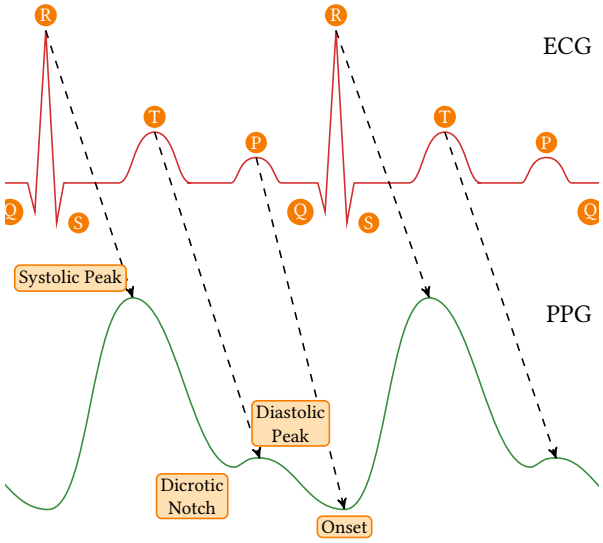

Figure 8: Correlation between PPG and ECG signals

Existing work uses a *dipole model* to calculate the effective voltage seen at the electrode due to the moving charged wavefront [2]. Equation (2) shows the equation used to calculate this induced voltage, where r is the distance between the electrode and the dipole, ψ is the angle of the dipole relative to the electrode, and C is the magnitude of the dipole.

$$V(r, \psi) = C \times \frac{cos(\psi)}{r^2} \tag{2}$$

We can make three key observations from this equation. Firstly, any dipole which is perpendicular to the electrode ($\psi = 90°$) will induce no voltage. Secondly, there is an inverse square relationship between the induced voltage and the distance. Thirdly, dipoles with a larger magnitude induce a larger voltage at the electrode.

3 ELECTROGRAM GENERATION

While the existing dipole model [2] serves as a good abstraction for how EGMs are generated, it is not perfect. Here, we will outline three challenges with the dipole model, particularly when used with an abstract cardiac network, and the steps that were taken to address them.

3.1 Challenge 1: Orientation Changes

Firstly, let's consider the impact that the angle of a dipole (ψ) has on the generated EGM. In a case where conduction pathways are smooth and continuous, the angle of the dipole will therefore also be continuous. However, in abstract cardiac models (such as [33]) the conduction network is approximated by a series of linear paths. As a result, whenever the wavefront moves from one path to another it experiences a sudden (discontinuous) change in its orientation, and hence also the induced voltage at the electrode.

Figure 9a shows an illustration of this problem where the wavefront makes a 90° bend while passing the electrode (E). Figure 9b shows the induced voltage at the electrode as the wavefront moves along this path, with the sudden discontinuity happening at 5 s.

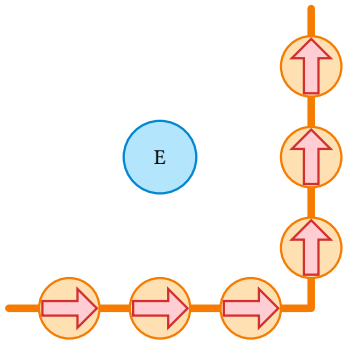

(a) Illustration of the problem

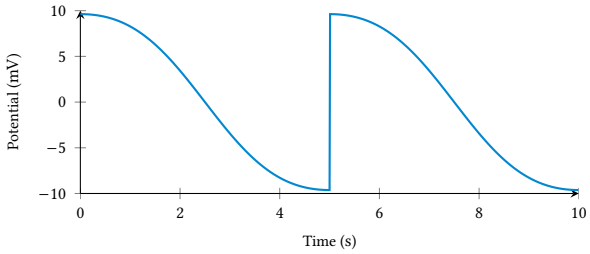

(b) Induced voltage at the electrode

Figure 9: Problems due to dipole orientation changes

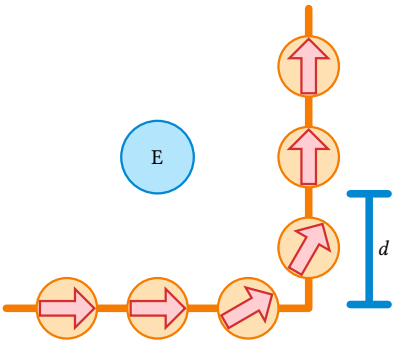

(a) Illustration of the solution

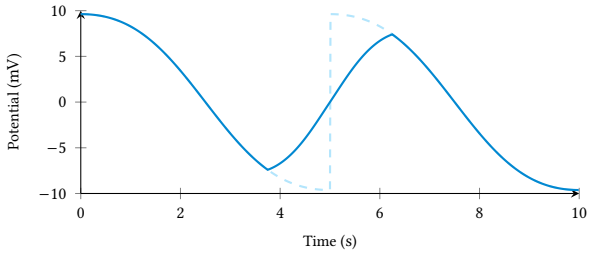

(b) Induced voltage at the electrode

Figure 10: Rotating dipoles as a solution

This problem could be solved by creating curved, rather than straight, paths between nodes and having the dipoles continuously align themselves with the path direction. However, this would incur a high computational cost and impact our ability to execute the model in real-time. Instead, we can achieve a similar effect by rotating the orientation of the dipoles as they progress along the path. This allows us to align the dipole orientations at each node and smooth out any sudden changes that would otherwise occur. We rotate the orientation of the dipoles only when they are within some distance d of a node, otherwise the dipole orientation is aligned with the path as normal.

Figure 10a shows this rotating dipole behaviour in practice as it gets within a distance d of the change in direction, where the dipole orientation would be exactly the midpoint (i.e. pointing to the top-right) while at this point. The resulting induced voltage at the electrode (E) is shown in Figure 10b, where the signal no longer experiences any discontinuity at 5 s, as was the case before (reproduced as a faded dashed line). The distance d determines how much smoothing is performed to the induced voltage, where a value of 0 would be the same as having non-rotating dipoles and larger values would tend towards a smooth line from positive to negative.

We note that this approach is not perfect, as there remains a discontinuity in the derivative of this signal as the dipole passes the distance d. However, for an abstract model, we see this as a fair-enough concession to make while providing improvements over the original technique.

Additionally, there are several nodes in the abstract cardiac network (Figure 2) which do not just contain a single input and output path that the dipole follows. For example, the LVA and RVA nodes each have four paths connected to them, each of which require aligned rotations at the node. How should this orientation be decided?

Instead of manually assigning the orientation at each node, we instead automatically calculate an "average" orientation under normal conditions. To do this, we decide the direction of *forwards* conduction for each path in the network, decided based on when the heart is behaving normally. Next, at each node we determine the average input and output vectors for the connected paths, and then take the midpoint between these average vectors as the resulting orientation.

Figure 11 shows this approach for a node with four connected paths — two inputs (in green) and two outputs (in orange). The average input and output vectors are determined from these paths, and are then shown in dark red. Finally, the orange arrow indicates the resulting orientation of the dipoles when at this node, ensuring that under normal conditions there is minimal rotation required on each path.

3.2 Challenge 2: Intensity Changes

The second challenge that we encounter is where the intensity of the dipole (C) suddenly changes. In the abstract model, this occurs whenever a single path splits into multiple parallel paths or multiple paths converge into a single one, such as through the BH. Looking back to eq. (2), we can see that if the intensity of the dipole was to double then the corresponding induced voltage at the electrode would also double.

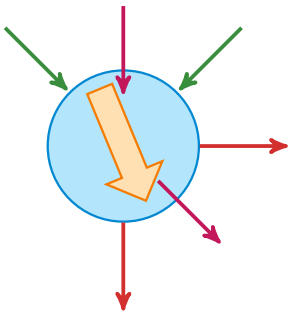

Figure 11: Determining a node's dipole orientation

Figure 12a shows an illustration of this problem where a single path is diverging into two parallel paths, each with the same intensity and resulting in twice the total charge. Figure 12b then shows the induced voltage from this setup, where the divergence happens at 5 s causing an instantaneous doubling in the induced voltage. Note that the electrode (E) is offset from the divergence point, as otherwise the induced potential when the discontinuity happens would be at 0 mV.

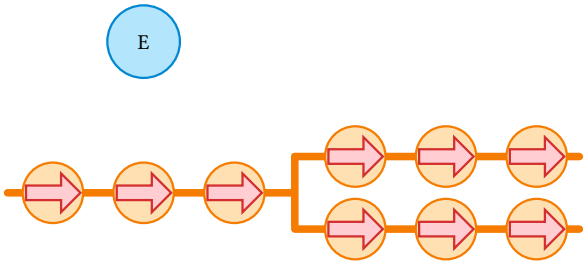

(a) Illustration of the problem

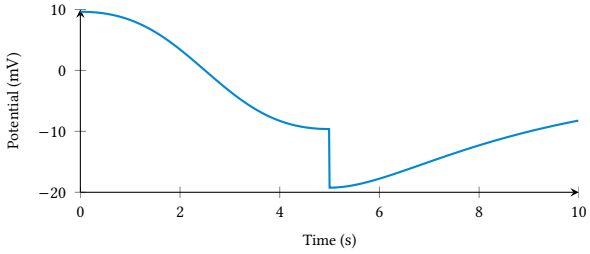

(b) Induced voltage at the electrode

Figure 12: Problems due to dipole intensity changes

To solve this problem, we again follow a similar approach to the first challenge, but rather than linearly varying the orientation of the dipole, we now vary its intensity. At each divergence, the outgoing dipoles have their intensity reduced as a proportion of the number of outgoing paths, and then have their intensities increased linearly over some distance d. Figure 13a shows an illustration of this solution, where the dipole intensities are immediately halved when encountering the divergence to ensure the total charge remains identical. Then, over a distance d, their intensities are increased up to their full charge. The resulting induced voltage is shown in Figure 13b which now no longer contains that discontinuity that was previously present (indicated by the faded dashed line).

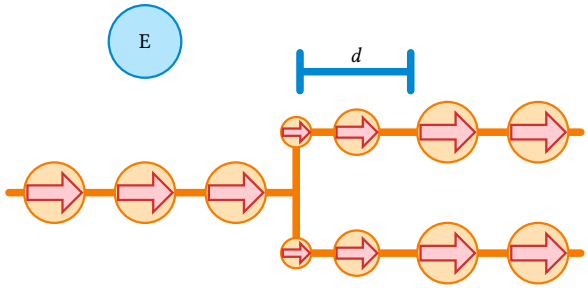

(a) Illustration of the solution

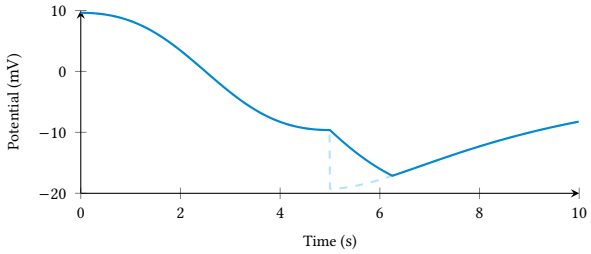

(b) Induced voltage at the electrode

Figure 13: Varying intensities as a solution

3.3 Challenge 3: Re-polarisation

The third and final challenge that we tackle here is that of re-polarisation. While the existing dipole models [2, 14] capture the instantaneous depolarisation of the AP (Figure 3), they do not capture the slower re-polarisation that occurs afterwards.

Let's start by simplifying the AP of the wavefront as a square wave. Such a wave moving past an electrode should produce a symmetrical induced voltage at an electrode. Figure 14a shows how this was previously captured, using only a single dipole representing the depolarisation (i.e. leading edge) of the wavefront. As a result, the induced voltage (Figure 14b) is asymmetrical as the square wave moves by, which is not correct. Additionally, we would expect to see a predominantly positive induced voltage due to the positively charged square wave, however this is also not the case with large negative induced voltages. Even though this challenge does not cause any obvious discontinuities, the re-polarisation effects are often an important aspect of physiological signals.

To address this, we simply add a corresponding dipole that captures the re-polarisation of the wavefront. This dipole should be opposite in direction to that of the depolarisation one, and follow behind it with some time delay. Figure 15a shows the addition of a re-polarisation dipole for the square wave example, using one which is equally charged but opposite in direction. The corresponding induced voltage at the electrode (Figure 15b) now shows a symmetric signal which is predominantly positive, as expected, when

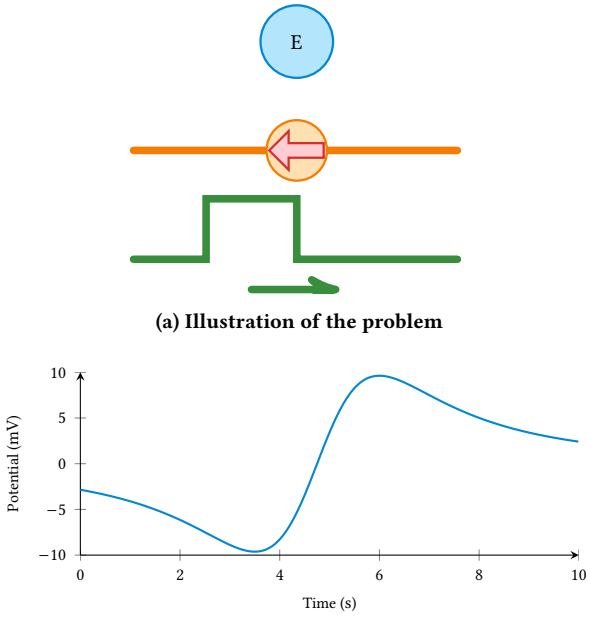

(a) Illustration of the problem

(b) Induced voltage at the electrode

Figure 14: Problems due to a lack of re-polarisation

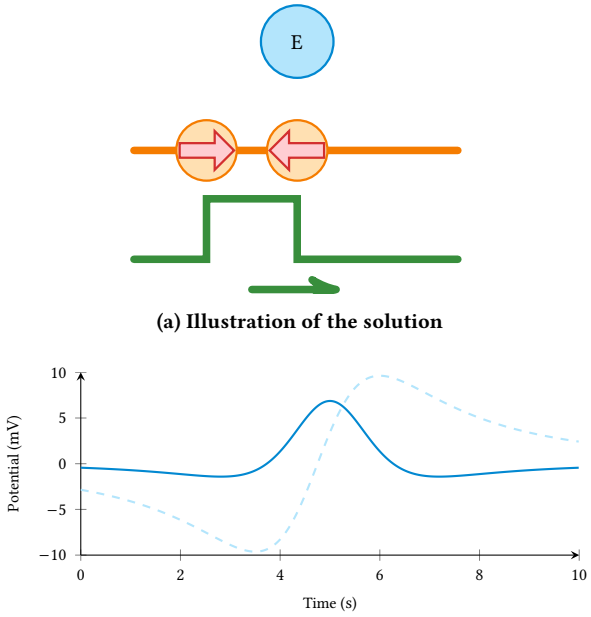

(a) Illustration of the solution

(b) Induced voltage at the electrode

Figure 15: Re-polarisation dipoles as a solution

compared to the original signal (faded dashed line). Furthermore, this modification makes the electrode much more sensitive to signals which are close to it, as the two dipoles roughly cancel out when further away. This is a much more realistic phenomenon, particularly in EGMs where the sensing range of electrodes is highly localised.

When capturing the wavefront of an AP through the abstract cardiac model, the depolarisation and re-polarisation dipoles should be separated by a time corresponding to its APD, namely the time between the upstroke and its RRP. This generally causes, in the cardiac context, the two dipoles to be further spaced apart than those shown in Figure 15a, with noticeable depolarisation and re-polarisation artefacts in the resulting induced voltages. The fact that these can be distinguished is an important part of EGM (and ECG) signal processing in order to calculate a person's APD.

This model can also be extended to more closely capture the shape of a cardiac AP, which is slower to re-polarise than it is to depolarise. A simple example of this using two re-polarisation dipoles is shown in Figure 16a, where each of the re-polarisation dipoles is half the magnitude of the depolarisation one. This captures a morphology as indicated by the green wavefront. By varying the number of dipoles and their relative intensities, this morphology can be adjusted to the desired level of accuracy for an AP. The resulting induced voltage is then shown in Figure 16b, with the faded dashed line showing the previous square-wave case, allowing us to see the minor changes that can be made.

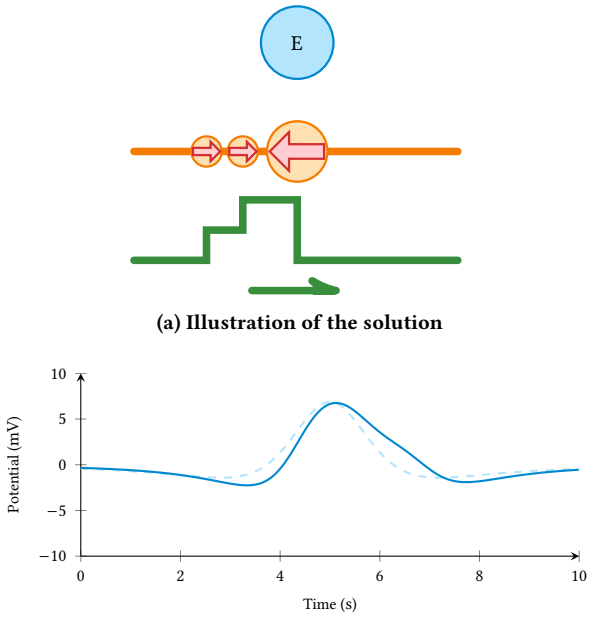

(a) Illustration of the solution

(b) Induced voltage at the electrode

Figure 16: Multiple re-polarisation dipoles to adjust morphology

3.4 Resulting Signals

Before making these improvements, the challenges presented in this section caused generated EGMs which, while generally good at a high level, contained many discontinuities as shown in Figure 17a. Additionally, some aspects of the generated signals were uncharacteristic of real EGMs, such as the crater in the atrial signal at around 120 ms or the flat section in the ventricular signal at around

340 ms to 380 ms. Such abnormalities can reduce the suitability of these generated signals for use in the device testing space, as they do not accurately represent what would be seen in reality.

After applying the corrective measures and improvements, the resulting atrial and ventricular EGMs are shown in Figure 17b which resolves many of the problems with the original signals. Now, there are no discontinuities in the induced potentials and the uncharacteristic "flat" sections have been removed.

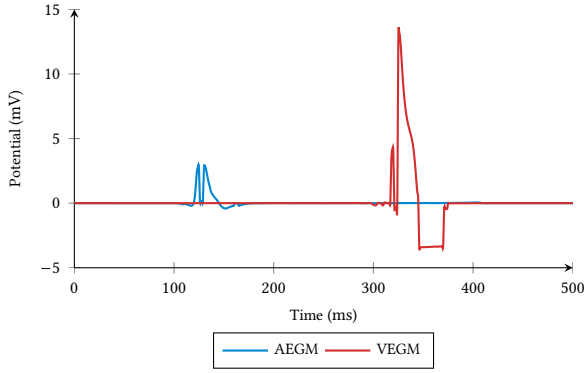

(a) Before improvements

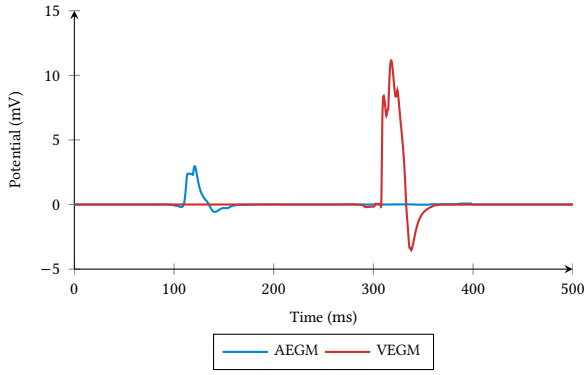

(b) After improvements

Figure 17: Generated Electrograms before and after the improvements

While this section introduced changes to the cardiac model from the point of view of EGM generation, the same problems (and associated solutions) apply to other physiological signal generation such as ECGs.

4 ELECTROCARDIOGRAM GENERATION

To generate ECGs from the abstract cardiac model, a similar approach to that for EGMs, namely using dipoles, can be used. The difference here is that the electrodes are now placed outside the heart, are further away, and there are more of them. The larger distance between the electrode and the conduction network means that the distance r is increased, and so the induced voltage from the dipole model (eq. (2)) is greatly reduced.

Firstly, the six precordial lead electrodes are placed relatively close to the heart in their typical locations, wrapping around the virtual heart in Three Dimensional (3D) space. The radial distance from the centre of the heart out to these electrodes ranges from around 3 cm for V1 and V2, placed around the sternum, up to around 20 cm for V5 and V6, placed on the left side of the torso. V3 and V4 are then placed between these two previous sets on the left side of the chest.

As mentioned in Section 2.3.2, the limb leads effectively capture the directions of AP propagation through 30° intervals. An abstraction of these limb leads can therefore be modelled as shown in Figure 18 using 12 electrodes equally spaced around the centre of the heart, where red and green vectors capture the standard and augmented limb leads respectively. For example, the aVR lead corresponds to the vector from R9 to R3, while the augmented lead I corresponds to R10 to R4. Using this abstraction allows us to decrease the distance between the heart and electrodes, which minimises any residual error due to the small induced voltages.

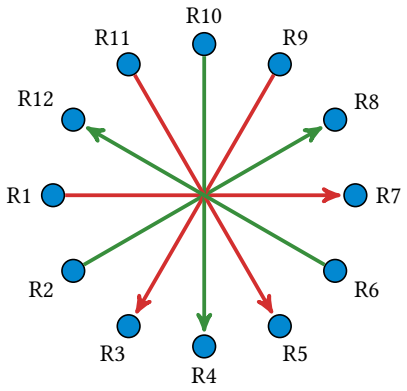

Figure 18: Set of vectors from virtual limb leads of an Electrocardiogram. Standard limb leads are shown in red, while augmented limb leads are in green.

Each lead of the ECG can then be generated using the same approach as for EGMs, making use of the improvements discussed in the previous section to improve their output.

Generally, limb lead II is the one which most people think of when they refer to an ECG, and is often the most useful from a diagnostic point of view. This lead physically uses electrodes from the right arm to left leg, and corresponds to the vector from R11 to R5 in our abstraction (Figure 18). Figure 19 shows the generated output for this lead across two heartbeats using the techniques discussed in this paper with our abstract conduction network. Here, the typical morphology of an ECG can be seen with an initial P wave, followed by the QRS complex, and finally the T wave.

4.1 Experimental Validation

To demonstrate the potential of our abstract cardiac model for ECG-based digital twin applications, ambulatory ECG data was collected using a single-lead Vivalink patch during walking activity, providing real-world validation data despite the model's lack of a torso volume conductor. Figure 20 presents the direct overlay comparison between experimental (blue) and simulated (red) ECGs, with time scales normalized for morphological comparison rather than

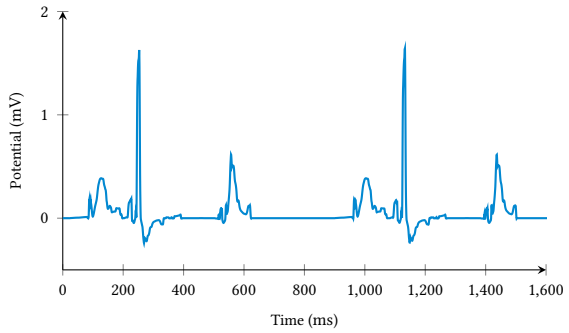

Figure 19: Generated Electrocardiogram for limb lead II

temporal synchronization, revealing similar P-QRS-T morphologies and rhythmic patterns, with automated wave detection successfully identifying corresponding cardiac events in both signals.

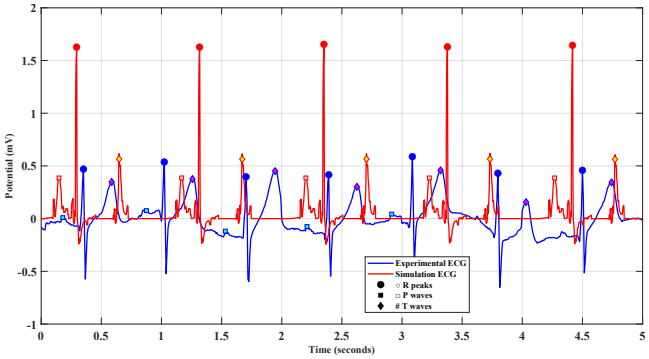

Figure 20: Overview of the experimental and simulated Electrocardiograms

Figure 21 provides detailed wave annotation over a 3 s window, clearly demonstrating the preservation of fundamental cardiac timing relationships between the two signals. Quantitative analysis reveals that while the simulation exhibits a more bradycardic rhythm (roughly 58 bpm, RR interval about (1030 ± 7) ms) compared to the ambulatory experimental ECG (86.8 bpm, RR interval about (691 ± 14) ms), both signals maintain physiologically consistent PR intervals ((149.0 ± 1.6) ms simulation vs (169 ± 12) ms experimental).

This fidelity of the purely electrophysiological model to generate ECG-comparable signals, without any precise torso modelling, validates the hypothesis that our HA framework contains sufficient cardiac dynamics to serve as a foundation for personalized digital twins. The next phase of this work will involve parameter optimization algorithms to systematically adjust the 44-node network's cellular and conduction parameters to minimize the difference between simulated and patient-specific ECG patterns, thereby creating individualized cardiac digital twins capable of predicting patient responses to pacemaker therapy and other cardiac interventions in a personalized medicine framework.

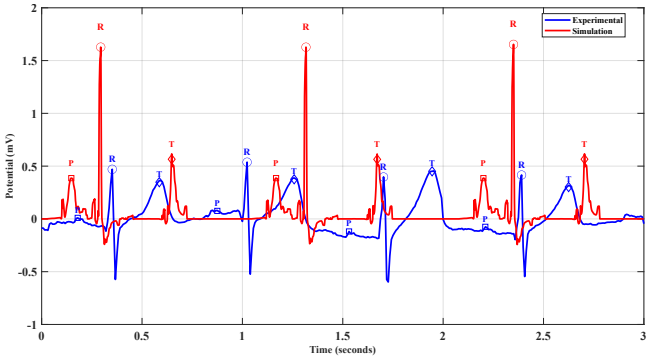

Figure 21: Detailed analysis of experimental and simulated Electrocardiograms

5 CONCLUSIONS

Wearable health monitoring applications have opened new avenues in healthcare which were previously unthinkable [21]. These devices can be seen as a key enabler for the Digital Twins (DTs) of human organs, as they provide a window into their activity. Body surface sensors can now sense not only the activity of the heart and brain but also that of the gut[1]. These have opened new avenues for healthcare, such as apps which could interact with DTs of your own organs in real-time to enable diagnostic and therapeutic improvements, improving the lives of millions of people.

To this end, we propose methods towards the first Cardiac Digital Twin (CDT) which combines our earlier work on real-time cardiac models and biomedical sensing, both inspired by formal methods. Our real-time heart model has been extensively published, and has been synthesised on a Field Programmable Gate Array (FPGA) using hardware-software co-design techniques to enable device validation [6]. Recently, MathWorks has worked on this model and could reproduce similar results using the Speedgoat hardware [4]. In this paper, we have discussed methods for improved Electrogram (EGM) and Electrocardiogram (ECG) generation, without using 3D torso modelling. We have also shown initial attempts at relating ambulatory ECGs with the generated ECG.

While this is an initial attempt at steps towards a CDT, we need to consider several future avenues of improvement. These include algorithms for adapting the timing parameters of the conduction system to match that of the individual using stateful modelling. We will also examine approaches such as Large Language Models (LLMs) and Kalman filtering to ensure that the CDT functions like a real human. However, a key would be the use of sensor fusion to ensure that by combining ECG and Photoplethysmogram (PPG) information, we are able to obtain the correct timing values for parametrisation. Hence, this research has much potential for cooperation with the formal methods community.

ACKNOWLEDGMENTS

The authors acknowledge Weiwei Ai and Mark Trew from the Auckland Bioengineering Institute, especially for their work on real-time heart modelling and dipole models for EGMs.

[1] https://www.alimetry.com/

REFERENCES

[1] Houssam Abbas, Kuk Jin Jiang, Zhihao Jiang, and Rahul Mangharam. 2016. Towards model checking of implantable cardioverter defibrillators. In *Proceedings of the 19th International Conference on Hybrid Systems: Computation and Control*. 87–92.

[2] W. Ai, N. Patel, P. Roop, A. Malik, N. Allen, and M. L. Trew. 2017. An intracardiac electrogram model to bridge virtual hearts and implantable cardiac devices. In *2017 39th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. 1974–1977. https://doi.org/10.1109/EMBC.2017.8037237

[3] Weiwei Ai, Nitish D. Patel, Partha S. Roop, Avinash Malik, and Mark L. Trew. 2020. Cardiac Electrical Modeling for Closed-Loop Validation of Implantable Devices. *IEEE Transactions on Biomedical Engineering* 67, 2 (2020), 536–544. https://doi.org/10.1109/TBME.2019.2917212

[4] Weiwei Ai, Mark Trew, and Partha Roop. 2024. A Real-Time Heart-in-the Loop: A Novel Method for Validation of Cardiac Devices. https://www.mathworks.com/videos/a-real-time-heart-in-the-loop-a-novel-method-for-validation-of-cardiac-devices-1730977844895.html.

[5] Nathan Allen, Aron Jeremiah, Robin Murphy, Rachael Sumner, Anna Forsyth, Nicholas Hoeh, David B Menkes, William Evans, Suresh Muthukumaraswamy, Frederick Sundram, et al. 2024. LSD increases sleep duration the night after microdosing. *Translational Psychiatry* 14, 1 (2024), 191.

[6] Nathan Allen and Partha S Roop. 2020. Semantics-directed hardware generation of hybrid systems. In *2020 ACM/IEEE 11th International Conference on Cyber-Physical Systems (ICCPS)*. IEEE, 259–268.

[7] Rajeev Alur. 2015. *Principles of cyber-physical systems*. MIT press.

[8] Sidharta Andalam, Nathan Allen, Avinash Malik, Partha S Roop, and Mark Trew. 2017. A novel emulation model of the cardiac conduction system. *ACM Transactions on Embedded Computing Systems (TECS)* 16, 5s (2017), 1–20.

[9] Travis M Austin, Mark L Trew, and Andrew J Pullan. 2006. Solving the cardiac bidomain equations for discontinuous conductivities. *IEEE Transactions on Biomedical Engineering* 53, 7 (2006), 1265–1272.

[10] Taolue Chen, Marco Diciolla, Marta Kwiatkowska, and Alexandru Mereacre. 2014. Quantitative verification of implantable cardiac pacemakers over hybrid heart models. *Information and Computation* 236 (2014), 87–101.

[11] Taolue Chen, Marco Diciolla, Marta tkowska, and Alexandru Mereacre. 2014. Quantitative verification of implantable cardiac pacemakers over hybrid heart models. *Information and Computation* 236 (2014), 87–101.

[12] Robert L Drury, Stephen Porges, Julian Thayer, and JP Ginsberg. 2019. Heart rate variability, health and well-being: A systems perspective. , 323 pages.

[13] Michael W Grieves. 2023. Digital twins: past, present, and future. In *The digital twin*. Springer, 97–121.

[14] WERNER IRNICH. 1985. Intracardiac Electrograms and Sensing Test Signals: Electrophysiological, Physical, and Technical Considerations. *Pacing and Clinical Electrophysiology* 8, 6 (1985), 870–888. https://doi.org/10.1111/j.1540-8159.1985.tb05907.x arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-8159.1985.tb05907.x

[15] Zhihao Jiang, Miroslav Pajic, Salar Moarref, Rajeev Alur, and Rahul Mangharam. 2012. Modeling and Verification of a Dual Chamber Implantable Pacemaker. In *Proceedings of the 18th International Conference on Tools and Algorithms for the Construction and Analysis of Systems* (Tallinn, Estonia) *(TACAS'12)*. Springer-Verlag, 188–203.

[16] Marta Z Kwiatkowska, Alexandru Mereacre, Volker Turau, Marta Kwiatkowska, Rahul Mangharam, and Christoph Weyer. 2014. Automated Verification of Quantitative Properties of Cardiac Pacemaker Software. In *MCPS*. 137–140.

[17] Lei Li, Julia Camps, Blanca Rodriguez, and Vicente Grau. 2024. Solving the Inverse Problem of Electrocardiography for Cardiac Digital Twins: A Survey. *arXiv preprint arXiv:2406.11445* (2024).

[18] Tatjana Loncar-Turukalo, Eftim Zdravevski, José Machado Da Silva, Ioanna Chouvarda, Vladimir Trajkovik, et al. 2019. Literature on wearable technology for connected health: scoping review of research trends, advances, and barriers. *Journal of medical Internet research* 21, 9 (2019), e14017.

[19] William Martin, Patrick Lincoln, and William Scherlis. 2022. Formal methods at scale. *IEEE Security & Privacy* 20, 03 (2022), 22–23.

[20] Erin L O'Callaghan, Renata M Lataro, Eva L Roloff, Ashok S Chauhan, Helio C Salgado, Edward Duncan, Alain Nogaret, and Julian FR Paton. 2020. Enhancing respiratory sinus arrhythmia increases cardiac output in rats with left ventricular dysfunction. *The Journal of physiology* 598, 3 (2020), 455–471.

[21] Aleksandr Ometov, Viktoriia Shubina, Lucie Klus, Justyna Skibińska, Salwa Saafi, Pavel Pascacio, Laura Flueratoru, Darwin Quezada Gaibor, Nadezhda Chukhno, Olga Chukhno, et al. 2021. A survey on wearable technology: History, state-of-the-art and current challenges. *Computer Networks* 193 (2021), 108074.

[22] Abhinandan Panda, Ayush Anand, Srinivas Pinisetty, and Partha Roop. 2024. An Explainable and Formal Framework for Hypertension Monitoring Using ECG and PPG. *IEEE Embedded Systems Letters* 16, 4 (2024), 405–408.

[23] Abhinandan Panda, Srinivas Pinisetty, and Partha Roop. 2022. A novel mapping of ECG and PPG to ensure the safety of health monitoring applications. *IEEE Embedded Systems Letters* 15, 1 (2022), 49–52.

[24] Abhinandan Panda, Srinivas Pinisetty, and Partha Roop. 2022. Policy-Based Diabetes Detection using Formal Runtime Verification Monitors. In *2022 IEEE 35th International Symposium on Computer-Based Medical Systems (CBMS)*. 333–338. https://doi.org/10.1109/CBMS55023.2022.00066 ISSN: 2372-9198.

[25] Stefan Sammito, Beatrice Thielmann, and Irina Böckelmann. 2024. Update: factors influencing heart rate variability–a narrative review. *Frontiers in physiology* 15 (2024), 1430458.

[26] Chenyu Tang, Wentian Yi, Edoardo Occhipinti, Yanning Dai, Shuo Gao, and Luigi G Occhipinti. 2024. A roadmap for the development of human body digital twins. *Nature Reviews Electrical Engineering* 1, 3 (2024), 199–207.

[27] Mark Trew, Ian Le Grice, Bruce Smaill, and Andrew Pullan. 2005. A finite volume method for modeling discontinuous electrical activation in cardiac tissue. *Annals of biomedical engineering* 33, 5 (2005), 590–602.

[28] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).

[29] Luman Wang, Avinash Malik, Partha S Roop, Leo K Cheng, and Niranchan Paskaranandavadivel. 2019. A formal approach for scalable simulation of gastric ICC electrophysiology. *IEEE Transactions on Biomedical Engineering* 66, 12 (2019), 3320–3329.

[30] Jim Woodcock, Peter Gorm Larsen, Juan Bicarregui, and John Fitzgerald. 2009. Formal methods: Practice and experience. *ACM computing surveys (CSUR)* 41, 4 (2009), 1–36.

[31] P. Ye, E. Entcheva, S. A. Smolka, and R. Grosu. 2008. Modelling excitable cells using cycle-linear hybrid automata. *IET Systems Biology* 2, 1 (1 2008), 24–32. https://doi.org/10.1049/iet-syb:20070001

[32] Eugene Yip, Sidharta Andalam, Partha S Roop, Avinash Malik, Mark L Trew, Weiwei Ai, and Nitish Patel. 2018. Towards the emulation of the cardiac conduction system for pacemaker validation. *ACM transactions on cyber-physical systems* 2, 4 (2018), 1–26.

[33] Eugene Yip, Sidharta Andalam, Partha S. Roop, Avinash Malik, Mark L. Trew, Weiwei Ai, and Nitish Patel. 2018. Towards the Emulation of the Cardiac Conduction System for Pacemaker Validation. *ACM Trans. Cyber-Phys. Syst.* 2, 4, Article 32 (July 2018), 26 pages. https://doi.org/10.1145/3134845