Towards Model Checking of Implantable Cardioverter Defibrillators

Houssam Abbas, Kuk Jin Jang, Zhihao Jiang, Rahul Mangharam
 Department of Electrical and Systems Engineering
 University of Pennsylvania, Philadelphia, PA, USA
 {habbas, jangkj, zhihaoj, rahulm}@seas.upenn.edu

ABSTRACT

Ventricular Fibrillation is a disorganized electrical excitation of the heart that results in inadequate blood flow to the body. It usually ends in death within a minute. A common way to treat the symptoms of fibrillation is to implant a medical device, known as an *Implantable Cardioverter Defibrillator* (ICD), in the patient's body. Model-based verification can supply rigorous proofs of safety and efficacy. In this paper, we build a hybrid system model of the human heart+ICD closed loop, and show it to be a STORMED system, a class of o-minimal hybrid systems that admit finite bisimulations. In general, it may not be possible to compute the bisimulation. We show that approximate reachability can yield a finite *simulation* for STORMED systems, and that certain compositions respect the STORMED property. The results of this paper are theoretical and motivate the creation of concrete model checking procedures for STORMED systems.

1. INTRODUCTION

Implantable Cardioverter Defibrillators (ICDs) are life-saving medical devices. An ICD is implanted under the shoulder, and connects directly to the heart muscle through two electrodes and continuously measures the heart's rhythm (Fig. 1). If it detects a potentially fatal accelerated rhythm known as Ventricular Tachycardia (VT), the ICD delivers a high-energy electric shock or sequence of pulses through the electrodes to reset the heart's electrical activity. Without this therapy, the VT can be fatal within seconds of onset. In the US alone, 10,000 people receive an ICD every month. Studies have presented evidence that patients implanted with ICDs have a mortality rate reduced by up to 31%.

Unfortunately, ICDs suffer from a high rate of *inappropriate therapy* due to poor detection of the current rhythm on the part of the ICD. In particular, a class of rhythms known as Supraventricular Tachycardias (SVTs) can fool the detection algorithms. Inappropriate shocks increase patient stress, reduce their quality of life, and are linked to

Figure 1: ICD connected to a human heart via two electrodes. The ICD monitors three electrical signals (known as electrograms) traversing the heart muscle.

increased morbidity. Depending on the particular ICD and its settings, the rates of inappropriate therapy can range from 46% to 62% of all delivered therapy episodes. Current practice for ICD verification relies heavily on testing and software cycle reviews. With the advent of computer models of the human heart, *Model-Based Design* (MBD) can supply rigorous evidence of safety and efficacy. This paper presents hybrid system models of the human heart and of the common modules of ICDs currently on the market, and shows that the closed loop formed by these models admits a finite bisimulation. The objective is to develop model checkers for ICDs to further their MBD process.

No work exists on ICD verification. Earlier work on verification of medical devices (formal or otherwise) focuses on pacemakers, which measure the timing of heart events. ICD algorithms are more complex than a pacemaker's, because an ICD also measures and processes the *morphology* of the electrical signal to distinguish many types of arrhythmias. This takes the model out of the realm of timed automata and into hybrid automata proper. Previous work on biological hybrid and/or nonlinear systems uses approximate reachability techniques to verify system invariants [7, 8], and demonstrates success in parameter space exploration.

The first contribution of this paper is to develop a hybrid system model of the heart and ICD measurement process (Section 3), of the ICD sensing process (Section 4), and of the algorithmic components of ICDs from most major manufacturers on the market (Section 5, see Fig. 2). We show that the composition of these three models admits a finite

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

HSCC '16, April 12–14, 2016, Vienna, Austria.

© 2016 ACM. ISBN 978-1-4503-3955-1/16/04...\$15.00

DOI: <http://dx.doi.org/10.1145/2883817.2883841>

Figure 2: The whole heart is modeled as a 2D mesh of cells (Section 3). The ICD electrodes are shown in the right atrium and ventricle. The electrogram signals measured through the electrodes are processed by the sensing module (see Section 4). The detection algorithm determines the current rhythm using the processed signal (Section 5).

bisimulation. The models presented here are the first formalization of ICD operation. To establish this result we use the theory of STORMED systems [15], a class of hybrid systems that have finite bisimulations. Our second contribution is two general results for STORMED systems. First we prove that parallel compositions of STORMED systems yield STORMED systems (Section 6). Secondly, we show that any definable over-approximate reach tubes can replace the exact trajectories of a STORMED system, yielding a system that still admits a finite *simulation* (Section 7). All proofs are in the online report [1].

2. HYBRID SYSTEMS AND SIMULATIONS

Definition 2.1. A hybrid automaton is a tuple

$$\mathcal{H} = (X, L, H_0, \{f_\ell\}, \text{Inv}, E, \{R_{ij}\}_{(i,j) \in E}, \{G_{ij}\}_{(i,j) \in E})$$

where $X \subset \mathbb{R}^n$ is the continuous state space equipped with the Euclidian norm $\|\cdot\|$, $L \subset \mathbb{N}$ is a finite set of modes, $H_0 \subset X \times L$ is an initial set, $\{f_\ell\}_{\ell \in L}$ determine the continuous evolutions with unique solutions, $\text{Inv} : L \rightarrow 2^X$ defines the invariants for every mode, $E \subset L^2$ is a set of discrete transitions, $G_{ij} \subset X$ is guard set for the transitions (so $\mathcal{H}$ transitions $i \rightarrow j$ when $x \in G_{ij}$), $R_{ij} : X \rightarrow X$ is an edge-specific reset function.

Set $H = L \times X$. Given $(\ell, x_0) \in H$, the flow $\theta_\ell(\cdot; x_0) : \mathbb{R}_+ \rightarrow \mathbb{R}^n$ is the solution to the IVP $\dot{x}(t) = f_\ell(x(t))$, $x(0) = x_0$.

The associated transition system is $T_{\mathcal{H}} = (H, E \cup \{\tau\}, \rightarrow, H_0)$ where H is the state set, $E \cup \{\tau\}$ is the label set for transitions, H_0 is the set of initial states, and $\rightarrow = (\bigcup_{e \in E} \xrightarrow{e} \cup \xrightarrow{\tau})$ where $(i, x) \xrightarrow{e} (j, y)$ iff $e = (i, j)$, $x \in G_{ij}$, $y = R_{ij}(x)$ and $(i, x) \xrightarrow{\tau} (j, y)$ iff $i = j$ and there exists a flow $\theta_i(\cdot; x)$ of $\mathcal{H}$ and $t \geq 0$ s.t. $\theta_i(t; x) = y$ and $\forall t' \leq t$, $\theta_i(t'; x) \in \text{Inv}(i)$. Let $\sim$ be an equivalence relation on H and $H/\sim$ the corresponding partition. Let $\mathcal{F}_t(H/\sim)$ be the coarsest bisimulation with respect to $\xrightarrow{\tau^1}$ respecting the partition $H/\sim$, and $\mathcal{F}_d(H/\sim) := \{(h_1, h_2) \mid (h_1 \xrightarrow{e} h'_1) \implies (\exists e' \in E, h'_2. h_2 \xrightarrow{e'} h'_2 \wedge h'_1 \sim h'_2)\} \cap H/\sim$ [15]. The iteration

$$W_0 = \mathcal{F}_t(H/\sim), \quad \forall i \geq 0, W_{i+1} = \mathcal{F}_d(\mathcal{F}_t(W_i)) \quad (1)$$

computes a bisimulation of $\mathcal{H}$. However it does not necessarily terminate for hybrid systems because the system's

¹I.e., $\mathcal{F}_t$ only considers the continuous transition relation: it is a bisimulation of $T_{\mathcal{H}}^c := (H/\sim, \{\tau\}, \xrightarrow{\tau}, H_0/\sim)$.

reach set might intersect a given block of $H/\sim$ an infinite number of times (see [11] for an example). The class of systems introduced in the next section has the property that the iteration does terminate for it and returns a finite $\mathcal{S}$.

Given a set of atomic propositions, if $\sim$ is s.t. $\eta \sim \eta'$ iff both states satisfy exactly the same atomic propositions, then model checking temporal logic properties can be done on the finite bisimulation instead of the possibly infinite $\mathcal{H}$.

2.1 O-minimality and STORMED systems

We give a very brief introduction to o-minimal structures. A more detailed introduction can be found in [11] and references therein. We are interested in sets and functions in $\mathbb{R}^n$ that enjoy certain finiteness properties, called order-minimal sets (o-minimal). These are defined inside structures $\mathcal{A} = (\mathbb{R}, <, +, -, \cdot, \exp, \dots)$. The subsets $Y \subset \mathbb{R}^n$ we are interested in are those that are *definable* using first-order formulas φ : $Y = \{(a_1, \dots, a_n) \in \mathbb{R}^n \mid \varphi(a_1, \dots, a_n)\}$. (First-order formulas use the boolean connectives and the quantifiers $\exists, \forall$). The atomic propositions from which the formulas are recursively built allow only the operations of the structure $\mathcal{A}$ on the real variables and constants, and the relations of $\mathcal{A}$ and equality. For example $2x - 3.6y < 3z$ and $x = y$ are valid atomic propositions of the structure $\mathcal{L}_{\mathbb{R}} = (\mathbb{R}, <, +, -, \cdot)$, while $\cosh(x) < 3z$ is not because $\cosh$ is not in the structure. These structures are already sufficient to describe a set of dynamics rich enough for our purposes and for various classes of linear systems.

Definition 2.2. A theory of $(\mathbb{R}, \dots)$ is o-minimal if the only definable subsets of $\mathbb{R}$ are finite unions of points and (possibly unbounded) intervals. A function $f : x \mapsto f(x)$ is o-minimal if its graph $\{(x, y) \mid y = f(x)\}$ is a definable set.

We use the terms o-minimal and definable interchangeably, and they refer to $\mathcal{L}_{\text{exp}} = (\mathbb{R}, <, +, -, \cdot, \exp)$ which is known to be o-minimal. The dot product between $x, y \in \mathbb{R}^n$ is denoted $x \cdot y$, and $d(Y, S) = \inf\{\|y - s\| \mid (y, s) \in Y \times S\}$.

Definition 2.3. [15]. A STORMED hybrid system (SHS) Σ is a tuple $(\mathcal{H}, \mathcal{A}, \phi, b_-, b_+, d_{\min}, \epsilon, \zeta)$ where $\mathcal{H}$ is a hybrid automaton, $\mathcal{A}$ is an o-minimal structure, $d_{\min}, \epsilon, \zeta$ are positive reals, $b_-, b_+ \in \mathbb{R}$ and $\phi \in X$ such that:

(S) The system is $d_{\min}$ -separable, meaning that for any $e = (\ell, \ell') \in E$ and $\ell'' \neq \ell'$, $d(R_e(G_{(\ell, \ell')}), G_{(\ell', \ell'')}) > d_{\min}$ ²

²The original definition of separability [15] required the guards themselves to be separated, which is insufficient to guarantee that if $\mathcal{H}$ flows, it flows a uniform minimum dis-

(T) The flows (i.e., the solutions of the ODEs) are Time-Independent with the Semi-Group property (TISG), meaning that for any $\ell \in L, x \in X$, the flow θ_ℓ starting at (ℓ, x) satisfies: 1) $\theta_\ell(0; x) = x$, 2) for every $t, t' \geq 0$, $\theta_\ell(t + t'; x) = \theta_\ell(t'; \theta_\ell(t; x))$

(O) All the sets and functions of $\mathcal{H}$ are definable in the o-minimal structure $\mathcal{A}$

(RM) The resets and flows are monotonic with respect to the same vector ϕ , meaning that

1) (Flow monotonicity) for all $\ell \in L, x \in X$ and $t, \tau \geq 0$, $\phi \cdot (\theta_\ell(t + \tau; x) - \theta_\ell(t; x)) \geq \epsilon \|\theta_\ell(t + \tau; x) - \theta_\ell(t; x)\|$, and
 2) (Reset monotonicity) for any edge $(\ell, \ell') \in E$ and any $x^-, x^+ \in X$ s.t. $x^+ = R_{\ell, \ell'}(x^-)$,

1. if $\ell = \ell'$, then either $x^- = x^+$ or $\phi \cdot (x^+ - x^-) \geq \zeta$
2. if $\ell \neq \ell'$, then $\phi \cdot (x^+ - x^-) \geq \epsilon \|x^+ - x^-\|$

(ED) Ends are Delimited: for all $e \in E$ we have $\phi \cdot x \in (b_-, b_+)$ for all $x \in G_e$

Intuitively, the above conditions imply the trajectories of the system always move a minimum distance along ϕ whether flowing or jumping, which guarantees that no area of the state space will be visited infinitely often. This is at the root of the finiteness properties of STORMED systems.

Theorem 2.1. [15] Let $\mathcal{H}$ be a STORMED hybrid system, and let $\mathcal{P}$ be an o-minimal partition of its hybrid state space. Then $\mathcal{H}$ admits a finite bisimulation that respects $\mathcal{P}$.

3. HEART MODEL

For the verification of ICDs, we adopt the cellular automata (CA)-based heart model developed in [14],[5]. This model lies in-between high spatial fidelity but slow to compute PDE-based whole heart models, and low spatial fidelity but very fast-to-compute automata-based models [12]. Ionic currents [9] and PDE-based models may be more accurate but are not currently amenable to formal verification (however see [7] for reachability analysis of discretized PDEs). CA-based models were used in [2] and [4]. This paper's model also has the important advantage of forming the basis of software used to train electrophysiologists, and allows interactive simulation of surgical procedures like ablation [13]. In particular, it can simulate tachycardias.

This paper's automata: All hybrid automata in this paper have the whole state space as invariants and transitions are urgent (taken immediately when the guard is enabled). The electrogram (EGM) voltage signal s has upper and lower bounds. We also observe that, as will be seen in Section 5, i) while observing a rhythm, the ICD will always reach a decision of VT or SVT in finite time ii) at which point it resets its controlled (software) variables so new values are computed for the next arrhythmia episode. So while the heart can beat indefinitely, for the purposes of ICD verification, there's a uniform upper bound on the length of time of any execution. Let $D \geq 0$ be this duration (D is on the order of 30sec depending on device settings. More recent ICD models might wait for longer for self-termination). Therefore, every mode of every automaton in what follows has a transition to mode End in which

tance along ϕ . Indeed assume the guards are separated. If $x \in G_{(\ell, \ell')}$ and $y = R_{(\ell, \ell')}(x)$, it can be that $y \in G_{(\ell', \ell'')}$ and thus a jump happens, even though $G_{(\ell, \ell')}$ and $G_{(\ell', \ell'')}$ are separated. Therefore we need $d(y, G_{(\ell', \ell'')}) > d_{min}$ for all $y \in R_e(G_e)$, which is the condition we use in Def. 2.3. The properties of SHS, in particular the existence of finite bisimulation, are therefore preserved by this change.

Figure 3: Hybrid model $\mathcal{H}_c$ of one cell of the heart model. AP figure from [6]. $V_{th,2} > V_{th}$, $V_{max,2} < V_{max}$

time does not progress. We don't show these transitions in the automata figures of this paper to avoid congestion.

3.1 Cellular automata model

The heart has two upper chambers called the *atria* and two lower chambers called the *ventricles* (Fig. 1) The synchronized contractions of the heart are driven by electrical activity. Under normal conditions, the SinoAtrial (SA) node (a tissue in the right atrium) spontaneously *depolarizes*, producing an electrical wave that propagates to the atria and then down to the ventricles (Fig.2) In this model, the myocardium (heart's muscle) is treated as a 2D surface (so it has no depth), and discretized into *cells*, which are simply regions of the myocardium (Fig. 2). Thus we end up with N^2 cells in a square N -by- N grid. A cell's voltage changes in reaction to current flow from neighboring cells, and in response to its own ion movements across the cell membrane. This results in an *Action Potential* (AP).

Fig. 3 shows how the AP is generated by a given cell [10]: in its quiescent mode (Phase 4), a cell (i, j) in the grid has a cross-membrane voltage $V(i, j, t)$ equal to $V_{min} < 0$. As it gathers charge, $V(i, j, t)$ increases until it exceeds a threshold voltage V_{th} . In Phase 0, the voltage then experiences a very fast increase (Phase 0), called the upstroke, to a level $V_{max} > 0$, after which it decreases (Phase 1) to a plateau (Phase 2). It stays at the plateau level for a certain amount of time **PD** then decreases linearly to below V_{th} (Phase 3 - ERP). Once below V_{th} it is said to be in the Relative Refractory Period (Phase 3 - RRP). In Phase 3 - RRP, the cell can be depolarized a second time, albeit at a higher threshold $V_{th,2}$, slower and to a lower plateau level $V_{max,2} < V_{max}$ (Upstroke 2). Otherwise, when the voltage reaches V_{min} again, the cell enters the quiescent stage again. This model is suitable for both pacemaker and non-pacemaker cells, the main differences being in the duration of the plateau (virtually non-existent for pacemaker cells), and the duration of phases 0 and 4 (both are shorter for pacemaker cells). In Fig. 3, $V(i, j) \in \mathbb{R}$ denotes the voltage in cell (i, j) of the grid, and $V = (V(1, 1), \dots, V(N, N))^T$ in $\mathbb{R}^{N^2}$ groups the cross-membrane voltages of all cells in the heart. The whole heart model $\mathcal{H}_{CA}$ is the parallel composition of these N^2 single-cell models. The $(i, j)^{th}$ cell's voltage at time t

Figure 4: $\mathcal{H}_{Sense}$. States not shown in a mode have a 0 derivative, e.g., $eF = 0$ in all modes.

in Phase 4 depends on that of its neighbors and its own as follows [14]

$$\begin{aligned} \dot{V}(i, j, t) &= \frac{[V(i-1, j, t) + V(i+1, j, t) - 2V(i, j, t)]}{R_h(i, j)} \\ &+ \frac{[V(i, j-1, t) + V(i, j+1, t) - 2V(i, j, t)]}{R_v(i, j)} \\ &= a(i, j)^T V(t), \quad a(i, j) \in \mathbb{R}^{N^2} \end{aligned} \quad (2)$$

where R_h, R_v are conduction constants that can vary across the myocardium. Thus V evolves according to a linear ODE $\dot{V} = AV$ where A is the matrix whose rows are the $a(i, j)$. The two states t and t_p are clocks. Clock t_p keeps track of the value of the last discrete jump. We will use this arrangement in all our models: it avoids resetting the clocks which preserves Reset Monotonicity.

ICDs observe the electrical activity through three channels (Fig. 1). Each signal is called an electrogram (EGM) signal. The signal read on a channel is given by [5]:

$$s(t) = \frac{1}{K} \sum_{i,j} \left(\frac{1}{\|p_{i,j} - p_0\|} - \frac{1}{\|p_{i,j} - p_1\|} \right) \dot{V}(i, j, t) \quad (3)$$

where $\|\cdot\|$ is the Euclidian norm, p_0 and p_1 are the electrodes' positions and $p_{i,j}$ is the position of the $(i, j)^{th}$ cell on the 2D myocardium ($p_0, p_1, p_{i,j} \in \mathbb{R}^2$). Positions p_0, p_1 should be chosen different from $p_{i,j}$ to avoid infinities.

Extensions. The Action Potential Duration (APD) restitution mechanism as modeled in [14] can be included in this model without changing its formal properties.

We now state the main result of this section.

Theorem 3.1. *Let $\mathcal{H}_{CA}$ be the whole heart cellular automaton model obtained by parallel composition of N^2 models $\mathcal{H}_c$ with state vector $x = [V, t, t_p, s] \in \mathbb{R}^{N^2} \times \mathbb{R}^3$. Assume that all executions of the system have a duration of $D \geq 0$. Then $\mathcal{H}_{CA}$ is STORMED.*

4. ICD SENSING

Sensing is the process by which cardiac signals s measured through the leads of the ICD are converted to timing events. The ICD declares events when the signal exceeds a dynamically-adjusted threshold Th .

Fig. 4 shows the model $\mathcal{H}_{Sense}$ of the sensing algorithm, and Fig. 5 illustrates its operation. The sensing takes place

Figure 5: Example of dynamic threshold adjustment in ICD sensing algorithm. The shown signal is rectified.

on the rectified EGM signal $y = |s|$. After an event is declared at the current threshold value ($y(t) \geq Th(t)$ in Fig. 4), the algorithm tracks the signal in order to measure the next peak's amplitude (Peak Tracking). For a duration $MinTP$ (min tracking period) the latest peak is saved in y_M . A variable f indicates that a peak was found. After a peak is found ($f == 1$) and after the end of the tracking period, the algorithm enters a fixed *Blanking Period* (Blanking), during which additional events are ignored. On the transition to Blanking, Th, Th_0 and the exponential factor of decay eF are updated. At the end of the blanking period, the algorithm transitions to the Exponential Decay mode in which Th decays exponentially from Th_0 to a minimum level (Exponential Decay), and stays there for at least a sampling period of $MinDecP$. Different manufacturers may use a step-wise decay instead of exponential, but the principle is the same. Local peak detection is modeled via the $\dot{y} = 0 \wedge \dot{y} < 0$ transition. While $y = |s|$ is non-differentiable at 0, the peak will occur away from 0, as shown in Fig. 5. States t, t_p are clocks and $minTh$ and TC are constant parameters.

Theorem 4.1. $\mathcal{H}_{Sense}$ is STORMED.

5. ARRHYTHMIA DETECTION

A sustained Ventricular Tachycardia (VT) (or Ventricular Fibrillation (VF)) can be fatal whereas a SupraVentricular Tachycardia (SVT) is usually not fatal, so the ICD's main task is to discriminate VT from SVT and deliver therapy to the former only [3]. Most VT/SVT detection algorithms found in ICDs today are composed of individual *discriminators*. A discriminator is a software function whose task is to decide whether the current arrhythmia is SVT or VT. No one discriminator can fully distinguish between SVT and VT. Thus a detection algorithm is often a decision tree built using a number of discriminators *running in parallel*. We have modeled each discriminator in Boston Scientific's detection algorithm as a hybrid automaton. **The ICD system is thus $\mathcal{H}_{ICD} = \mathcal{H}_{Sense} || \mathcal{H}_{Detection- Algo}$ where $\mathcal{H}_{Detection- Algo}$ is the parallel composition of the discriminator automata.** We now illustrate the models we created with three discriminators and prove they are SHS.

5.1 Three Consecutive Fast Intervals

Our first module simply detects whether three consecutive fast intervals have occurred, where 'fast' means the interval length, measured between 2 consecutive peaks on the EGM signal, is shorter than some pre-set amount. See Fig. 6. States t and t_p are clocks as before. The vector L_3 is three-dimensional, and stores the values of the last three intervals. The event $VEvent?$ is shorthand for the transition $y(t) \geq Th$

Figure 6: Three Consecutive Fast Intervals $\mathcal{H}_{TCFI}$

Figure 7: EGMs of different origin have different morphologies, while EGMs of similar origins have very similar morphologies.

being taken by the $\mathcal{H}_{Sense}$ automaton. In other words, it indicates a ventricular event. Then L_3 gets reset to $L_3^+ = (z_1, z_2, z_3)^+ := \text{Circulate}(L_3, t - t_p)$ where

$$L_3^+ = \begin{pmatrix} z_2 \\ z_3 \\ t - t_p \end{pmatrix} = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix} L_3 + \begin{pmatrix} 0 \\ 0 \\ t - t_p \end{pmatrix} \quad (4)$$

Lemma 5.1. $\mathcal{H}_{TCFI}$ is STORMED.

5.2 Vector Timing Correlation

It has been clinically observed that a depolarization wave originating in the ventricles (as produced during VT for example) will in general produce a different EGM morphology than a wave originating in the atria (as produced during SVT) [3]. See Fig. 7. A morphology discriminator measures the correlation between the morphology of the current EGM and that of a stored *template* EGM acquired during normal sinus rhythm. If the correlation is above a pre-set threshold for a minimum number of beats, then this is an indication that the current arrhythmia is supraventricular in origin. Otherwise, it might be of ventricular origin.

Boston Scientific's implementation of a morphology discriminator is called Vector and Timing Correlation (VTC). VTC first samples 8 *fiducial* points $s_i, i = 1, \dots, 8$ on the current EGM s at pre-defined time instants. Let $s_{m,i}$ be the corresponding points on the template EGM. A simple 0-shift correlation ρ_{new} is calculated between the two sequences. If 3 out of the last 10 calculated correlation values exceed the threshold, then SVT is decided and therapy is withheld.

The system of Fig. 8 implements the VTC discriminator. As before, t is a local clock. μ accumulates the values of the current EGM, α accumulates the product $s_i s_{m,i}$, β accumulates s_i^2 . State w is an auxiliary state we need to establish the STORMED property. $\vec{v}$ is a 10D binary vector: $v_i = -1$ if the i^{th} correlation value fell below the threshold, and is $+1$ otherwise. L_3 is the state of $\mathcal{H}_{TCFI}$: the guard condition $L_3 \leq th$ indicates that all its entries have values less than the tachycardia threshold, which is when $\mathcal{H}_{VTC}$ starts computing. *WindowEnds* indicates the 'end' of an EGM, measured as a window around the peak sensed by $\mathcal{H}_{Sense}$.

Figure 8: VTC calculation. iT_s is the sampling time for the i th fiducial point, $i = 1, \dots, 8$. $R_{21}, \dots, R_{28}$ are the corresponding resets. For clarity of the figure, 8 transitions are represented on the same edge R_{2i} .

Figure 9: Stability discriminator.

Lemma 5.2. $\mathcal{H}_{VTC}$ is STORMED.

5.3 Stability discrimination

Stability refers to the variability of the peak-to-peak cycle length. A rhythm with large variability (above a pre-defined threshold) is said to be *unstable*. The Stability discriminator is used to distinguish between atrial fibrillation, which is usually unstable, and VT, which is usually stable.

The Stability discriminator shown in Fig. 9 simply calculates the variance of the cycle length over a fixed period called a Duration (measured in seconds). Let $DL \geq 0$ be the Duration length. The events *DurationBegins?* and *DurationEnds?* indicate the transitions of a simple system that measures the lapse of one Duration (not shown here). State t is a clock, L_1 accumulates the sum of interval lengths (and will be used to compute the average length), L_2 accumulates the squares of interval lengths, and κ is a counter that counts the number of accumulated beats. σ_2 is assigned the value of the variance given by $\frac{1}{\kappa}[L_2 - L_1^2/\kappa]$

Lemma 5.3. $\mathcal{H}_{Stab}$ is STORMED.

Now that each system was shown to be STORMED, it remains to establish that their parallel composition is STORMED. This result does not hold in general - Thm. 6.1 gives conditions under which parallel composition respects the STORMED property. Intuitively, we require that whenever a sub-collection of the systems jumps, the remaining systems that did not jump are separated from all of their respective guards by a uniform distance. This is a requirement that can be shown to hold for our systems by modeling various minimal delays in the systems' operation. We may now state:

Theorem 5.1. Consider the collection of systems $\mathcal{H}_{CA}$,

$\mathcal{H}_{ICD} = \mathcal{H}_{Sense} || \mathcal{H}_{Detection- Algo}$ where the latter is the parallel composition of the discriminator systems. This collection satisfies the hypotheses of Thm. 6.1 (Section 6) and therefore the parallel system $\mathcal{H}_{CA} || \mathcal{H}_{ICD}$ is STORMED and has a finite bisimulation.

6. COMPOSING STORMED SYSTEMS

The results in this section and the next apply to SHS in general, including those with time-unbounded operation. We write $[m] = \{1, \dots, m\}$. In this section given hybrid systems $\mathcal{H}_1, \dots, \mathcal{H}_m, x^i, G^i, \theta^i, \dots$ etc refer to a state, guard, flow ... of system $\mathcal{H}_i$. Recall that $\theta_\ell(t; x)$ is the flow starting at (ℓ, x) . The parallel composition $\mathcal{H} = \mathcal{H}_1 || \dots || \mathcal{H}_m$ is defined in the usual way: $\mathcal{H}.X = \Pi_i X^i$, $\mathcal{H}.L = \Pi_i L^i$, $\mathcal{H}.H_0 = \Pi_i H_0^i$, $Inv(\ell) = \Pi_i Inv^i(\ell^i)$, and the flow $\theta_\ell(x, t) = [\theta_{\ell^1}^1(x^1, t), \dots, \theta_{\ell^m}^m(x^m, t)]^T$. The system jumps if any of its subsystems jumps. When a guard of a subsystem is satisfied, the state of that subsystem is reset according to its reset map. The guards are disjoint to avoid non-determinism.

We show that the parallel composition of SHS is still a SHS. In general $\mathcal{H}$ is not separable: indeed for any candidate value of d_{min} , one could find a transition (i, j) of $\mathcal{H}$ due to, say, a jump of $\mathcal{H}_1$, s.t. at that moment x^2 is closer than d_{min} to one of its own guards, say $G_{(j^2, k^2)}^2$. This causes $\mathcal{H}$ to further jump $j \rightarrow k$ without having traveled the requisite minimum distance, thus violating the separability of $R_{ij}(G_{ij})$ and G_{jk} . Therefore we need to impose an extra condition on minimum separability across sub-systems.

Theorem 6.1. *Let $\Sigma_i = (\mathcal{H}_i, \mathcal{A}, \phi^i, b^{i,-}, b^{i,+}, d_{min}^i, \varepsilon^i, \zeta^i)$, $i = 1, \dots, m$ be deterministic SHS defined using the same underlying o-minimal structure, and where each state space X^i is bounded by B_{X^i} .*

Define parallel composition $\Sigma = (\mathcal{H}, \mathcal{A}, \phi, b^-, b^+, d_{min}, \varepsilon, \zeta)$ where $\mathcal{H} = \mathcal{H}_1 || \dots || \mathcal{H}_m$, $\phi = (\phi^1, \dots, \phi^m)^T \in \mathbb{R}^{mn}$, $b^{i,-} = \inf_{x \in X} \phi \cdot x$, $b^{i,+} = \sup_{x \in X} \phi \cdot x$, $\varepsilon = \min(\min_i \varepsilon^i, \min_i \frac{\zeta^i}{B_{X^i}})$, $\zeta = \min_i \zeta^i$ and

$$d_{min} = \min_{I \subset [m]} \left(\min_{i \in I} d_{min}^i, \min_{i \in I, j \in [m] \setminus I} d_{min}^{ij} \right)$$

Assume that the following Collection Separability condition holds: for all $i, j \leq m, i \neq j$ there exists $d_{min}^{ij} > 0$ s.t. if $x \in X$ is in the reachable set of $\mathcal{H}$ and $x^i \in G_e^i \wedge x^j \notin G_{e'}^j, \forall e' \in E^j$ then $d(x^j, G_{e'}^j) > d_{min}^{ij}$ for all $e' \in E^j$ where E^j is the edge set of Σ_j and $G_{e'}^j$ is a guard of Σ_j on edge $e' \in E^j$. Then Σ is STORMED.

7. FINITE STORMED SIMULATION

In general it is not possible to compute the reach sets required by the iteration (1) exactly unless the underlying theory is decidable. The $\mathcal{H}_{ICD} || \mathcal{H}_{CA}$ closed loop is definable in $\mathcal{L}_{exp}$, and the latter is not known to be decidable. Here we show that if an approximate reachability tool with definable over-approximations is available for the continuous dynamics, it can be used in (1) to yield a finite simulation. Since we only have a simulation, counter-examples on the abstraction should be validated in a CEGAR-like fashion.

Lemma 7.1. *Let $\Sigma = (\mathcal{H}, \dots)$ be a SHS and $\sim$ an equivalence relation on X . For any mode ℓ of $\mathcal{H}$, the dynamical system $\mathcal{D}$ with state space $X = \mathcal{H}.X$ and set-valued flow $\Theta(t; x) = \{y \in \mathbb{R}^n \mid \|y - \theta_\ell(t; x)\|^2 \leq \epsilon^2\}$ admits a finite simulation $\mathcal{S}_\ell$ that respects $\sim$.*

Let $\mathcal{F}_t^\epsilon(\mathcal{P}) := \cap_{\ell} \mathcal{S}_{\ell \in L}$ where $\mathcal{P} = X / \sim$. $\mathcal{F}_t^\epsilon$ refines all the $\mathcal{S}_\ell$'s, and it is a finite simulation of $\mathcal{H}$ by itself w.r.t. the continuous transition $\xrightarrow{\tau}$.

Theorem 7.1. *Let $\mathcal{H}$ be a STORMED hybrid system, and $\mathcal{P}$ be a finite definable partition of its state space. Define*

$$W_0 = \mathcal{F}_t^\epsilon(\mathcal{P}), \quad \forall i \geq 0, W_{i+1} = \mathcal{F}_t^\epsilon(\mathcal{F}_d(W_i)) \quad (5)$$

Then there exists $U \in \mathbb{N}$ s.t. $W_{U+1} = W_U$ and $\mathcal{F}_t^\epsilon(W_U)$ is a simulation of $\mathcal{H}$ by itself.

This paper has presented the first formal models of ICD operation and shown that they admit finite bisimulations by proving new results in the theory of STORMED systems.

8. REFERENCES

- [1] H. Abbas, K. J. Jang, Z. Jiang, and R. Mangharam. Model checking implantable cardioverter defibrillators. 2016. Arxiv 1512.08083.
- [2] E. Bartocci, F. Corradini, M. D. Berardini, E. Entcheva, S. Smolka, and R. Grosu. Modeling and simulation of cardiac tissue using hybrid I/O automata. *Th. Com. Sci.*, 410(33), 2009.
- [3] Boston Scientific Corporation. The Compass - Technical Guide to Boston Scientific Cardiac Rhythm Management Products. *Device Documentation*, 2007.
- [4] T. Chen, M. Diciolla, M. Kwiatkowska, and A. Mereacre. Quantitative verification of implantable cardiac pacemakers over hybrid heart models. *Information and Computation*, 236:87 – 101, 2014.
- [5] D. D. Correa de Sa, N. Thompson, J. Stinnett-Donnelly, P. Znojkwicz, N. Habel, J. G. Muller, J. H. Bates, J. S. Buzas, and P. S. Spector. Electrogram fractionation. *Circ Arrhythm Electrophysiol*, 55:909 – 916, Dec 2011.
- [6] R. Hood. The EP Lab. Accessed 10/20/2015.
- [7] Z. Huang, C. Fan, A. Mereacre, S. Mitra, and M. Kwiatkowska. Invariant verification of nonlinear hybrid automata networks of cardiac cells. In A. Biere and R. Bloem, editors, *CAV*. 2014.
- [8] M. A. Islam, R. DeFrancisco, C. Fan, R. Grosu, S. Mitra, and S. Smolka. Model checking tap withdrawal in c. elegans. In *HSB*. 2015.
- [9] M. A. Islam, A. Murthy, A. Girard, S. A. Smolka, and R. Grosu. Compositionality results for cardiac cell dynamics. *HSCC*, 2014.
- [10] R. Klabunde. *Cardiovascular electrophysiology concepts*. Lippincott-Williams, 2 edition, 2011.
- [11] G. Lafferriere, G. J. Pappas, and S. Sastry. O-minimal hybrid systems. *Mathematics of Control, Signals and Systems*, 13(1):1–21, 2000.
- [12] M. Pajic, Z. Jiang, I. Lee, O. Sokolsky, and R. Mangharam. Safety-critical medical device development using the upp2sf model translation tool. *ACM Trans. Embed. Comput. Syst.*, 13(4), 2014.
- [13] P. S. Spector. Visible EP. Accessed 10/20/2015.
- [14] P. S. Spector, N. Habel, B. E. Sobel, and J. H. Bates. Emergence of complex behavior: An interactive model of cardiac excitation provides a powerful tool for understanding electric propagation. *Circulation: Arrhythmia and Electrophysiology*, 4(4):586–591, 2011.
- [15] V. Vladimerou, P. Prabhakar, M. Viswanathan, and G. Dullerud. Stormed hybrid systems. In *Automata, Languages and Programming*. 2008.