Closing the Loop for Medical CPS: From Verified Models to Verified Code and Beyond

Houssam Abbas

University of Pennsylvania

habbas@seas.upenn.edu

From Verified Models to Verified Code

PART I
From Verified Models to Verified Code for Medical Devices
(NSF CPS Large 2010-2015)

PART II
Computer-Aided Clinical Trials
(NSF Frontiers 2015-2020)

Part III
Bringing formal and approximate approaches to cardiology

Medical device recalls due to software

More problems…

1996: 10% of all medical device recalls were caused by software-related issues.

2008-12: **15% of all** the medical device recalls (Class I, II & III) due to software

Medical device recalls due to software

More problems…

1996: 10% of all medical device recalls were caused by software-related issues.

2008-12: **15% of all** the medical device recalls (Class I, II & III) due to software

To more people…

Every month: 10,000 new patients implanted with a defibrillator in the US

2005-2011: Virtually 60 countries saw increases in implant numbers

Open-Loop Testing

1996: 10% of all medical device recalls were caused by software-related issues.

2008-12: **15% of *all*** the medical device recalls (Class I, II & III) due to software

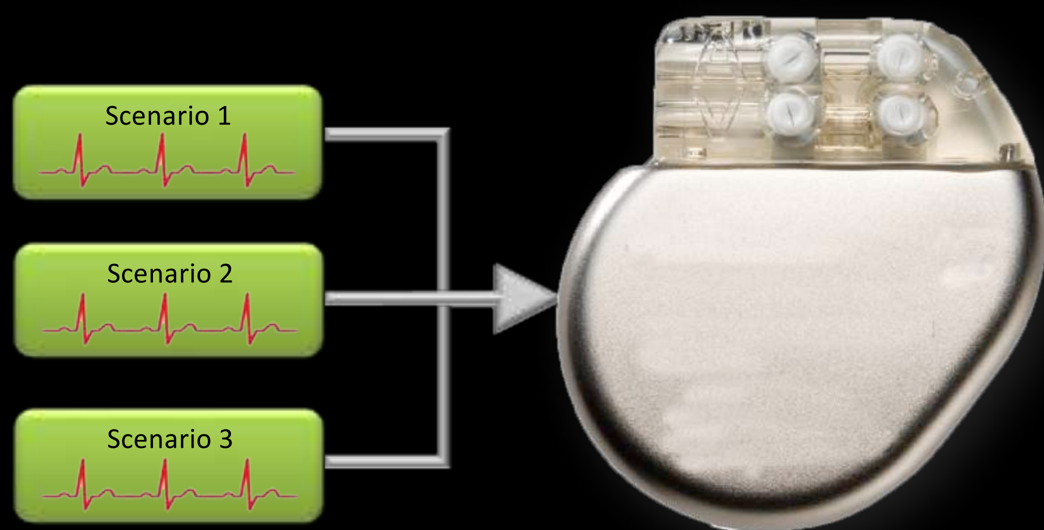

MUST TEST THE CLOSED LOOP

1996: 10% of all medical device recalls were caused by software-related issues.

2008-12: **15% of *all*** the medical device recalls (Class I, II & III) due to software

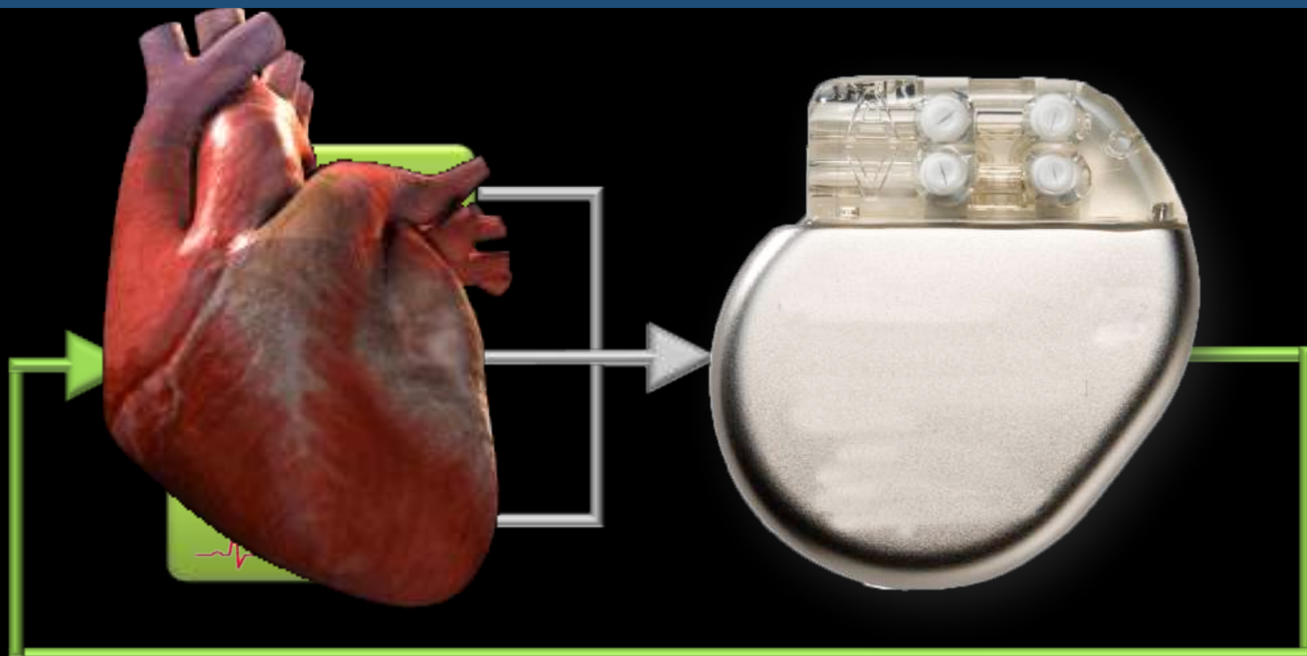

Use physiological models

1996: 10% of all medical device recalls were caused by software-related issues.

2008-12: **15% of *all*** the medical device recalls (Class I, II & III) due to software

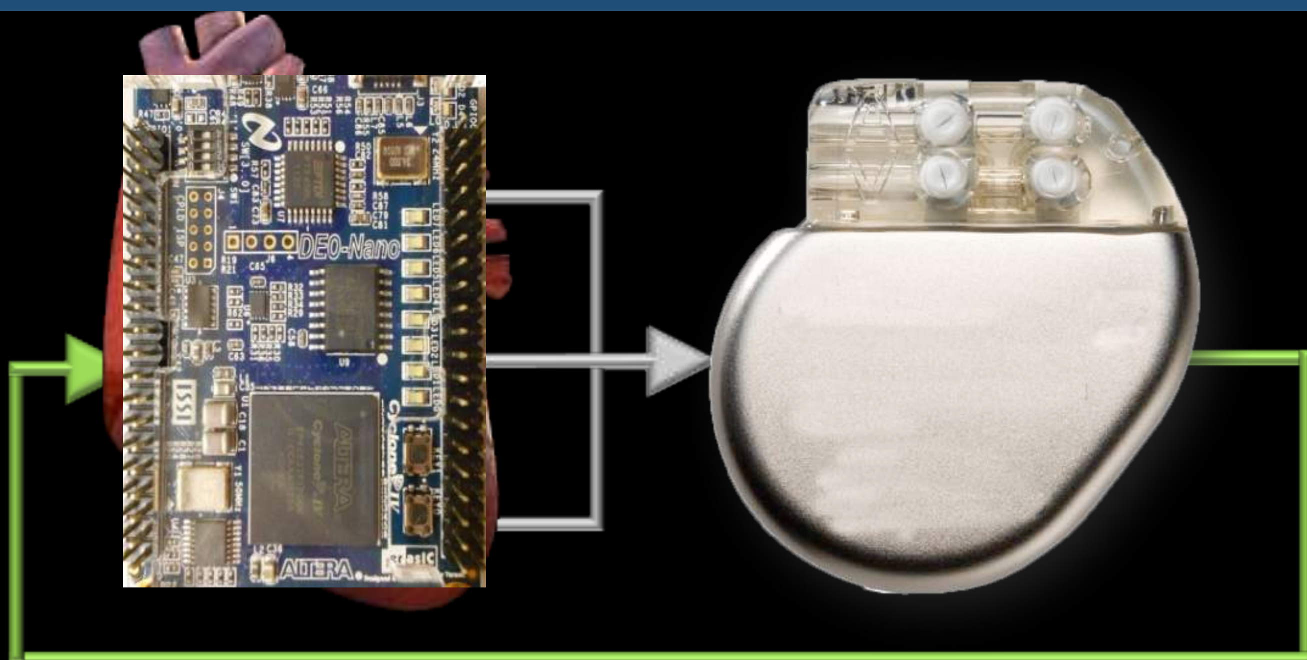

Ingredient 1: Heart model

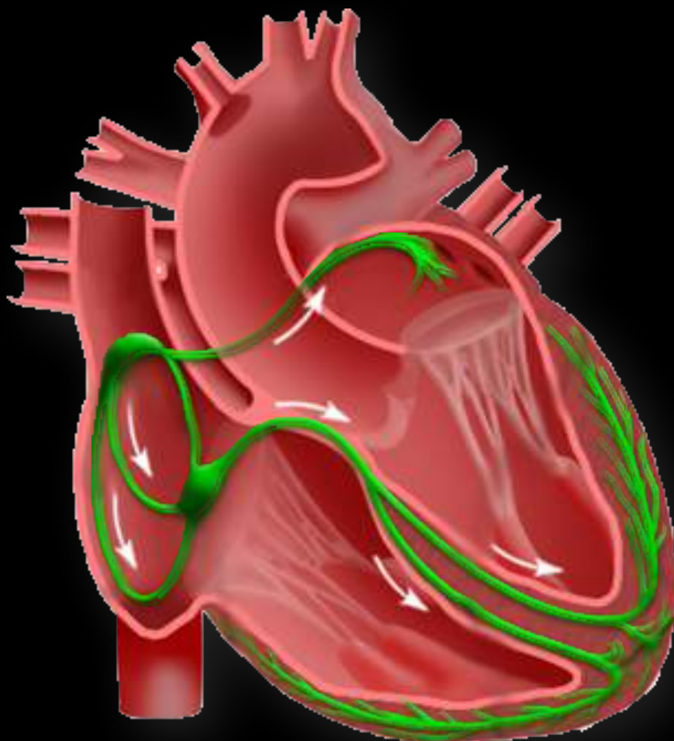

Ingredient 1: Heart model

Ingredient 1: Heart model

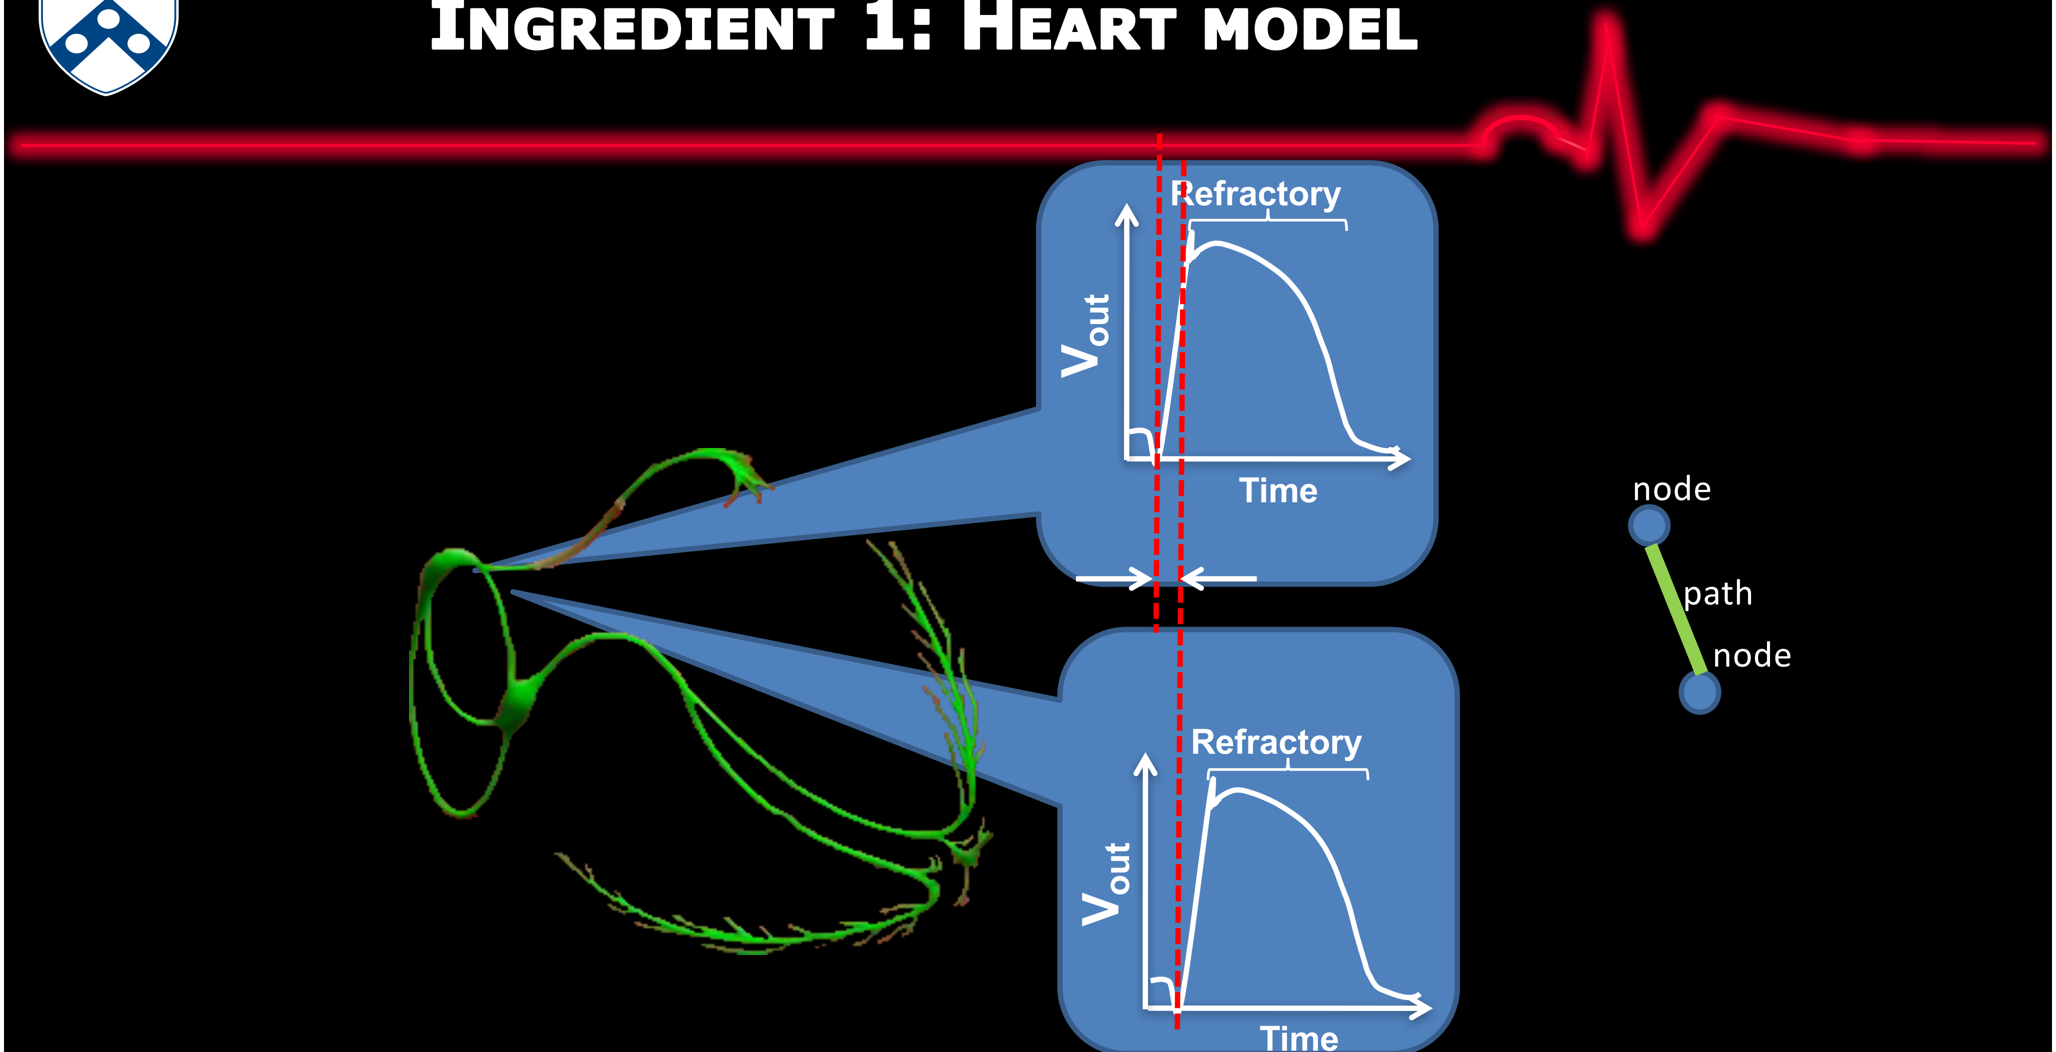

Timed Automata Heart Model

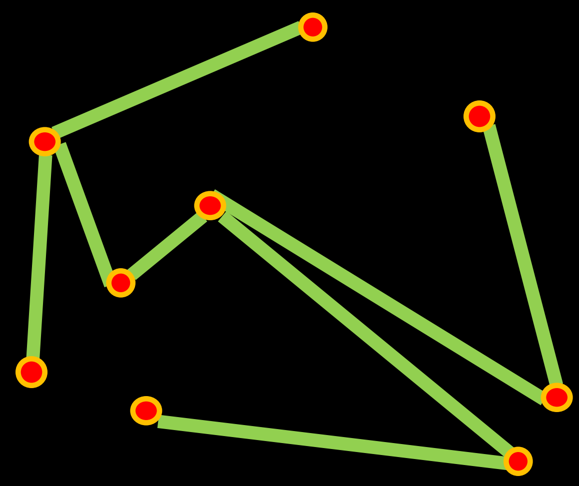

Timed Automata Heart Model

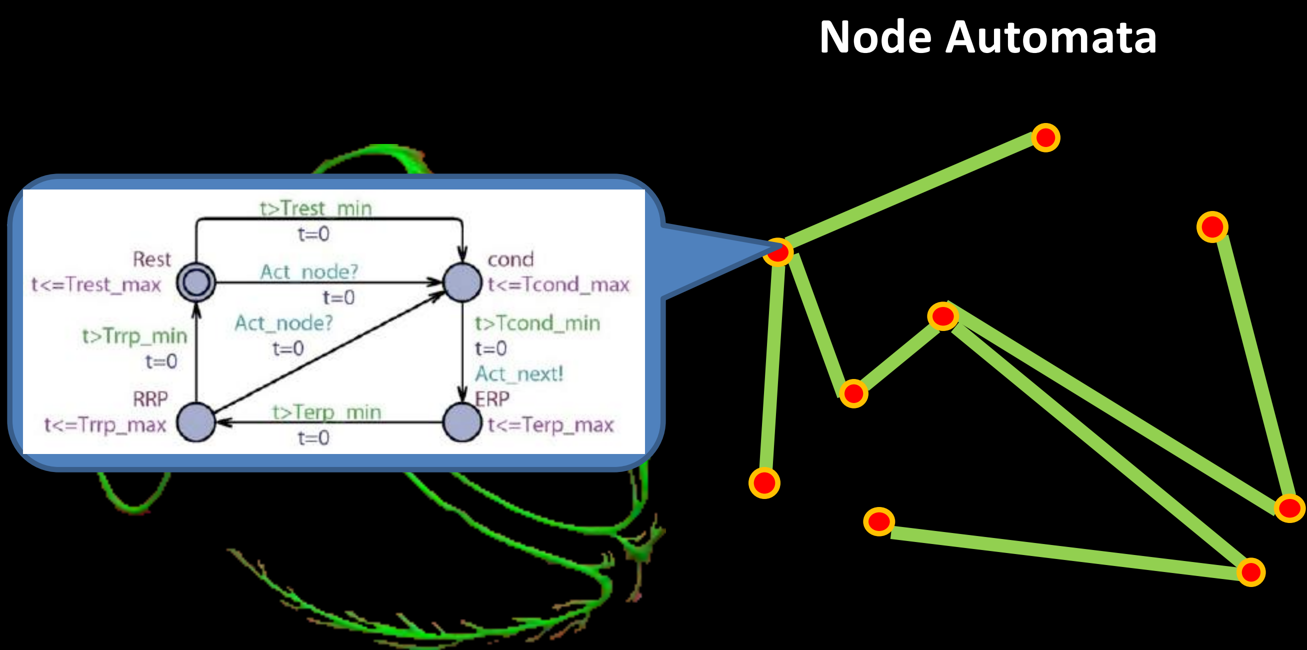

Timed Automata Heart Model

Path Automata

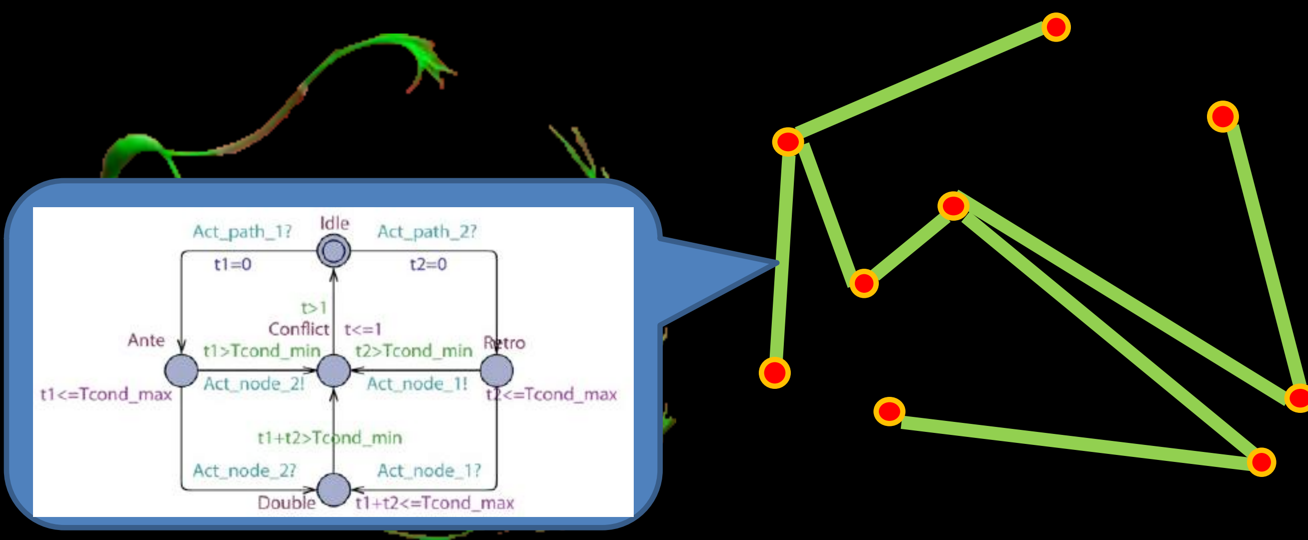

Timed automata Heart Model

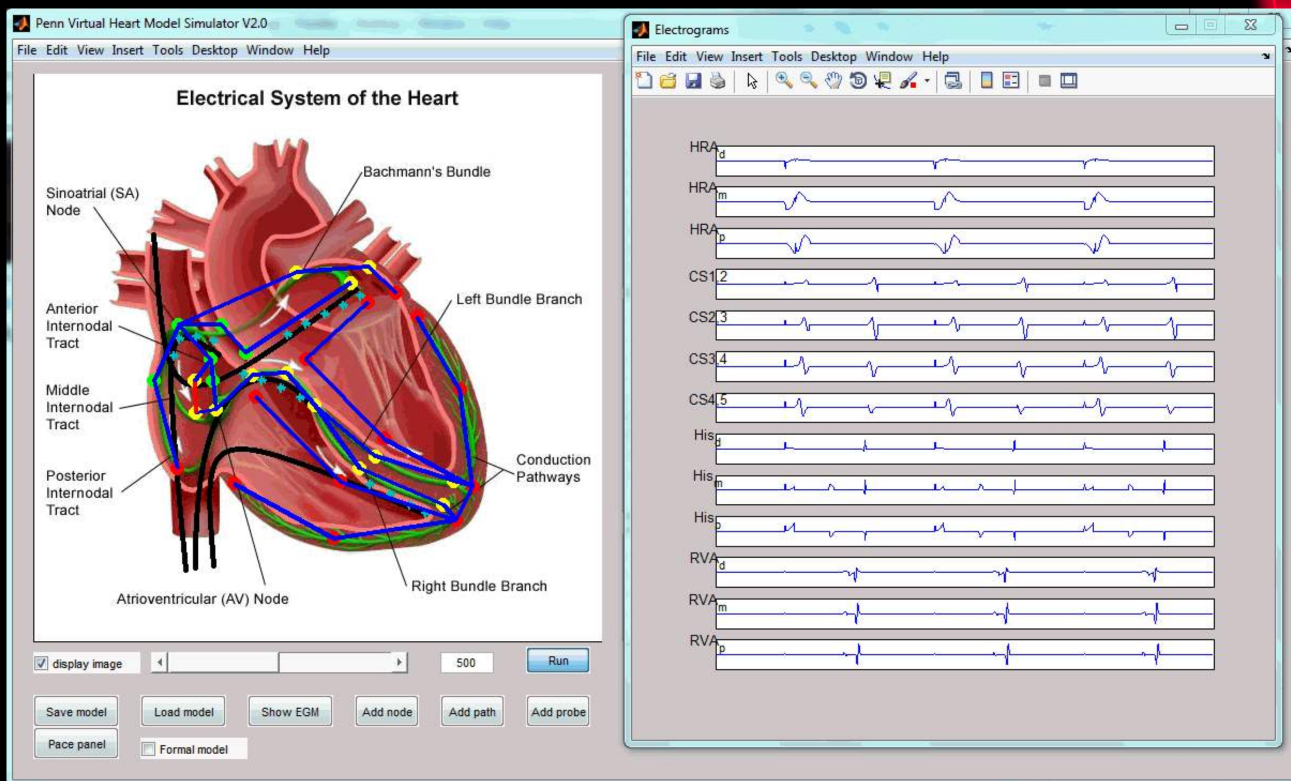

Ingredient 2: Pacemaker Model

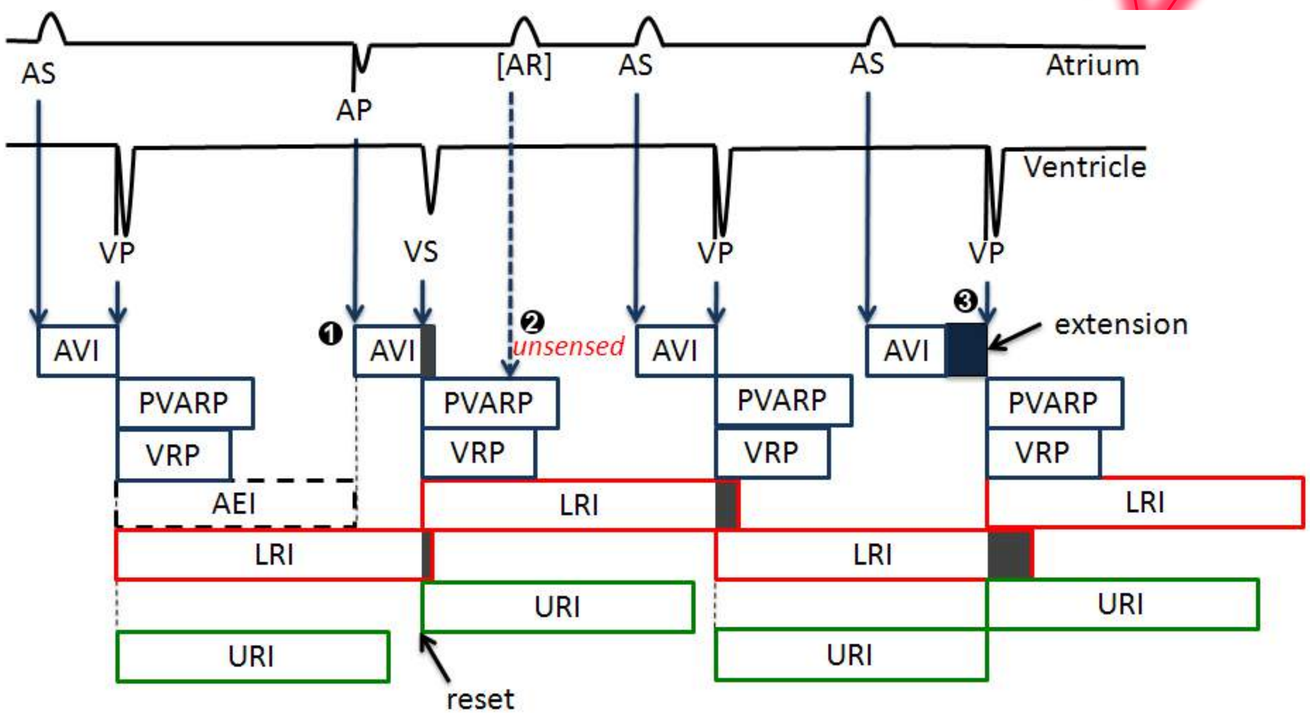

The timed automata model of the closed-loop system

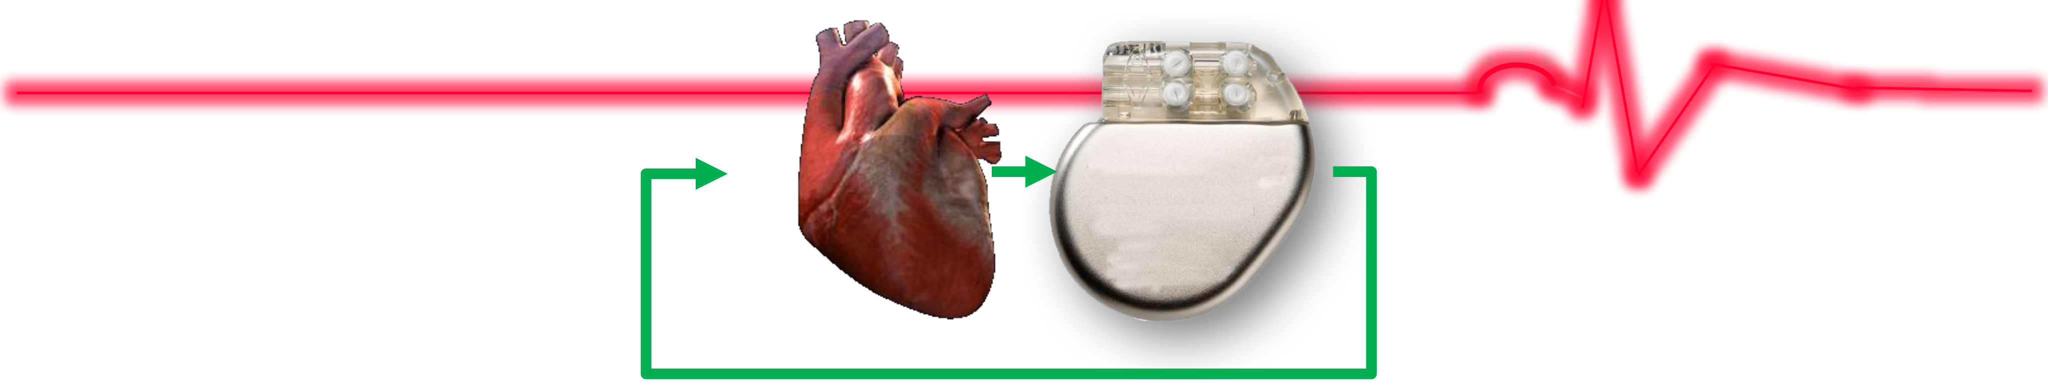

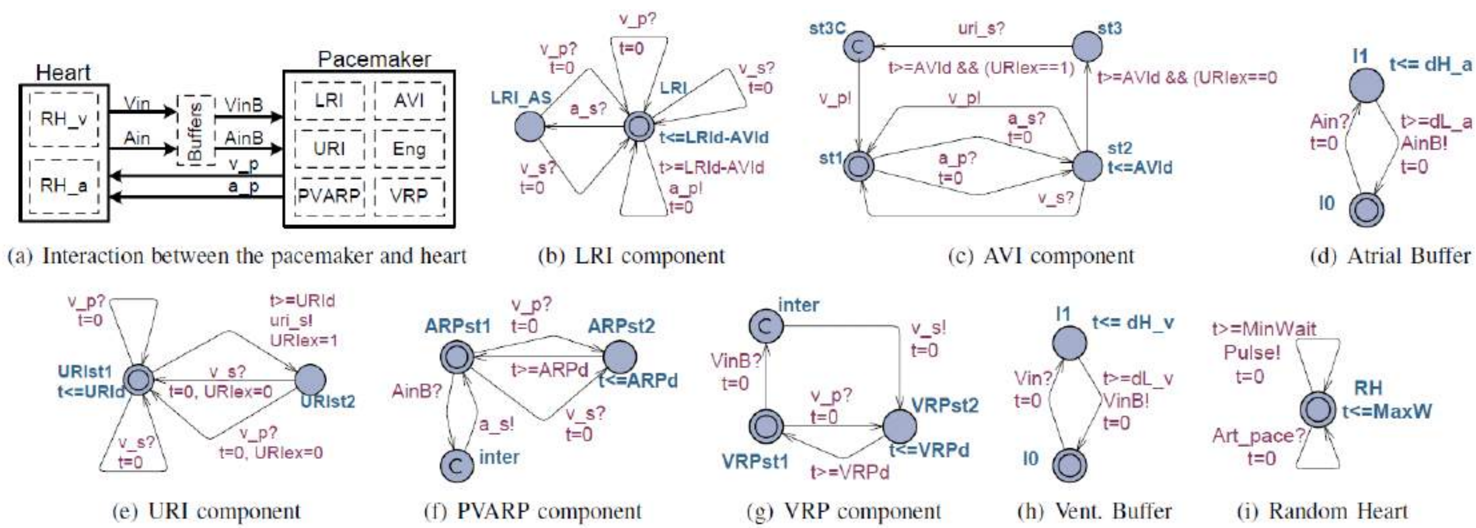

(a) Interaction between the pacemaker and heart (b) LRI component (c) AVI component (d) Atrial Buffer

(e) URI component (f) PVARP component (g) VRP component (h) Vent. Buffer (i) Random Heart

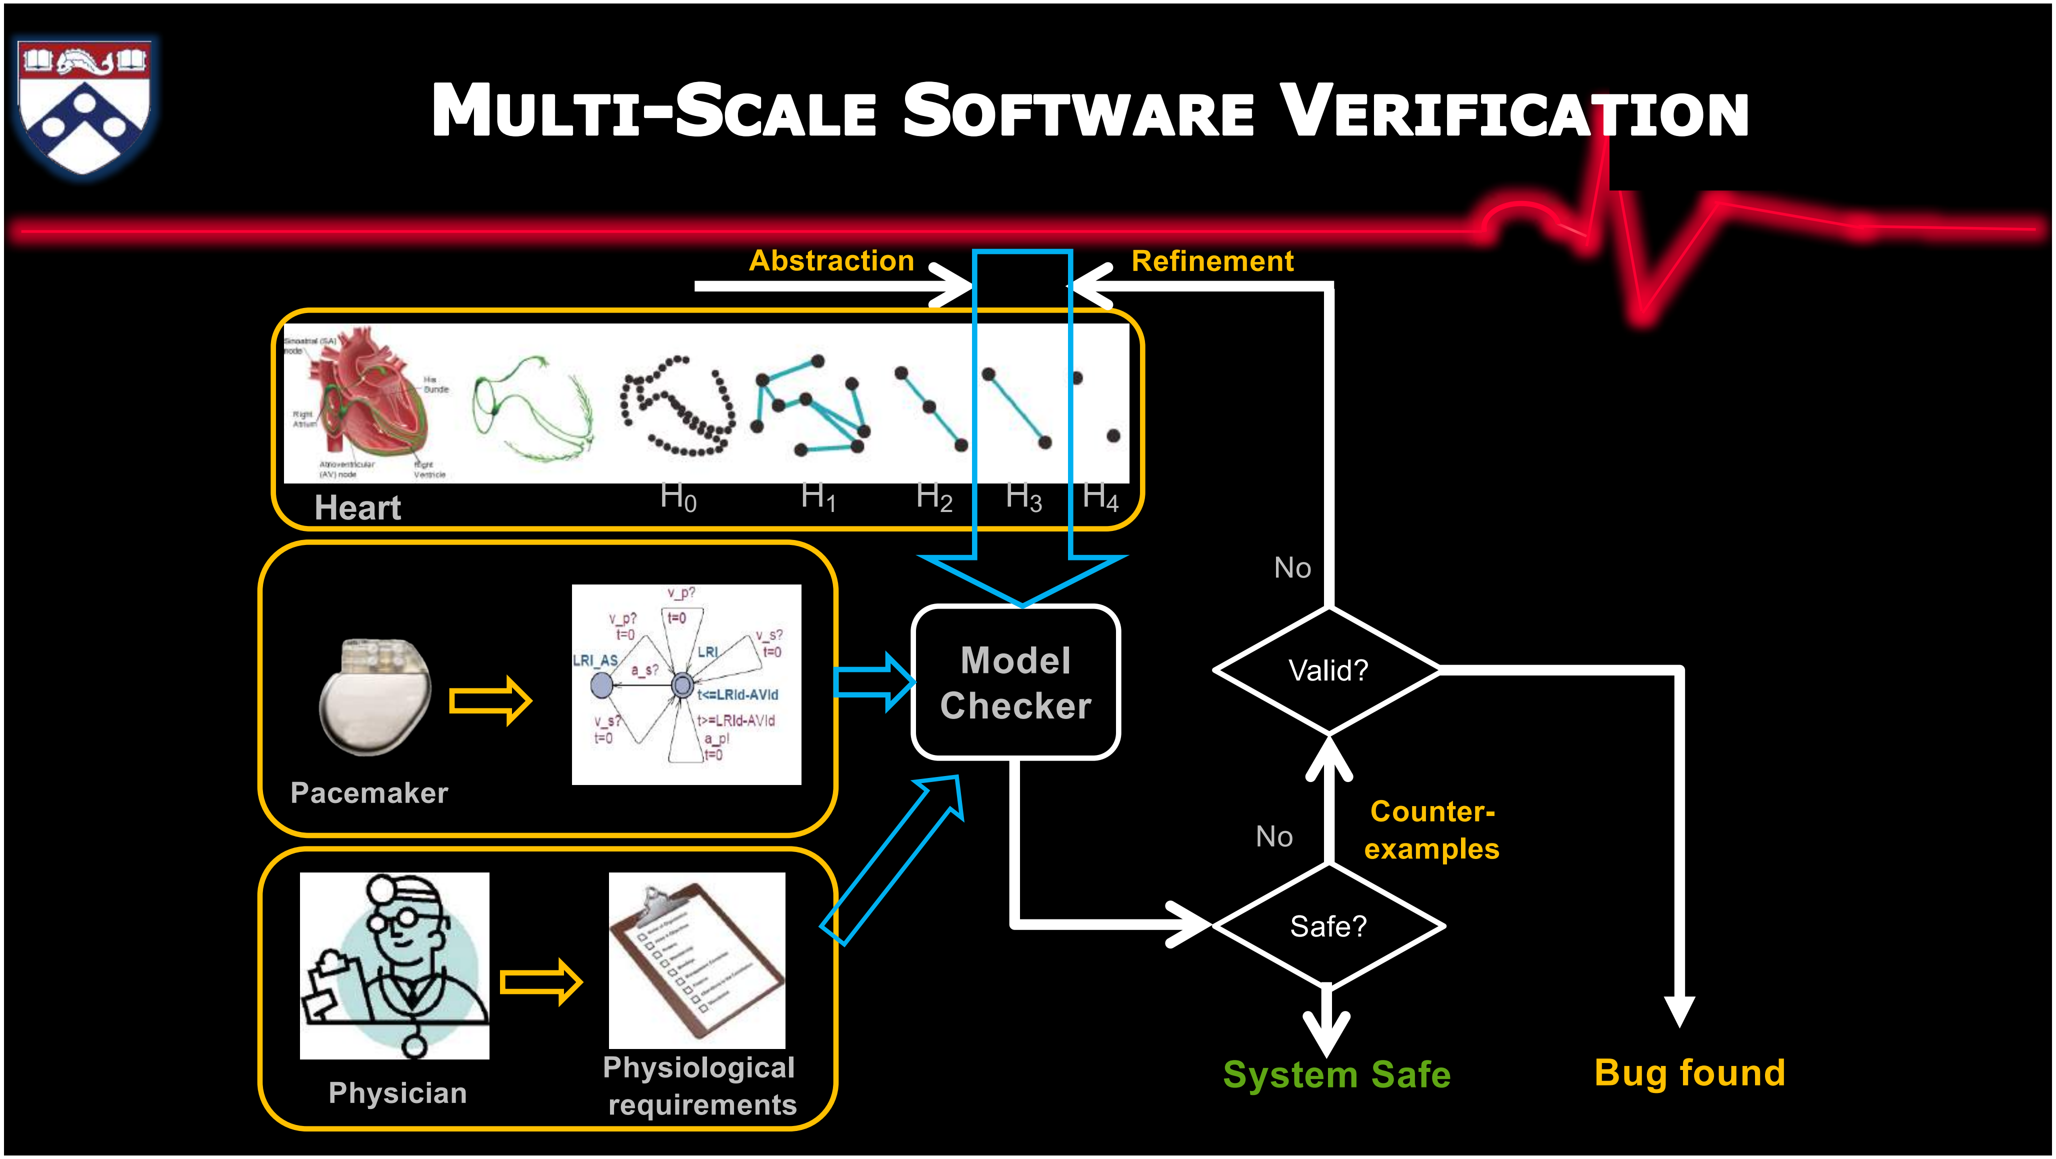

Multi-Scale Software Verification
Abstraction
Refinement
Heart
H0
H1
H2
H3
H4
Pacemaker
v_p?
t=0
LRI_AS
a_s?
LRI
v_s?
t<=LRId-AVId
t>=LRId-AVId
a_p!
Physician
Physiological requirements
Model Checker
No
Valid?
Counter-examples
Safe?
System Safe
Bug found

COUNTER-EXAMPLE-GUIDED ABSTRACTION & REFINEMENT

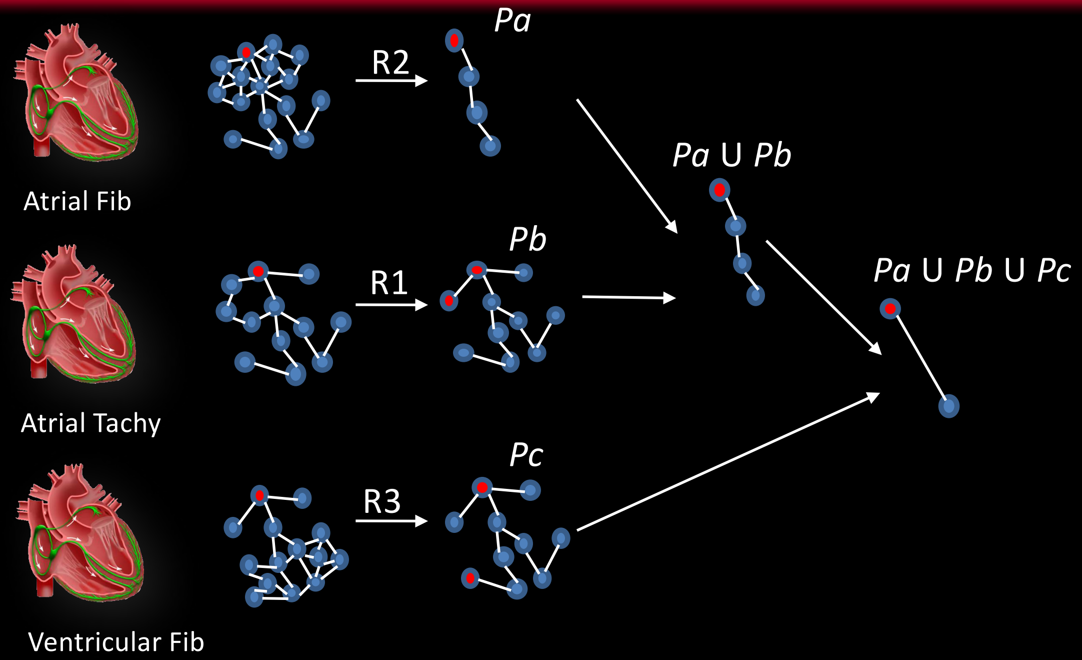

Least coverage
Least nb of invalid counterexample
Least ambiguous counterexample

Most coverage
Largest nb of invalid counterexamples
Most ambiguous counterexample

MBD Toolchain: UPP2SF Model translation UPPAAL → Stateflow → Generated code

The goal is to integrate:

- System modeling
- Verification
- Model-based WCET analysis
- Simulation
- Code generation
- Testing

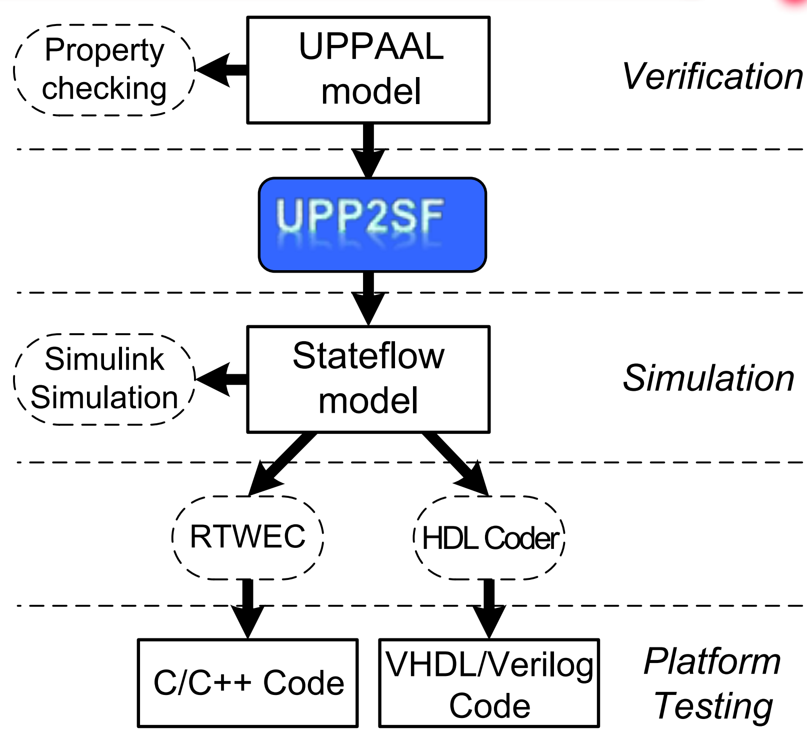

MBD Toolchain: UPP2SF Model translation UPPAAL → Stateflow → Generated code

The goal is to integrate:

- System modeling
- Verification
- Model-based WCET analysis
- Simulation
- Code generation
- Testing

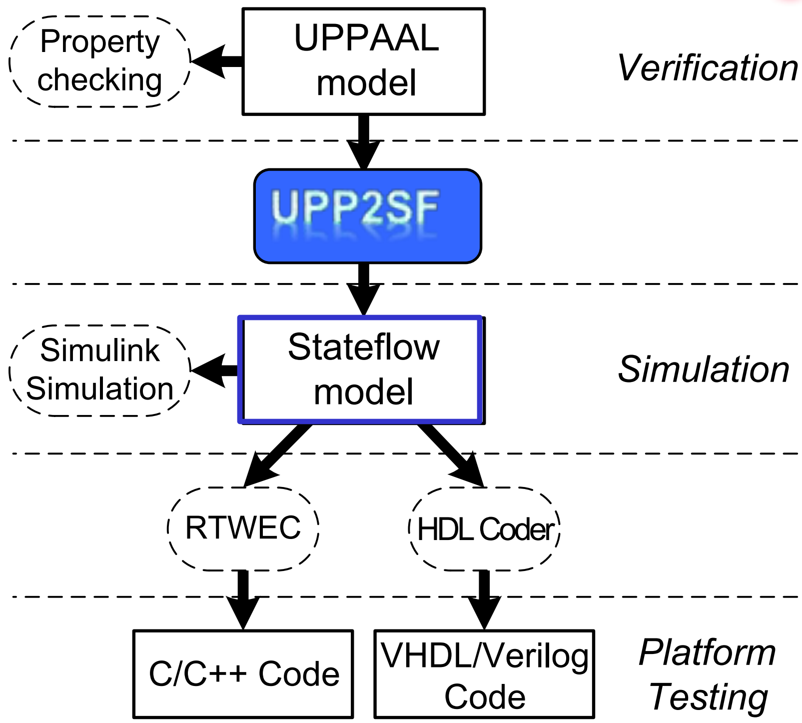

MBD Toolchain: UPP2SF Model translation UPPAAL → Stateflow → Generated code

The goal is to integrate:

- System modeling
- Verification
- Model-based WCET analysis
- Simulation
- Code generation
- Testing

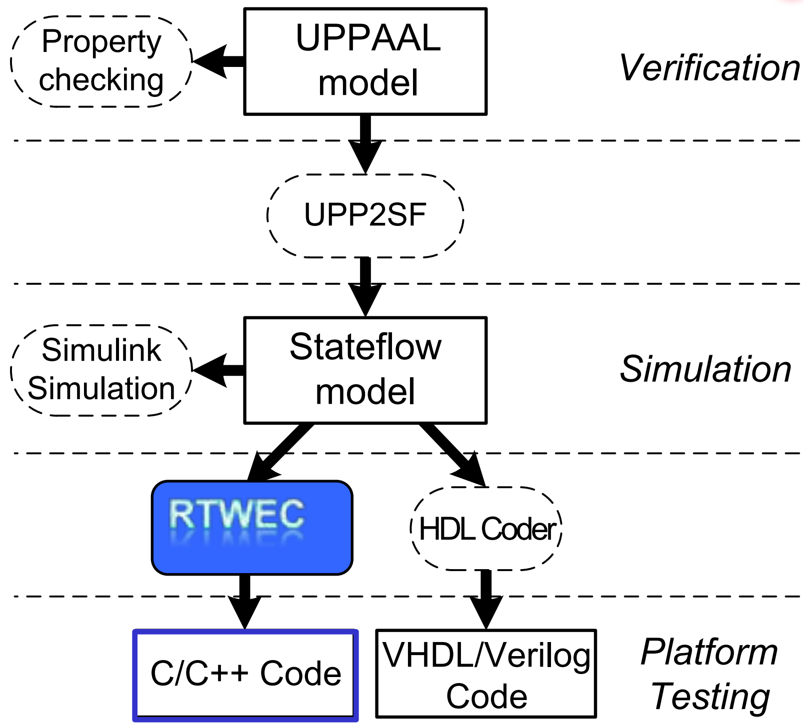

Listing 1. `bitsForTID0` definition

```
struct {
    uint_T is_AVI:3;
    uint_T is_LRI:2;
    uint_T is_PVARP:2;
    uint_T is_VRP:2;
    uint_T is_URI:2;
    uint_T is_active_AVI:1;
    uint_T is_active_LRI:1;
    uint_T is_active_PVARP:1;
    uint_T is_active_VRP:1;
    uint_T is_active_URI:1;
    uint_T is_active_Eng:1;
    uint_T is_Eng:1;
    uint_T URI_ex:1;
} bitsForTID0;
```

MBD Toolchain: UPP2SF Model translation
UPPAAL → Stateflow → Generated code

The goal is to integrate:

- System modeling
- Verification
- Model-based WCET analysis
- Simulation
- Code generation
- Testing

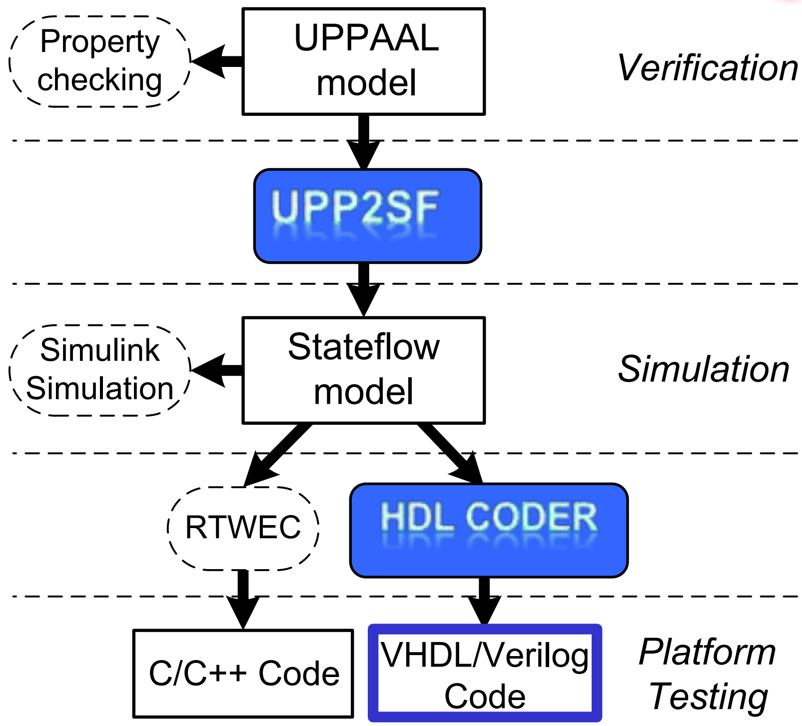

Published in: IEEE ECRTS'10, EMBC'10, Proceedings of IEEE'11, ICCPS'11, EMBC'11, TACAS'12, RTAS'12, STTT'13, BMES'14, Frontiers of EDA'15, IEEE Computer '16

Heart-on-Chip Platform for Closed-loop Testing

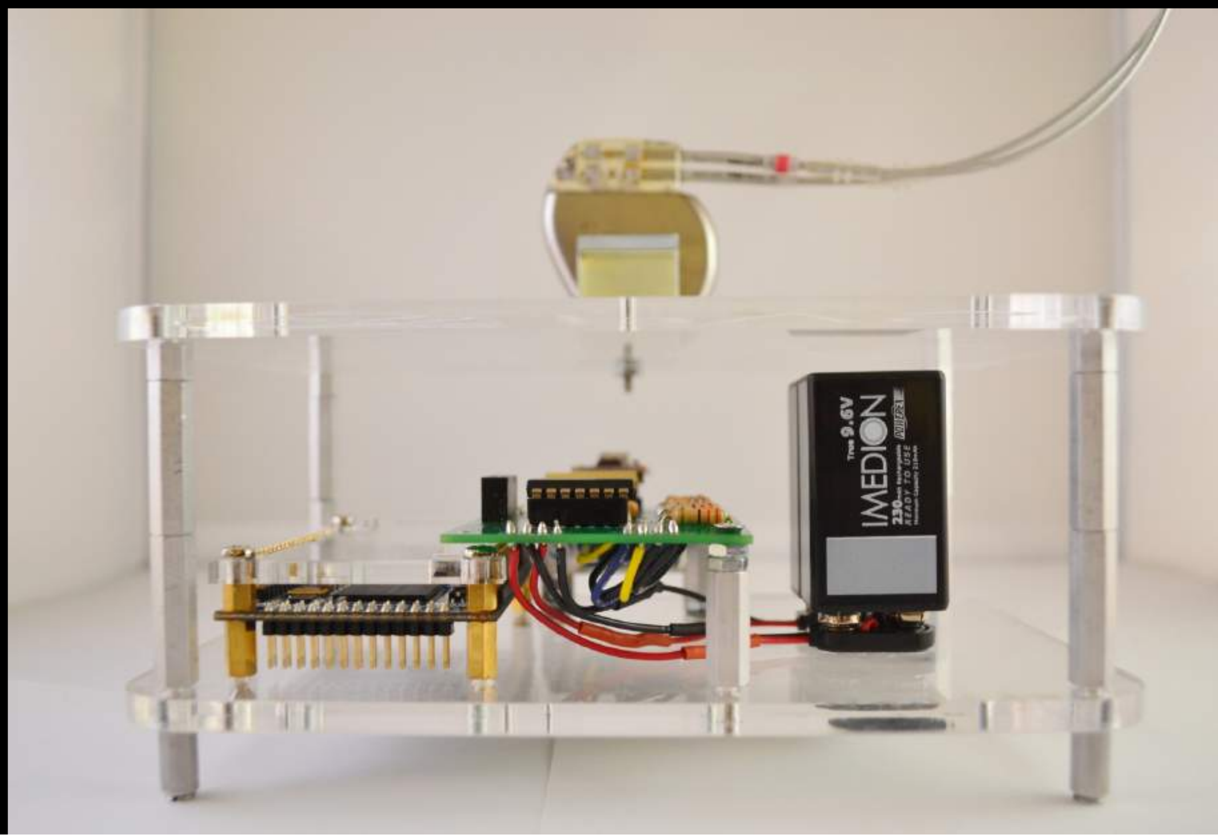

From Verified Models to Verified Code

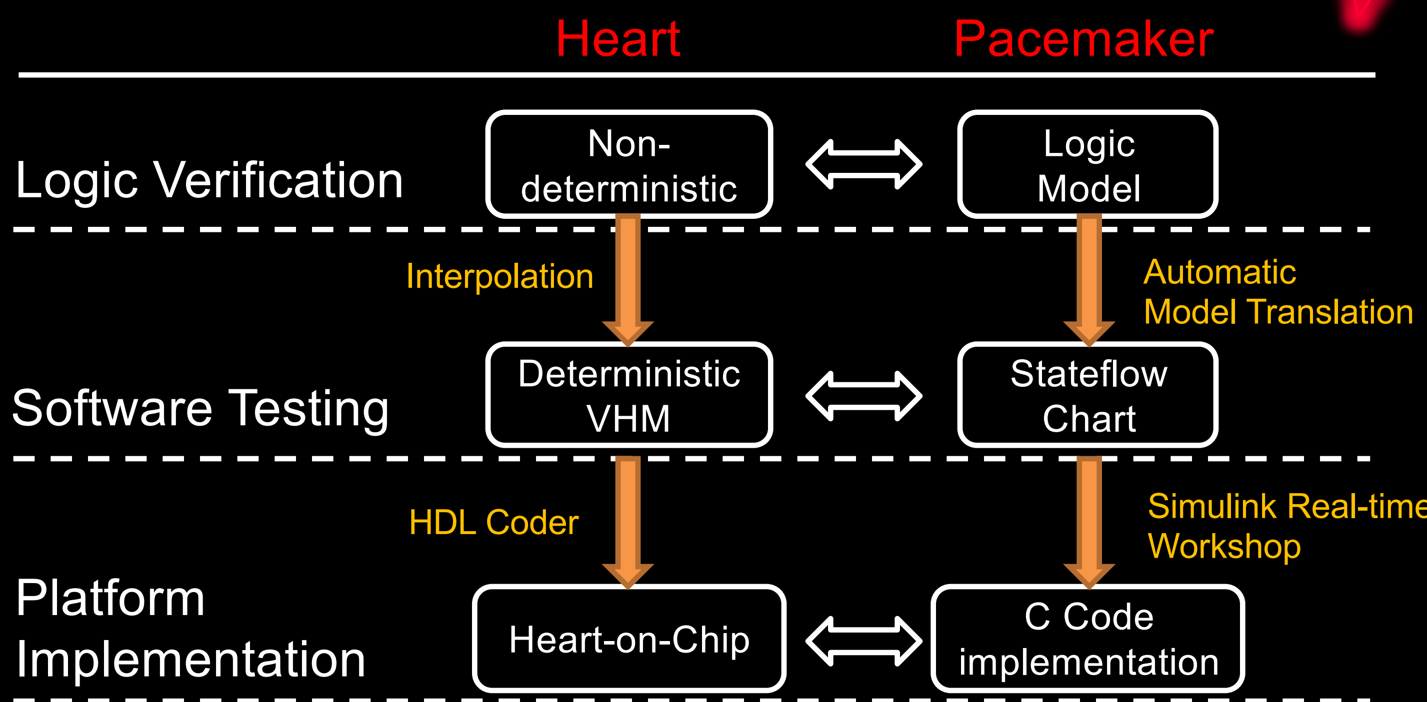

Published in: IEEE ECRTS'10, EMBC'10, Proceedings of IEEE'11, ICCPS'11, EMBC'11, TACAS'12, RTAS'12, STTT'13, BMES'14, Frontiers of EDA'15, IEEE Computer '16

Medical Devices vs Consumer Electronics

PART I
From Verified Models to Verified Code for Medical Devices
(NSF CPS Large 2010-2015)

PART II
Computer-Aided Clinical Trials
(NSF Frontiers 2015-2020)

Part III
Bringing formal and approximate approaches to cardiology

The clinical trial

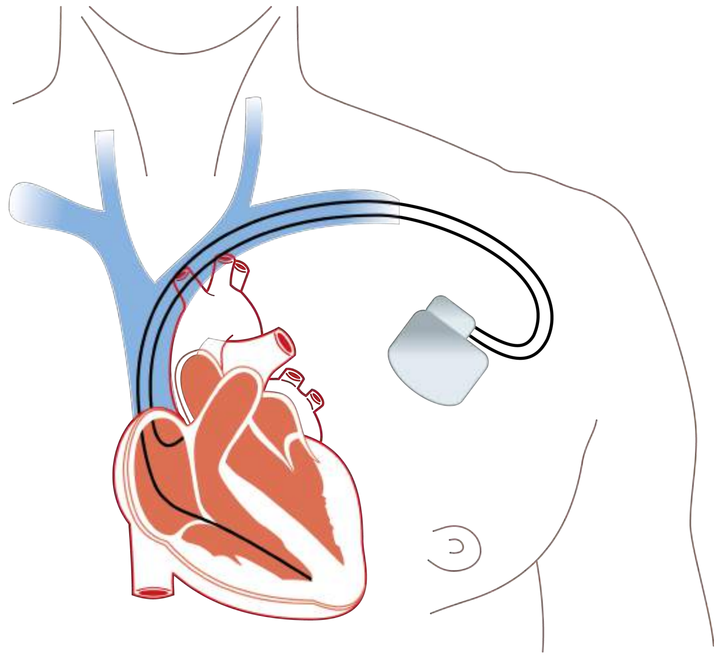

The ultimate closed-loop test

Trials are costly

- Device trial Costs can be $10-20 million
- Trial Time and effort: 4-6 years
- Ethical burden: putting patients at risk
- High percentage of failure

A clinical trial is a hypothesis test

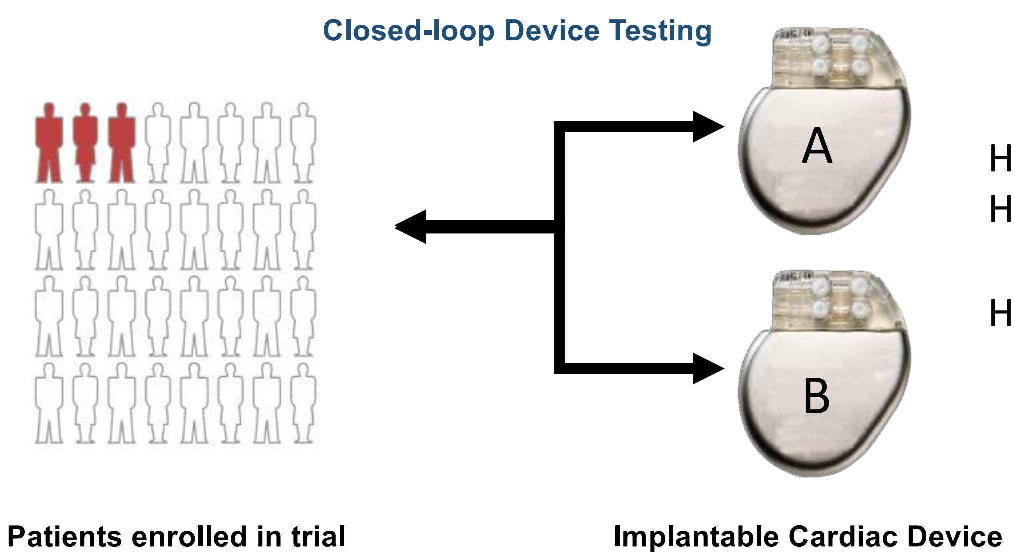

Hypothesis test:

H_0: Performance(A) = Performance(B)

H_a: Performance(A) ≠ Performance(B)

A computer-aided clinical trial is a hypothesis test

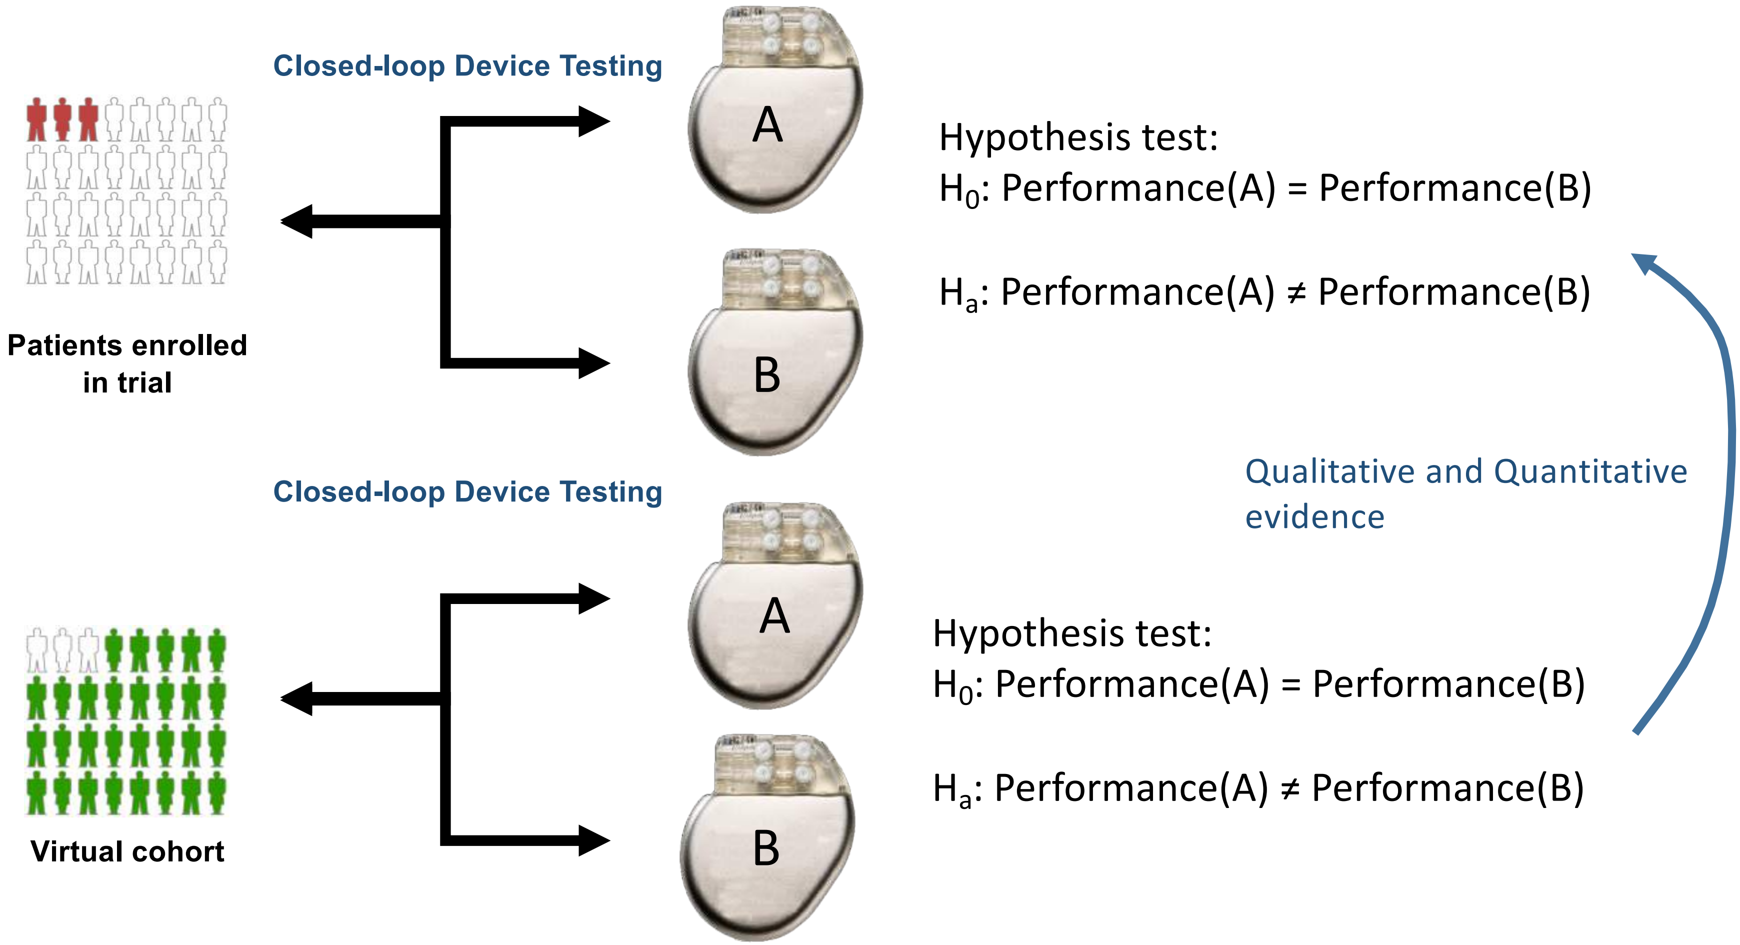

The RIGHT trial

The Rhythm ID Going Head to Head Trial*

Do patients on the two devices experience different time-to-first inappropriate therapy?

~2,000 patients, 4 years

Medtronic ICD
(**the control arm**)

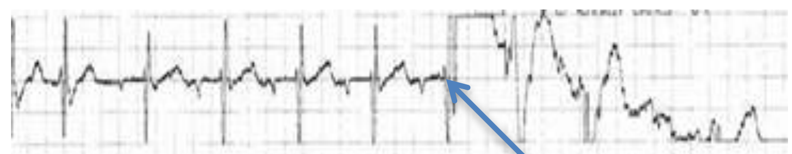

Vitality II ICD
(**the treatment arm**)

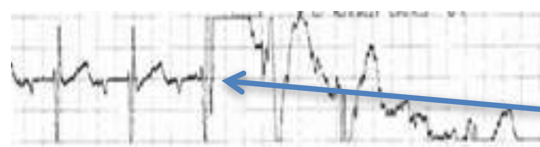

Inappropriate Therapy

*Berger et al., "The Rhythm ID Going Head to Head Trial", Journal of Cardiovascular EP, Vol. 17, No. 7, July 2006

RIGHT Trial Results – Inappropriate Therapy

Gold et al RIGHT of Inappropriate ICD Therapy 373

Table 2 Adjudication summary of spontaneous episodes where therapy was delivered

	n episodes (% of total events)			
Adjudicated rhythm	VITALITY 2	Selected Medtronic	Overall	*P* value
Artifact	23 (1.1)	90 (4.6)	113 (2.8)	.0094
Ventricular tachycardia	705 (34.9)	994 (51.0)	1699 (42.8)	.2490
Ventricular fibrillation	59 (2.9)	61 (3.1)	120 (3.0)	.4265
Sinus tachycardia	506 (25.0)	220 (11.3)	726 (18.3)	<.0001
Atrial fibrillation	431 (21.3)	101 (5.2)	532 (13.4)	<.0001
Atrial flutter	66 (3.3)	19 (1.0)	85 (2.1)	.0076
Atrial tachycardia	20 (1.0)	100 (5.1)	120 (3.0)	.0001
AVNRT	17 (0.8)	39 (2.0)	56 (1.4)	.5956
Other supraventricular tachycardia/unknown	178 (8.8)	325 (16.7)	503 (12.7)	.4436
Sinus rhythm with premature ventricular complexes	18 (0.9)	1 (0.1)	19 (0.5)	NE
Total events	2023	1950	3973	

NE = nonestimable; AVNRT = Atrioventricular nodal re-entry tachycardia.

Inappropriate Therapy

VITALITY 2: 62.2%

Medtronic: 54.1%

*Michael R. Gold, Primary results of the Rhythm ID Going Head to Head Trial, Heart Rhythm, Vol 9, No 3, March 2012

Computer-aided trial

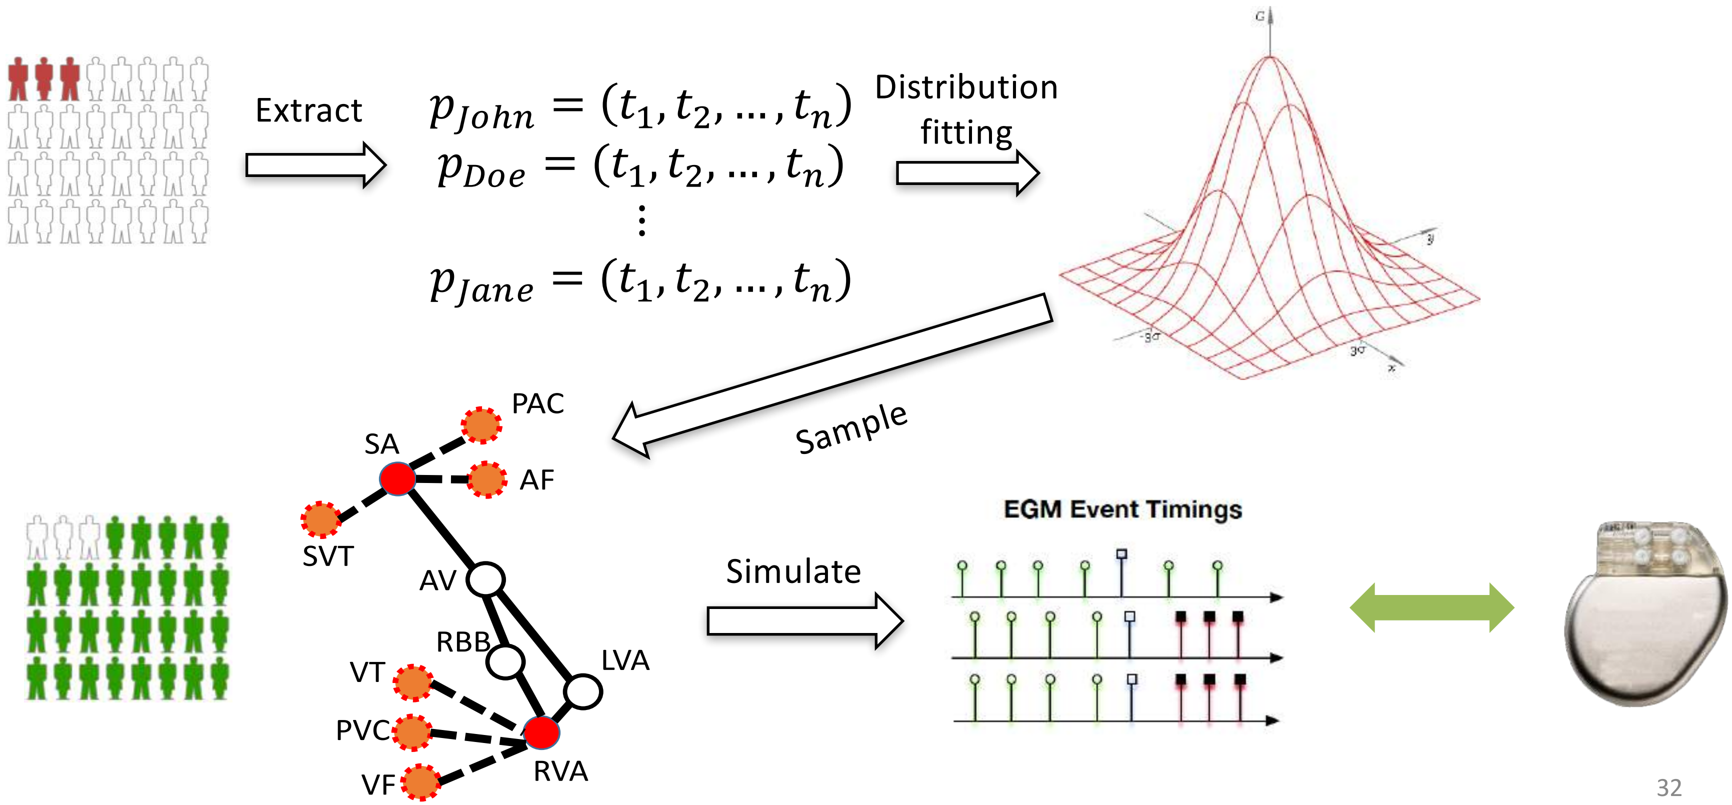

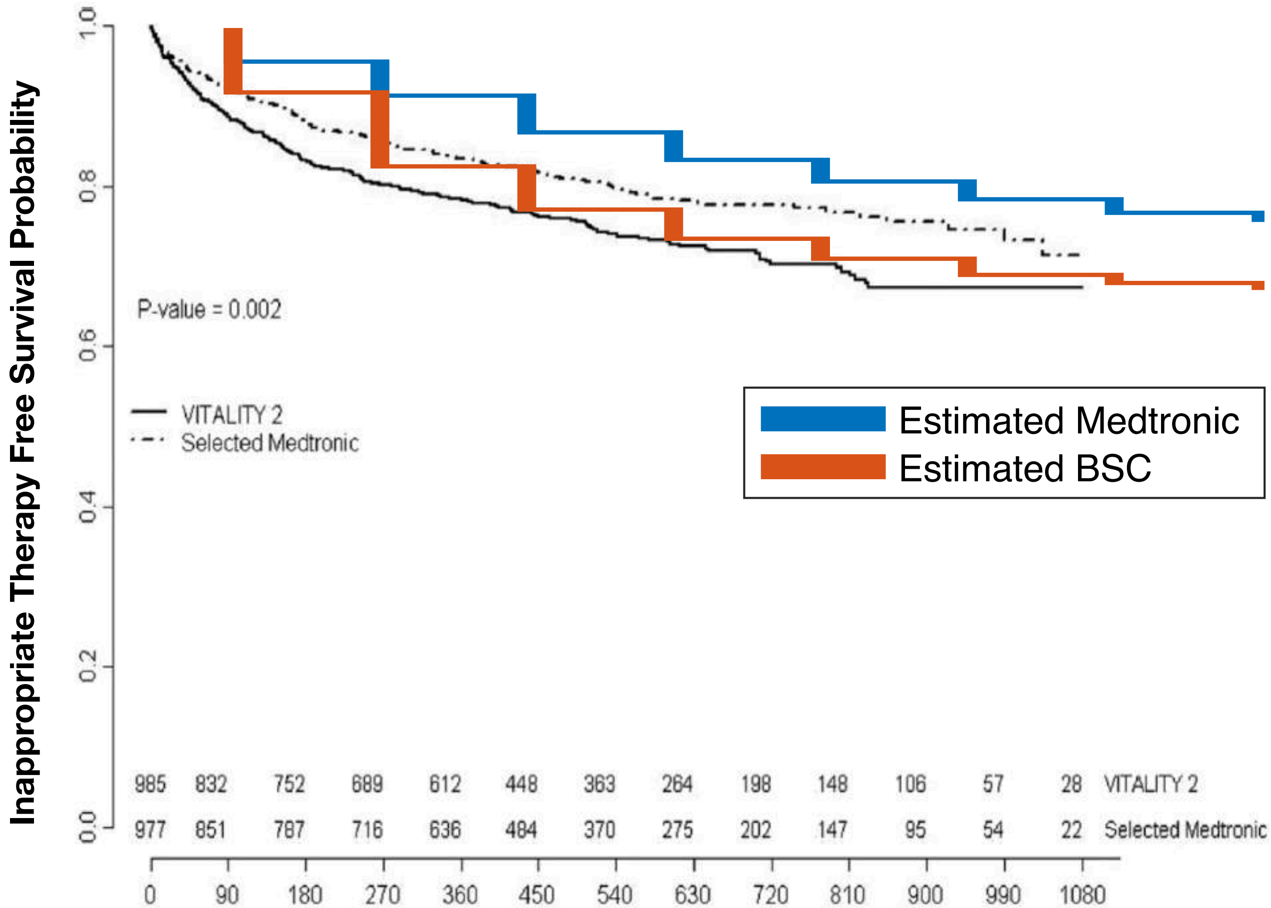

Kaplan-Meier curve — Inappropriate therapy-free survival

CPS Challenges

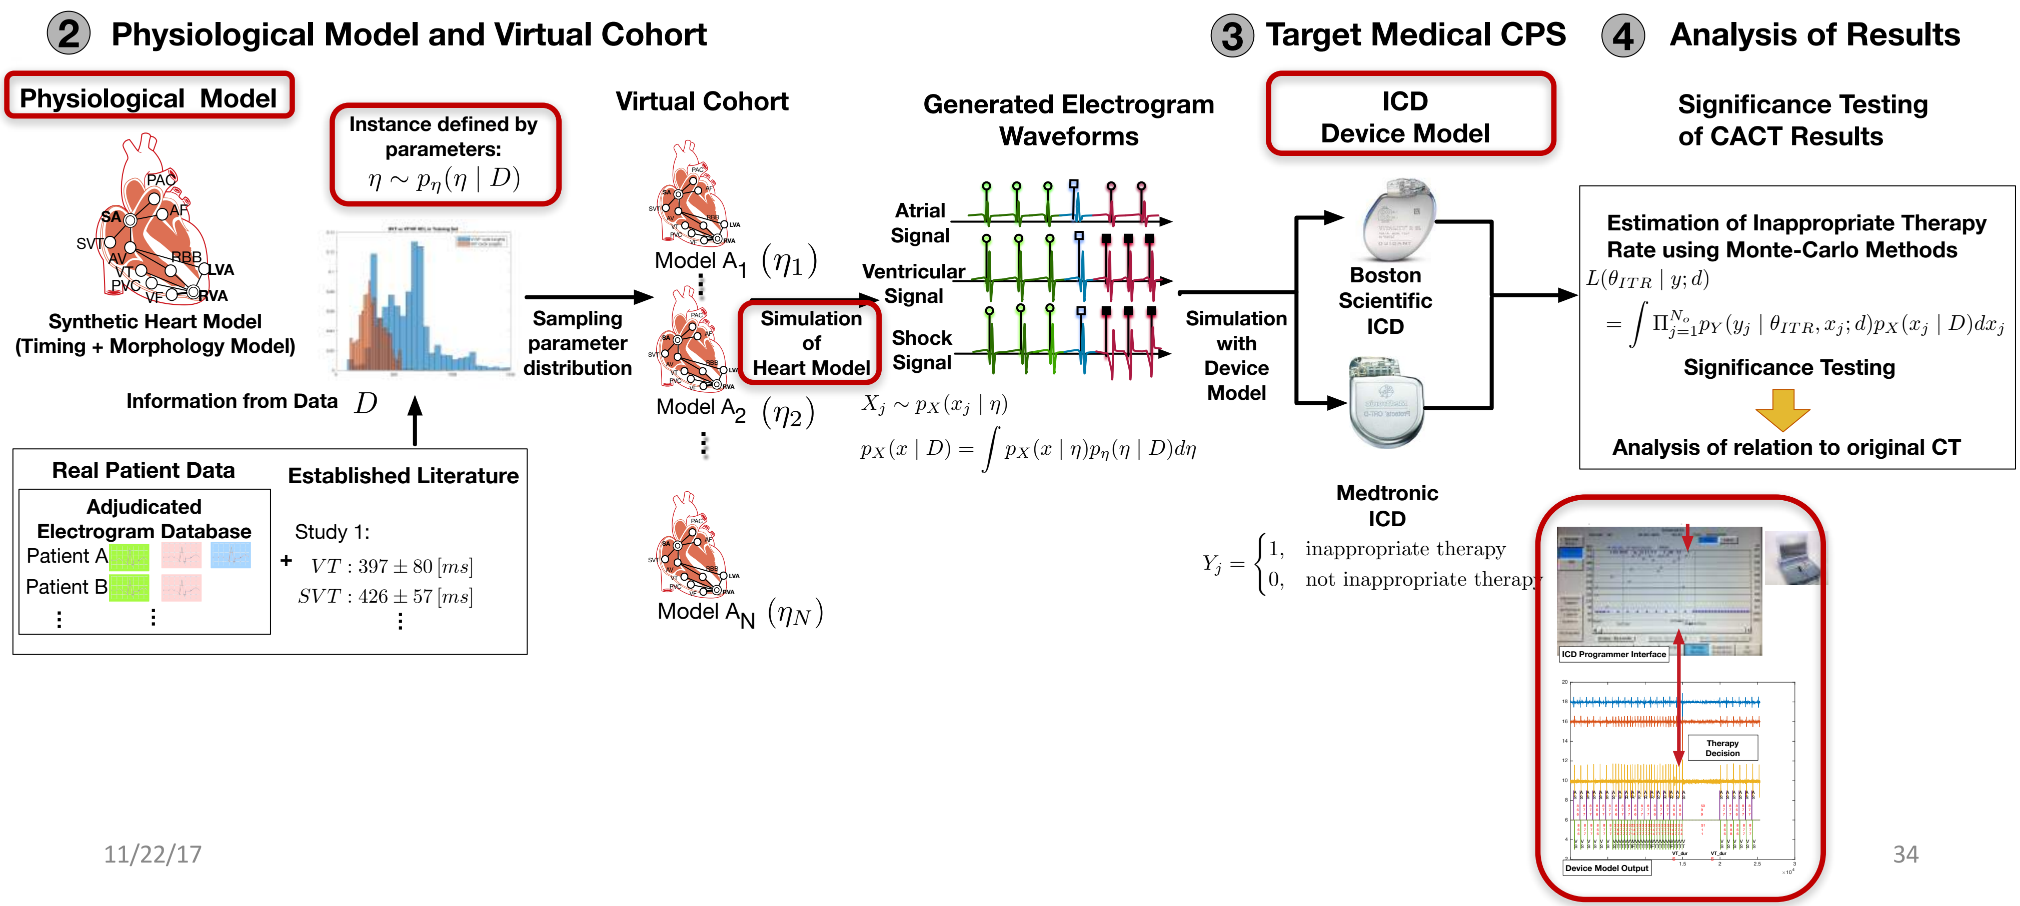

Roadmap: Increasing the use of CM&S evidence

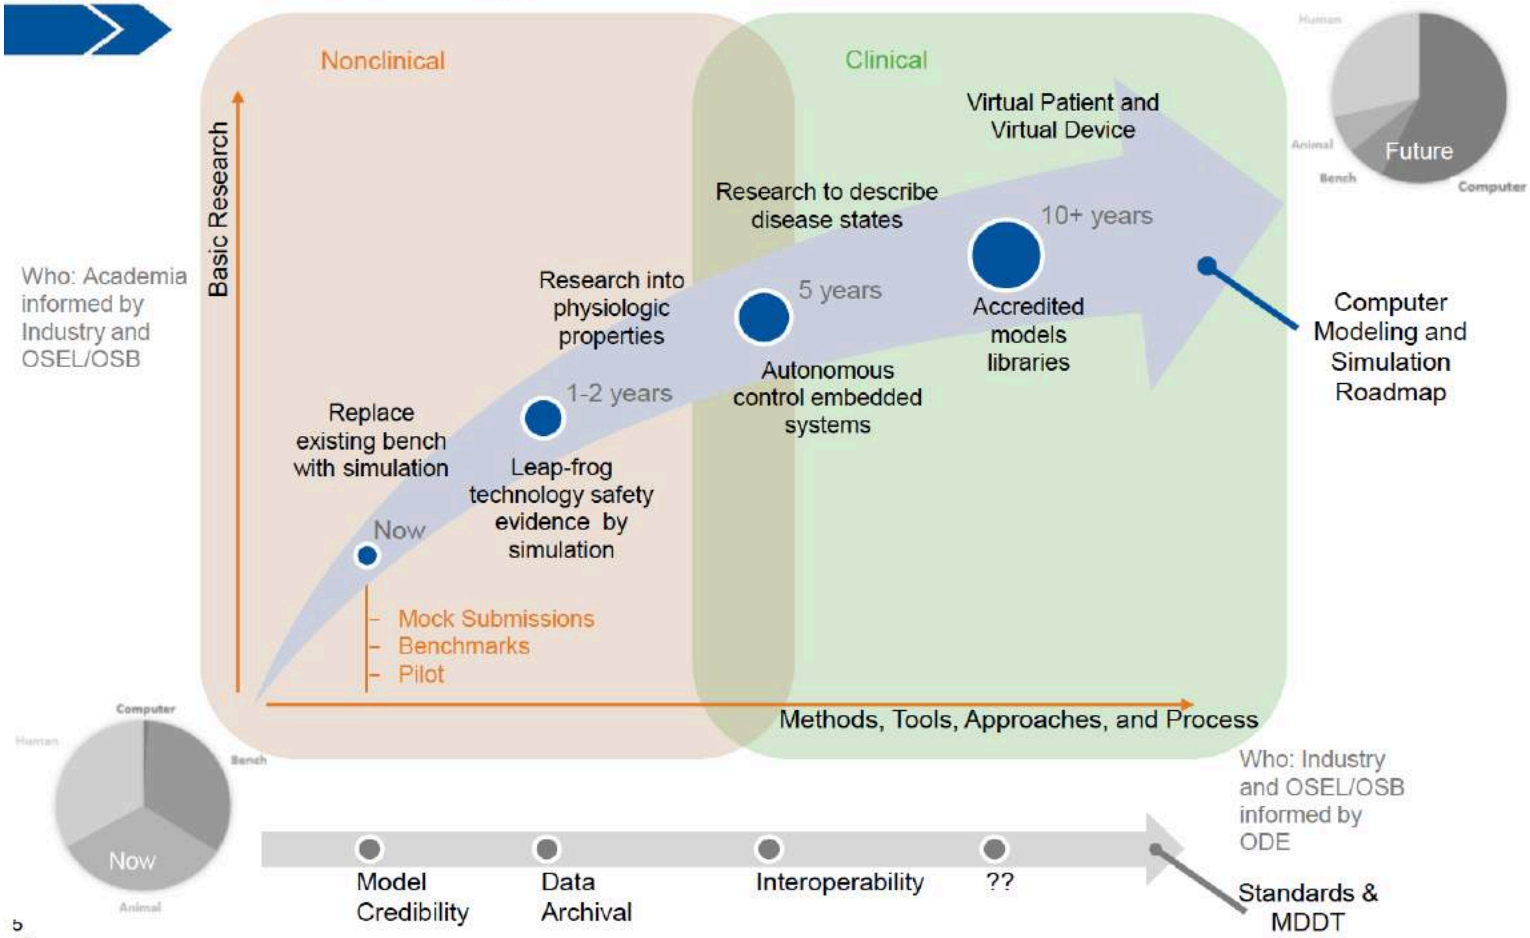

Slide credit: Dawn Bardot, MDIC

Medical Devices vs Consumer Electronics

PART I
From Verified Models to Verified Code for Medical Devices
(NSF CPS Large 2010-2015)

PART II
Computer-Aided Clinical Trials
(NSF Frontiers 2015-2020)

Part III
Bringing formal and approximate approaches to cardiology

Robust Artificial Pancreas

CyberCardia

Stony Brook University

Data-driven robust control of insulin therapy

- Artificial pancreas (AP): automated treatment of type 1 diabetes (T1D) through control algorithms integrating insulin pump and glucose sensor
- **Fully closed-loop therapy is challenging:** blood glucose (BG) depends on **disturbances related to the patient's behavior**, mainly **meals and physical activity**
- To account for uncertainties, we **construct data-driven models** of meal and exercise behavior, and develop a **robust model-predictive control (MPC) algorithm**

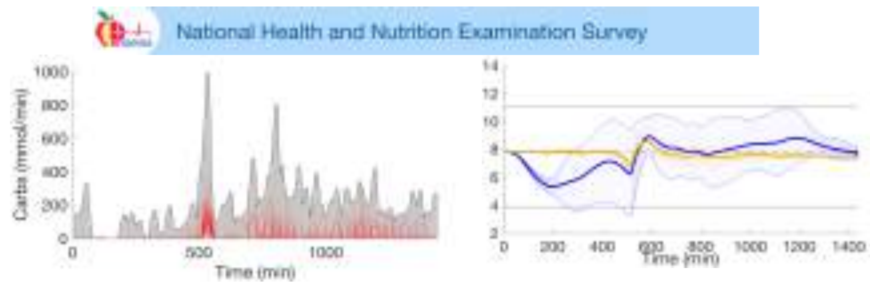

Left: ***uncertainty sets*** *constructed from* ***data*** *(CDC NHANES database) with probabilistic coverage guarantees. The robust MPC controller minimizes the worst-case performance wrt these sets.*
Right*: BG comparison between* ***our controller*** *and an* ***ideal controller*** *that has exact knowledge of plant state and disturbances.*

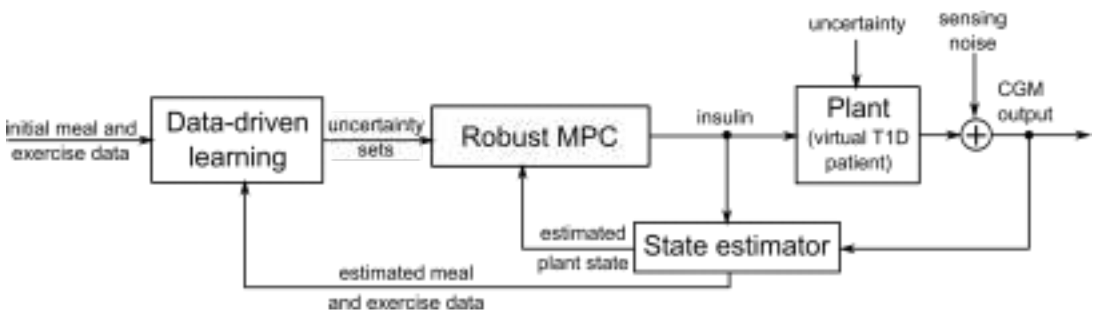

Paoletti, N., Liu, K.S., Smolka, S.A., Lin, S. (2017) Data-Driven Robust Control for Type 1 Diabetes Under Meal and Exercise Uncertainties. Computational Methods in Systems Biology. LNCS 10545, pp. 214–232.

Robust Artificial Pancreas

CyberCardia Stony Brook University

SMT-based synthesis of safe and robust PID controllers

- New method for the **automated synthesis of PID controllers with safety and performance guarantees for hybrid systems with stochastic and nonlinear dynamics**.
- Controllers are **robust by design** since they minimize the probability of reaching an unsafe state under random disturbances.
- We leverage **SMT solvers over the reals and nonlinear differential equations** (e.g. dReal, iSAT) to provide formal guarantees that the controller satisfies a given probabilistic bounded reachability property.
- Application to insulin regulation for T1D

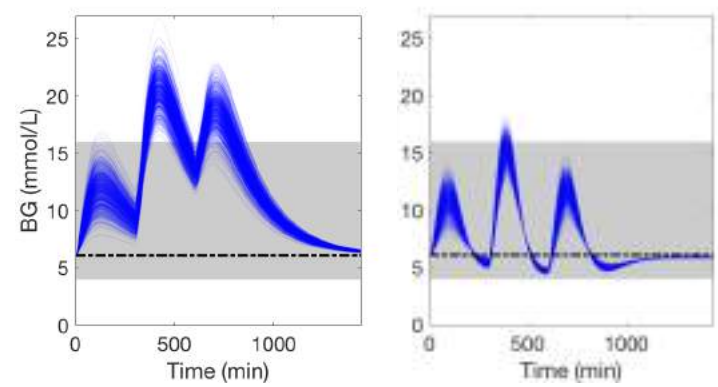

BG for basal (left) vs synthesized (right) insulin controllers

Shmarov, F., Paoletti, N., Bartocci, E., Lin, S., Smolka, S., Zuliani, P. (2017) SMT-based Synthesis of Safe and Robust PID Controllers for Stochastic Hybrid Systems. Haifa Verification Conference (to appear).

Focus at UMD in CyberCardia

- Foundations, tools for reasoning about CPS
  - Formal modeling of CPS
  - Formal specification, verification
- This year: Specification reconstruction
  - Given model M, infer temporal properties that M (likely) satisfies
  - Motivations
    - Model understanding
    - Specification updating
    - Means for "jump-starting" formal specifications in often unfamiliar notations
- See poster (48-50)!

Specific Results in 2017

- Linear temporal-logic query checking
  - Problem
    - Given Kripke structure M, LTL “template” phi[x]
    - Find most general solution phi’ for missing formula x so that M satisfies phi[x:=phi’]
  - Algorithmic solution based on model checking developed, implemented, evaluated
  - Work presented at AVoCS/FMICS 2017
- Invariant mining from test data
  - Problem
    - Given (Simulink) model M, state variables of interest
    - Propose invariants describing relationships among variables
  - Approach: use data-mining on test data coupled with retesting to generate likely invariants
  - Evaluation used 11 models from automotive, medical-device domain
  - Work presented at EMSOFT 2017

Reachability Analysis of Cardiac Alternans (CMSB'16 and TCS'18)

CyberCardia

Stony Brook University

- Alternans is a phenomenon in cardiac cells that can contribute to fibrillation.
- Want to detect initial conditions that lead to alternans.
- Model as hybrid automata and use delta-reachability and statistical sampling technique

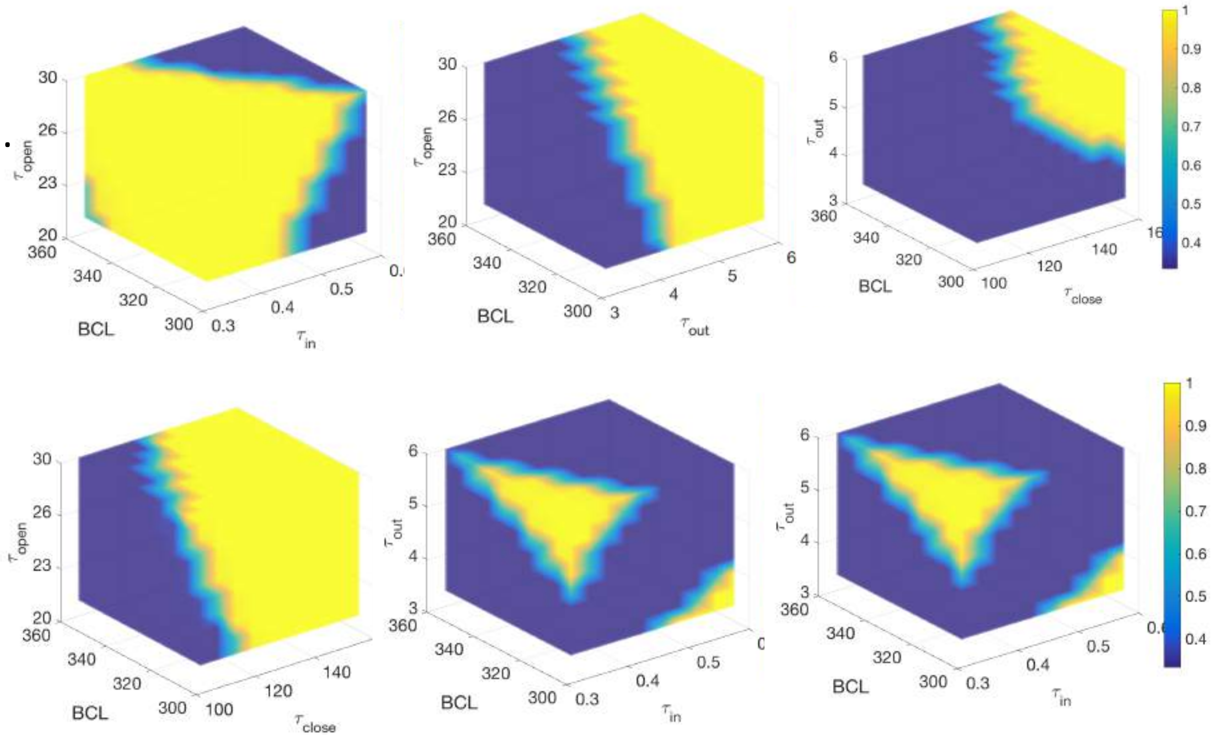

Yellow: Alternans region
Dark blue: Non-alternans region
Light-blue: Bifurcation hypersurface.

Lagrangian Reachability Analysis [CAV'17]

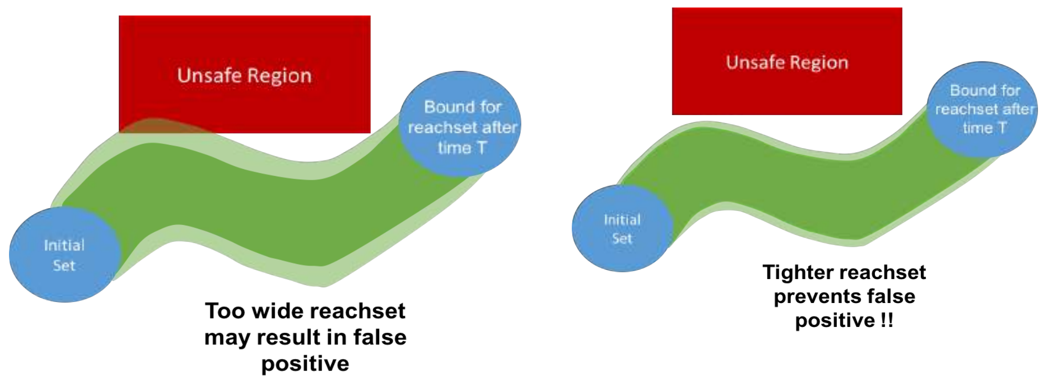

Lagrangian Reachability

- Compute over-estimate for the gradient of the solution-flows
- Compute over-estimate for Cauchy-Green (CG) deformation tensor from the gradient
- Optimize for positive-definite symmetric matrix M_1, defining the weighted norm in which the CG stretching factor is minimized
- Compute an upper bound for the CG stretching factor Λ, then the ball over-estimate at time t_1 is $B_{M_1}(\phi_{t_0}^{t_1}(x_0), \Lambda . \delta_0)$

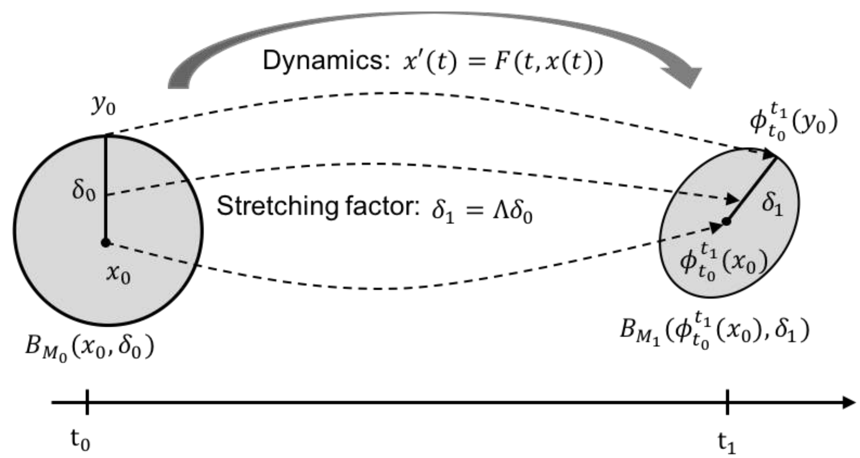

Shock-induced Virtual Electrode Formation

CyberCardia

- Understand the physical mechanisms of defibrillation shocks
- Approach: use high-fidelity numerical solutions of governing equations (several hours to simulate a 400ms heartbeat)
- Finding: large blood vessels act as *virtual electrodes,* that are favored paths for defibrillation shock to travel through and propagate from.

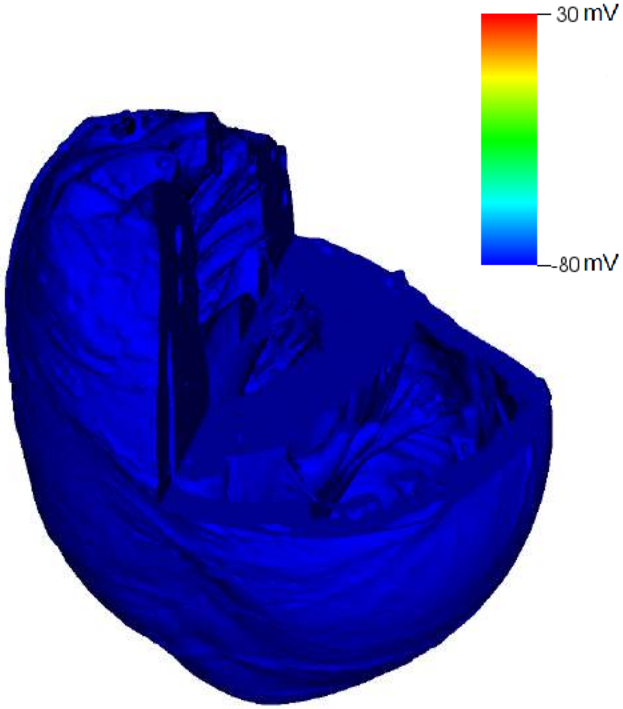

Stony Brook University

Shock-induced Virtual Electrode Formation

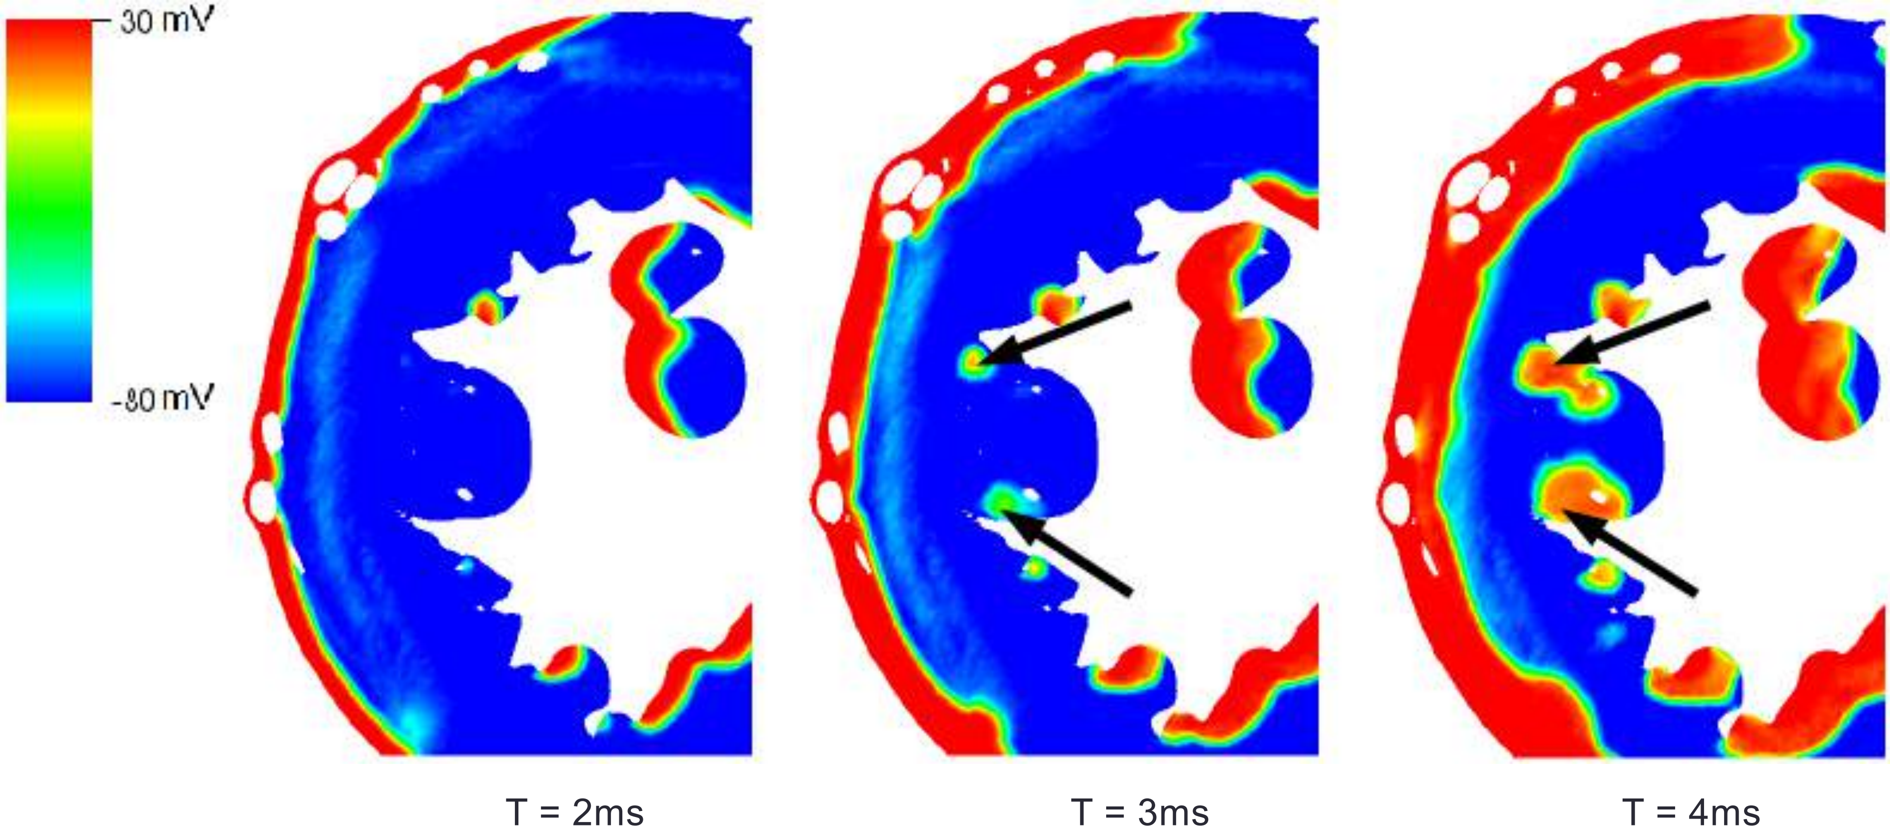

Re-thinking the basics: Distance functions

[Abbas et al., Heart Rhythm Sessions 2017]

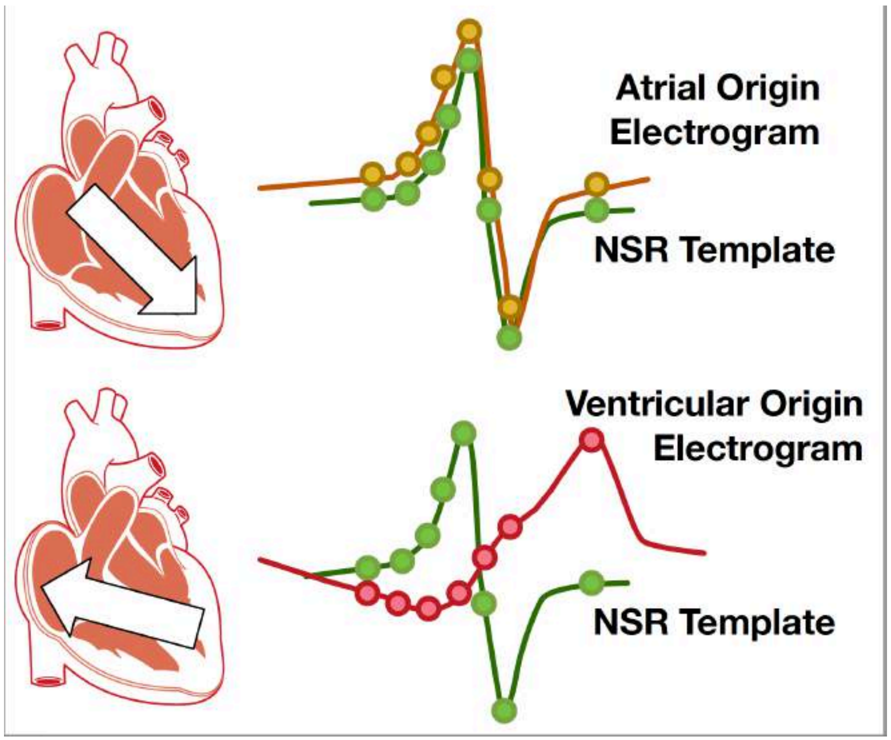

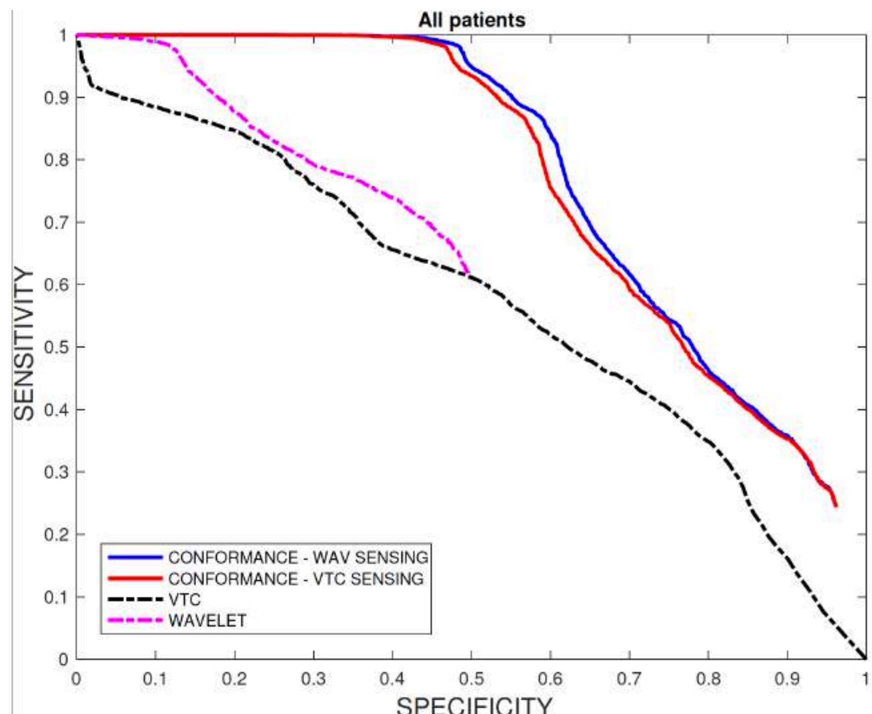

Penn Engineering

Re-thinking the basics: Programming languages

[Abbas, Rodionova, Bartocci, Smolka, Grosu, CMSB 2017]

DEVICE	Signal Processing	Decision logic
 DEFIBRILLATOR	Detect peaks (local maxima) in input signal Correlate two real-valued signals 	Is the average heart rate above a threshold? Do we see an “A(V+)A pattern” with a given delay between events? → Is the heart in fatal arrhythmia?
 PACEMAKER	Detect peaks in input signal	Do the ventricles always beat within 150-250ms of the atria? → Is the heart in bradychardia?

The number of heartbeats in a one-minute time interval is between 120 and 150.

→ Quantitative Regular Expression →

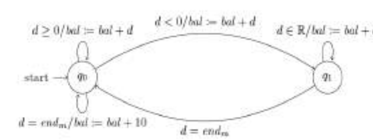

ESE 680-004: Digital Twins - Model-Based Embedded Systems

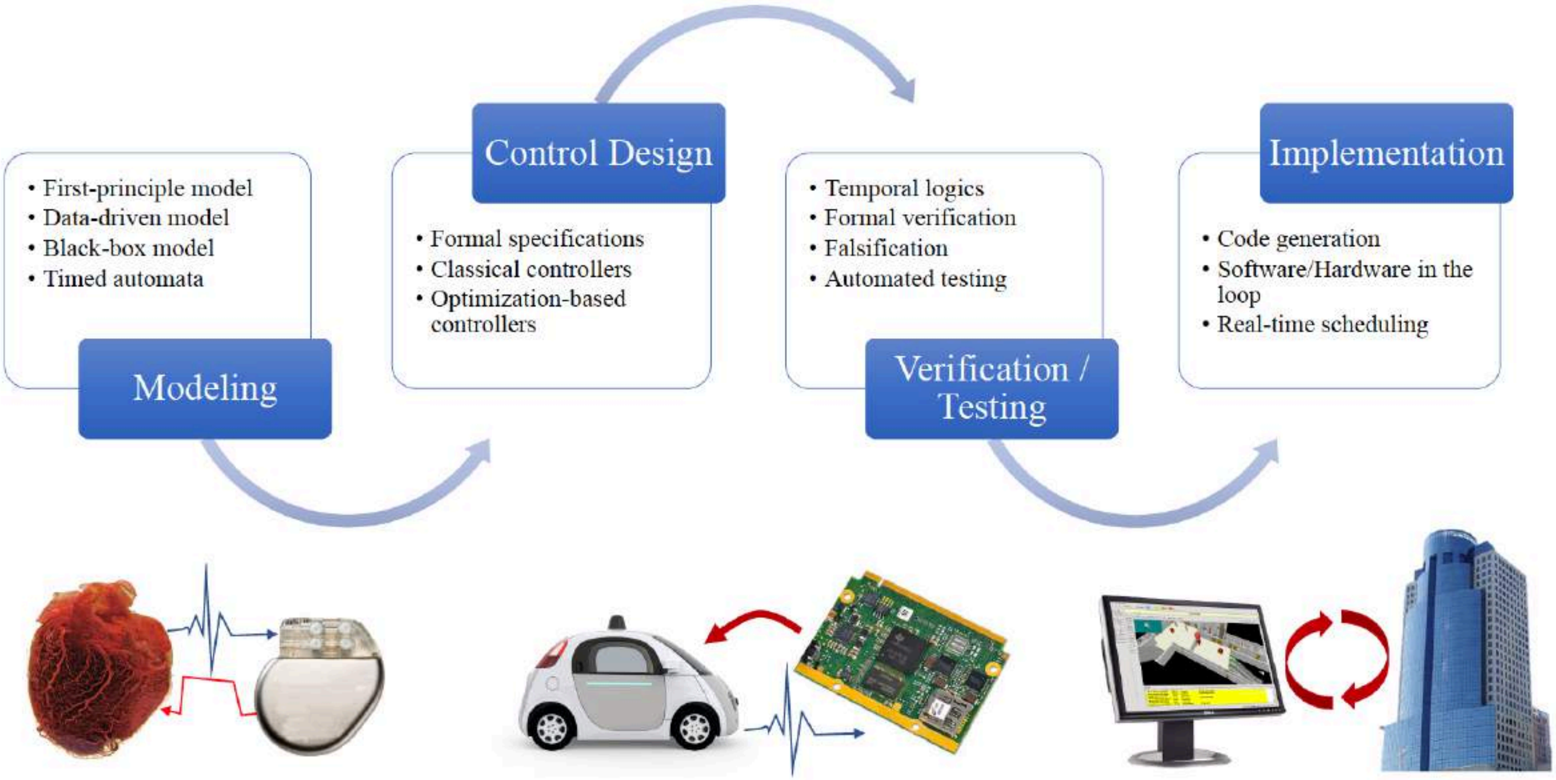

Module A
Life-critical medical devices

Module B
Safety-critical automotive control

Module C
Building modeling and control

What’s in it for you?

Set yourself apart from ”regular” embedded systems engineers by knowing when and how to apply formal methods, complemented with simulation and testing

This is a valuable and rare skill

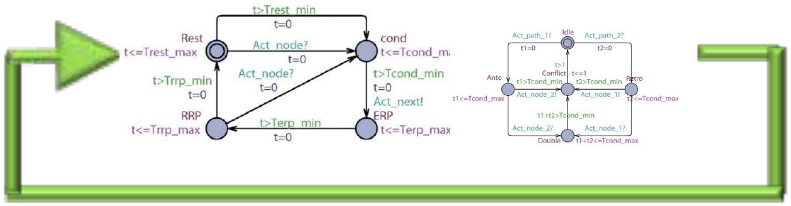

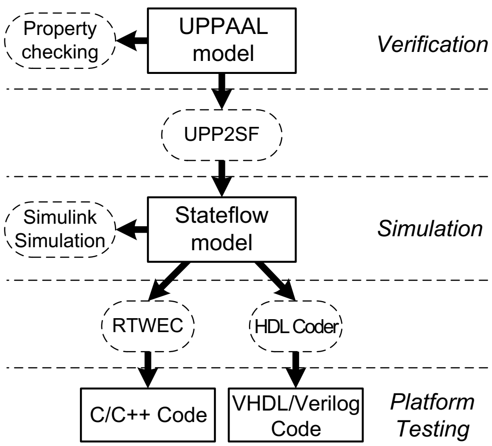

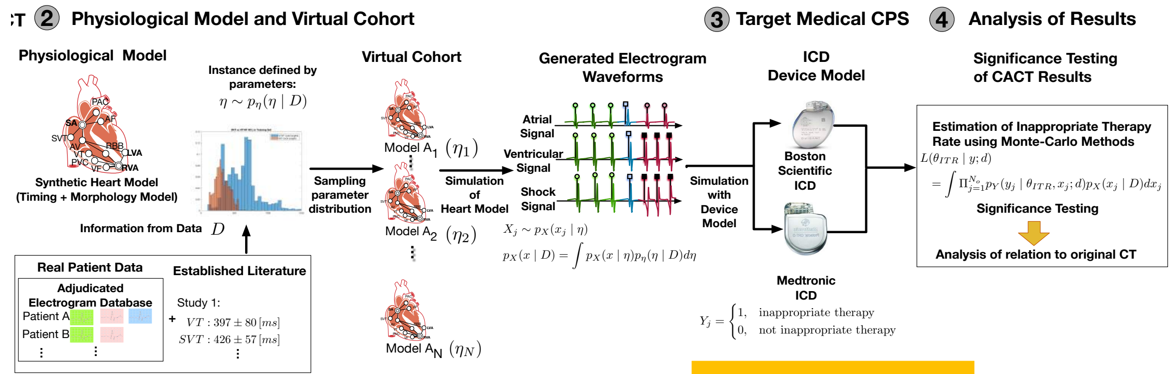

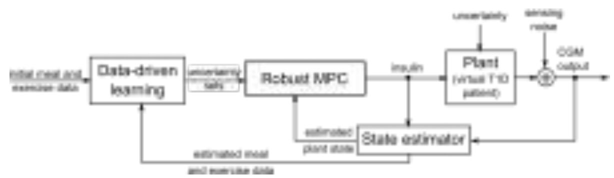

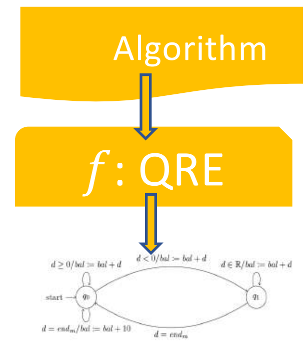

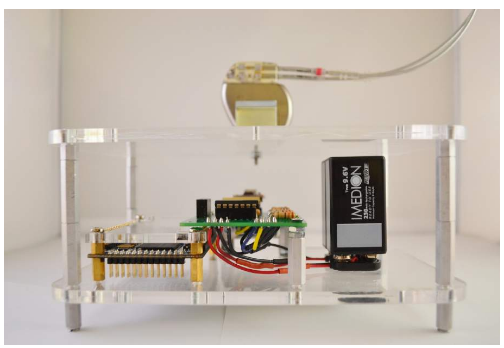

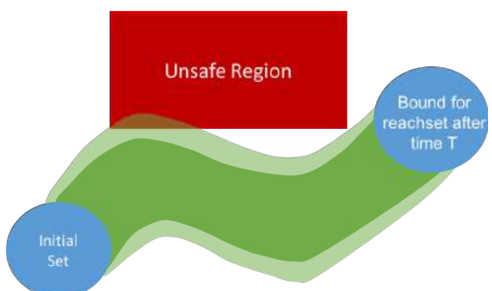

Too wide reachset
may result in false
positive

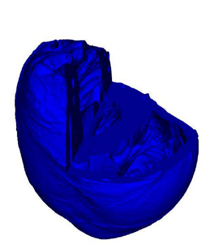

ESE 680-004: Digital Twins - Model-Based Embedded Systems