Compositional, Approximate, and Quantitative Reasoning for Medical Cyber-Physical Systems with Application to Patient-Specific Cardiac Dynamics and Devices

Radu Grosu[1], Elizabeth Cherry[2], Edmund M. Clarke[3], Rance Cleaveland[4], Sanjay Dixit[5], Flavio H. Fenton[6], Sicun Gao[3], James Glimm[1], Richard A. Gray[7], Rahul Mangharam[5], Arnab Ray[8], and Scott A. Smolka[1]

[1] Stony Brook University
[2] Rochester Institute of Technology
[3] Carnegie Mellon University
[4] University of Maryland
[5] University of Pennsylvania
[6] Georgia Institute of Technology
[7] U.S. Food and Drug Administration
[8] Fraunhofer USA Center for Experimental Software Engineering

Abstract. The design of bug-free and safe medical device software is challenging, especially in complex implantable devices that control and actuate organs who's response is not fully understood. Safety recalls of pacemakers and implantable cardioverter defibrillators between 1990 and 2000 affected over 600,000 devices. Of these, 200,000 or 41%, were due to firmware issues that continue to increase in frequency. According to the FDA, software failures resulted in 24% of *all* medical device recalls in 2011. There is currently no formal methodology or open experimental platform to test and verify the correct operation of medical-device software within the closed-loop context of the patient.

The goal of this effort is to develop the foundations of modeling, synthesis and development of *verified medical device software* and systems *from verified closed-loop models* of the device and organ(s). Our research spans both implantable medical devices such as cardiac pacemakers and physiological control systems such as drug infusion pumps which have multiple networked medical systems. These devices are physically connected to the body and exert direct control over the physiology and safety of the patient. The focus of this effort is on (a) Extending current binary safety properties to quantitative verification; (b) Development of patient-specific models and therapies; (c) Multi-scale modeling of complex physiological phenomena and compositional reasoning across a range of model abstractions and refinements; and (d) Bridging the formal reasoning and automated generation of safe and effective software for future medical devices.

1 Introduction

Between 1992-1998, less than 10% of medical devices were recalled due to software issues. This rate has more than doubled in 2011 with software failures accounting for

T. Margaria and B. Steffen (Eds.): ISoLA 2014, Part II, LNCS 8803, pp. 356–364, 2014.

24% of all medical device recalls. There is currently no formal design methodology or open experimental platforms that can be used to ensure the correct operation of medical devices *within the physiological closed-loop context*. Furthermore, the present approach of ad hoc and open-loop testing of medical device software and the design process significantly increase the time and cost for validation and do not provide strong guarantees on the safety and efficacy of the closed-loop system of the device and the patient. Given the increasing complexity and features built in medical devices, the rate and volume of devices recalled will continue on its current trajectory, unless a systematic approach for medical device software verification, validation and testing, within clinical and physiological relevant contexts, is adopted.

The focus of this proposal is on the development of a model-based design framework for medical devices to verify and test the safety and efficacy of device software for implantable cardiac devices such as pacemakers and defibrillators. This will be accomplished in three phases:

(a) Integrated Functional and Formal Modeling: We propose a multi-scale modeling approach where abstract physiological and device models are used to prove basic safety closed-loop properties and progressively refined models automatically prove more complex properties. We are particularly interested in cases where the device may drive the heart into unsafe states, such as in Pacemaker Mediated Tachycardia. To accomplish this, we will develop approximate and probabilistic physiological models for quantitative verification for competitive analysis of new cardiac rhythm therapies.

(b) Patient-specific Modeling: Using the generalized modeling approaches we will now employ patient data to develop patient-specific tuned heart models and conduct sensitivity and parametric analysis for model-based clinical trials of implantable cardiac devices.

(c) Pre-Clinical Validation and Platforms The modeling effort will be directed and supported by clinical validation with evaluation of therapies on animal models and organs. The heart and device models and the therapies developed in this effort will be implemented in closed-loop testing platforms to standardize the toolchains for low-cost and efficient medical device software evaluation. With collaboration with the US FDA, the proposed framework, models, platforms and toolchain will be harmonized into the current regulatory guidelines for development of high-confidence medical device software and systems.

1.1 From Verified Models to Verified Code for Medical Devices

Model-based approaches are revolutionizing the development of cyber-physical systems in general, and embedded control systems in particular. In these paradigms, which are variously called Model-Based Development (MBD) or Model-Driven Engineering (MDE), engineers first build models of the components of the system under development. They then use simulations to verify that the system exhibits desired properties [3, 4, 9, 17, 19], synthesis techniques (“autocoding”) to generate portions of the implementation automatically, and hybrid simulation/hardware-test infrastructure (“hardware-in-the-loop testing”) to verify implementations of components as they become available. The motivations for MBD/MDE approaches stem from time and cost

efficiencies in engineering processes: the "virtual prototyping" enabled by computer-based modeling permits much more thorough analysis of a design, at much lower cost, than does traditional physical prototyping.

The use of MBD/MDE is especially advanced in the automotive and aerospace control domains, where detailed simulation models for the physics of controlled systems ("plants") have been developed and serve as the basis for assessing models of control strategies proposed by engineers building these vehicles. In other domains, such as medical-device design, these techniques have yet to achieve much headway, due in part to a lack widely accepted behavioral models for human biological systems, but also due to the wide variability observed in individual patient's biological functions.

The goals of this proposal are to develop the theoretical and practical underpinnings of a new verification framework for cyber-physical systems that would support compositional, highly parameterizable, approximate, and quantitative reasoning; to build the tool support for conducting the deep analysis this framework will allow; and to use these tools and techniques to advance the state of the art in cardiac therapy devices. We are particularly interested in *closed-loop verification* of cardiac device software and therapies [15, 19]. In this setting, a computational model of the heart (the biological plant) is under closed-loop control of a computational model of the cardiac device (the controller), and verification is conducted on this closed-loop system. Moreover, we will develop a *multi-scale* formal modeling approach, in which simpler properties are verified using more abstract models for the heart and device, while more complex properties require progressively refined plant and controller models.

We are also very interested in developing *patient-specific* heart models, which we plan to obtain from ablation and other cardiac-related procedures. Having such models in the verification loop will improve the level of confidence in the safety and efficacy of the device, thereby potentially reducing the expense of failed clinical trials.

An architectural overview of our proposed framework, which we call HyRes, is given in Figure 1.[1] In what follows, we summarize our proposed work on the verification technology needed to support closed-loop verification of medical CPSs, its application to cardiac devices, especially the recent proposed Low-Energy Anti-Fibrillatory Pacing (LEAP) [16] approach of PIs Fenton and Cherry and its interaction with more traditional pacing and anti-arrhythmic therapies and our planned education and outreach activities.

2 Computational Foundations for Medical CPSs

Compositional Reasoning (Plug and Play). If two components (subsystems) are approximately equivalent (they can simulate each other's behavior up to a small error δ), then it would be highly desirable if you could replace one with the other in a larger context in a guaranteed *safe* manner; i.e., the resulting total behaviors are approximately equivalent up to a small error ϵ, which is a function of δ. For this to be the case, one needs to provide appropriate *proof rules* (based most likely on a *small-gain condition*), since in most interesting cases the larger context is nonlinear.

[1] The name HyRes derives from *Hybrid systems*, a modeling formalism for CPSs amenable to formal verification, and the *Resolution* or precision at which the verification is carried out.

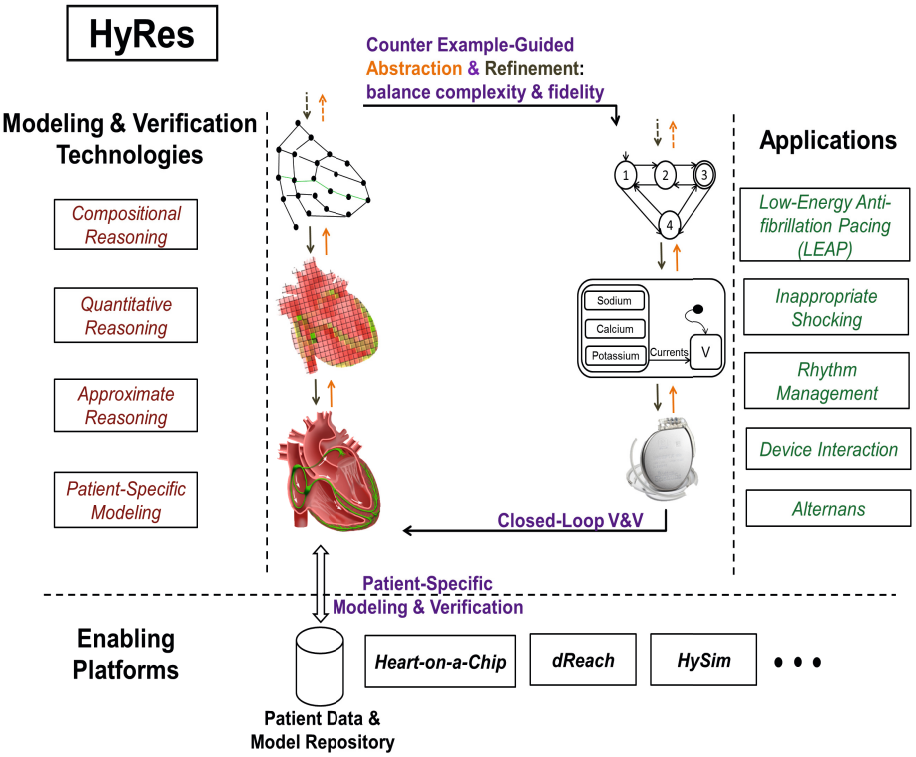

Fig. 1. The HyRes framework for closed-loop verification of Medical CPS. The verification technologies we propose to develop are shown on the left, the intended applications on the right, and the supporting computational platforms and repositories along the bottom of the figure. A hierarchy of models, capturing the electrophysiology of the heart at varying levels of complexity, will be devised using abstraction and refinement techniques. The figure shows a highly detailed model at the base of the hierarchy, which is spatially abstracted to obtain a grid-based computational model, which is further abstracted to obtain a network of Timed automata for reasoning about timing-related properties.

PIs Grosu, Smolka and others have recently used this kind of reasoning to show that the 13-variable sodium-channel component of the 67-variable IMW cardiac-cell model (Iyer-Mazhari-Winslow) can be replaced by an approximately bisimilar, 2-variable HH-type (Hodgkin-Huxley) abstraction [12–14, 18] . Moreover, this substitution of (approximately) equals for equals is safe in the sense that the approximation error between sodium-channel models is not amplified by the feedback-loop context in which it is placed.

Being able to reason about dynamical systems compositionally [1, 13, 14] is important for two reasons: the plug-and-safely-play nature of compositional reasoning is highly *efficient*, as it avoids the state-explosion problem that bedevils automated verification; and composition of subsystems can be used to uncover *bad interactions* between subsystems, AKA the *feature interaction* problem.

Compositional reasoning can also be used as basis for *synthesis*: not just controller synthesis (well-studied), but plant synthesis. That is, given a controller, infer the plant for which it works.

Approximate Logical Reasoning (Maximum Precision). When proving that a system satisfies a particular property, there is a "Plank discretization" (precision) limit. For example, suppose a curve (given by an analytic function) separates the plane, and that there is a small grid of "Plank size". For the grid squares cut by the (zero-width) curve, we cannot say whether they are on one side or the other of the curve (that is, satisfy or do not satisfy the property). For all other squares, one has a definitive answer.

Approximate verification can also be used to turn an undecidable decision problem over the reals into a decidable decision problem, and efficiently at that. In a multi-time-scale approach to verification, choose the level of approximation that matches the granularity of the time-scale under consideration.

The team brings expertise in this approach to the proposed effort in the form of the dReal and dReach *reachability analysis platform for nonlinear hybrid systems*, cultivated by PIs Gao and Clarke during the course of the CMACS NSF Expedition in Computing. In the spirit of this proposal, Gao and Clarke have applied dReal/dReach to the analysis of a highly nonlinear cardiac-cell model [7–9].

Quantitative Logical Reasoning (How Good). Classical temporal-logic model checking provides a boolean yes/no answer to the question "Does a system Σ satisfy a temporal logic formula φ?" When Σ is a dynamical system such as a CPS, one can demand a more quantitative assessment of how well Σ does or does not satisfy φ. If φ is satisfied by Σ, then how *robustly* is it satisfied? If φ is violated, then how badly is it violated? How many (abstract) points in the state space violate the property? If a point violates the property, then by how much does it violate it? Quantitative reasoning [2, 10, 11] can be seen as lifting the model checking problem from a boolean setting to one in which the results are interpreted over a metric space.

With quantitative reasoning, once can also augment temporal logic with *quantitative operators*. For example, consider the following *convergence* property $FG(x \leq \tau)$, which states that eventually the value of x is always less than or equal to threshold τ. In the quantitative setting, one can also measure the *speed* at which the G-subformula eventually becomes true, and the *average* value of x, once x always $\leq \tau$.

Quantitative reasoning can also play a role *diagnostically*. Consider the safety property: an ICD should not deliver an inappropriate shock, or the occurrence of one should be minimized. Quantitatively, one can compute e.g. the average amount of energy consumed by an ICD every time an inappropriate shock is delivered to the patient.

Adversarial Reasoning (Games, and Open Systems). The controller and the plant do not always represent a closed system. They may be in a game-like situation with the environment from which they receive additional *adversarial* input. A winning strategy for the controller + system is one for which they behave *safely* regardless of the the moves the environment makes. The environment may be nondeterministic (making it difficult to compete against), or stochastic (making it somewhat easier to deal with

as one can then model it with belief states and partially observable Markov decision processes).

Closed-loop Verification with Automated Model Abstraction and Refinement. While complex physiological models of the heart with over 4 million finite elements or 100K ODEs exist, they do not provide a suitable level of interaction with a device such as a pacemaker which only observes the state of the heart from two or three points. We propose a multi-scale formal modeling approach to verify a set of closed-loop properties (i.e., where the heart can affect the device and, more importantly, where the device can drive the heart into safe/unsafe states).

In this approach, simpler properties are verified with more abstract models of the heart/device, while more complex properties require progressively refined plant models. In support of this approach, we will develop automated a Counter-Example-Guided Abstraction and Refinement (CEGAR) framework to balance model complexity and fidelity in accordance with increasingly complex closed-loop issues such as Pacemaker Mediated Tachycardia, where the pacemaker drives the heart into an unsafe state.

3 Application to Patient-Specific Cardiac Models, Therapies, and Devices

Patient-Specific Modeling The construction of patient-specific heart models will enable:

- Improved level of confidence in the safety and efficacy of the device with a patient-model in the loop. This will reduce the expense of failed clinical trials and potentially reduce the extent of clinical trials, in general.
- Physicians to maintain actionable patient records between operations and perform pre-op evaluations on these models.
- Semi-automatic tuning of device parameters to the specific patient requirements.
- Model-based training of EP fellows and medical students.

The requisite data will be obtained from ablation and other medical, cardiac-related procedures. We will use this data to learn/personalize heart models, device settings, etc. We refer to this process as *Patent-Specific (P-S) Modeling*. In order to have P-S heart models, we will need to incorporate patient data in our models and tool chain. The most accurate patient data comes from the electro-physiology study before implantation. Catheters with probes are inserted into the patient's heart and local electrical activities are recorded as Electrogram (EGM) signals. From the EGM signals, we can extract timing delays between different heart locations. As the VHM and EP studies use the same parameters, we can incorporate patient EGM data to form a P-S heart model. We will pursue this in two steps:

(1) Model Construction Using Synthetic EGMs: The VHM is able to generate synthetic EGM signals. Since EGM signals mainly carry timing information, the synthetic EGMs are comparable to realistic EGMs - however with known probability distributions. As the VHM is a more controlled environment than a real patient, it is much

easier to evaluate for quantitative verification for patient-specific conditions.
(2) Model Construction using Realistic EGMs: We will use EGMs from a real patient to construct our model. This will require noise filtering, determining the catheter positioning and benchmark analysis for the constructed model.

We are also interested in *Property-Based Modeling*. If one is only interested in timing-related aspects of patient therapy, as may be the case with a pacemaker, learn a Timed Automaton (TA) model of the patient's heart and of the device. If voltage is of interest, for example in the treatment of VT and VFib, learn a voltage-based Hybrid Automaton (HA) model.

A key aspect of property-based modeling will be to ensure that the models we derive are related to one another in the ways we intend them to be. For example, is the TA model an abstraction of the HA model? We can ensure this is the case by following a process of *abstraction refinement* in deriving e.g. the HA model from the TA one.

Such a framework will allow for enforcement of property priorities, where under certain physiological conditions, some properties may be violated while higher-priority properties remain enforced. This will allow for verification of multi-scale and multi-mode systems, whose properties must adapt to the mode of the patient.

Closed-Loop Verification of Cardiac Therapies and their Interactions. We will put a P-S cardiac model in the loop with a cardiac device with P-S parameter settings, and apply the analysis techniques developed in Part I of the proposal to the resulting systems. Recent work by the PIs in compositional verification [12–14, 18] (Grosu and Smolka), approximate verification [5–8] (Clarke and Gao), and quantitative reasoning [10] (Grosu, Smolka, and others) on which this proposal will build makes us confident that we will be successful.

We will consider both pacemakers and ICDs and their interactions. Some devices combine a pacemaker and ICD in one unit for persons who need both functions, and this is becoming more and more common. Thus, the need to carefully analyze their interactions is on the rise.

Low-Energy Anti-Fibrillatory Pacing (LEAP). PIs Fenton, Cherry, and others have developed a new approach to eradicating life-threatening arrhythmias. Instead of one large jolt of electricity to the heart, the new approach, called low-energy anti-fibrillatory pacing (LEAP), uses a series of smaller electrical pulses. An article describing this breakthrough appeared in a recent issue of Nature [16]. The goal of LEAP is not to eliminate the arrhythmia at once, but rather to synchronize the electrical state of the heart gradually. In this way, undesirable side effects can be avoided while still restoring the heart to its normal condition.

Computational modeling, initially using simple models and then more complex models [3, 17], validated this approach and provided guidance for a series of preclinical experimental trials that demonstrated LEAP's effectiveness. The modeling and analysis techniques put forth in this proposal will be used to further optimize the method so that it can be used in human clinical trials. We will also study its interactions with other pacing-based therapies.

References

1. Bartocci, E., Bortolussi, L., Nenzi, L.: A temporal logic approach to modular design of synthetic biological circuits. In: Gupta, A., Henzinger, T.A. (eds.) CMSB 2013. LNCS (LNBI), vol. 8130, pp. 164–177. Springer, Heidelberg (2013)
2. Bartocci, E., Bortolussi, L., Nenzi, L., Sanguinetti, G.: On the robustness of temporal properties for stochastic models. In: Proc. of HSB 2013: The 2nd Intern. Workshop on Hybrid Systems and Biology. EPTCS, vol. 125, pp. 3–19 (2013)
3. Bartocci, E., Cherry, E.M., Glimm, J., Grosu, R., Smolka, S.A., Fenton, F.H.: Toward real-time simulation of cardiac dynamics. In: Proceedings of the 9th International Conference on Computational Methods in Systems Biology, CMSB 2011, pp. 103–112. ACM, New York (2011)
4. Bartocci, E., Singh, R., von Stein, F.B., Amedome, A., Caceres, A.J., Castillo, J., Closser, E., Deards, G., Goltsev, A., Ines, R.S., Isbilir, C., Marc, J.K., Moore, D., Pardi, D., Sadhu, S., Sanchez, S., Sharma, P., Singh, A., Rogers, J., Wolinetz, A., Grosso-Applewhite, T., Zhao, K., Filipski, A.B., Gilmour, R.F., Grosu, R., Glimm, J., Smolka, S.A., Cherry, E.M., Clarke, E.M., Griffeth, N., Fenton, F.H.: Teaching cardiac electrophysiology modeling to undergraduate students: Laboratory exercises and GPU programming for the study of arrhythmias and spiral wave dynamics. Adv. Physiol. Educ. 35(4), 427–437 (2011)
5. Gao, S., Avigad, J., Clarke, E.M.: Delta-complete decision procedures for satisfiability over the reals. In: Proceedings of the 6th International Joint Conference on Automated Reasoning (IJCAR), pp. 286–300 (2012)
6. Gao, S., Avigad, J., Clarke, E.M.: Delta-decidability over the reals. In: Proceedings of the 27th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), pp. 305–314 (2012)
7. Gao, S., Kong, S., Clarke, E.M.: dReal: An SMT solver for nonlinear theories over the reals. In: Bonacina, M.P. (ed.) CADE 2013. LNCS (LNAI), vol. 7898, pp. 208–214. Springer, Heidelberg (2013)
8. Gao, S., Kong, S., Clarke, E.M.: Satisfiability modulo ODEs. In: Proceedings of the 13th International Conference on Formal Methods in Computer Aided Design, FMCAD (2013)
9. Grosu, R., Batt, G., Fenton, F.H., Glimm, J., Le Guernic, C., Smolka, S.A., Bartocci, E.: From cardiac cells to genetic regulatory networks. In: Gopalakrishnan, G., Qadeer, S. (eds.) CAV 2011. LNCS, vol. 6806, pp. 396–411. Springer, Heidelberg (2011)
10. Grosu, R., Peled, D., Ramakrishnan, C.R., Smolka, S.A., Stoller, S.D., Yang, J.: Compositional branching-time measurements. In: Bensalem, S., Lakhneck, Y., Legay, A. (eds.) FPS 2014 (Sifakis Festschrift). LNCS, vol. 8415, pp. 118–128. Springer, Heidelberg (2014)
11. Grosu, R., Peled, D., Ramakrishnan, C.R., Smolka, S.A., Stoller, S.D., Yang, J.: Using statistical model checking for measuring systems. In: Margaria, T., Steffen, B. (eds.) ISoLA 2014, Part II. LNCS, vol. 8803, pp. 223–238. Springer, Heidelberg (2014)
12. Islam, M.A., Murthy, A., Bartocci, E., Girard, A., Smolka, S., Grosu, R.: Compositionality results for cardiac cell dynamics. In: Gupta, A., Henzinger, T.A. (eds.) CMSB 2013. LNCS, vol. 8130, pp. 242–244. Springer, Heidelberg (2013)
13. Islam, M.A., Murthy, A., Bartocci, E., Cherry, E.M., Fenton, F.H., Glimm, J., Smolka, S.A., Grosu, R.: Model-order reduction of ion channel dynamics using approximate bisimulation. Theoretical Computer Science (in press, 2014)
14. Islam, M.A., Murthy, A., Girard, A., Smolka, S.A., Grosu, R.: Compositionality results for cardiac cell dynamics. In: Proc. of HSCC 2014: The 17th International Conference on Hybrid Systems: Computation and Control, HSCC 2014, pp. 243–252. ACM, New York (2014)
15. Jiang, Z., Pajic, M., Alur, R., Mangharam, R.: Closed-loop verification of medical devices with model abstraction and refinement. STTT 16(2), 191–213 (2014)

16. Luther, S., Fenton, F.H., Kornreich, B.G., Squires, A., Bittihn, P., Hornung, D., Zabel, M., Flanders, J., Gladuli, A., Campoy, L., Cherry, E.M., Luther, G., Hasenfuss, G., Krinsky, V.I., Pumir, A., Gilmour, R.F., Bodenschatz, E.: Low-energy control of electrical turbulence in the heart. Nature 475(7355), 235–239 (2011)
17. Murthy, A., Bartocci, E., Fenton, F., Glimm, J., Gray, R., Cherry, E., Smolka, S., Grosu, R.: Curvature analysis of cardiac excitation wavefronts. IEEE/ACM Transactions on Computational Biology and Bioinformatics 10(2), 323–336 (2013)
18. Murthy, A., Islam, M.A., Bartocci, E., Cherry, E.M., Fenton, F.H., Glimm, J., Smolka, S.A., Grosu, R.: Approximate bisimulations for sodium channel dynamics. In: Gilbert, D., Heiner, M. (eds.) CMSB 2012. LNCS, vol. 7605, pp. 267–287. Springer, Heidelberg (2012)
19. Pajic, M., Jiang, Z., Lee, I., Sokolsky, O., Mangharam, R.: From verification to implementation: A model translation tool and a pacemaker case study. In: Proceedings of IEEE 18th Real Time and Embedded Technology and Applications Symposium, Beijing, China, April 16-19, pp. 173–184 (2012)