

Three Challenges in Cyber-Physical Systems

Rahul Mangharam, Houssam Abbas, Madhur Behl, Kuk Jang, Miroslav Pajic and Zhihao Jiang

School of Engineering and Applied Sciences

University of Pennsylvania

Philadelphia, Pennsylvania, USA 19104

Email: {rahulm, habbas, mbehl, jangkj, pajic, zhihaoj}@seas.upenn.edu

Abstract—The tight coupling of computation, communication and control with physical systems such as actuation of closed-loop medical devices within the human body, peak power minimization by coordination of controllers across large industrial plants, and fast life-critical decision making by autonomous vehicles, present a set of fundamental and unique challenges. Each of these require new approaches at the intersection of multiple scientific, human and systems disciplines. We discuss five such challenges which require creative insights and application of model-based design, control systems, scheduling theory, formal methods, statistical machine learning and domain-specific experimentation. We ask the following questions: (1) An autonomous medical device is implanted to control your heart over a period of 5-7 years. How do you guarantee the software in the device provides safe and effective treatment under all physiological conditions? (2) Electricity prices in the US have summer peaks that are over 32x their average prices and winter peaks that are 86x. How can buildings respond to massive swings in energy prices at fast time scales? (3) While wireless has been successfully used for open-loop monitoring and tracking, how can we operate closed-loop control systems over a network of wireless controllers. Furthermore, how can we ensure robust, optimal and secure control in the presence of node/link failures and topology changes?

I. INTRODUCTION

We describe five challenge problems at the foundations of Cyber-Physical Systems (CPS). CPS are the new generation of time-critical and safety-critical Real-Time Embedded Systems where computation and communication are tightly coupled to control large, complex and “messy” plants. Unlike classical Real-Time Systems where the system is designed in a constrained manner to limit the complexity for predictability, in current CPS, the plants are difficult to model precisely because they are non-deterministic, interactive and often scale to thousands of controllers. We outline research approaches to address these structural concerns in the design of future Cyber-Physical Systems through closed loop modeling, architectures, algorithms and platforms. The eventual goal of efforts such as this is to develop CPS to transform how we interact with and manipulate the physical world, just as the Internet transformed how we interact with information systems.

A key aspect of this cross-cutting work is that each challenge bridges two or more well-established scientific fields such as scheduling theory, control systems, formal methods, statistical machine learning, and experimentation. We illustrate this through five domain-specific problems spanning the safety of closed-loop medical devices, control over wireless, data-driven control of energy systems in volatile price markets, safe decision making in autonomous systems and co-design

of computation and control for future automotive systems. We address the foundations of CPS across four themes: Modeling, Algorithms and Architectures.

- **CPS Modeling:** From verified models into verified code for life-critical systems. We focus on the formal modeling, synthesis and certification of high-confidence medical device software and systems. We describe efforts on both implantable medical devices and physiological control systems to ensure that both the functional and formal aspects are verified, validated and tested within the closed-loop context of their physiological systems.
- **CPS Algorithms:** Data-driven Controller Synthesis. In highly volatile electricity pricing markets, we required buildings and built environments to rapidly respond to peak pricing spikes. We describe a new data-driven approach for synthesis of control strategies to achieve real-time demand response while ensuring custom climate conditions are maintained.
- **CPS Architectures:** Distributed Control over Wireless Networks. We describe radical architectural approaches for completely in-network computation for robust, optimal and secure control over unreliable infrastructure.

II. THE DESIGN OF SAFE AND EFFECTIVE CLOSED-LOOP MEDICAL DEVICES

The design of safe, bug-free, and efficient medical device software is challenging, especially in implantable devices that directly control and actuate partially understood organs. However, such design is also essential; safety recalls of pacemakers and implantable cardioverter defibrillators between 1990 and 2000 affected over 600,000 devices [1]. Importantly, 41% of these recalls were due to software issues [2]. According to the US Food and Drug Administration (FDA), in 1996, 10% of *all* medical device recalls were caused by software-related issues. This percentage rose to an average of 15% of recalls from 2008 to 2012. And, surprisingly given these numbers, there is currently no formal methodology or open experimental platform to test the correct operation of medical device software within the *closed-loop* context of the patient [3]. Unlike other industries such as aviation and automotive where the safety concern is focused on a well-defined physical plant [4], [5], in the medical device domain patient response is complex and nondeterministic. As a result there are no well-established standards for the development of medical device software that directly control and actuate the patient.

For device manufacturers, this has prompted recent interest in applying formal modeling [6], [7], [8] and verification techniques in medical devices software development [9], [10].

How do you guarantee the device software will not adversely affect the patient under all physiological conditions?

An effective software verification methodology is therefore needed for the risk analysis and certification of medical device software during the FDA’s pre-market submission phase. Testing medical device software currently is ad hoc, open loop, and very expensive [11], [12]. Test generation must be interactive and adaptive and must consider the current state of the patient when generating the next input in a way that advances the purpose of the test. The problem, therefore, becomes one of controller synthesis and cannot be addressed by an off-the-shelf model checker [13]. The key challenge is in the generation of *physiologically relevant* software that does not provide inappropriate therapy or adversely affect the patient. This requires validated patient models of the appropriate abstraction levels and that testing is conducted within the closed-loop context of the patient model.

Consequently, there is a need for fundamentally novel approaches to the closed-loop modeling, analysis, and design of medical cyber-physical systems, as well as the development of holistic, heterogeneous physiological models, approaches, and tools that address the many different physical, functional and logical aspects of the device and patient interaction. The scientific agenda of this research on medical cyber-physical systems is to unify theories of formal methods, control of physiological systems, communication and computing systems. Our early efforts are conducted with implantable medical devices, such as cardiac pacemakers and defibrillators, and physiological control systems such as infusion pumps with networks of discrete sub-systems. Both feature tight coupling with the patient-in-the-loop, exposing safety and efficacy risks of autonomous and automatic control of the body.

1. Implantable Medical Devices - Cardiac Pacemakers: The human heart is perhaps the most important real-time system, generating electrical impulses that determine the heart’s rhythm and proper function. Irregularities with timing, i.e. cardiac arrhythmias, cause inefficient and unsafe function of the blood-oxygen system, necessitating the maintenance of the heart rate artificially. The cardiac pacemaker is a rhythm management device that maintains the minimum heart rate and synchrony between its chambers, thereby improving the condition of patients with cardiac arrhythmias. Cardiac rhythm management devices have grown in complexity with over 80,000 to 100,000 lines of code [14]. The primary approach to system-level testing of medical devices is unit testing using a playback of pre-recorded electrogram and electrocardiogram signals. This tests if the input signal triggers a particular response by the pacemaker but cannot evaluate if the response was appropriate for the patient condition. Furthermore, this approach of *open loop* “tape testing” is unable to check for safety violations due to inappropriate stimulus by the pacemaker. Pacemaker Mediated Tachycardia (PMT), a condition where the pacemaker inappropriately drives the intrinsic heart-




Fig. 1. Pacemaker operating in a closed-loop with the heart. The leads sense cardiac electrophysiological activity from inside the heart tissue (AS/VS = Atrial/Ventricular Sense event) and actuate the heart (AP/VP = Atrial/Ventricular Pacing event) to maintain the heart rate. rate toward the maximum rate, is a strong example of why we need an interactive and adaptive *closed-loop* verification and testing of such systems. With a tape test, PMT would not be observed and the response of the pacemaker could be classified as appropriate therapy.

Our proposed model-based design (MBD) for Medical CPS begins with developing integrated functional and formal heart models that interact with real and modeled pacemakers for closed-loop verification and testing [15]. As shown from the top of Fig. 2, the heart-pacemaker closed-loop systems is first modeled abstractly to facilitate verification of the basic pacemaker design with maximum coverage [16]. In our case, we use timed automata [17] and the UPPAAL model checker [18], [19], [20] at this design stage. Next, the models are translated to more detailed models that take into account the complex dynamics of the heart and interaction with more detailed pacemaker model [21], [22], [23]. We use Stateflow and Simulink [24] at this design stage. These models are validated by physicians for their clinical relevance. The automatic model translation procedure, from UPPAAL to Stateflow, ensures that abstract models used for verification over-approximate the more detailed models used downstream [25]. Once the detailed models pass simulation-based testing with closed-loop dynamics, they are automatically generated into code and are subject to platform-level integration testing [26] as shown in Fig. 3. This MBD approach ensures the closed-loop safety properties are retained through the design toolchain and facilitates the development of verified software from verified models.

2. Physiological Control Systems: The understanding

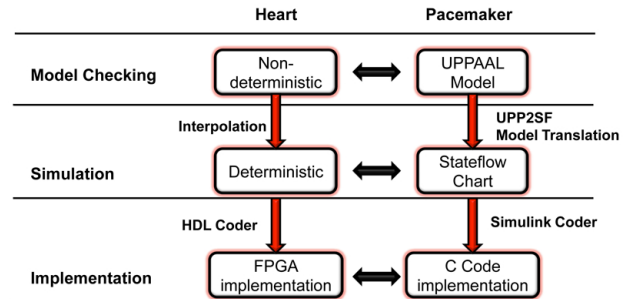


Fig. 2. From closed-loop verification to simulation-based testing, code generation and platform evaluation

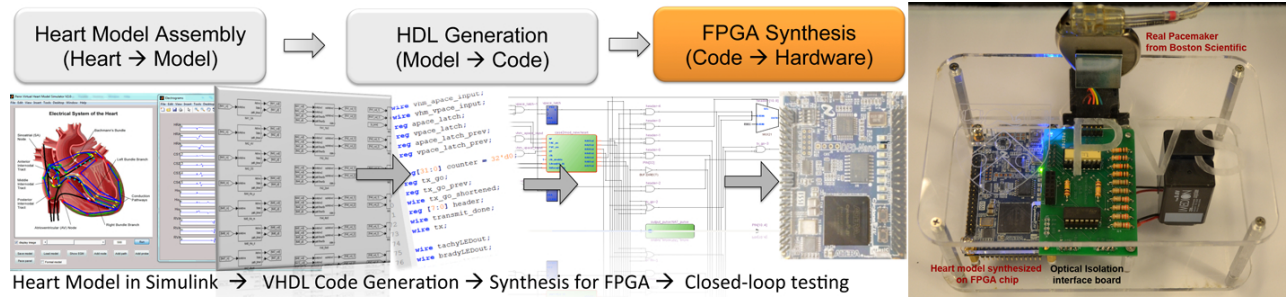


Fig. 3. From verified models to verified code: Translating models to a heart-on-a-chip platform for closed-loop testing of cardiac devices

of closed-loop safety analysis of single-system devices such as pacemakers has naturally broadened our work to physiological control systems featuring multiple networked devices with the patient-in-the loop [27]. Such a clinical scenario is viewed as a control system, in which the patient is the plant, bedside monitors are sensors and drug infusion pumps are actuators [6]. In this setup, caregivers traditionally perform the role of the controller. In many cases automatic controllers for drug infusion can reduce the burden on the caregiver and avoid human errors. However, vendors of medical equipment continue to avoid closed-loop scenarios due to an insufficient understanding of the human body's response to treatment. Furthermore, a particular challenge arises from the complex interplay between the continuous dynamics of the patient's reaction to treatment, and discrete controller and network. Consequently, there is a need for model driven safety analysis of closed-loop medical systems within uncertain parameters. Both the abstract, formal model and the detailed, informal model are needed in the process of verification, validation, and regulatory approval of closed-loop medical device systems. Formal models allow us to exhaustively explore the possible behaviors of the system and prove its safety, detailed models allow us to use high-fidelity simulation that take real system dynamics into account. Both kinds of results can be used to make the case for regulatory approval if the abstract model is guaranteed to over-approximate the patient's dynamics with respect to the control algorithms used.

3. Model-based Clinical Trials: Regulatory authorities require that the safety and efficacy of a new high-risk medical device be proven in a Clinical Trial (CT), in which the effects of the device on a group of patients are compared to the effects

of the current standard of care. Phase III trials can run for several years, cost millions of dollars, and pose an inherent risk to the patients by exposing them to an unproven device. With the use of computational modeling and simulation, we would like to investigate how to use a large model-based synthetic group of patients and device models to improve the planning and execution of a CT so as to increase the chances of a successful trial.

As an example, we apply our initial efforts are in applying it to a real CT that compares two algorithms within implantable cardioverter defibrillators (ICDs) for the detection of potentially fatal cardiac arrhythmias [28]. In 2011, a CT posited that one algorithm (Boston Scientific's) would be better than the other (Medtronic's) but the results of the trial were opposite to this hypothesis [29]. We begin by modeling the heart and processing 100's of real patients' electrogram signals, mapping the timing and morphology components of the signals to the heart model. This is followed by generating a population of 10,000+ synthetic heart models, by perturbing the parameters of the initial heart models for different arrhythmias, and implementing diagnostic algorithms of two very commonly used ICD platforms in the USA, i.e. Boston Scientific and Medtronic. We conducted conformance testing to validate our device models against real ICDs. Now, using the closed-loop of the device models and synthetic patient population we conducted multiple trials to compare the performance of the two algorithms to appropriately discriminate between potentially fatal ventricular tachycardias (VT) and non-fatal SupraVentricular Tachycardias (SVTs).

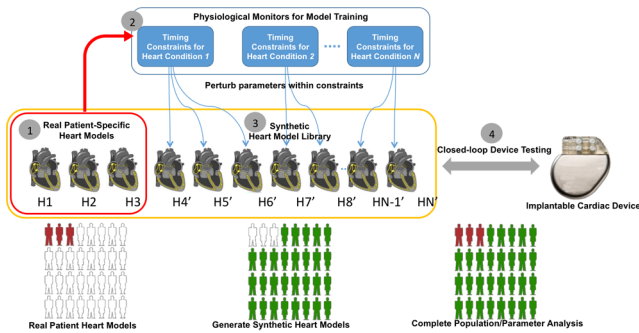


Fig. 4. Model-based Clinical Trials

The results of our model-based clinical trials (MBCT) indicate we could have accurately predicted this with our model, i.e. that Boston Scientific's algorithm was less able to discriminate between SVT and VT and so may lead to inappropriate therapy. We further demonstrated that the result continues to hold if we vary the characteristics of the synthetic population and device parameters. While MBCTs do not seek to replace a CT, they may provide early insight into the factors which affect the outcome at a fraction of the cost and duration and without the ethical issues. This effort is an early step towards using computer modeling as regulatory-grade evidence for medical device certification.




Fig. 5. Majority of DR today is manual and rule-based. The fixed rule based DR is inconsistent and could under-perform compared to the required curtailment, resulting in DR penalties. Using data-driven models DR-Advisor uses DR strategy evaluation and DR strategy synthesis for a sustained and sufficient curtailment.

III. DATA-DRIVEN MODELING, CONTROL AND TOOLS FOR CYBER-PHYSICAL ENERGY SYSTEMS

In 2013, a report by the U.S. National Climate Assessment provided evidence that the most recent decade was the nation's warmest on record [30] and 2015 is very likely to become the hottest year on record since the beginning of weather recording in 1880 [31]. Heat waves in summer and polar vortexes in winter are growing longer and pose increasing challenges to an already over-stressed electric grid.

Furthermore, with the increasing penetration of renewable generation, the electricity grid is also experiencing a shift from predictable and dispatchable electricity generation to variable and non-dispatchable generation. This adds another level of uncertainty and volatility to the electricity grid as the relative proportion of variable generation vs. traditional dispatchable generation increases. The volatility due to the mismatch between electricity generation and supply further leads to volatility in the wholesale price of electricity. For e.g., the polar vortex triggered extreme weather events in the U.S. in January 2014, which caused many electricity customers to experience increased costs. Parts of the Pennsylvania-New Jersey-Maryland (PJM) electricity grid experienced a 86 fold increase in the price of electricity from \$31/MWh to \$2,680/MWh in a matter of a few minutes [32]. Similarly, the summer price spiked 32 fold from an average of \$25/MWh to \$800/MWh in July of 2015. Energy industry experts are now considering the concept that extreme weather, more renewables and resulting electricity price volatility, could become the new norm.

Across the United States, electric utilities and independent service operators (ISOs or grid coordinators) are devoting increasing attention and resources to demand response (DR) [33]. It is considered as a reliable means of mitigating the uncertainty and volatility of renewable generation and extreme weather conditions and improving the grid's efficiency and reliability. The potential demand response resource contribution from all U.S. demand response programs is estimated to be nearly 72,000 megawatts (MW), or about 9.2 percent of U.S. peak demand [34] making DR the largest virtual generator in the U.S. national grid. The annual revenue to end-users from DR markets with PJM ISO alone is more than \$700 million [35]. Global DR revenue is expected to reach nearly \$40 billion from 2014 through 2023 [36].

In order to shield themselves from the volatility and risk of high prices, such consumers must be more flexible in their

electricity demand. Consequently, large commercial, industrial and institutional customers are increasingly looking to demand response programs to help manage their electricity costs. As shown in Fig. 5, DR programs involve a voluntary response of a building to a price signal or a load curtailment request from the utility or the curtailment service provider. Upon successfully meeting the required curtailment level the end-users are financially rewarded, but may also incur penalties for under-performing and not meeting a required level of load curtailment.

A. Challenges

On the surface demand response may seem simple. Reduce your power when asked to and get paid. However, in practice, one of the biggest challenges with end-user demand response for large scale consumers of electricity is the following: *Upon receiving the notification for a DR event, what actions must the end-user take in order to achieve an adequate and a sustained DR curtailment?* This is a hard question to answer because of the following reasons:

1. Modeling complexity and heterogeneity: Unlike the automobile or the aircraft industry, each building is designed and used in a different way and therefore, it must be uniquely modeled. The user expertise, time, and associated sensor costs required to develop a model of a single building (e.g., with EnergyPlus [37]) is very high [38] - often taking 5-12 months to model and tune the model. This is because usually a building modeling domain expert requires the geometry of a building from the building design and equipment layout plans, detailed information about material properties, and of equipment and operational schedules.

2. Control complexity and scalability: Upon receiving a notification for a DR event, the building's facilities manager must determine an appropriate DR strategy to achieve the required load curtailment. These control strategies can include adjusting zone temperature set-points, supply air temperature and chilled water temperature set-point, dimming or turning off lights, decreasing duct static pressure set-points and restricting the supply fan operation etc. For a large building, it is difficult to assess the effect of one control action on other sub-systems and on the building's overall power consumption because the building sub-systems are tightly coupled. Therefore, it is extremely difficult for a human operator to accurately gauge the building's or a campus's response.

3. Interpretability of modeling and control: Predictive models for buildings, regardless how sophisticated, lose their effectiveness unless they can be interpreted by human experts and facilities managers in the field. For e.g., artificial neural networks (ANN) obscure physical control knobs and interactions and hence, are difficult to interpret by building facilities managers. Therefore, the required solution must be transparent, human centric and highly interpretable.

B. Real-Time Data-driven Demand Response

We have developed a method called DR-Advisor (Demand Response-Advisor) [39], which acts as a recommender system

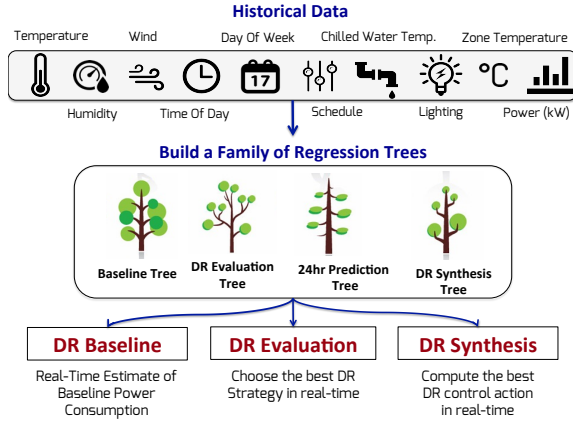


Fig. 6. DR-Advisor Architecture for real-time demand response

for the building’s facilities manager and provides the power consumption prediction and control actions for meeting the required load curtailment and maximizing the economic reward. Using historical meter and weather data along with set-point and schedule information, DR-Advisor builds a family of interpretable regression trees to learn non-parametric data-driven models for predicting the power consumption of the building (Figure 8). DR-Advisor can be used for real-time demand response baseline prediction, strategy evaluation and, most importantly, for control synthesis, without having to learn first principles based models of the building:

1. DR Baseline Prediction: A baseline is an estimate of the electricity that would have been consumed by a customer in the absence of a demand response event. Typical demand response programs rely upon financial incentive for customers based on the extent to which they reduce their energy consumption and therefore require a reliable system to measure the energy reduction. For this reason the measurement and verification of demand response is the most critical component of any DR program. Using regression trees based approaches, DR-Advisor achieves a prediction accuracy of 93% to 98.9% for baseline estimates of eight buildings on the Penn campus by just analyzing data as shown in Figure 8.

2. DR Strategy Evaluation: A DR strategy refers to what sequence of control actions, and at what times, a system (lighting, HVAC or plug loads) will actuate. Furthermore, there could be several of such fixed DR strategies, but only one specific strategy can be used at a time. This brings us to our question, *how can we choose good DR strategies from a pre-determined set of strategies?* Instead of predicting the baseline power consumption Y_{base} , in this case we want the ability to predict the actual response of the building Y_{kW} due to any given strategy. At the beginning of the DR event we use the auto-regressive tree for predicting the response of the building due to each rule-based strategy and choose the one which performs the best over the predicted horizon. The prediction and strategy evaluation is re-computed periodically throughout the event.

3. DR Control Synthesis: While black-box and data-driven machine learning approaches are suitable for prediction, our

primary contribution is making them capable of controller synthesis. Unlike rule-based DR, which does not account for building state and external factors, in DR synthesis the optimal control actions are derived based on the current state of the building, forecast of outside weather and electricity prices. We introduce a novel model based control with regression trees (mbCRT) algorithm to enable control with regression trees for real-time DR synthesis. Using the mbCRT algorithm, we can optimally trade off thermal comfort inside the building against the amount of load curtailment. While regression trees are a popular choice for prediction based models, this is the first time regression tree based algorithms have been used for controller synthesis with applications in demand response. Our synthesis algorithm outperforms rule based DR strategy by 17% (based on guidelines from Siemens) while maintaining bounds on thermal comfort inside the building.

We have evaluated the performance of DR-Advisor using a mix of real data from 8 buildings (1.2 million square feet) on the campus of the University of Pennsylvania, a large office building in Philadelphia and data-sets from a virtual building test-bed for the Department of Energy’s (DoE) large commercial reference building. We also compared the performance of DR-Advisor against other data-driven methods using a bench-marking data-set from AHRAE’s great energy predictor shootout challenge [40] and rank second. The first place was achieved by the use of neural networks which are not interpretable, while DR-Advisor with regression trees is highly interpretable by facilities managers and provides them both guidance and provenance for decisions made to control their infrastructure.

IV. CLOSED-LOOP CONTROL OVER WIRELESS NETWORKS

The current generation of embedded wireless systems has largely focused on open-loop sensing and monitoring applications. To address actuation in closed-loop wireless control systems there is a strong need to re-think the communication architectures and protocols for reliability, coordination and control [41]. Wireless networked control systems, or Networked Cyber-Physical Systems (Networked-CPS), fundamentally differ from standard distributed systems in that the dynamics of the *network* (variable channel capacity, probabilistic connectivity, topological changes, node and link failures) can change the operating points and physical dynamics of the *closed-loop system* [42], [43]. The most important objective of control in Networked-CPS is to provide stability of the closed-loop system. It is therefore necessary for the network (along with its interfaces to sensors and actuators) to be able to provide some form of guarantee of the control system’s stability in the face of the non-idealities of the wireless links and the communication constraints of the wireless swarm network. A secondary goal in Networked-CPS is to allow for composition of additional controllers and plants within the same network without requiring reconfiguration of the entire network operation.

The most common approach to incorporating Networked-CPS into the feedback loop is to use it primarily as a com-

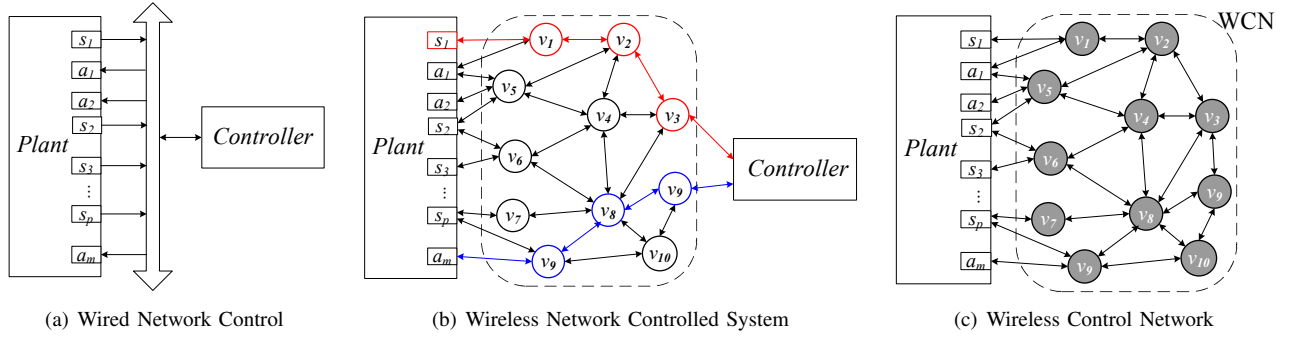


Fig. 7. Standard architectures for Networked Control Systems; (a) Wired system with a shared bus and dedicated controller; (b) Red links/nodes - routing data from the plant's sensors to the controller; Blue links/nodes - routing data from the controller to the plant's actuators; (c) A multi-hop wireless control network used as a distributed controller.

munication medium: the nodes in the network simply route information to and from one or more *dedicated controllers*, which are usually specialized CPUs capable of performing computationally expensive procedures (see Fig. 7(b)). The use of dedicated controllers imposes a routing requirement along one or more fixed paths through the network, which must meet the stability constraints, encapsulated by end-to-end delay requirements [44], [45]. However, this assignment of routes is a static setup, which commonly requires global reorganization for changes in the underlying topology, node population and wireless link capacities.

Routing couples the communication, computation and control problems [46], [47], [48]. Therefore, when a new route is required due to topological changes, the computation and control configurations must also be recalculated. Merely inserting a WNCS into the standard network architecture “sensor \rightarrow channel \rightarrow controller/estimator \rightarrow channel \rightarrow actuator” requires the addition of significant software support [45], [49], as the overhead of completely recomputing the computation and control configurations, due to topological changes or packet drops, is too expensive and does not scale.

While providing a review of classical and recent approaches for control over wireless networks, we present two complementary approaches on maintaining stability in the presence of environment and network disturbances. The first approach adopts a “computer systems” perspective on the design of robust architectures for embedded wireless control and actuation. We call this scheme Embedded Virtual Machines (see Fig. 8) which provides software mechanisms to decouple controller functionality from the physical node - thus providing resilience to node, link and topology changes. The second approach adopts a “control theoretic” perspective on distributed control within the network (see Fig. 7(c)). This provides control mechanisms to remove controller functionality from a dedicated node to all nodes in the network - thus eliminating the need for routing and guaranteeing stability and optimal control in the presence of link, node and topology changes.

A. Embedded Virtual Machines

Current approaches for robust networked control [43] require the underlying network to satisfy a minimal set of requirements (e.g. guaranteed packet deliver rate, upper bound

on network induced delay) and reduce the network model to that of a single channel with random delays. In addition, they do not address the spatial aspects of the network, i.e., how changes in the network topology affect the closed-loop system performance.

As the links, nodes and topology of wireless systems are inherently unreliable, such time-critical and safety-critical applications require programming abstractions where the tasks are assigned to the sensors, actuators and controllers as a *single component*, rather than statically mapping a set of tasks to a specific physical node at design time (as shown in Fig. 8). Such wireless controller grids are composed of many nodes that share a common sense of the control application but without regard to physical node boundaries. Our approach, is to *decouple* the functionality (i.e., tasks) from the inherently unreliable physical substrate (i.e., nodes) and allow tasks to migrate/adapt (Fig. 8(d)) to changes in the topology.

To this end, we introduced the Embedded Virtual Machine (EVM), a powerful and flexible programming abstraction where a Virtual Component (VC) and its properties are maintained across node boundaries [45], [50], as shown in Fig. 2(c). EVMs differ from classical system virtual machines. In the enterprise or on PCs, one (powerful) physical machine may be partitioned to host multiple virtual machines for higher resource utilization. On the other hand, in the embedded domain, an EVM is composed across multiple physical nodes with the goal to maintain correct and high-fidelity operation even under changes in the physical composition of the network. The goal of the EVM is to maintain a set of *functional invariants*, such as a control law and *para-functional invariants* such as timeliness constraints, fault tolerance and safety standards across a set of controllers given the spatio-temporal changes in the physical network. Thus, the EVM introduces new degrees of freedom, task migration and routing which facilitates, at runtime, the network configuration (operating point, conditions) to meet the requirements of the networked control algorithms. However, the EVM does not provide explicit guarantees but only finds the optimal operation configuration in terms of routing and task assignment.

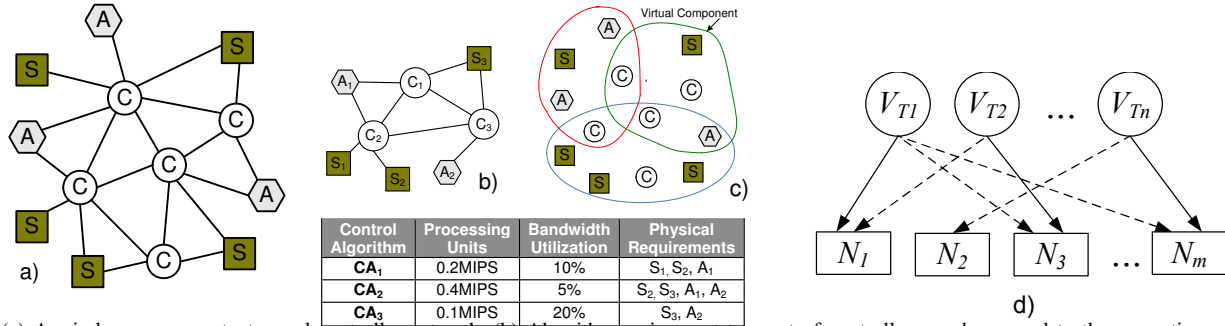


Fig. 8. (a) A wireless sensor, actuator and controller network. (b) Algorithm assignment to a set of controllers, each mapped to the respective nodes. (c) Three Virtual Components, each composed of several network elements. (d) Decoupled virtual tasks and physical nodes with runtime task mapping.

B. Distributed Control over Wireless Networks

Let us consider the problem of stabilizing a plant with a multi-hop network of resource constrained wireless nodes. In [51], we introduced the concept of a *Wireless Control Network* (WCN), which is a paradigm change for distributed control over a wireless network. In a WCN the entire network *itself* acts as a controller, as the computation is spread over the whole network, instead of assigning a particular node with the execution of the control procedure. We have devised a numerical design procedure that produces the coefficients of the linear combinations for each node and actuator to apply in order to stabilize the plant. The radio connectivity between nodes in the network induces topological constraints to the control algorithm, and this topology determines whether it is even possible to stabilize the system with the use of linear iterative strategies. In addition, a method to synthesize an optimal WCN, with respect to the standard cost functions, has been developed [52].

Given the fundamental unreliability of wireless communication, the WCN method handles topological constraints while maintaining mean square stability for packet drop rates **up to 20%** for a specific network topology and plant. This bridges the gap between the basic WCN and the theoretical upper bound of robustness to packet drops [53]. We also show a method to synthesize a WCN robust to a certain level of node failures, and then extend the synthesis procedures to allow for the use of the WCN for control of continuous-time plants.

While in the past efforts, we consider scenarios where the network topology is already set, in recent efforts [54], [55] we have investigated a dual problem, “how to synthesize the network so that a stable WCN configuration exists?” The topological conditions from [54], along with the results from [51] provide the essential building blocks for an integrated decentralized wireless control network design framework. Early experiments in an industrial process control case study of a distillation column in a process-in-the-loop test-bed demonstrate optimal control of continuous-time physical processes which maintain system stability under the presence of node and link failures.

1) *Advantages of the WCN*: The WCN introduces very low communication and computation overhead. The linear iterative runtime procedure is computationally very inexpensive as each node only computes a linear combination of its value and

values of its neighbors. This makes it suitable for resource constrained, low-power wireless nodes (e.g., Tmote). Furthermore, the communication overhead is also very small, as each node needs to transmit only its own state once per frame. In the case when a node maintains a scalar state it transmits only 2 bytes in each message, making it suitable to combine this scheme with periodic message transmissions in existing wireless systems.

The WCN utilizes a simple transmission schedule where each node is active only once during a TDMA cycle and the control-loop does not impose end-to-end delay requirements. This allows the network operator to decouple the computation schedule from the communication schedule, which significantly simplifies closed-loop system design and enables compositional design and analysis.

CONCLUSION

We present three CPS challenges across medical devices, energy systems and control over wireless. While these are three distinct domains, the fundamental problems of closed-loop operation across one or more controllers interfacing with messy, interactive and non-deterministic plants introduces fundamental challenges in ensuring the safety, stability and efficacy of these systems under all possible plant conditions. For each domain, we describe efforts at the intersection of formal methods, control theory, and data-driven systems to tackle the common themes of limited plant observability and high variability across different plants.

REFERENCES

- [1] List of Device Recalls, U.S. Food and Drug Admin., (last visited Jul. 19, 2010).
- [2] W. H. Maisel, M. O. Sweeney, W. G. Stevenson, K.E. Ellison, and L. M. Epstein. Recalls and safety alerts involving pacemakers and implantable cardioverter-defibrillator generators. *J. American Med. Ass.*, 286:793–799, 2001.
- [3] K. Sandler, L. Ohrstrom, L. Moy, and R. McVay. Killed by Code: Software Transparency in Implantable Medical Devices. *Software Freedom Law Center*, 2010.
- [4] AUTOSAR website: www.autosar.org/.
- [5] AVSI website: <http://www.avsi.aero>.
- [6] I. Lee, George J. Pappas, Rance Cleaveland, John Hatcliff, Bruce H. Krogh, Peter Lee, Harvey Rubin, and Lui Sha. High-Confidence Medical Device Software and Systems. *IEEE Computer*, 39(4):139–148, 2006.
- [7] A. O. Gomes and M. V. Oliveira. Formal Specification of a Cardiac Pacing System. In *Proceedings of the 2nd World Congress on Formal Methods*, FM ’09, pages 692–707. Springer-Verlag, 2009.

- [8] E. Jee, I. Lee, and O. Sokolsky. Assurance Cases in Model-Driven Development of the Pacemaker Software. In *Leveraging Applications of Formal Methods, Verification, and Validation*, volume 6416 of *LNCS*, pages 343–356. 2010.
- [9] R. Alur, D. Arney, E. L. Gunter, I. Lee, J. Lee, W. Nam, F. Pearce, S. Van Albert, and J. Zhou. Formal Specifications and Analysis of the Computer-Assisted Resuscitation Algorithm (CARA) Infusion Pump Control System. *Intl. Journal on Software Tools for Technology Transfer (STTT)*, 5:308–319, 2004.
- [10] Annette ten Teije et. al. Improving medical protocols by formal methods. *Artificial Intelligence in Medicine*, 36(3):193 – 209, 2006.
- [11] J. M. Cortner. Testing Implantable Medical Devices. *Global Healthcare Medical Device Manufacturing Technology*, pages 2–4, 2003.
- [12] Medtronic ViP-II Virtual Interactive Patient: User's Manual Software v1.5. Rivertek Medical Systems, 2006.
- [13] John Rushby. Verified software: Theories, tools, experiments. chapter Automated Test Generation and Verified Software, pages 161–172. Springer-Verlag, 2008.
- [14] Personal communication with Paul L. Jones, Senior Systems/Software Engineer, Office of Science and Engineering Laboratories, Center for Devices and Radiological Health, US FDA. August, 2010.
- [15] Zhihao Jiang, Miroslav Pajic, and Rahul Mangharam. Cyber-Physical Modeling of Implantable Cardiac Medical Devices. *Proceeding of IEEE, Special Issue on Cyber-Physical Systems*, 2011.
- [16] Z. Jiang, M. Pajic, S. Moarref, R. Alur, and R. Mangharam. Modeling and Verification of a Dual Chamber Implantable Pacemaker. In *TACAS'12: 18th Conf. on Tools and Algorithms for the Construction and Analysis of Systems*. 2012.
- [17] R. Alur and D. L. Dill. A Theory of Timed Automata. *Theoretical Computer Science*, 126:183–235, 1994.
- [18] G. Behrmann, A. David, and K.G. Larsen. A tutorial on uppaal. In *Formal Methods for the Design of Real-Time Systems (revised lectures)*, volume 3185 of *LNCS*, pages 200–237, 2004.
- [19] K.G. Larsen, Paul Pettersson, and Wang Yi. Uppaal in a Nutshell. *International Journal on Software Tools for Technology Transfer (STTT)*, pages 134–152, 1997.
- [20] Gerd Behrmann, Alexandre David, and Kim G. Larsen. A Tutorial on Uppaal. *Formal Methods for the Design of Real-Time Systems, Lecture Notes in Computer Science*, pages 200–236, 2004.
- [21] Zhihao Jiang, M. Pajic, A. Connolly, S. Dixit, and R. Mangharam. Real-Time Heart Model for Implantable Cardiac Device Validation and Verification. In *22nd Euromicro Conference on Real-Time Systems (ECRTS)*, pages 239 –248, July 2010.
- [22] Z. Jiang and R. Mangharam. Modeling Cardiac Pacemaker Malfunctions with the Virtual Heart Model. *33rd Intl. Conf. IEEE Engineering in Medicine and Biology Society*, 2011.
- [23] Z. Jiang, M. Pajic, and R. Mangharam. Model-based Closed-loop Testing of Implantable Pacemakers. In *ICCPs'11: ACM/IEEE 2nd Intl. Conf. on Cyber-Physical Systems*, 2011.
- [24] Matlab R2011a Documentation → Stateflow. <http://www.mathworks.com/help/toolbox/stateflow>.
- [25] M. Pajic, Z. Jiang, O. Sokolsky, I. Lee, and R. Mangharam. From Verification to Implementation: A Model Translation Tool and a Pacemaker Case Study. In *18th IEEE Real-Time and Embedded Technology and Applications Symposium (IEEE RTAS)*, 2012.
- [26] Virtual Heart Model website - <http://medcps.org>.
- [27] David Arney, Miroslav Pajic, Julian M. Goldman, Insup Lee, Rahul Mangharam, and Oleg Sokolsky. Toward patient safety in closed-loop medical device systems. In *ACM/IEEE International Conference on Cyber-Physical Systems*, pages 33–38, 2010.
- [28] Ronald D Berger et al. The Rhythm ID Going Head to Head Trial (RIGHT): Design of a Randomized Trial Comparing Competitive Rhythm Discrimination Algorithms in Implantable Cardioverter Defibrillators. *Journal of Cardiovascular Electrophysiology*, 17(7):749–753, 2006.
- [29] Michael R. Gold, Saleem Ahmad, Kevin Browne, Kellie Chase Berg, Lisa Thackeray, and Ronald D. Berger. Prospective comparison of discrimination algorithms to prevent inappropriate ICD therapy: Primary results of the Rhythm ID Going Head to Head Trial. *Heart Rhythm*, 9(3):370 – 377, 2012.
- [30] Jerry M Melillo, TC Richmond, and Gary W Yohe. Climate change impacts in the united states: the third national climate assessment. *US Global change research program*, 841, 2014.
- [31] NOAA National Centers for Environmental Information. State of the climate: Global analysis for august 2015. [published online September 2015, retrieved on October 15, 2015 from <http://www.ncdc.noaa.gov/sotc/global/201508>.]
- [32] PJM Interconnection Michael J. Kormos. Pjm response to consumer reports on 2014 winter pricing. 2014.
- [33] Charles Goldman. Coordination of energy efficiency and demand response. *Lawrence Berkeley National Laboratory*, 2010.
- [34] Federal Energy Regulatory Commission et al. Assessment of demand response and advanced metering. 2012.
- [35] PJM Interconnection. 2014 demand response operations markets activity report. 2014.
- [36] Navigant Research. Demand response for commercial & industrial markets market players and dynamics, key technologies, competitive overview, and global market forecasts. 2015.
- [37] Drury B Crawley, Linda K. Lawrie, et al. Energyplus: creating a new-generation building energy simulation program. 33(4):319 – 331, 2001.
- [38] D. Sturzenegger, D. Gyalistras, M. Morari, and R.S. Smith. Model predictive climate control of a swiss office building: Implementation, results, and cost-benefit analysis. *Control Systems Technology, IEEE Transactions on*, 2015.
- [39] Madhur Behl and Rahul Mangharam. Sometimes, money does grow on trees: Data-driven demand response with dr-advisor. In *Proceedings of the 2Nd ACM International Conference on Embedded Systems for Energy-Efficient Built Environments*, BuildSys '15, pages 137–146, New York, NY, USA, 2015. ACM.
- [40] J. F Kreider and J. S Haberl. Predicting hourly building energy use: The great energy predictor shootout-overview and discussion of results. Technical report, Am. Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc., 1994.
- [41] A. Willig, K. Matheus, and A. Wolisz. Wireless technology in industrial networks. *Proceedings of the IEEE*, 2005.
- [42] W. Zhang, M.S. Branicky, and S.M. Phillips. Stability of networked control systems. *IEEE Control Systems Magazine*, 21(1):84–99, 2001.
- [43] J. P. Hespanha, P. Naghshtabrizi, and Y. Xu. A survey of recent results in networked control systems. *Proceedings of the IEEE, Special Issue on Technology of Networked Control Systems*, 95(1):138 – 162, 2007.
- [44] A. Saifullah, Y. Xu, C. Lu, and Y. Chen. Real-Time Scheduling for WirelessHART Networks. In *31st IEEE Real-Time Systems Symposium*, pages 150 –159, 2010.
- [45] M. Pajic and R. Mangharam. Embedded virtual machines for robust wireless control and actuation. In *RTAS'10: 16th IEEE Real-Time and Embedded Technology and Applications Symposium*, pages 79–88, 2010.
- [46] R. Alur, A.D'Innocenzo, K. H. Johansson, G. J. Pappas, and G. Weiss. Compositional modeling and analysis of multi-hop control networks. *IEEE Transactions on Automatic Control*, 56(10):2345–2357, Oct. 2011.
- [47] G. Fiore, V. Ercoli, A.J. Isaksson, K. Landernäs, and M. D. Di Benedetto. Multi-hop Multi-channel Scheduling for Wireless Control in WirelessHART Networks. In *IEEE Conference on Emerging Technology & Factory Automation*, pages 1 – 8, 2009.
- [48] A. D'Innocenzo, G. Weiss, R. Alur, A.J. Isaksson, K.H. Johansson, and G.J. Pappas. Scalable scheduling algorithms for wireless networked control systems. In *CASE'09: IEEE International Conference on Automation Science and Engineering*, pages 409–414, 2009.
- [49] S. Graham, G. Baliga, and P.R. Kumar. Abstractions, architecture, mechanisms, and a middleware for networked control. *IEEE Transactions on Automatic Control*, 54(7):1490–1503, 2009.
- [50] M. Pajic and R. Mangharam. robust architectures for embedded wireless network control and actuation.. In *ACM Transactions of Embedded Computing Systems (TECS)*, 2011.
- [51] M. Pajic, S. Sundaram, G. J. Pappas, and R. Mangharam. The Wireless Control Network: A New Approach for Control over Networks. *IEEE Transactions on Automatic Control*, 56(10):2305–2318, 2011.
- [52] Rahul Mangharam and Miroslav Pajic. Distributed Control for Cyber-Physical Systems. *Journal of the Indian Institute of Science*, 93(3):353–387, 2013.
- [53] C. N. Hadjicostis and R. Touri. Feedback control utilizing packet dropping network links. In *Proceedings of the 41st IEEE Conference on Decision and Control*, pages 1205–1210, 2002.
- [54] M. Pajic, R. Mangharam, G. J. Pappas, and S. Sundaram. Topological Conditions for In-Network Stabilization of Dynamical Systems. *IEEE Journal on Selected Areas in Communications*, 31(4):794–807, 2013.
- [55] M. Pajic, S. Sundaram, G. J. Pappas, and R. Mangharam. Topological Conditions for Wireless Control Networks. In *Proceedings of the 50th IEEE Conference on Decision and Control*, pages 2353–2360, 2011.