10-2015

Parameter-Invariant Design of Medical Alarms

James Weimer

University of Pennsylvania, weimerj@seas.upenn.edu

Radoslav Ivanov

University of Pennsylvania, rivanov@seas.upenn.edu

Alexander Roederer

University of Pennsylvania, roederer@seas.upenn.edu

Sanjian Chen

University of Pennsylvania, sanjian@seas.upenn.edu

Insup Lee

University of Pennsylvania, lee@cis.upenn.edu

Follow this and additional works at: https://repository.upenn.edu/cis_papers

Part of the [Analytical, Diagnostic and Therapeutic Techniques and Equipment Commons](#), [Computer Sciences Commons](#), and the [Other Medicine and Health Sciences Commons](#)

Recommended Citation

James Weimer, Radoslav Ivanov, Alexander Roederer, Sanjian Chen, and Insup Lee, "Parameter-Invariant Design of Medical Alarms", *IEEE Design & Test* 32(5), 9-16. October 2015. <http://dx.doi.org/10.1109/MDAT.2015.2451083>

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/cis_papers/788
For more information, please contact repository@pobox.upenn.edu.

Parameter-Invariant Design of Medical Alarms

Abstract

The recent explosion of low-power low-cost communication, sensing, and actuation technologies has ignited the automation of medical diagnostics and care in the form of medical cyber physical systems (MCPS). MCPS are poised to revolutionize patient care by providing smarter alarm systems, clinical decision support, advanced diagnostics, minimally invasive surgical care, improved patient drug delivery, and safety and performance guarantees. With the MCPS revolution emerges a new era in medical alarm systems, where measurements gathered via multiple devices are fused to provide early detection of critical conditions. The alarms generated by these next generation monitors can be exploited by MCPS to further improve performance, reliability, and safety.

Currently, there exist several approaches to designing medical monitors ranging from simple sensor thresholding techniques to more complex machine learning approaches. While all the current design approaches have different strengths and weaknesses, their performance degrades when underlying models contain unknown parameters and training data is scarce. Under this scenario, an alternative approach that performs well is the parameter-invariant detector, which utilizes sufficient statistics that are invariant to unknown parameters to achieve a constant false alarm rate across different systems. Parameter-invariant detectors have been successfully applied in other cyber physical systems (CPS) applications with structured dynamics and unknown parameters such as networked systems, smart buildings, and smart grids; most recently, the parameter-invariant approach has been recently extended to medical alarms in the form of a critical shunt detector for infants undergoing a lung lobectomy. The clinical success of this case study application of the parameter-invariant approach is paving the way for a range of other medical monitors.

In this tutorial, we present a design methodology for medical parameter-invariant monitors. We begin by providing a motivational review of currently employed medical alarm techniques, followed by the introduction of the parameter-invariant design approach. Finally, we present a case study example to demonstrate the design of a parameter-invariant alarm for critical shunt detection in infants during surgical procedures.

Keywords

medical computing, paediatrics, surgery, critical shunt detection, infants, medical alarm techniques, medical parameter-invariant monitors, parameter-invariant alarm design, parameter-invariant design approach, surgical procedures, biomedical monitoring, data models, detectors, lungs, mathematical model, medical diagnostic imaging, monitoring

Disciplines

Analytical, Diagnostic and Therapeutic Techniques and Equipment | Computer Sciences | Other Medicine and Health Sciences

Parameter Invariant Design of Medical Alarms

James Weimer, Radoslav Ivanov, Alexander Roederer,
Sanjian Chen, and Insup Lee

The recent explosion of low-power low-cost communication, sensing, and actuation technologies has ignited the automation of medical diagnostics and care in the form of medical cyber physical systems (MCPS). MCPS are poised to revolutionize patient care by providing smarter alarm systems, clinical decision support, advanced diagnostics, minimally invasive surgical care, improved patient drug delivery, and safety and performance guarantees. With the MCPS revolution emerges a new era in medical alarm systems, where measurements gathered via multiple devices are fused to provide early detection of critical conditions. The alarms generated by these next generation monitors can be exploited by MCPS to further improve performance, reliability, and safety.

Currently, there exist several approaches to designing medical monitors ranging from simple sensor thresholding techniques to more complex machine learning approaches. While all the current design approaches have different strengths and weaknesses, their performance degrades when underlying models contain unknown parameters and training data is scarce. Under this scenario, an alternative approach that performs well is the parameter-invariant detector, which utilizes sufficient statistics that are invariant to unknown parameters to achieve a constant false alarm rate across different systems [1]. Parameter-invariant detectors have been successfully ap-

plied in other cyber physical systems (CPS) applications with structured dynamics and unknown parameters such as networked systems, smart buildings, and smart grids; most recently, the parameter-invariant approach has been recently extended to medical alarms in the form of a critical shunt detector for infants undergoing a lung lobectomy [2]. The clinical success of this case study application of the parameter-invariant approach is paving the way for a range of other medical monitors.

In this tutorial, we present a design methodology for medical parameter-invariant monitors. We begin by providing a motivational review of currently employed medical alarm techniques, followed by the introduction of the parameter-invariant design approach. Finally, we present a case study example to demonstrate the design of a parameter-invariant alarm for critical shunt detection in infants during surgical procedures.

COMMON MEDICAL ALARM DESIGNS

Design of medical alarms has its roots in classical fault/anomaly detection and identification (see [3], [4], [5], [6] and citations within). While a complete review of the related literature is beyond the scope of this tutorial, this section provides a high-level overview of common medical alarm designs, including sensor thresholding, model-based, data-driven-based, and hybrid approaches, as motivation for the parameter-invariant design.

Sensor Threshold Alarms

The most common detectors currently in widespread use in hospitals are threshold

• All authors are with the PRECISE center in the School of Engineering and Applied Sciences, University of Pennsylvania, Philadelphia, PA, 19104.
E-mail: {weimerj, rivanov, roederer, sanjian}@seas.upenn.edu, lee@cis.upenn.edu

alarms. Threshold alarms monitor a single physiologic signal as measured by a biosensor [7]. When the value of the signal crosses a predefined fixed threshold, the system produces an alarm that can be observed by a clinician. The goal of a threshold alarm is to alert a clinician whenever the patient transitions from a normal, healthy state to a dangerous, unhealthy state based on the physiologic signal being monitored.

Threshold alarms are popular because they are simple; they are easy to implement, and easy for humans to understand. However, extensive research has shown that single sensor threshold alarms are severely limited. Various studies have documented sensor threshold false alarm rates ranging from 57% to 99% of all alarms, causing alarm fatigue in caregivers [8]. Thus threshold alarms can ultimately fail to provide clinicians with a reliable understanding of their patients' state.

Model-Based Alarms

Model-based fault detection is a vast and sophisticated area in control engineering, where varying model fidelities are needed across applications [3]. In medicine, physiological models can be leveraged to provide personalized alarms and treatment advice. Models of the target physiological process can be used to interpret clinical measurements and estimate the patient's state in real-time. The information can then be used to generate predictive alarms or advice treatment actions [9].

Developing a robust model-based alarm system requires a high-fidelity model of the physiological process that is being monitored or controlled. To apply the model-based approach to physiological systems in general, the model must be expressive enough to be able to track all relevant states, and at the same time, it must still be identifiable from measurement data so that it is usable in a practical clinical setting. The trade-off between complexity and usability is one of the fundamental grand challenges in physiological modeling [10].

Data Driven Alarms

The unreliability of underlying physiological processes has led towards the data-driven design of *smart alarms*. The data-driven approach

leverages large amounts of patient data to learn a model directly and employ the model for clinical detection. Data-driven approaches have been successfully applied in many medical applications and have resulted in better monitoring systems [11].

However, there are often practical challenges to data-driven approaches in the medical domain. Machine learning algorithms usually require rich training data with accurate event annotations [6]. In the medical domain, such annotations are often rare or unavailable [8]. Moreover, temporal reasoning over clinical data using machine learning techniques is still an open and vibrant area of research [11], where challenges arise due to the often large number of possible features which must be learned using possibly restrictive data sets.

Hybrid Alarms

In an attempt to leverage the strengths of multiple alarm designs, various hybrid approaches to medical monitoring have been proposed (e.g. [12]). In general, these techniques have shown promise by leveraging the strengths of various approaches to overcome their collective shortcomings. However, these hybrid approaches require some combination of accurate models and rich training data [6], [10]. As an alternative to the designs presented in this section, the following section describes the parameter-invariant alarm design approach, which has been shown to perform well in situations where model parameters are unknown and training data is scarce [2].

PARAMETER-INVARIANT ALARM DESIGN

In applications where patient or condition variability precludes accurate models and rich training data, the parameter-invariant design provides an alternative approach. A parameter-invariant design represents a guided approach to the removal of unknown or corrupted signals from the measurement. To design a parameter-invariant test, we assume there exists a real-world process $\mathcal{P}$ which can be modeled as $\mathcal{M}_{\mathcal{P}}$ and there exists a test, $\mathcal{T}_{\mathcal{H}}$, for the hypothesis (question) $\mathcal{H}$. In designing a parameter-invariant alarms, three design requirements must be satisfied:

- $\mathcal{T}_{\mathcal{H}}$ is invariant to $\mathcal{M}_{\mathcal{P}}$
- $\mathcal{M}_{\mathcal{P}}$ is accurate for $\mathcal{P}$
- $\mathcal{T}_{\mathcal{H}}$ is accurate for $\mathcal{H}$

The remainder of this section discusses the above requirements, while the subsequent section provides a step-by-step tutorial example for designing a parameter-invariant alarm.

Foundations of Parameter-Invariant Design

At the core of the parameter-invariant alarm design is the requirement that the test, $\mathcal{T}_{\mathcal{H}}$, is invariant to the unknown model parameters, $\mathcal{M}_{\mathcal{P}}$. In this subsection, we first introduce the high-level mathematical foundations for parameter-invariant design, then provide examples of three parameter-induced transformations and their respective invariant tests common in real-world medical monitoring applications namely: bias, scale, and rotation.

To begin, let us assume that the unknown model parameters induce a *group of transformations*, $G_{\mathcal{M}}$, on the measurements, $\mathbf{y}$. A test statistic $t(\mathbf{y})$ is invariant to the group of transformations, $G_{\mathcal{M}}$, if and only if

$$\forall g \in G_{\mathcal{M}} \quad t(\mathbf{y}) = t(g(\mathbf{y}))$$

where, in words, we say the statistic is invariant to the group of transformation if the statistic has the same value regardless of the unknown parameters. There are many invariant statistics (e.g. $t(\mathbf{y}) = 0$ is trivially invariant to all parameters, but is useless as a test statistic.), thus we wish (if possible) to choose an invariant statistic that is *maximally invariant*, namely the statistic which only removes the effect of the unknown parameters. This concept is captured mathematically by the following implication: for any two measurement vectors $\mathbf{y}'$ and $\mathbf{y}''$,

$$t(\mathbf{y}') = t(\mathbf{y}'') \longrightarrow \exists g \in G_{\mathcal{M}}, \mathbf{y}' = g(\mathbf{y}'').$$

Except in rare cases, a maximally invariant statistic cannot be realized; however, designing the statistic to be *near* maximally invariant improves the test performance. To present different groups of transformations and their corresponding maximally invariant statistics, the remainder of this section utilizes Figure 1 illustrating transformations induced by bias, scale, and rotation.

Bias occurs when the outputs are (partially) generated through a known process with unknown parameters (e.g., the glucose response of an individual is governed by the individual's insulin sensitivity, which varies across the population). A subspace bias is illustrated in Figure 1a, where outputs $\mathbf{y}$ are summed with a vector $\mathbf{H}\boldsymbol{\theta}$ of unknown magnitude and direction (in a known subspace, $\langle \mathbf{H}^{\parallel} \rangle$) such that $\mathbf{y} + \mathbf{H}\boldsymbol{\theta}$ results. Thus, we say that $\boldsymbol{\theta}$ induces a group of transformations $G_{\boldsymbol{\theta}} = \{g \mid g(\mathbf{y}) = \mathbf{y} + \mathbf{H}\boldsymbol{\theta}\}$. From the Figure, we observe that the vector statistic $t(\mathbf{y}) = \mathbf{H}^{\perp} \mathbf{y}$ is invariant to any bias in the subspace $\langle \mathbf{H}^{\parallel} \rangle$, as $\mathbf{H}^{\perp} \mathbf{H} = 0$, and also maximal since only the subspace $\langle \mathbf{H}^{\parallel} \rangle$ is removed from the measurements.¹

Scaling occurs when the outputs are generated by a process which multiplies them by an unknown magnitude. All robust alarms are normalized by the noise variance. In situations where the noise variance is unknown (e.g. the variance of metabolism in a diabetic population), the normalization process induces an unknown scaling of the measurements. The effect of scaling is captured in Figure 1b, where an output $\mathbf{y}$ is scaled according by σ such that $\sigma \mathbf{y}$ results, (i.e. $G_{\sigma} = \{g \mid g(\mathbf{y}) = \sigma \mathbf{y}\}$). From the Figure, we observe that the direction of the measurement, $t(\mathbf{y}) = \mathbf{y} / \|\mathbf{y}\|$, is maximally invariant to scaling.

A *rotation* occurs when a translation is performed on a signal in a subspace without affecting the magnitude of the signal. In practical applications, rotations manifest as the eigenvectors of the measurement covariance matrix, in a subspace. The effect of a rotation is illustrated in Figure 1c, where an output $\mathbf{y}$ is rotated in the subspace $\langle \mathbf{H}^{\parallel} \rangle$ by applying a unitary operation, $\mathbf{Q} \in \mathcal{U}^2$, in the subspace of $\langle \mathbf{H}^{\parallel} \rangle$, denoted as $\mathbf{U}\mathbf{Q}\mathbf{U}^{\top}$ when $\mathbf{U}\mathbf{U}^{\top} = \mathbf{H}^{\parallel}$. The corresponding group of transformations induced by a rotation in $\mathbf{H}^{\parallel}$ is written as $G_{\mathbf{Q}}(\mathbf{y}) = \{g \mid g(\mathbf{y}) = (\mathbf{H}^{\perp} + \mathbf{U}\mathbf{Q}\mathbf{U}^{\top}) \mathbf{y}\}$. From the Figure, we observe that the angle between the cone of rotation and the subspace $\langle \mathbf{H}^{\perp} \rangle$, namely $t(\mathbf{y}) = \cos^{-1} \left(\left\| \frac{\mathbf{H}^{\perp} \mathbf{y}}{\|\mathbf{y}\|} \right\| \right)$ is invariant to an arbitrary rotation in $\langle \mathbf{H}^{\parallel} \rangle$.

1. See [1] for proofs of invariance and maximality for all transformations considered herein.

2. $\mathcal{U}$ denotes the set of unitary matrices.

Fig. 1: Types of transformations included in the transformation group $\mathcal{G}$.

While bias, scaling, and rotation represent three types of translations for which individual tests can be designed invariant, combinations of these transformations can also be considered. A detailed description of the compositional translation and the resulting parameter-invariant statistics is provided in [1]. Moreover, the set of transformations presented herein does not represent a complete set of transformations; however we have found bias, scaling, and rotation sufficient to design powerful detectors in many real-world applications.

Physiological Process Modeling

To develop a model, $\mathcal{M}_{\mathcal{P}}$, which is accurate for a physiological process, $\mathcal{P}$, we consider classical compartment models of physiological phenomena [10]. A compartmental model is (typically) a set of differential or difference equations in which the state variables represent the quantities of the target substances within each compartment and the equations represent the interconnections among the compartments. The rates at which the substances enter or leave a compartment can be captured by a set of model parameters. Since most compartmental models are derived from first-principle physics, those parameters usually have physiological meanings: for example, a parameter may represent how fast a chemical diffuses from one body compartment (e.g., an organ or a type of tissue) to another.

In general, compartment models can be designed to arbitrary accuracy through a natural grouping (or ungrouping) of physiological effects. For instance, in the glucose-insulin system one can model the cumulative effect of carbohydrate ingestion as a single first-order differential equation representing collectively the stomach, intestine, and blood compartments as the *glucose system*, or one can model each of the three physical compartments individually as interactive systems. When designing a parameter-invariant alarm, the choice in model fidelity affects the performance of the resulting detector. As a rule of thumb, the designer should utilize a model which captures the *general trend* of the system while simultaneously ensuring that the unknown parameters of the resulting model yield an invariant statistic. This represents a design tradeoff which varies with the monitoring application.

Regardless of the model fidelity, the parameter-invariant design requires two parameterized models: null model ($\mathcal{H}_0$) and event model ($\mathcal{H}_1$). The null model describes the normal physiology, while the event model describes the physiology in the presence of the event or condition. After discussing the design of a hypothesis test in the following subsection, we will demonstrate the modeling approach in a case study example.

Hypothesis Testing

Lastly, we consider that a test, $\mathcal{T}_{\mathcal{H}}$, must be accurate for the hypothesis problem, $\mathcal{H}$. While the parameter-invariant approach utilizes two models (null and event), it is unlikely that these two models always accurately explain all possible scenarios. Thus, we utilize two statistics, $t_0(\mathbf{y})$ and $t_1(\mathbf{y})$, under to test the different hypotheses independently. The first statistic, $t_0(\mathbf{y})$, assumes that the null hypothesis, $\mathcal{H}_0$, is true, and aims to test the event hypothesis, $\mathcal{H}_1$. The second statistic, $t_1(\mathbf{y})$ assumes that the event hypothesis is true and aims to test the null hypothesis. To ensure a minimum level of performance, the parameter-invariant design bounds the following probabilities:

$$P[t_0(\mathbf{y}) \notin \phi_0] \leq \alpha \quad \text{and} \quad P[t_1(\mathbf{y}) \notin \phi_1] \leq \beta$$

where α and β are the probability of false alarm (type I error) and probability of miss (type II error), respectively, and ϕ_0 and ϕ_1 define when each tests accepts the assumed true hypothesis (event or null). Employing these acceptance regions, the parameter-invariant test decides to sound an alarm based on the following table.

TABLE 1: Test Decision Space for Alarm System

	$t_0(\mathbf{y}) \notin \phi_0$	$t_0(\mathbf{y}) \in \phi_0$
$t_1(\mathbf{y}) \notin \phi_1$	Warning (inaccurate model)	No alarm
$t_1(\mathbf{y}) \in \phi_1$	Alarm	Warning (indecision)

In Table 1, the parameter invariant test makes a definitive decision (i.e., alarm or no alarm) when both statistic tests agree. When the tests do not agree, the parameter-invariant alarm generates one of two types of warnings. An indecision warning occurs when neither test rejects its assumption, indicating that there is not enough power in the test to disambiguate between the two hypotheses to the level of accuracy specified. A model inaccuracy warning occurs when both tests reject their assumptions, indicating neither model accurately describes the measurements. The benefit of this two-sided testing approach is that the event (or null) hypothesis won't be accepted/rejected just because there wasn't enough power or because the models were inaccurate.

EXAMPLE: CRITICAL SHUNT ALARM

As a tutorial for the parameter-invariant alarm design, this section develops a parameter-invariant alarm to provide early detection of critical-pulmonary shunts in infants during lung lobectomy surgeries, as described in detail in [2].

Hypothesis Formulation

In intubated patients, a pulmonary shunt can result in single lung ventilation through the insertion of the endotracheal tube into the main stem bronchus of the healthy lung. A critical shunt occurs when the single lung ventilation is insufficient to maintain the blood oxygen content. Currently, the level of oxygen in the blood is monitored at the peripheral capillaries (e.g. finger tip) through a pulse oximeter, which does provide an accurate measure of the arterial blood oxygen content, but is time-delayed and once the measurement decreases the patient is already in a critical state. The aim of this monitor is to use respiratory measurements, sampled at 15 second intervals, corresponding to respiratory rate, end-tidal CO_2 (i.e. amount of CO_2 exhaled), and tidal volume (i.e. inhaled air volume) to provide accurate early detection of the critical pulmonary shunts. Thus the null hypothesis, in words, is *no shunt has occurred* and the event hypothesis is *a shunt has occurred*.

Physiological Model

Recalling that the physiological model must be accurate for the process while simultaneously allowing for a test statistic to be designed invariant to any unknown parameters, we aim to develop a model which captures the general trends in the blood oxygen content under the shunt and no shunt scenarios. The process which relates the hypotheses (i.e. the absence or presence of a shunt) to the respiratory measurements is governed by the circulatory and respiratory dynamics. A simplified schematic model of the gas partial pressures in the circulatory and respiratory systems is illustrated in Figure 2³, where the compartments of interest are the airways (denoted by $P_e\text{O}_2$, $P_e\text{CO}_2$ and

3. For a complete model explanation, including supporting physiological evidences, see our previous work [2].

P_iO_2 , P_iCO_2 for expiration and inspiration, respectively), the alveoli (P_AO_2 , P_ACO_2), arteries (P_aO_2 , P_aCO_2), pulmonary capillaries (P_pO_2 , P_pCO_2), veins (P_vO_2 , P_vCO_2) and pulmonary veins ($P_v^RO_2$, $P_v^RCO_2$ and $P_v^LO_2$, $P_v^LCO_2$ for the right and left side, respectively), also referred to as “lung-specific partial” pressures. For completeness, the end-tidal airway CO_2 partial pressure (i.e. P_eCO_2) is measured via an anesthesia machine along with the tidal volume and respiratory rate.

Under no shunt, alveolar diffusion governs the transfer of gases between the blood and air in the lungs. We recall from [2] that due to the diffusive capacity of CO_2 and O_2 , the change in blood CO_2 content is inversely proportional to

In the presence of a shunt, only one lung participates in diffusion as indicated by Figure 3b, which tends to increase CO_2 blood content (and decrease the O_2 blood content). Lastly, we observe that the alveolar CO_2 partial pressure is approximately equal to the expiratory CO_2 partial pressure such that an increase in expiratory CO_2 indicates an increase blood CO_2 content (corresponding to a decrease in blood O_2 content).

Parameter-Invariant Test

$$t_0(\mathbf{y}) = \frac{\mathbf{y}^\top \mathbf{P}_{01} \mathbf{y}}{\mathbf{y}^\top (\mathbf{P}_0 - \mathbf{P}_{01}) \mathbf{y}} \quad (1)$$

Fig. 3: Model of the respiratory and cardiovascular partial pressures with and without a shunt.

and

$$t_1(\mathbf{y}) = \frac{\mathbf{y}^\top \mathbf{P}_{10} \mathbf{y}}{\mathbf{y}^\top (\mathbf{P}_1 - \mathbf{P}_{10}) \mathbf{y}} \quad (2)$$

where,

$$\mathbf{P}_i = \mathbf{I} - \mathbf{F}_i (\mathbf{F}_i^\top \mathbf{F}_i)^{-1} \mathbf{F}_i^\top$$

$$\mathbf{P}_{ij} = \mathbf{P}_i \mathbf{F}_j (\mathbf{F}_j^\top \mathbf{P}_i \mathbf{F}_j)^{-1} \mathbf{F}_j^\top \mathbf{P}_i.$$

For $t_0(\mathbf{y})$, invariance is established with respect to G_0 by first projecting the measurements onto the null space of $\mathbf{F}_0$ (eliminating the bias induced by $\boldsymbol{\theta}$) using $\mathbf{P}_0$. Then, a ratio is constructed where the numerator is the energy of the remaining measurements explained by signal under the event hypothesis, $\mathbf{P}_0 \mathbf{F}_1$, and the denominator is the energy not explained by the signal, thus eliminating the scaling induced by σ .

DISCUSSION AND ONGOING WORK

The parameter invariant design approach overview in this paper addresses a design problem commonly encountered in medical applications. Based on an evaluation (detailed in [2]), the critical shunt alarm evaluated on real-patient data available from the Children's Hospital of Philadelphia (CHOP) from lobectomy surgeries over the last decade, for 26 patients with shunt-induced hypoxia, the detector gave early predictions for about 80% of the cases while maintaining a false alarm rate of about 2 false alarms per hour as tested on 172

open surgeries that were known to not have a shunt. Ongoing work in designing parameter invariant medical alarms include the developing alarms for hypovolemia and hypoglycemia conditions. The design of detectors for these conditions will give insight into the robustness of the parameter-invariant design.

ACKNOWLEDGMENTS

This research was supported in part by NSF CNS-1035715.

REFERENCES

- [1] L. L. Scharf, *Statistical signal processing*. Addison-Wesley Reading, MA, 1991, vol. 98.
- [2] R. Ivanov, J. Weimer, A. Simpaio, M. Rehman, and I. Lee, "Early detection of critical pulmonary shunts in infants," in *Proceedings of the 6th International Conference on Cyber-Physical Systems*, 2015, pp. 110–119.
- [3] A. S. Willsky, "A survey of design methods for failure detection in dynamic systems," *Automatica*, vol. 12, no. 6, pp. 601–611, 1976.
- [4] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," *ACM Computing Surveys (CSUR)*, vol. 41, no. 3, p. 15, 2009.
- [5] H. L. Van Trees, *Detection, estimation, and modulation theory*. John Wiley & Sons, 2004.
- [6] C. M. Bishop et al., *Pattern recognition and machine learning*. Springer New York, 2006, vol. 1.
- [7] B. R. Eggins, *Biosensors: an introduction*. Wiley Chichester, UK, 1996.
- [8] J. Edworthy and E. Hellier, "Fewer but better auditory alarms will improve patient safety," *Quality and Safety in Health Care*, vol. 14, no. 3, pp. 212–215, 2005.
- [9] Z. Jiang, M. Pajic, and R. Mangharam, "Cyber-physical modeling of implantable cardiac medical devices," in *Proceedings of the IEEE*, 2012, pp. 122–137.

- [10] C. Cobelli and E. Carson, *Introduction to modeling in physiology and medicine*. Academic Press, 2008.
- [11] M. Stacey and C. McGregor, "Temporal abstraction in intelligent clinical data analysis: A survey," *Artificial intelligence in medicine*, vol. 39, no. 1, pp. 1–24, 2007.
- [12] M. Ghorbani and P. Bogdan, "A cyber-physical system approach to artificial pancreas design," in *Proceedings of the ninth IEEE/ACM/IFIP international conference on hardware/software codesign and system synthesis*. IEEE Press, 2013, p. 17.

James Weimer received a B.S. degree in Electrical Engineering from Purdue University, West Lafayette in 2005 and M.S. and Ph.D. degrees in Electrical and Computer Engineering from Carnegie Mellon University in 2007 and 2010, respectively. He is currently a Postdoctoral Researcher in the Department of Information and Computer Science at the University of Pennsylvania

and was previously a Postdoctoral Researcher in the Department of Automatic Control at the Royal Institute of Technology KTH, Stockholm. His research interests focus on the design and analysis of cyber-physical systems with application to medical devices/monitors, networked systems, building energy management, and smart grids.

Sanjian Chen received his B.E. degree from Tsinghua University, Beijing and M.S.E. degree in Computer and Information Science from the University of Pennsylvania. He is currently a Ph.D. student in the Department of Computer and Information Science at the University of Pennsylvania. His research interests include model-based analysis of cyber-physical systems with application to medical and automotive systems. He received IEEE RTSS 2012 Best Paper Award.

Insup Lee is the Cecilia Fidler Moore Professor of Computer and Information Science and serves as the Director of the PRECISE Center at the University of Pennsylvania. He holds a secondary appointment in the Department of Electrical and Systems Engineering. His research interests include cyber-physical systems, real-time and embedded systems, run-

time assurance and verification, trust management, and high-confidence medical systems. He received a PhD in Computer Science from the University of Wisconsin, Madison. He is IEEE fellow. He received IEEE TC-RTS Outstanding Technical Achievement and Leadership Award in 2008.

Radoslav Ivanov received his B.A. in Computer Science and Economics from Colgate University, New York and is currently pursuing a doctoral degree in the Computer and Information Science Department at the University of Pennsylvania. His main research focus is on the security and control of cyber-physical systems (CPS), in particular automotive and medical CPS.

He is also interested in predictive and retrospective analysis of medical patient data.

Alexander Roederer is a Ph.D. student in Computer and Information Science at the University of Pennsylvania. He received his B.S. in Computer Science and Mathematics from the University of Miami and his M.S.E. in Computer and Information Science from the University of Pennsylvania. His research interests include the application of machine learning to high-frequency,

multi-source physiologic data to develop clinical decision support systems.