Research Statement

James Weimer
weimerj@seas.upenn.edu

My research develops foundations for the control of cyber-physical systems (CPS) – primarily targeting ***healthcare*** applications. CPS represent a new era of safety-critical embedded systems that feature tight coupling between communication and computation used to control complex, uncertain, and potentially adverse physical plants. In critical infrastructures (e.g. healthcare, energy, transportation, manufacturing) existing robust engineering paradigms do not address the combination of challenges imposed by closed-loop CPS: strict safety and performance requirements, constrained sensing and actuation capabilities, unreliable communication and computation platforms, and potentially malicious feedback information. In ***medical CPS***, these challenges also include unidentifiable non-linear time-varying physiological processes that result in significant model uncertainty – an unfavorable scenario for closed-loop control.

My work utilizes *parameter-invariant* (PAIN) designs [1--8] to address model uncertainty. For some objective (e.g. monitoring whether an event happened), rather than employing estimated models, PAIN designs remove the effect of unknown nuisance modeling variables from its solution. Theoretically, PAIN designs have constant true/false alarm rates regardless of model uncertainty. Practically, PAIN designs provide near-constant true/false alarm rates in medical and non-medical applications –a PAIN medical monitor implemented in an area hospital reliably alerts clinicians to a hard-to-diagnose life-threatening surgical condition, without requiring patient tuning.

My research focuses on control systems engineering for CPS, encompassing ***control system design, analysis, and implementation***. The long-term goal enables ***high-assurance control*** of CPS through new algorithm *designs*, which consider *implementation* constraints, such that the resulting real-world system facilitates *analysis*. Achieving this goal – especially in healthcare – requires control system designs that conform to the platform resources while also providing predictable performance in spite of complex cyber-physical interactions. Quantifying the trade-offs between performance, safety, reliability, and security in the control designs requires new analysis tools and techniques. Validating control system design assumptions requires real-world implementations that provide evidentiary support and enable identifying unforeseen challenges and future research.

My work bridges control theory and its practical application. The importance of both theory and practice in my research agenda stems from my early work on wireless sensor networks [9, 10] and energy-efficient buildings [3]. In these applications, I developed theoretical solutions to resource-constrained problems and implemented the solutions in real-world testbeds. These implementations revealed, to me, that generating accurate models for CPS is difficult at best and economically infeasible at worst. This observation initiated my development of PAIN designs, which in turn motivated me to consider application domains with significant uncertainty to highlight utility. Hence, my ongoing and future research concerns ***engineering human health*** and ***cyber-physical security*** –wherein both theory and practice play important roles.

Engineering Human Health - *Developing High-Assurance Medical CPS*

Medicine's evolution into an information- and technology-driven discipline stands to revolutionize human health. My vision of the next generation healthcare systems include advanced wearable medical diagnostics, nearly full automation of inpatient care, and accurate prediction of patient trends and outcomes. Towards this vision, my work aims to develop techniques for high-assurance medical CPS – targeting physiological monitoring and control systems. Specifically, this research contains two thrusts: monitoring life-critical physiological events, and the design and analysis of high-assurance medical control systems.

Research Thrust I: *Monitoring Life-Critical Physiological Events.* Clinical decision and support systems can alert overloaded clinicians to life-critical situations; however, these systems face fundamental challenges stemming from complex cyber-human interfaces and patient variability. Threshold-based alarming on vital signs have false alarm rates between 57% and 99% [2]. Restricted sensing and actuation capabilities coupled with complicated physiological dynamics, makes standard model-based monitoring impractical in many scenarios. Additionally, variability between patient physiologies and co-morbidities can prevent reliable data-driven design. Thus, this work targets monitor development for life-critical events that, by design, provides low false alarm rates across all patients.

Towards this goal, I have developed PAIN algorithms for surgical, critical, and outpatient care to monitor life-threatening events, namely to *predict critical pulmonary shunts in infants* [5], *detect early-stage hypovolemia* [6], and *detect meals in type I diabetics* [4]. In all applications, real-patient data evaluations of the PAIN monitors outperformed existing monitors while exhibiting near-constant and low false positive rates across all patients – without individualization. Notably, clinicians at the Children's Hospital of Philadelphia (CHOP) currently use an implementation of the shunt prediction monitor – a best paper finalist at ICCPS'15 – during infant lung lobectomy surgeries. These early results motivate my future research into physiological monitor design, analysis, and implementation.

As future work, I plan to integrate data-driven (machine learning) techniques into the PAIN design, while retaining its theoretical properties. Incorporating data-driven approaches will provide insight towards efficient extraction of critical physiological information from electronic health records using modified PAIN designs. Evaluating these designs requires support from analysis tools. As a first step towards analysis, I developed a *physiological benchmark* for surgical glucose control to evaluate potential tools [12].

Leveraging the design and analysis techniques, my vision for future implementations includes building/improving monitors for hard-to-predict critical conditions such as hypovolemia, sepsis, and opiate overdose. Additionally, I plan to recall my early work on wireless sensor networks [9, 10] to extend the physiological monitoring technology to wearable platforms. Outcomes of this work – beyond deployed monitors for surgical, critical, and outpatient care – will be theoretical foundations for general CPS monitoring and wearable examples for demonstration and outreach.

Research Thrust II: ***Design and Analysis of High-Assurance (Medical) Control Systems.*** Closed-loop patient care exists in very constrained contexts, such as cardiac implants, and relies on extensive tuning by a clinician before deployment. Increasing the presence of closed-loop medical technologies lightens the clinician's workload – enabling improved doctor-patient interaction and high-level patient monitoring. However, designing and analyzing closed-loop systems including complex physiological dynamics presents a significant challenge to medical control. In this thrust, I build on the my early physiological monitoring work, discussed in the previous research thrust, to perform closed-loop physiological control.

Due to the convergent nature of physiological systems, many patient care scenarios involve first identifying a life-critical situation, then taking a prescribed clinical action. Thus, my vision for medical cyber-physical control systems adopts a hierarchal control architecture where a *context-aware* controller, likely incorporating PAIN technologies, supervises optimized control strategies (i.e. data-driven, adaptive, etc.), in real-time. However, system design and analysis challenges arise due to unknown physiological dynamics. Addressing these challenges requires algorithm designs that can provide performance guarantees despite physiological uncertainty, such as PAIN designs. Based on the control solution's theoretical capabilities, other CPS domains will be considered (energy-efficient buildings and robotics) to broaden the impact of this work and reveal new research directions.

Practical medical applications will be investigated to evaluate the efficacy of the theoretical results. Likely candidates include diabetic glucose control (outpatient) and ventilator weaning (intensive care) based on the imposed control challenge and accessibility to data. A high-fidelity FDA-accepted model and simulator exists for the diabetic glucose control, providing a unique bridge between theory and practice to study medical closed-loop control. Ventilator weaning – the process of taking a patient off mechanical ventilation – has clinical guidelines and primarily interacts with the cardio-pulmonary system for which I have previous experience. To integrate multiple devices into a single testbed, I plan to develop an open-access *Integrated Clinical Environment* (ICE) compliant testbed for networked medical systems, leveraging my CHOP implementation experiences. This work is partly contained in my *NSF CPS Frontier* proposal – *Rigorous Design and Development of Closed-Loop Medical CPS* – for which I am a Co-PI.

Cyber-Physical Security - *Defenses for Cyber-Physical Exploits*

New control system security vulnerabilities arise when malicious attackers exploit the physical environment – not protected by cybersecurity defenses – to execute an attack (e.g. GPS spoofing). Although small disturbance attenuation is a centerpiece of all robust control systems, stability and safety claims become invalid when model and measurement deviations violate the design assumptions. I envision future secure control solutions that defend the cyber-physical surface through co-design and systems integration allowing cyber defenses to leverage physical design features to yield unprecedented resilience. Towards this future, my research aims to design secure cyber-physical control systems with minimal environmental (physical) assumptions that tests the boundary of cyber-physical security.

Research Thrust III: ***Context-Aware Attack-Resilient Control.*** While most security-related research focuses on cybersecurity, recent events demonstrate that attacks on the control environment (e.g. sensors, actuators, and communication media) can be equally disruptive. Attack-resilient control strategies have emerged as a viable solutions in restricted contexts. This work aims to enable wide-spread application of attack-resilient strategies in broad contexts.

Towards this goal, we have designed multiple attack-resilient state estimation approaches, under different attacker capability assumptions, and implemented them on robotic platforms and in an American-built automobile [11, 14]. For analysis, we proved resilience for an attack-resilient sate estimator with bounded modeling errors [13] – winning the best paper award at ICCPS'14 – and developed an attack-resilient assurance-case for its robotic implementation [15]. These results motivate my future cyber physical system security research.

Most approaches to secure cyber-physical control utilize a nominal model of the physical world for the purposes of comparing predicted and actual measurements. The reliance on a nominal model results in restrictive operating conditions requiring supplemental sensing systems to validate model assumptions. This can introduce additional attack surfaces and increases the system complexity – prime conditions for malicious exploits. To address this issue, my future work aims to design attack-resilient control strategies that expand the operational envelope by requiring minimal model assumptions. My vision includes an attack monitor that can detect malicious sensors over a broad range of operating conditions. When an attack is detected, a *safe mode* controller takes over with restricted operating capabilities. Through this project, I anticipate extensions to distributed networked systems where only a local model is known. These results will be evaluated on ground and air robotic platforms initially, and applied to healthcare systems when mature.

Funding Potential:

My proposed research agenda on engineering human health and security has several funding sources. NSF funding potential presents through the CAREER award as well as CPS and SCH programs. Security and soldier health funding has been available through DARPA and ONR. Additionally, utilizing my clinical relationships, my work can be partially funded through NIH, particularly BISTI.

References

[1] J. Weimer, A. Roederer, R. Ivanov, S. Chen, and I. Lee. Learning robust classifiers for physiological systems. *50 pages*, 2015 *(preprint available by request)*.

[2] J. Weimer, R. Ivanov, A. Roederer, S. Chen, and I. Lee. Parameter-invariant design of medical alarms. *Design Test, IEEE*, 32(5):9--16, Oct 2015.

[3] J. Weimer, J. Araujo, M. Amoozadeh, S. Ahmadi, H. Sandberg, and K. Johansson. Parameter-invariant actuator fault diagnostics in cyber-physical systems with application to building automation. In *Control of Cyber-Physical Systems*, volume 449 of *Lecture Notes in Control and Information Sciences*, pages 179--196. Springer International Publishing, 2013.

[4] S. Chen, J. Weimer, M. Rickels, A. Peleckis, and I. Lee. Physiology-invariant meal detection for type i diabetes. *Diabetes Technology and Therapeutics*, 2015 *(accepted)*.

[5] R. Ivanov, J. Weimer, A. Simpao, M. Rehman, and I. Lee. Prediction of critical pulmonary shunts in infants. *Transaction on Control Systems Technology, IEEE*, 2015 *(accepted)*.

[6] A. Roederer, J. Weimer, J. Dimartino, J. Gutsche, and I. Lee. Robust monitoring of hypovolemia in intensive care patients using photoplethysmogram signals. In *37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, EMBC '15, 2015.

[7] J. Weimer, D. Varagnolo, M. Stankovic, and K. Johansson. Parameter-invariant detection of unknown inputs in networked systems. In *IEEE Conference on Decision and Control (CDC)*, pages 4379--4384, Dec 2013.

[8] J. Weimer, N. Bezzo, M. Pajic, G. Pappas, O. Sokolsky, and I. Lee. Resilient parameter-invariant control with application to vehicle cruise control. In *Control of Cyber-Physical Systems*, volume 449 of *Lecture Notes in Control and Information Sciences*, pages 197--216. Springer International Publishing, 2013.

[9] J. Weimer, J. Araujo, and K. Johansson. Distributed event-triggered estimation in networked systems. In *Analysis and Design of Hybrid Systems*, pages 178--185, 2012.

[10] J. Weimer, B. Krogh, M. Small, and B. Sinopoli. An approach to leak detection using wireless sensor networks at carbon sequestration sites. *International Journal of Greenhouse Gas Control*, 9:243 -- 253, 2012.

[11] N. Bezzo, J. Weimer, M. Pajic, O. Sokolsky, G. Pappas, and Insup L. Attack resilient state estimation for autonomous robotic systems. In *Intelligent Robots and Systems (IROS 2014)*, pages 3692--3698, Sept 2014.

[12] S. Chen, M. O'Kelly, J. Weimer, O. Sokolsky, and I. Lee. An intraoperative glucose control benchmark for formal verification. In *Analysis and Design of Hybrid Systems*, ADHS'15, 2015.

[13] M. Pajic, J. Weimer, N. Bezzo, P. Tabuada, O. Sokolsky, Insup Lee, and G. Pappas. Robustness of attack-resilient state estimators. In *ACM/IEEE International Conference on Cyber-Physical Systems (ICCPS), 2014* .

[14] J. Weimer, N. Bezzo, M. Pajic, O. Sokolsky, and Insup Lee. Attack-resilient minimum mean-squared error estimation. In *American Control Conference (ACC), 2014*, pages 1114--1119, June 2014.

[15] J. Weimer, O. Sokolsky, N. Bezzo, and I. Lee. Towards assurance cases for resilient control systems. In *Cyber-Physical Systems, Networks, and Applications (CPSNA), 2014 IEEE International Conference on*, pages 1--6, Aug 2014.