On Threat Modeling and Mitigation of Medical Cyber-Physical Systems

Hussain Almohri[*‡], Long Cheng[†], Danfeng (Daphne) Yao[†], Homa Alemzadeh[‡]
[*]Department of Computer Science, Kuwait University, Kuwait
[†]Department of Computer Science, Virginia Tech, USA
[‡]Department of Electrical and Computer Engineering, University of Virginia, USA

***Abstract*—Medical Cyber Physical Systems (MCPS) are life-critical networked systems of medical devices. These systems are increasingly used in hospitals to provide high-quality healthcare for patients. However, MCPS also bring concerns about security and safety and new challenges to protect patients from acts of theft or malice. In this paper, we focus our investigation on a thorough understanding of threat modeling in MCPS. We examine the roles of stakeholders and system components and sketch an abstract architecture of a MCPS to demonstrate various threat modeling options. We also discuss possible security techniques and their applicability and utility for the design of secure MCPS. This work forms a basis for understanding threatening conditions in MCPS, and embarks on promising state-of-the-art research trends for addressing MCPS security concerns.**

***Index Terms*—Medical information systems, Medical control systems, Safety management, Security management**

I. Introduction

Medical cyber-physical systems (MCPS) aim to improve patient treatment effectiveness, provide intelligent information to the caregiver, and ensure patient safety [1], [2]. These systems enable unprecedented usage scenarios considering a broad spectrum of the patient's health. For examples, health-monitoring systems continuously monitor patients' various body parameters in real time to improve patient treatment effectiveness [3]. Robotic surgical systems aid surgical procedures by performing actions with smooth and feedback-controlled motions [4]. MCPS are increasingly used in hospitals to provide high-quality healthcare, and have emerged as promising platforms for monitoring and controlling multiple aspects of patient health.

Similar to other systems involving automated decisions with impacts on human lives, MCPS impose numerous security and privacy challenges that are either shared with classical distributed systems or unique to MCPS as a result of the nature of connected components and system requirements [5]. For example, there is no doubt that MCPS must maintain patient's privacy information, avoiding leakage through direct exposure to unauthorized parties, side-channel information, or poor system implementation practices. Securing MCPS against malicious attacks is of paramount importance because a compromise of the system can easily impact patient's health and safety [6].

The more interesting aspect of security challenges imposed by medical cyber-physical systems lies within their unique application requirements and characteristics, in particular their inevitable interoperability requirement with components that may not have been originally designed to work in an environment with untrustworthy parties. Among the security issues that MCPS face is that operators of such systems are likely to have limited knowledge in systems security and privacy. As a result, a sequence of unstudied consequent executions demanded from the system may simply jeopardize the patient's privacy or safety. Dedicated protocols designed specifically for MCPS constitute another class of potential security threats. For instance, Medical Device Plug-and-Play (MD PnP) Interoperability initiative [2], if used without proper threat modeling as a forefront issue, can cause a sweet spot for attacks.

MCPS can choose to accommodate smarter and more centralized control environment that have absolute authority over the various components of the system. The alternative is to inject intelligent algorithms and decision making capabilities within individual components, creating a more dynamic and less centralized authority. In either case, one important question needs careful investigation: *what are the various levels of threat modeling and security requirements that a MCPS should entertain?* The importance of threat modeling and analysis has been also emphasized in recent FDA guidance documents on pre-market submission [7] and post-market management [8] of cyber-security in medical devices.

This work investigates a thorough understanding of threat modeling in MCPS as a step towards patient security and privacy. We examine the roles of stakeholders and system components and sketch an abstract architecture of a MCPS to demonstrate various threat modeling options. We also comment on the role of major security techniques that have been well established in the state-of-the-art and investigate their applicability and utility for the design of MCPS.

II. Related Work

In this section, we briefly summarize the prior work on threat modeling in cyber-physical systems (CPS).

Threat modeling is an approach for analyzing potential threats of a system at the design level and could possibly provide recommendations to system designers to address these identified threats with appropriate countermeasures. It can be performed from different perspectives, *e.g.*, attacker-centric model or system-centric model [9]. Attacker-centric model begins from identifying possible attackers, then evaluating their goals as well as the knowledge and resources available, and finally predicting how these goals might be achieved by adversaries. For example, researchers proposed to use game theory to assess security of CPS and improve the system's survivability in the face of strategic adversaries [10]. System-centric model focuses on identifying all possible attacks that target each of the system elements. Another commonly used approach is the attack tree based modeling [11]. It represents potential threats against a system using a tree structure, where the root node represents the overall threat and leaf nodes representing the different ways of attacks.

Though threat modeling has been widely explored in conventional IT systems, existing threat analysis in legacy IT systems can not be directly applied to CPS due to its complexity and the human in the loop. Javaid *et al.* [11] analyzed various security threats to an unmanned aerial vehicle (UAV) system using the attack tree approach, and proposed a cyber-security threat model showing possible attack paths in UAV. Martins *et al.* [9] adopted the Microsoft SDL Threat Modeling Tool [12] for software-related threat modeling in CPS domain, and used a railway temperature monitoring system as the case study to validate the proposed approach. Nourian *et al.* [13] analyzed the threats exploited by Stuxnet attack, which targeted nuclear centrifuges at the Iranian uranium enrichment plant. The authors demonstrated that the vulnerabilities exploited by Stuxnet could have been addressed if adopting a systems theoretic threat modeling approach at the design phase. Manikas *et al.* [14] applied multiple-valued logic decision diagrams to threat trees to assess the risk of medication delivery to a patient within the traditional hospital environment.

Several recent research efforts have focused on safety analysis of medical systems [15]–[17]. Alemzadeh *et al.* [15] conducted an in-depth study of public FDA recall data to characterize safety-critical computer failures in medical devices. They found that, in many instances of recalled devices, their safety mechanisms were not designed/implemented correctly or they were designed without identifying and handling the safety issues at all. Pajic *et al.* [16] studied the safety of a medical device system for the physiologic closed-loop control of drug infusion. They utilized a timed automaton model to express the safety property of a medical system, and modelled timing relationship between system components to prove safety of the system.

Halperin *et al.* [18] presented a general framework for evaluating the security and privacy of wireless implantable medical devices. Burleson *et al.* [19] discussed design principles for securing implantable medical devices. Rushanan *et al.* [20] emphasized the need of achieve trustworthy communication, trustworthy software, and trustworthy hardware and sensor interfaces in implantable medical devices and body area networks. Zhang *et al.* [21] discussed trustworthiness concerns of medical devices and possible countermeasures against these threats in MCPS. The authors summarized major trustworthiness requirements of a MCPS, including reliability and availability of medical devices, confidentiality and integrity of patient data. Tamara *et al.* [22] analyzed cyber-security attacks against teleoperated robotic surgery system, with a focus on denial-of-service (DoS) attacks.

The prior work on security and threat modeling for cyber-physical systems also shed light on similar problems with medical cyber-physical systems. Cyber-physical systems have common design characteristics, such as integrating components from multiple vendors with various design goals, and are under similar general safety requirements, such as maintaining human lives. However, to the best of our knowledge, a systematic analysis of security threats for modern medical CPS systems isn't a well studied subject in the literature.

III. Trust and Threat Models

A. MCPS Stakeholders

A medical cyber-physical system can be designed in various forms. We assume a centralized design model in which a control component is in place. In this model, the control component has two subsystems. One subsystem interacts directly with the practitioners and the second subsystem manages other components with direct human interactions. However, the second subsystem does need a network interface to a technical team to interfere when the system needs recovery from an urgent failure.

Practitioners are the main stakeholders in MCPS. Practitioners (including doctors and nurses) can have various roles and responsibilities and may differ in their usage patterns. For example, during an operation, an assistant nurse may be in charge of heartbeat monitoring. She can control the electrocardiography (ECG) device, which transmits current heart conditions to another unit through a connected component. The data may be viewed on a mobile device. In this case, simple protocols define how users provide input to the ECG device for only retrieving the data.

System administrators (or the technical team) are another important group of stakeholders that maintain the system's reliability and availability. The system administration team may have various responsibility levels. In one wide open case, system administrators can be as powerful as having access to the source code of individual system components, can temper with any operation, and can replace executable binaries on the devices with updated ones. A less ambitious role model for the system administrator demands to limit the capabilities of system administrators by mainly restricting tasks such as having full access to the device, and thus being able to modify the software, and by disallowing access to initiating medical commands, such as starting (or stopping) heartbeat monitoring.

Non-medical staff can interact with individual components in a MCPS. Office staff in a unit play important roles and can directly work with the system. An office staff typically inputs and retrieves patient data. Depending on hospital policies, however, the roles of non-medical staff can differ. In some environments, non-medical staff may not directly interact with a MCPS component that closely interacts with patients' body. In this scenario, non-medical staff manage patient's data through interaction with a demilitarized system component, *i.e.*, exercising limited privileges. However, the interconnected nature of MCPS may demand an indirect link with the interfaces provided to non-medical staff.

B. Trust Models

Trust modeling of MCPS is a particularly important and difficult task. Trusting a particular component, software, or stakeholder can have catastrophic consequences involving patient's health conditions and lives. We categorize the trust model with respect to individuals interacting with a MCPS as follows.

1) *Trustworthy* users that represent relatively trustworthy individuals. The system may not completely function and achieve its goals without defining a set of fully trusted users. This is despite the fact that any individual, regardless of skills and career background, may intentionally or unintentionally commit mistakes that can affect patients' lives.
2) *Trusted but error-prone* users are similar to fully trusted users in their intentions but are expected to commit mistakes due to lack of knowledge or training. Thus, this group of users cannot handle all the tasks handled by the fully trusted users without proper supervision or time limitation.
3) *Untrustworthy* individuals exist in any environment. They do include general users that are not authorized to perform any medical actions or individuals that are either stakeholders or work closely with them but are not particularly authorized to work in a team that is responsible for the treatment of a patient.
4) *Temporarily trustworthy* individuals that may need access to a part of the system with specific operation authorization for a limited period of time. Such individuals may not be allowed operations that can have consequences beyond their time period. For example, a temporarily trustworthy individual may not create a permanent account in the system.

Individual software and hardware components of the system can have various levels of trust, which are categorized as:

1) *Trustworthy* hardware or software that provide the trusted computing base for other parts of the system. The trustworthy hardware or software is assumed to be well designed such that they guard against basic security threats. Also, they should have very limited surfaces and controlled interfaces, making the potential attacks particularly difficult to succeed.
2) *Trustworthy but vulnerable* components are trusted in the general sense that the components were designed by a trusted engineering team and the intentions of the software are not malicious. However, the security and reliability of these components are not established, and the attack surface on them is assumed to be substantially wider than the trustworthy components.
3) *Untrustworthy* components are designed by teams that are not trusted within the MCPS environment. They may be either highly vulnerable or have mild to severe malicious intentions. They can attempt to collect patient data or inject malicious signals. They may temper with the normal execution path in other components.
4) *External components* are any software that can interact with the system over an open network. Such an interaction is highly unexpected to happen directly. However, external network request can penetrate a hospital system, aiming to reach a component within MCPS. A naive design of MCPS may allow direct access through Internet.

C. Threat Models

Threat models differ depending on the trust model instantiated in the system. A weak and open trust model involves moving most of the individuals to the trustworthy category. Similarly, the trust model can become weak if arbitrary systems from unverified sources and vendors set to be trusted or systems become easily accessible from Internet.

We analyze the MCPS threat models from a motivation perspective. Threat motivations are closely related to the functions provided in MCPS. In general, attackers are motivated to either

1) *Breach the privacy* of patients by passively or actively collecting precise data from the system. The motivation may either be to use the data for purely commercial purposes (directing commercials to the patient or the patient's family) or to influence the patients' life by means of blackmailing or more severe attacks. For example, an attacker may derive the existence of certain disease of a patient by intercepting the communications of the patient's medial device via wireless hacking tools [18].
2) *Direct influence* on a patient's health conditions is another threat, that is often politically or criminally motivated and targets specific individuals. For example, a criminal hacker might inject false commands by using wireless tools to change a medical devices state to harm a patient's health condition [18]. In Section IV, we present a new architecture and a policy set that can significantly improve the safety and security of MCPS.

D. Threat Analysis

Figure 1 lists the potential threats posed to MCPS in terms of three security properties of the system: confidentiality, integrity, and availability. We classify the vulnerabilities in MCPS into four types according to their sources: 1) communication links, 2) software, 3) platform/hardware; and 4) users. The causes of these vulnerabilities include isolation assumption (*i.e.*, neglecting security-by-design), increased connectivity, and heterogeneity [23].

Short-range wireless communications are usually used in MCPS. The confidentiality of MCPS can be compromised by eavesdropping due to the lack of proper encryption. An attacker could exploit software vulnerabilities such as the buffer overflow vulnerability to steal sensitive information [24]. It is also possible the attacker physically accesses to medical devices to intercept sensitive information. Insider attacks can affect all the three security properties of MCPS. In addition,

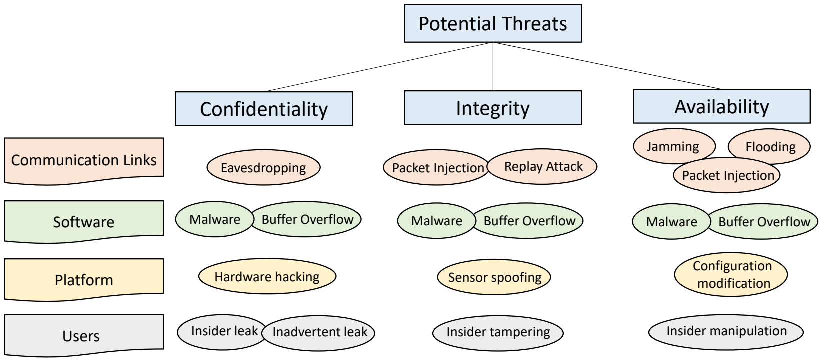

Fig. 1. Potential threats in MCPS

medical information may be leaked to unauthorized parties by users' inadvertent activities [25].

Integrity of a system is violated by modifying existing information or fabricating new information [11]. Attacks infiltrating the integrity of MCPS can be the false packet/command injection or replaying outdated measures in the communication layer. Software based exploits are a common way of compromising the integrity, such as overwriting sensor values or critical control decision variables through memory corruption attacks [26]. MCPS are also vulnerable to sensor data spoofing attacks in the physical layer, *e.g.*, resonant acoustic injection attacks can disable the function of MEMS-based gyroscopes [27].

DDoS/DoS attacks are the main threats to the availability of MCPS, *e.g.*, traffic jamming that disrupts communication through interference or collision, overflowing the buffer memory of network cards, and broadcasting spoofed network packets. By physically accessing to a medical device, it is possible an attacker modifies system configurations or corrupt data that impact the availability of the system.

IV. Example Architecture

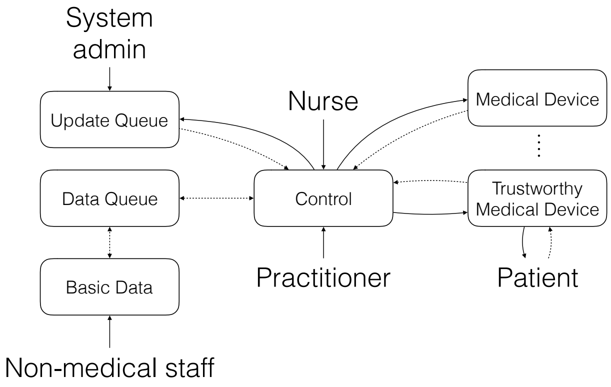

Fig. 2. A proposed architecture in which a number of untrustworthy components are decoupled from the control center in a medical cyber-physical system. The dotted lines capture flow of medical data while the solid lines capture the action initiation direction (according to general access control policies). The boxes represent system components and stakeholders are mentioned in text.

Traditional clinical scenarios can be modelled as loosely coupled "closed-loop control systems", where patients are plants, caregivers are the controllers, and individual medical devices act as sensors and actuators [3]. Caregivers together with medical devices affect the state of a patient. However, the controller (*i.e.*, caregiver) may not be able to continuously monitor a large number of sensors and apply timely treatments to the patient, thus, accidents caused by human errors are inevitable. With the advancement of digital technologies, MCPS has emerged as a promising platform to bridge the gap, by embracing the potential of embedded software and network connectivity in medical systems.

An architecture that considers the threat and trust models in Section III is presented in Figure 2. The main idea of the architecture is to decouple sensitive operations from untrustworthy components and individuals. Since some potentially untrustworthy stakeholders (or components) must interact with the system, we aim to increase the security of the system by a mandatory *first-check-then-update* method. For example, if a software failure occurs in a component, a system administrator must provide an update. This update enters an *update queue* and does not take effect immediately. The updates will only pass through, if a trusted individual views it, and verifies its intentions, and receives a test confirmation from the control component. Since this architecture involves medical operations, our assumption is that the trusted individuals are practitioners that directly work with the patient.

The central part of the architecture is the control component. The control component is the trusted computing base for the entire system. It must be thoroughly verified and tested before entering operations. The core part of the control component is a command center that receives and monitors medical commands. The control component also contains several queues of requests from untrustworthy parts including updates from system administrators, non-medical staff, and untrustworthy medical devices, which are all decoupled from direct interaction with the patient. Further, the control component includes testing capabilities. Dynamic and static testing tools, that can work efficiently in real-time, must be included in the control component for performing checks, especially on the updates proposed by the system administrators.

V. Security Remedies

In this section, we discuss promising approaches for securing MCPS against possible threats and attacks.

A. Anomaly Detection

Runtime monitoring of MCPS is an effective countermeasure against various attacks. There has been considerable research activity on attack detection for CPS [28]. The majority of existing works in this field are behavior model-based, which can be further divided into two lines of research based on physical process models or cyber models. Physics-based models define normal operations in CPS for anomaly detection, where system states must follow immutable laws of physics in CPS. Cyber-based models characterize the expected program/system behaviors to recognize potential attacks [29], [30]. Since CPS are application-specific, most existing works are designed to detect specific attacks for specific applications, such as smart grid [31], unmanned aerial vehicles [32], and industrial control process [33]. Recent studies [26], [34] have shown that CPS may suffer from a variety of runtime attacks, including code-reuse attacks [35], malicious code injection [36], non-control data attacks [37], and false data injection attacks [38]. C-FLAT [26] instruments target programs running on CPS devices to achieve the remote attestation of execution paths of these programs. Zimmer *et al.* [39] exploited worst-case execution time information obtained through static analysis of application code to detect code injection attacks in CPS. In particular, Mitchell *et al.* [40] analyzed a behavior-rule specification-based technique for intrusion detection in MCPS. They proposed to transform behavior rules to a state machine, then at runtime check against the transformed state machine for deviation from a medical device's behavior specification.

B. Cryptographic Measures

Cryptography is a commonly used approach for securing the communication channel from unauthorized access. However, most of the traditional cryptographic primitives that have been employed in general-purpose IT systems, both ciphers and hash functions, cannot be directly applied to MCPS due to the size, real-time, and power constraints of medical devices. For example, the high energy and implementation overhead of asymmetric cryptography pose significant challenges for encrypting sensitive data in MCPS. To mitigate this problem, compression techniques may be used before encryption to reduce the overhead [21]. Lightweight cryptography has recently received considerable attention and many lightweight block ciphers [41]–[44] have been proposed in the literature. These low-cost low-latency encryption techniques are proposing to provide cryptographic building blocks for resource constrained medical devices. Kocabas *et al.* [6] surveyed conventional and emerging encryption schemes that might be used to provide secure storage, data sharing, and computation in MCPS.

C. System Hardening

Secure execution environment can be used to defend a wide range of threats in MCPS. Isolating security-critical applications from untrusted OS is a promising technique to enhance MCPS security, such as by the hardware security support of Intel's TrustLite or ARM's TrustZone technologies. Shepherd *et al.* [45] analyzed different technologies of trust computing and their applications to the emerging domains of CPS, including Trusted Platform Module (TPM), Secure Elements (SE), Hypervisors and Virtualisation, Intel TXT, Trusted Execution Environments (TEE) and Encrypted Execution Environment (E3). Increasing the integrity of the underlying operating system is a crucial step for MCPS. Previous work such as [46] achieves higher security when untrustworthy components are present. MCPS can benefit from inter-authentication of components to improve the system's integrity. Big picture analysis of the overall distributed environment for MCPS can be achieved by graph-theoretic methods such as the one presented in [47]. Graph-based optimization, as suggested in [47], coupled with parameters with MCPS can provide a basis for reasoning about the overall integrity of the system.

VI. Conclusions

A thorough understanding of threat modeling in MCPS is a step towards improving patient security and privacy, while benefiting from efficiency provided by MCPS. Through sketching an abstract architecture of a MCPS, as we examined the roles of stakeholders and components, and demonstrated various threat modeling options in MCPS, we envision future MCPS to enable clean security models that are verifiable. The discussion of major security techniques to mitigate the threats in MCPS can further enhance design decisions made in future systems. Enhancing the security and privacy in medical cyber-physical systems remains a serious challenge demanding careful considerations and joint efforts by the industry, the health systems, and the research community.

Acknowledgement

This work has been supported in part by Security and Software Engineering Research Center (S^2ERC), an NSF sponsored multi-university Industry/University Cooperative Research Center (I/UCRC).

References

[1] D. Arney, R. Jetley, P. Jones, I. Lee, and O. Sokolsky, "Formal methods based development of a PCA infusion pump reference model: Generic infusion pump (gip) project," in *Proceedings of the 2007 Joint Workshop on High Confidence Medical Devices, Software, and Systems and Medical Device Plug-and-Play Interoperability*, ser. HCMDSS-MDPNP '07, 2007, pp. 23–33.

[2] C. Kim, M. Sun, S. Mohan, H. Yun, L. Sha, and T. F. Abdelzaher, "A framework for the safe interoperability of medical devices in the presence of network failures," in *ICCPS '10*, 2010.

[3] D. Arney, M. Pajic, J. M. Goldman, I. Lee, R. Mangharam, and O. Sokolsky, "Toward patient safety in closed-loop medical device systems," in *PICCPS '10*, 2010, pp. 139–148.

[4] H. Alemzadeh, R. K. Iyer, Z. T. Kalbarczyk, N. Leveson, and J. Raman, "Adverse events in robotic surgery: A retrospective study of 14 years of FDA data," *CoRR*, vol. abs/1507.03518, 2015.

[5] I. Lee, O. Sokolsky, S. Chen, J. Hatcliff, E. Jee, B. Kim, A. King, M. Mullen-Fortino, S. Park, A. Roederer, and K. K. Venkatasubramanian, "Challenges and research directions in medical cyber physical systems," *Proceedings of the IEEE*, vol. 100, no. 1, pp. 75–90, 2012.

[6] O. Kocabas, T. Soyata, and M. K. Aktas, "Emerging security mechanisms for medical cyber physical systems," *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, vol. 13, no. 3, pp. 401–416, May 2016.

[7] "Content of premarket submissions for management of cybersecurity in medical devices," https://www.fda.gov/downloads/medicaldevices/deviceregulationandguidance/guidancedocuments/ucm356190.pdf, [Online; accessed 20-June-2017].

[8] "Postmarket management of cybersecurity in medical devices," https : / / www . fda . gov / downloads / MedicalDevices / DeviceRegulationandGuidance / GuidanceDocuments / UCM482022 . pdf, [Online; accessed 20-June-2017].

[9] G. Martins, S. Bhatia, X. Koutsoukos, K. Stouffer, C. Tang, and R. Candell, "Towards a systematic threat modeling approach for cyber-physical systems," in *2015 Resilience Week (RWS)*, 2015, pp. 1–6.

[10] H. Sandberg, S. Amin, and K. H. Johansson, "Cyberphysical security in networked control systems: An introduction to the issue," *IEEE Control Systems*, vol. 35, no. 1, pp. 20–23, 2015.

[11] A. Y. Javaid, W. Sun, V. K. Devabhaktuni, and M. Alam, "Cyber security threat analysis and modeling of an unmanned aerial vehicle system," in *2012 IEEE Conference on Technologies for Homeland Security (HST)*, 2012, pp. 585–590.

[12] "Microsoft SDL threat modelling tool," *Network Security*, vol. 2009, no. 1, pp. 15 – 18, 2009.

[13] A. Nourian and S. Madnick, "A systems theoretic approach to the security threats in cyber physical systems applied to Stuxnet," *IEEE Transactions on Dependable and Secure Computing*, 2015.

[14] T. W. Manikas, D. Y. Feinstein, and M. A. Thornton, "Modeling medical system threats with conditional probabilities using multiple-valued logic decision diagrams," in *2012 IEEE 42nd International Symposium on Multiple-Valued Logic*, 2012, pp. 244–249.

[15] H. Alemzadeh, R. K. Iyer, Z. Kalbarczyk, and J. Raman, "Analysis of safety-critical computer failures in medical devices," *IEEE Security Privacy*, vol. 11, no. 4, pp. 14–26, 2013.

[16] M. Pajic, R. Mangharam, O. Sokolsky, D. Arney, J. Goldman, and I. Lee, "Model-driven safety analysis of closed-loop medical systems," *IEEE Transactions on Industrial Informatics*, vol. 10, no. 1, pp. 3–16, 2014.

[17] P. L. Wu, L. Sha, R. B. Berlin, and J. M. Goldman, "Safe workflow adaptation and validation protocol for medical cyber-physical systems," in *2015 41st Euromicro Conference on Software Engineering and Advanced Applications*, 2015, pp. 464–471.

[18] D. Halperin, T. S. Heydt-Benjamin, K. Fu, T. Kohno, and W. H. Maisel, "Security and privacy for implantable medical devices," *IEEE Pervasive Computing*, vol. 7, no. 1, pp. 30–39, 2008.

[19] W. Burleson, S. S. Clark, B. Ransford, and K. Fu, "Design challenges for secure implantable medical devices," in *DAC Design Automation Conference 2012*, 2012, pp. 12–17.

[20] M. Rushanan, A. D. Rubin, D. F. Kune, and C. M. Swanson, "Sok: Security and privacy in implantable medical devices and body area networks," in *Proceedings of the 2014 IEEE Symposium on Security and Privacy*, ser. SP '14, 2014, pp. 524–539.

[21] M. Zhang, A. Raghunathan, and N. K. Jha, "Trustworthiness of medical devices and body area networks," *Proceedings of the IEEE*, vol. 102, no. 8, pp. 1174–1188, 2014.

[22] T. Bonaci, J. Yan, J. Herron, T. Kohno, and H. J. Chizeck, "Experimental analysis of denial-of-service attacks on teleoperated robotic systems," in *ICCPS '15*, 2015, pp. 11–20.

[23] A. Humayed, J. Lin, F. Li, and B. Luo, "Cyber-physical systems security – a survey," *IEEE Internet of Things Journal*, 2017.

[24] S. Hanna, R. Rolles, A. Molina-Markham, P. Poosankam, K. Fu, and D. Song, "Take two software updates and see me in the morning: The case for software security evaluations of medical devices," in *Proceedings of the 2Nd USENIX Conference on Health Security and Privacy*, ser. HealthSec'11, 2011.

[25] L. Cheng, F. Liu, and D. D. Yao, "Enterprise data breach: causes, challenges, prevention, and future directions," *WIREs Data Mining and Knowledge Discovery*, 2017.

[26] T. Abera, N. Asokan, L. Davi, J. Ekberg, T. Nyman, A. Paverd, A. Sadeghi, and G. Tsudik, "C-FLAT: control-flow attestation for embedded systems software," in *CCS*, 2016.

[27] T. Trippel, O. Weisse, W. Xu, P. Honeyman, and K. Fu, "WALNUT: Waging doubt on the integrity of MEMS accelerometers with acoustic injection attacks," in *Proceedings of the 2nd Annual IEEE European Symposium on Security and Privacy*, 2017.

[28] R. Mitchell and I.-R. Chen, "A survey of intrusion detection techniques for cyber-physical systems," *ACM Comput. Surv.*, vol. 46, no. 4, pp. 55:1–55:29, Mar. 2014.

[29] X. Shu, D. Yao, and N. Ramakrishnan, "Unearthing stealthy program attacks buried in extremely long execution paths," in *CCS*, 2015.

[30] K. Xu, K. Tian, D. Yao, and B. G. Ryder, "A sharper sense of self: Probabilistic reasoning of program behaviors for anomaly detection with context sensitivity," in *2016 46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, 2016, pp. 467–478.

[31] S. Sridhar, A. Hahn, and M. Govindarasu, "Cyber physical system security for the electric power grid," *Proceedings of the IEEE*, vol. 100, no. 1, pp. 210–224, 2012.

[32] R. Mitchell and I.-R. Chen, "Adaptive intrusion detection of malicious unmanned air vehicles using behavior rule specifications," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 44, no. 5, pp. 593–604, 2014.

[33] D. I. Urbina, J. A. Giraldo, A. A. Cardenas, N. O. Tippenhauer, J. Valente, M. Faisal, J. Ruths, R. Candell, and H. Sandberg, "Limiting the impact of stealthy attacks on industrial control systems," in *CCS*, 2016.

[34] D. D. Chen, M. Egele, M. Woo, and D. Brumley, "Towards automated dynamic analysis for Linux-based embedded firmware," in *ISOC NDSS*, 2016.

[35] R. Roemer, E. Buchanan, H. Shacham, and S. Savage, "Return-oriented programming: Systems, languages, and applications," *ACM Trans. Info. & System Security*, vol. 15, no. 1, Mar. 2012.

[36] A. Francillon and C. Castelluccia, "Code injection attacks on harvard-architecture devices," in *CCS*, 2008.

[37] S. A. Z. L. C. P. S. Z. L. Hong Hu, Shweta Shinde, "Data-oriented programming: On the expressiveness of non-control data attacks," in *IEEE Security and Privacy*, 2016.

[38] H. Alemzadeh, D. Chen, X. Li, T. Kesavadas, Z. T. Kalbarczyk, and R. K. Iyer, "Targeted attacks on teleoperated surgical robots: Dynamic model-based detection and mitigation," in *2016 46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, 2016, pp. 395–406.

[39] C. Zimmer, B. Bhat, F. Mueller, and S. Mohan, "Time-based intrusion detection in cyber-physical systems," in *ICCPS'10*, 2010.

[40] R. Mitchell and I. R. Chen, "Behavior rule specification-based intrusion detection for safety critical medical cyber physical systems," *IEEE Transactions on Dependable and Secure Computing*, vol. 12, no. 1, pp. 16–30, Jan 2015.

[41] A. Bogdanov, L. R. Knudsen, G. Leander, C. Paar, A. Poschmann, M. J. B. Robshaw, and Y. Seurin.

[42] C. Cannière, O. Dunkelman, and M. Knežević, "KATAN and KTANTAN – a family of small and efficient hardware-oriented block ciphers," in *Proceedings of the 11th International Workshop on Cryptographic Hardware and Embedded Systems*, ser. CHES '09, 2009, pp. 272–288.

[43] J. Guo, T. Peyrin, A. Poschmann, and M. Robshaw, "The LED block cipher," in *Proceedings of the 13th International Conference on Cryptographic Hardware and Embedded Systems*, ser. CHES'11, 2011, pp. 326–341.

[44] J. Borghoff, A. Canteaut, T. Güneysu, E. B. Kavun, M. Knezevic, L. R. Knudsen, G. Leander, V. Nikov, C. Paar, C. Rechberger, P. Rombouts, S. S. Thomsen, and T. Yalc, "Prince: A low-latency block cipher for pervasive computing applications," in *Proceedings of the 18th International Conference on The Theory and Application of Cryptology and Information Security (ASIACRYPT'12)*, 2012, pp. 208–225.

[45] C. Shepherd, G. Arfaoui, I. Gurulian, R. P. Lee, K. Markantonakis, R. N. Akram, D. Sauveron, and E. Conchon, "Secure and trusted execution: Past, present, and future - a critical review in the context of the internet of things and cyber-physical systems," in *2016 IEEE Trustcom*, 2016, pp. 168–177.

[46] H. M. J. Almohri, D. Yao, and D. G. Kafura, "Process authentication for high system assurance," *IEEE Trans. Dependable Secure Computing*, vol. 11, no. 2, pp. 168–180, 2014. [Online]. Available: http://dx.doi.org/10.1109/TDSC.2013.29

[47] H. M. J. Almohri, L. T. Watson, D. Yao, and X. Ou, "Security optimization of dynamic networks with probabilistic graph modeling and linear programming," *IEEE Trans. Dependable Secure Computing*, vol. 13, no. 4, pp. 474–487, 2016. [Online]. Available: http://dx.doi.org/10.1109/TDSC.2015.2411264