This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMSCS.2017.2720660, IEEE Transactions on Multi-Scale Computing Systems

DISASTER: Dedicated Intelligent Security Attacks on Sensor-triggered Emergency Responses

Arsalan Mosenia, *Student Member, IEEE,* Susmita Sur-Kolay, *Senior Member, IEEE,* Anand Raghunathan, *Fellow, IEEE,* and Niraj K. Jha, *Fellow, IEEE*

Abstract—Rapid technological advances in microelectronics, networking, and computer science have resulted in an exponential increase in the number of cyber-physical systems (CPSs) that enable numerous services in various application domains, e.g., smart homes and smart grids. Moreover, the emergence of the Internet-of-Things (IoT) paradigm has led to the pervasive use of IoT-enabled CPSs in our everyday lives. Unfortunately, as a side effect, the number of potential threats and feasible security attacks against CPSs has grown significantly.

In this paper, we introduce a new class of attacks against CPSs, called *dedicated intelligent security attacks against sensor-triggered emergency responses (DISASTER)*. DISASTER targets safety mechanisms deployed in automation/monitoring CPSs and exploits design flaws and security weaknesses of such mechanisms to trigger emergency responses even in the absence of a real emergency. Launching DISASTER can lead to serious consequences for three main reasons. First, almost all CPSs offer specific emergency responses and, as a result, are potentially susceptible to such attacks. Second, DISASTER can be easily designed to target a large number of CPSs, e.g., the anti-theft systems of all buildings in a residential community. Third, the widespread deployment of insecure sensors in already-in-use safety mechanisms along with the endless variety of CPS-based applications magnifies the impact of launching DISASTER.

In addition to introducing DISASTER, we describe the serious consequences of such attacks. We demonstrate the feasibility of launching DISASTER against the two most widely-used CPSs: residential and industrial automation/monitoring systems. Moreover, we suggest several countermeasures that can potentially prevent DISASTER and discuss their advantages and drawbacks.

Index Terms—Automation system, cyber-physical system, emergency response, Internet-of-Things, safety, security attack, smart thing, wireless sensors.

This paper can be cited as: A. Mosenia, S. Sur-Kolay, A. Raghunathan and N. K. Jha, "DISASTER: Dedicated intelligent security attacks on sensor-triggered emergency responses," in IEEE Trans. Multi-Scale Computing Systems (TMSCS), 2017, DOI: 10.1109/TMSCS.2017.272066

The latest version of this manuscript is available on http://ieeexplore.ieee.org/abstract/document/7959602/

1 Introduction

Cyber-physical systems (CPSs) offer a transformative approach to automation and monitoring through integration of processing, networking, and control. This combination and active collaboration of computational elements, e.g., powerful base stations, and small embedded devices, e.g., sensors, enable CPSs to reliably and efficiently control physical entities.

Recent and ongoing advances in microelectronics, networking, and computer science have resulted in significant CPS growth. Such systems facilitate automation and monitoring in various application domains, e.g., smart manufacturing lines, smart homes, smart cities, smart grids, and smart vehicles. Moreover, with the emergence of the Internet-of-Things (IoT) paradigm and IoT-enabled CPSs in the last decade, it has become clear that the economic and societal potential of such systems is far beyond what may have been imagined. Thus, major investments have been made worldwide to design and develop CPSs.

As a side effect of the rapid development and pervasive use of CPSs, the number of potential threats and possible attacks against the security of such systems is increasing drastically, while, unfortunately, their security needs are not yet well-recognized [1]–[3].

An essential component of a majority of CPSs is a safety mechanism, which aims to minimize harm to users' well-being or damage to equipment upon the detection of risks, hazards, or unplanned events. The security of safety mechanisms is an emerging research topic that is attracting increasing attention in academic, industrial as well as governmental research. A few real-world attacks and recent research efforts have demonstrated that generic classes of security attacks, e.g., computer worms, man-in-the-middle attacks, and denial of service (DoS), which have been extensively studied in the network/computer security domains, can be modified to *disable* the safety mechanisms of CPSs. For example, in 2003, the SQL Slammer worm infected the Davis-Besse nuclear power plant in Ohio, USA, and disabled the plant's safety parameter display system and plant process computer for several hours [4]. Stuxnet [5], [6], a real-world high-impact man-in-the-middle attack,

Acknowledgments: This work was supported by NSF under Grant no. CNS-1219570.

Arsalan Mosenia is with the Department of Electrical Engineering, Princeton University, Princeton, NJ 08544, USA (e-mail: arsalan@princeton.edu).

Susmita Sur-Kolay is with the Advanced Computing and Microelectronics Unit, Indian Statistical Institute, Kolkata 700108, India (e-mail: ssk@isical.ac.in).

Anand Raghunathan is with the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907, USA (e-mail: raghunathan@purdue.edu).

Niraj K. Jha is with the Department of Electrical Engineering, Princeton University, Princeton, NJ 08544, USA (e-mail: jha@princeton.edu).

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMSCS.2017.2720660, IEEE Transactions on Multi-Scale Computing Systems

was launched against safety mechanisms employed in thousands of industrial CPSs in 2012. Stuxnet faked industrial process control sensor signals so that safety mechanisms of infected systems were disabled and, as a result, their emergency responses were not activated even in the presence of a real emergency. Furthermore, Lamb developed a DoS attack against residential intrusion detection systems in which the attacker continuously jams the communication channel between motion sensors and the base station to suppress the system's alarm that is supposed to be trigerred in the presence of an intruder [7].

In this paper, we present a new class of attacks against CPSs, called *dedicated intelligent security attacks against sensor-triggered emergency responses (DISASTER)*. DISASTER targets *safety systems* utilized in CPSs in two-fold fashion. First, safety systems are often vulnerable to security attacks since they do not use strong cryptographic operations due to domain-specific requirements (in particular, low latency, low cost, and long battery lifetime). Second, they have the capability to *override normal operations of the whole CPS*. As opposed to the previously-proposed attacks that mainly aim to **disable** emergency responses of safety mechanisms in an emergency situation, DISASTER attempts to **trigger** the systems' emergency responses in the absence of a real emergency.

Our key contributions can be summarized as follows:

- We introduce DISASTER and discuss potential attackers who may be motivated to launch such security attacks.
- We discuss the impact of DISASTER by describing the consequences of launching such attacks.
- We examine common design flaws and security weaknesses of safety mechanisms and their components, which may be exploited by an attacker to launch DISASTER.
- We demonstrate the feasibility of launching DISASTER in realistic scenarios, e.g., residential and industrial automation/monitoring systems.
- We suggest several countermeasures to proactively address DISASTER, and discuss their advantages and drawbacks.

The remainder of the paper is organized as follows. Section II describes the threat model. Section III describes the typical architecture of CPSs that we consider in this paper. Then, it discusses different components, design flaws, and security weaknesses of their safety mechanisms. Section IV describes potential consequences of launching DISASTER. Section V demonstrates how it is feasible to launch the proposed attacks against two real CPSs. Section VI suggests several countermeasures to prevent DISASTER and describes why proactive countermeasures might not always be able to provide sufficient protection against the proposed attacks. Section VII briefly describes related work. Section VIII discusses the feasibility of IoT-enabled DISASTER and describes why even state-of-the-art cryptographic operations are not yet efficient enough to be deployed in the context of safety operations. Finally, Section IX concludes the paper.

2 Threat Model

In this section, we first describe what enables DISASTER and makes CPSs susceptible to such attacks, and discuss why launching DISASTER can be disastrous in real-word scenarios. Second, we discuss who the potential attackers may be who exploit vulnerabilities of CPSs to launch the proposed security attacks, and what their motivations may be.

2.1 Problem definition

As described later in Section 3, in a typical CPS, a centralized processing unit (commonly referred to as the base station) obtains a description of the environment based on the data that it collects from different sensors, and processes the sensory data along with user inputs to control physical objects. Given the need for direct interactions of such a system with both the environment and users, safety and security become fundamental requirements for it. Safety mechanisms employed in CPSs protect users from undesirable outcomes, risks, hazards, or unplanned events that may result in death, injury, illness, or other harm to individual's well-being, damage to equipment or harm to organizations, while security protocols are focused on protecting the system from intentional attacks [8].

Although safety and security seem to share very similar goals at first glance, a close examination of various safety and security requirements demonstrates that ensuring both safety and security of CPSs is not always possible due to the existence of unavoidable safety-security conflicts [9], [10]. When ensuring both safety and security is not feasible, safety is typically given preference and safety mechanisms willingly sacrifice security of the system to ensure users' safety. For example, modern vehicles commonly support an automatic door unlocking mechanism [11], which opens the vehicle's doors upon the detection of a collision. This safety mechanism ensures passengers' safety by completely disabling the car's security system after detecting an accident.

The unavoidable safety-security conflicts along with different design flaws and security weaknesses of components, e.g., sensors and base stations, used in safety mechanisms can, unfortunately, facilitate DISASTER. In DISASTER, the attacker exploits such conflicts/weaknesses to *fool the safety mechanism under attack into falsely labeling a normal situation as an emergency in an attempt to activate emergency responses even though they are not needed*. As discussed later in Section 4, activating emergency responses in the absence of an emergency can lead to catastrophic situations, ranging from system shutdown to life-threatening conditions.

Launching DISASTER can have severe negative consequences in real-world scenarios for three reasons. First, since the majority of CPSs offer emergency responses, they are susceptible to such attacks. Second, as demonstrated later in Section 5.2, DISASTER can be implemented to simultaneously target a large number of CPSs, e.g., the anti-theft systems of all buildings in a residential community. Third, the widespread use of vulnerable sensors along with the endless variety of CPS-based applications magnifies negative consequences of launching DISASTER.

Despite the fact that safety mechanisms are designed to control hazards and emergency situations and minimize

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMSCS.2017.2720660, IEEE Transactions on Multi-Scale Computing Systems

their associated risks, a careless design of a safety mechanism endangers both users' safety and system's security.

2.2 Potential attackers

Next, we discuss potential attackers who may target CPSs, and what their motivations might be.

As discussed earlier, DISASTER is widely applicable since CPSs used for automation/monitoring are in widespread use in our everyday lives. Such systems may manage a huge amount of information and be used for many services, ranging from industrial management to residential monitoring. This has made such CPSs targets of interest for a multitude of attackers, including, but not limited to, cyberthieves, hacktivists, occasional hackers, and cyberterrorists. Unfortunately, as described later in Section 5.2, an attacker with very limited resources, e.g., a very cheap radio transmitter such as HackRF [12], can easily launch powerful large-scale attacks against CPSs.

As extensively discussed later in Section 4, the attackers might launch DISASTER to access restricted areas, cause economic damage to companies or individuals, trigger life-threatening operations, or halt automation/monitoring processes. Moreover, they might try to make CPS use so inconvenient to the user that he is forced to shut down the whole system.

3 Typical Components and Weaknesses of Safety Mechanisms

In this section, we first describe the typical architecture of CPSs that we consider in this paper, and discuss different components of their safety mechanisms, and two main types of emergency responses that they provide. Second, we discuss design flaws and security weaknesses, which are commonly present in widely-used safety mechanisms.

3.1 Typical CPS architecture

Fig. 1 illustrates a common CPS architecture that includes safety mechanisms. A typical CPS consists of: (i) a base station, which collects and processes environment-related data and controls other components, (ii) wireless sensors that continuously collect data and transmit them to the base station, and (iii) physical objects that are controlled by the base station. State-of-the-art CPSs may also allow the user to remotely control, configure, or access the system over the Internet. The base station gathers data from different sources, e.g., sensors, cloud servers, and user inputs, and processes them to control different physical objects. Furthermore, a majority of modern CPSs have a safety mechanism, which typically needs two extra components: a safety unit and warning devices, e.g., speakers. The safety unit is usually integrated into the base station. When it detects an emergency, e.g., a fire or an accident, it activates warning devices or overrides control signals of physical objects to minimize safety risks associated with the situation. In each application domain, certain conditions and states are defined as emergency situations in which safety risks are present and should be actively and aggressively addressed.

Although emergency responses vary from one application to another, there are two types of emergency responses:

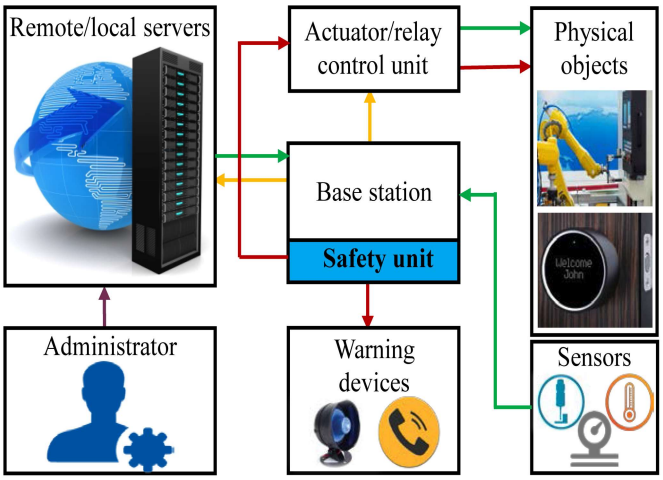

Fig. 1. Common architecture of CPS. Upon the detection of an emergency, the safety unit directly controls the physical objects or warns the users by activating the passive components.

active and passive. Next, we briefly describe each.

1. Active response: Such a response actively attempts to minimize the safety risks by controlling various actuators and components in the system. When the human operator is unable to provide a sufficiently fast response, it is required that the CPS offer an active response to minimize damages associated with the emergency situation. Typically, when a CPS provides an active response, it initiates a specific emergency procedure or halts the system's normal operation. For example, consider an insulin delivery system that continuously monitors blood glucose and injects insulin into the patient's body in order to regulate the blood glucose level. If the system detects a life-threatening low level of blood sugar, it immediately halts the injection procedure to ensure patient safety.

2. Passive response: A passive emergency response is provided by the system to warn human operators, proximate people, e.g., residents of a building, or emergency departments, e.g., fire department, about the need to take immediate action. Unlike an active response, it does not directly control physical entities. For example, consider a simple fire detection system. Upon detecting a fire, the system provides a passive response that activates various notification appliances, e.g., flashing lights and electromechanical horns.

3.2 Common design flaws and security weaknesses

Due to the inevitable complexities of CPSs, heterogeneity of entities that form them, and limitation of on-sensor resources (e.g., small amount of on-sensor storage), designing completely secure safety mechanisms and finding perfect safety design strategies are very daunting. As a result, it is easy to find several already-in-market products that are vulnerable to DISASTER. Next, we discuss three design flaws and security weaknesses of safety mechanisms that may enable an attacker to launch such attacks.

3.2.1 Using insecure sensors

As mentioned earlier, wireless sensors continuously collect and transmit data that are needed for detecting emergency

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMSCS.2017.2720660, IEEE Transactions on Multi-Scale Computing Systems

situations. A fundamental consideration in the design of wireless sensors is choosing the communication protocol. Many manufacturers have decided to design and use customized wireless protocols in an attempt to minimize expenses and maximize battery lifetime. We examined 70 already-in-market sensors from 10 different manufacturers (we reviewed documentations of 62 sensors, and as discussed later in Section 5.2, we closely inspected and attacked eight sensors under realistic scenarios) that are widely used in safety mechanisms. We noticed that the communication protocol of each sensor has at least one of the following security weaknesses.

1. Lack of data encryption or obfuscation: Unfortunately, a great number of customized wireless communication protocols used in already-in-market sensors support neither strong encryption nor obfuscation mechanisms and are, hence, susceptible to various forms of security attacks. In particular, due to the specific requirements of safety systems (e.g., fast response time and monitorability) and common limitations of sensing platforms (e.g., limited on-sensor energy, small amount of available memory, and limited computation power), the majority of sensors utilized in safety mechanisms transmit unencrypted non-obfuscated packets over the communication channel. As a result, an attacker can reverse-engineer the system's communication protocol and generate illegitimate packets (packets not created by legitimate sensors) using his equipment.

2. Lack of timestamp: A great number of sensors do not include a timestamp (i.e., a sequence of encrypted information identifying when the transmission occurred) in their packets. As a result, the system's base station is unable to distinguish new legitimate packets generated by the sensor from old seemingly legitimate ones, which are recorded and retransmitted by an attacker.

3. Using default passwords: Setting non-random default passwords at the manufacturing/installation time is a very common mistake that can lead to severe security attacks against the system even when strong encryption mechanisms are utilized to secure the communication channel. It is very common for system administrators or user to forget to change the system's default password at installation time. A recent article provided a list of more than 73,000 cameras that use standard communication protocols and a strong encryption mechanism, yet are not immune to security attacks because they use a default password for encrypting their communications [13].

4. Using short sensor identifiers: In order to distinguish different sensors from each other and enhance the security of the communication protocol, most sensors include their identifier (also referred to as the pin code or identification number) in all packets they transmit over the communication channel. The base station uses the sensor's identification code to ensure that the incoming packet comes from one of the already-registered sensors, which are known to the system. However, as demonstrated later in Section 5.2, several in-market sensors from well-known manufacturers use very short sensor identifiers (4-8 bits). As a result, they are susceptible to brute-force attacks (i.e., attacks consisting of systematically checking all possible sensor identifiers until the correct one is found).

3.2.2 *Offering inessential sacrifices*

As mentioned earlier, upon the detection of an emergency situation, automation/monitoring CPSs willingly sacrifice some of their security mechanisms to ensure users' safety. As an example, fire evacuation systems open all doors to enable firefighters to access different rooms and allow the occupants to safely leave the building. Although this evacuation mechanism seems essential to ensuring occupant safety, it might enable an attacker to access restricted areas by triggering an emergency response. This example demonstrates that designers should take both safety and security considerations into account, when designing emergency responses of a CPS.

3.2.3 *Relying on a single sensor type*

In order to provide a reactive emergency response when required, the automation/monitoring CPS must be able to correctly distinguish abnormal situations from normal ones. In fact, the most important steps in minimizing safety risks is detecting emergency situations. Therefore, before providing any response, the CPS needs to collect sufficient sensory data to obtain a clear description of its environment. If insufficient information is given to the system, it might fail to correctly recognize emergency situations. Unfortunately, in order to minimize costs, the majority of the already-in-use automation/monitoring CPSs only process a single environmental attribute. For example, consider a fire evacuation system that only relies on smoke detection sensors. Such a system provides an emergency response when at least one of the sensors detects the existence of smoke. Thus, an attacker can easily trigger the emergency response by only targeting a single vulnerable smoke detection sensor.

4 Potential Consequences of Launching DISASTER

As mentioned earlier, CPSs are in widespread use and handle sensitive tasks in various application domains. Hence, launching tailored attacks, like DISASTER, that are applicable to various forms of automation/monitoring CPSs, can lead to serious consequences. Such consequences depend on the type of emergency response activated by the attack. Generally, the negative impact of triggering an active response is more significant than the impact of triggering a passive response due to the fact that the former can actively control various critical operations and even bypass human operators' decisions. Next, we describe possible consequences associated with launching DISASTER.

4.1 Life-threatening conditions

Triggering the emergency response of a CPS that handles critical operations, e.g., medical or industrial automation tasks, can lead to serious life-threatening conditions. This can range from conditions affecting an individual to those affecting a large number of people. For example, consider an insulin delivery system that is equipped with a safety mechanism, which monitors the blood glucose and immediately stops the injection procedure when it detects the patient has hypoglycemia, i.e., a life-threatening low level of blood glucose. Triggering the active emergency response

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMSCS.2017.2720660, IEEE Transactions on Multi-Scale Computing Systems

of the insulin delivery system immediately shuts down the device. An attacker might be able to trigger such an active response even when the blood glucose level is normal/high to halt the delivery system and cause hyperglycemia, i.e., a life-threatening high level of blood glucose.

4.2 Economic collateral damage

Economic damages refer to monetary losses, including, but not limited to, loss of property, machinery, equipment, and business opportunities, costs of repair or replacement, the economic value of domestic services, and increased medical expenses. Almost all emergency responses have associated costs, even when they are triggered by a real emergency situation. For example, consider a vehicular CPS that is designed to inflate the vehicle's airbags to provide protection in the event of a collision. If a minor collision that results in the deployment of airbags, the whole dashboard panel, steering wheel, and all airbags have to be replaced. In such cases, the active emergency response provided by the CPS is quite costly. Therefore, if an attacker can activate such an emergency response, he will be able to cause collateral economic damage.

The cost associated with triggering emergency responses varies significantly from one application domain to another. For example, the economic collateral damage that results from triggering an emergency response of a vehicular CPS is much less than that of an industrial CPS that controls a manufacturing line.

4.3 Overriding access control mechanisms

In the presence of an emergency situation, a CPS might also be able to command the access control systems that control which users are authorized to access different restricted areas. Such a control is important for two reasons. First, the CPS can facilitate the evacuation procedure in the case of an emergency. Second, the system can lock down particular areas to prevent an intruder from escaping. For example, a residential CPS, which is able to control door locks, may open the main entrance to ensure that firefighters can easily access restricted areas and residents can safely evacuate the building. However, if an attacker can trigger this safety mechanism in the absence of an emergency situation, he might be able to bypass physical security mechanisms and access restricted areas.

4.4 Unintended ignorance

As mentioned earlier, CPSs provide both passive and active responses to minimize damage associated with an emergency situation. A majority of emergency responses could be extremely annoying to the proximate people if they are activated in the absence of an emergency situation. For example, upon the detection of an emergency situation, a great number of CPSs activate notification appliances, e.g., electromechanical horn and speaker, that generate a high-pitched noise to inform the nearby people about the need to take immediate action. If an attacker launches DISASTER that activates notification appliances several times in a short time frame, the system administrator might be convinced that the system is faulty and turn off the emergency responses. This might lead to serious safety/security risks for the duration the emergency responses remain off.

5 Launching DISASTER

In this section, we demonstrate the feasibility of launching DISASTER in realistic scenarios, e.g., residential and industrial automation/monitoring systems. As mentioned in Section 3.2, communication protocols utilized in wireless sensors commonly have various security weaknesses. Next, we first briefly describe two well-known types of attacks that exploit security weaknesses of communication protocols to create and transmit illegitimate packets. Second, we demonstrate how an attacker can tailor these generic forms of attacks to trigger the emergency responses of safety mechanisms and endanger both user safety and system security.

5.1 Creating and transmitting illegitimate packets

Here, we briefly describe two attacks against sensors that enable the attacker to send illegitimate packets to the base station. In both attacks, we use GNURadio [14], a development toolkit that can be used along with an external radio frequency (RF) hardware, e.g., HackRF [12], to implement various software-defined transmitters/receivers that are implemented as software programs to control external RF devices.

Attack 1: Retransmitting recorded packets

In this approach, an attacker aims to record data packets and retransmit them to the base station without processing or modifying their contents. To do so, the attacker first builds an RF receiver that listens to the communication channel between sensors and the base station and records the transmitted packets. Then, he uses an RF transmitter, which can retransmit previously-recorded packets on the same communication channel. If the frequency of the communication channel is known to the attacker, he can implement the above-mentioned receiver/transmitter using the built-in libraries of GNURadio.

The frequency of the communication channel can be extracted from documents submitted to Federal Communications Commission (FCC). FCC is an independent U.S. government agency that tests all wireless products sold in the U.S. and provides a public database, which includes test reports and documentations of the products [15]. In order to find the frequency of the communication channel used by a sensor, the attacker only needs to find the sensor's documentations by searching its FCC code (i.e., an identification code that specifies the sensor's manufacturer and type) in FCC's public database. FCC codes are commonly written on the sensor's cover.

Attack 2: Reverse engineering the communication protocol

In this approach, the attacker records several data packets and processes them to explore how the communication protocol sends digital data over the communication channel. We describe next how an attacker can reverse-engineer the communication protocol of an arbitrary sensor using HackRF and GNURadio.

1. The attacker first obtains the communication frequency of the sensor and records several packets from the sensor using the method discussed in the previous approach.

2. Then, the attacker finds the modulation type of the communication protocol. The two most commonly-used

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMSCS.2017.2720660, IEEE Transactions on Multi-Scale Computing Systems

communication protocols used in wireless sensors are on-off keying (OOK) and binary frequency-shift keying (BFSK). OOK is the simplest form of amplitude-shift keying modulation in which digital data are presented as the presence/absence of a carrier wave, and BFSK is the simplest form of frequency-shift keying (FSK) that uses a pair of discrete frequencies to transmit binary digital data. The modulation type of a communication protocol can be easily detected by examining the Fourier transform of a packet received by HackRF. A single peak (two discrete peaks) in the Fourier transform represents an OOK (BFSK) modulation.

3. After finding the modulation type of a communication protocol, the attacker implements a software-defined demodulator in GNURadio that extracts the transmitted digital data from the recorded analog signal. Fig. 2 demonstrates the implementation of an OOK demodulator in GNURadio.

4. If the sensor does not support any encryption mechanism, the attacker can easily examine the digital data to determine what each bit represents, and how he can generate seemingly legitimate packets with arbitrary content.

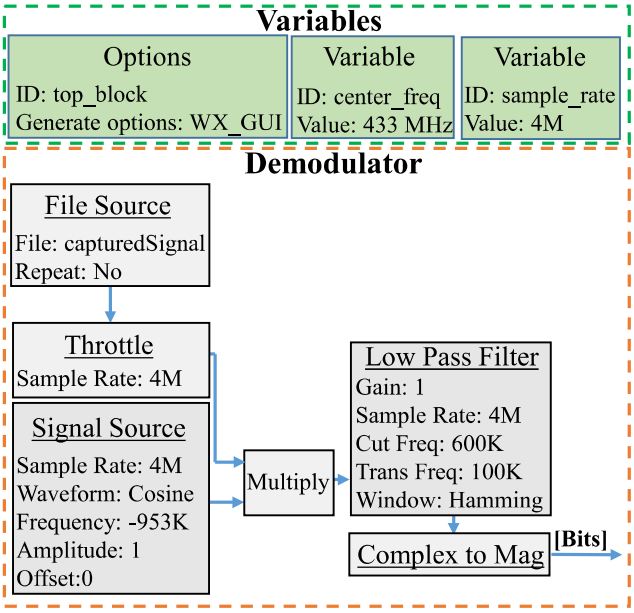

Fig. 2. The implementation of an OOK demodulator in GNURadio

This approach provides two fundamental advantages over the previous approach. First, the attacker does not need to capture any packets from the actual sensors that he targets. In fact, he can conduct several experiments in a test environment, e.g., a laboratory, to discover potential vulnerabilities in the system. Second, the attacker can improve the quality of the transmitted signal and increase its signal-to-noise ratio in an attempt to launch an attack against the system from a large distance.

5.2 DISASTER case studies

In this section, we demonstrate the feasibility of launching DISASTER against the two most widely-used automation/monitoring CPSs: residential and industrial.

5.2.1 Case I: Residential automation/monitoring CPSs

We first briefly discuss what a typical residential CPS does, what types of sensors are commonly used in such a system, and what emergency responses it offers in the presence of an emergency. Second, we demonstrate how an attacker can use the two well-known generic types of attacks discussed in Section 5.1 to trigger the emergency responses of residential CPSs in real-world scenarios.

Residential CPSs, their services and emergency responses

A residential CPS is mainly designed to offer physical security mechanisms, enhance residents' convenience, and minimize energy consumption. It can also offer increased quality of life to the residents who need special assistance, e.g., the elderly and disabled people. It processes the data collected by various sensors to control lighting, heating, and cooling, and to monitor/command the security locks of gates and doors. In addition, the base station continuously collects and processes real-time environment-related data gathered by sensors to detect emergency situations. Table 1 includes different sensors that are commonly used in residential CPSs, a short description of each sensor, and services that rely on each sensor.

State-of-the-art residential CPSs are able to detect two common emergency situations: fire and ongoing burglary, and provide three typical emergency responses: two passive and one active. Next, we elaborate on these responses and discuss the negative consequences of activating each response.

Passive response I: Activating warning devices: In the presence of an emergency situation, notification appliances, e.g., flashing lights, electromechanical horns, or speakers, are activated to warn the proximate people about the need to take immediate action. A majority of notification appliances generate a high-pitched sound to attract the attention of those nearby. The generated sound could be extremely annoying if it is activated in the absence of an emergency situation. Hence, if a potential attacker can trigger this emergency response several times in the absence of an emergency situation, residents might be convinced that the system is faulty and turn off the emergency response. This might lead to serious safety/security risks and concerns. For example, a burglar might try to trigger the anti-theft alarm several times in a short period of time, e.g., in an hour, in the hope of convincing the user to turn off the monitoring system.

Passive response II: Informing police/fire department: Requesting immediate help from the police/fire department, when a real threat is not present, puts firefighters, police officers, as well as the public at risk by needlessly placing heavy, expensive equipment on the streets while wasting fuel and causing traffic jams. Moreover, if an attacker can initiate a help request several times in a short period of time, he might be able to persuade firefighters, police officers, and occupants to believe that when an alarm goes off it is likely a false alarm. As demonstrated later, DISASTER can be launched from a large distance (e.g., over 100 m from the base station). Therefore, an attacker might be able to design a large-scale attack (e.g., he can launch DISASTER using HackRF, while driving in a residential community, to

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMSCS.2017.2720660, IEEE Transactions on Multi-Scale Computing Systems

TABLE 1
Different sensors used in a typical residential CPS, their descriptions, and services

Sensor type	Description	Service
Humidity sensor	measures moisture content in the environment	Heating/cooling
Temperature sensor	measures the current temperature of the room	Heating/cooling
Light sensor	measures the luminance in the environment	Lighting
Motion sensor	detects the presence of a person in the environment	Lighting and anti-theft mechanism
Door sensor	checks if a door has been opened	Anti-theft mechanism
Smoke detector sensor	detects the presence of smoke/fire in the environment	Fire detection

trigger the alarm systems of all houses in the community) to impose significant additional cost on both residents and the responsible governmental department, and convince the residents to turn off their security/safety alarms.

Active response: Controlling door locks: As mentioned earlier, a residential automation/monitoring system may be able to automatically control the locks upon the detection of a fire or burglary. In the presence of a fire, it opens the main doors/entrances to ensure that firefighters can enter the affected areas and residents can safely evacuate the building. Moreover, in the presence of an ongoing burglary, it locks the main entrances to ensure that the thief is not able to leave the crime scene until police officers arrive. Although this emergency response is offered to minimize potential safety risks, triggering this response by an attacker in the absence of an emergency situation could lead to serious security issues. For example, if the attacker triggers the fire evacuation procedure, he will be able to bypass the physical security mechanism of the building by unlocking main entrances. Similarly, the attacker might be able to confine the residents inside the house by initiating the anti-theft lock-down procedure.

Demonstration of DISASTER against residential CPSs

In order to examine the feasibility of launching DISASTER against residential CPSs, we developed two experimental scenarios using the approaches described in this section. In both scenarios, we targeted three types of sensors (highlighted in red in Table 1). The residential CPS processes the data gathered by these sensors to detect emergency situations (fire or burglary). In our experimental setup, we closely inspected six already-in-market sensors (two motion detectors, three smoke detectors, one door sensor) made by well-known manufacturers that cater to the home automation industry. Since these sensors are deployed in numerous already-in-use systems, we choose not to disclose their brand and model number in this paper.

Next, we describe how the two previously-mentioned generic attacks can be used to launch DISASTER and activate emergency responses of the system.

Experimental scenario 1: Retransmitting packets: In our experimental setup, we captured and retransmitted 20 packets from each sensor (120 packets in all) using the software-defined transmitter/receiver described in Section 5.1. Fig. 3 demonstrates a packet generated by the door sensor and captured by HackRF. The base station of all the under-experiment sensors accepted previously-recorded packets. This indicates that the packets generated by these sensors include neither a timestamp nor a sequence number. Thus, an attacker can record a packet from each of these sensors and retransmit it to the base station of the CPS that utilizes the sensor in an attempt to trigger emergency responses.

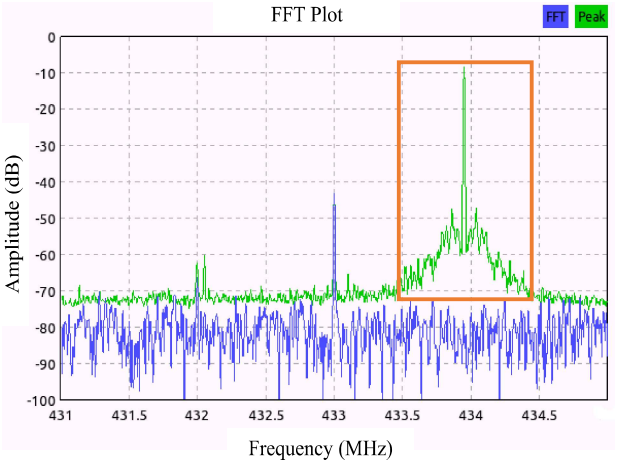

Fig. 3. The door sensor generates a packet as soon as it detects the door is open. The spike in the fast Fourier transform of the analog signal shows a single transmission using OOK modulation.

In this experimental scenario, we placed HackRF at different distances from each under-experiment sensor to find the maximum recording distance from which the attacker can record a packet using HackRF. Moreover, we increased the distance between HackRF and the base station to find the maximum retransmission distance for each sensor from which a previously-captured packet can be received and accepted by the sensor's base station. Table 2 summarizes results of this experiment for different sensors.

TABLE 2
Maximum recording distance and maximum retransmission distance for each sensor in Experimental scenario 1

Sensor	Maximum recording distance (m)	Maximum retransmission distance (m)
Motion sensor I	58	54
Motion sensor II	110	105
Smoke detector I	67	50
Smoke detector II	52	55
Smoke detector III	54	48
Door sensor	56	54

Sensors that enable the anti-theft mechanism (motion and door sensors) transmit data very frequently even when

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMSCS.2017.2720660, IEEE Transactions on Multi-Scale Computing Systems

the mechanism is completely disabled. The motion sensor transmits a packet when it detects a moving object, and the door sensor transmits a packet when it detects an open door. For such sensors, an attacker can simply capture a packet when the mechanism is disabled, e.g., when the residents are inside, and retransmit the packet when it is enabled. Unlike motion/door sensors, smoke sensors rarely transmit a packet to the base station since their event-driven transmission protocol only transmits a packet to the base station when an actual threat, e.g., a fire, is present. Thus, capturing and retransmitting a packet that is generated by smoke sensors are difficult and potentially very time-consuming for the attacker. In the second experimental scenario, we discuss how the attacker can reverse-engineer communication protocols deployed in smoke detectors to easily launch DISASTER against the system.

Experimental scenario 2: Reverse engineering: As mentioned earlier in this section, in this approach, the attacker records and demodulates transmitted packets in a test environment, e.g., a laboratory, to examine how a sensor transfers digital data over the communication channel.

We examined the communication protocol used in the six under-experiment sensors. The examination revealed that all sensors share a common security weakness: in order to provide a cost-effective solution, the manufacturers used very simple non-standard transmission protocols that do not provide any cryptographic mechanism. Indeed, the packets transmitted from these sensors to their base stations are neither cryptographically protected nor completely obfuscated. Fig. 4 demonstrates the bitstream transmitted by a door sensor to the base station of a residential CPS. The analog signal (Fig. 3) captured by HackRF is demodulated using the OOK demodulator (Fig. 2). This sensor repeatedly (20 times) transmits a single static packet (that does not change over time), which includes its 4-bit pin number (a very short sensor identification code), to its base station.

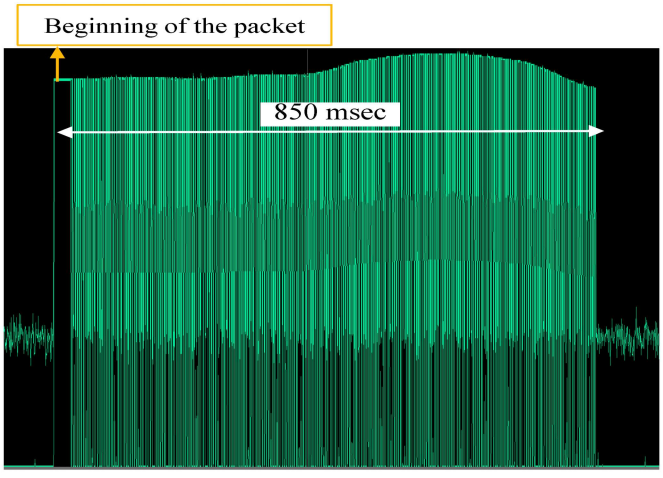

Fig. 4. The bitstream transmitted by the door sensor to the base station of the residential CPS. The door sensor repeatedly transmits a single static packet, which includes its 4-bit pin number, to its base station.

We were able to completely reverse-engineer the six communication protocols used in these sensors and determine that all of them generate static packets, which include fixed pin numbers. In other words, the only unknown field in the bitstream of an arbitrary packet generated by the sensors was the device's pin number. Table 3 specifies the communication frequency, modulation type, and pin length of each under-experiment sensor.

TABLE 3
Communication frequency, modulation type, and pin length of each residential sensor

Sensor	Communication frequency (MHz)	Modulation	Pin length (bits)
Motion sensor I	433	OOK	4
Motion sensor II	433	OOK	8
Smoke detector I	920	OOK	8
Smoke detector II	920	OOK	8
Smoke detector III	433	OOK	8
Door sensor	433	OOK	4

After completely reverse engineering the protocol, we implemented a brute-force attack using different possible values of the pin numbers of the sensors. We were able to find the actual pin numbers of all sensors in less than five seconds. Then, we placed HackRF at different distances from the base station of the under-experiment sensors and used the maximum transmission power of HackRF to determine the maximum transmission distance from which this attack is possible. Table 4 summarizes the results of this experiment for different sensors.

TABLE 4
Maximum transmission distance for each residential sensor in Experimental scenario 2

Sensor	Maximum transmission distance (m)
Motion sensor I	75
Motion sensor II	85
Smoke detector I	75
Smoke detector II	80
Smoke detector III	80
Door sensor	75

5.2.2 Case II: Industrial automation/monitoring CPS

We first briefly discuss different services and emergency responses offered by a typical industrial CPS. Second, we demonstrate how an attacker can trigger the emergency responses of industrial CPSs.

Industrial CPSs, their services and emergency responses

A typical industrial automation/monitoring CPS offers various automatic mechanisms to operate the equipment, e.g., machinery and boilers, with minimal human intervention, and several approaches that enable remote monitoring of the industrial environment. Generally, industrial automation/monitoring CPSs deal primarily with the automation of manufacturing, quality control, and material-handling processes. In addition, almost all modern industrial automation/monitoring CPSs continuously monitor the environment to detect emergency situations. These situations need to be aggressively addressed due to the fact they can be catastrophic in a large industrial setting.

State-of-the-art industrial CPSs are able to detect a variety of emergency situations, e.g., a tank overflow,

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMSCS.2017.2720660, IEEE Transactions on Multi-Scale Computing Systems

system failure, or a fire. Upon the detection of an emergency situation, they commonly provide four emergency responses, including two passive responses and two active responses. The two passive responses are similar to the ones provided by residential CPSs. Hence, we discuss the two active responses that are commonly offered by industrial CPSs and discuss the negative consequences of activating each response.

Active response I: Halting normal operation

A halting procedure is initiated to shut down a part, e.g., a plant or control unit, of the industrial setting or the whole production line when necessary. Upon the detection of an emergency situation, the centralized base station responds by placing the controllable elements, e.g., valves and pumps, into a safe state. For example, a halting procedure controls valves to stop the flow of a hazardous fluid or external gases upon the detection of a dangerous event. This provides protection against possible harm to people, equipment or the environment. Launching DISASTER against an industrial CPS, which activates the halting procedure, may lead to two consequences: production loss and profit penalty. A halting procedure shuts down specific units or the entire facility. This can lead to a significant production loss in chemical industries, e.g., gasoline-centric refinery, where shutting down a unit may stop chemical reactions from completing. Moreover, several time-consuming safety checks need to be done before restarting the normal operation. Thus, the facility might need to be shut down for a substantial amount of time. This could cause a significant impact on profits. For example, an average-sized U.S. Gulf Coast oil refinery loses 68,000 dollars a day for a downed unit [16].

Active response II: Initiating a damage control mechanism

Damage control mechanisms include any prudent action aimed at preventing/reducing any expected damage to the industrial setting, stabilize the situation caused by the damage or alleviate the effects of damage. The main purpose of damage control is to offer a way to return the production line to its normal operation with minimal loss of property or life. A common damage control mechanism in an industrial environment is automatic fire suppression, which employs a combination of dry chemicals and wet agents to extinguish a fire. It applies an extinguishing agent to a three-dimensional enclosed space in order to achieve a concentration of the agent that is sufficient to suppress the fire. A fire suppression system that primarily injects gases into enclosed spaces presents a risk of suffocation. Numerous incidents have been documented where individuals in such spaces have been killed by carbon dioxide agent release [17], [18]. Moreover, the positive pressure caused by these gases may be sufficient to break windows and even walls and destroy the surrounding equipment. Thus, launching DISASTER against an industrial CPS that triggers its fire suppression mechanism may lead to serious consequences, ranging from severe damage to the equipment to life-threatening conditions.

Demonstration of DISASTER against industrial CPSs

In order to investigate the feasibility of launching DISASTER against industrial CPSs, we targeted two industrial systems that use level sensors to monitor liquid level changes in storage tanks (Fig. 5). Level monitoring-based safety mechanisms are commonly used in various industrial environments, e.g., the oil industry, to detect an emergency situation that is called tank overflow.

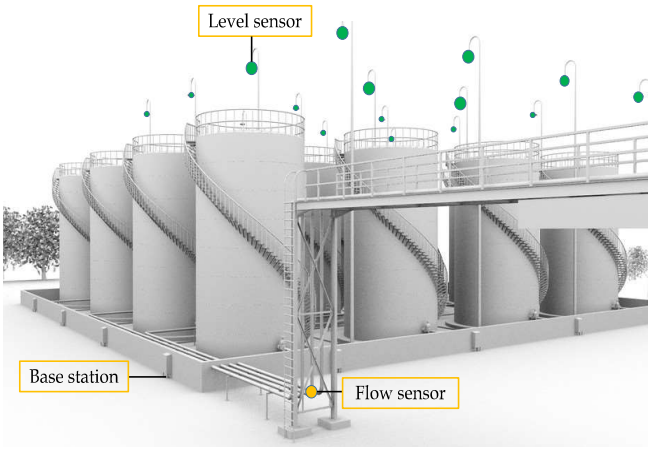

Fig. 5. A simple industrial automation/monitoring CPS

To activate the emergency responses of such a system, the attacker can transmit an illegitimate packet to the base station that indicates that the storage tank is full. Using each of the two approaches described in Section 5, an attacker can generate such packets. However, capturing and retransmitting level sensors' regular data packets (i.e., the packets that are periodically transmitted to the base station to report the level of liquid) cannot activate the system's emergency responses. In fact, the attacker needs to record a packet in the presence of a real emergency situation, which is extremely rare in real-world industrial environments, and use it later. Thus, the first approach may not be practical for launching DISASTER against industrial systems described above.

A close examination of two commonly-used industrial level monitoring-based CPSs revealed that none of their sensors uses a secure transmission protocol. Indeed, communications between the sensors and their corresponding base stations are not cryptographically-protected. Therefore, we were able to completely reverse-engineer the communication protocols used in these sensors. To do this, we captured and demodulated 40 data packets (80 packets in all) generated by each sensor. Fig. 6 demonstrates the bitstream transmitted by one of the level sensors.

We found that one of the level sensors simply transmits an unecrypted packet that includes a 10-bit pin number and a data field that represents the liquid level in the tank. In order to trigger the emergency response of the CPS that uses this sensor, we transmitted 1024 packets with different pin numbers. We were able to trigger the emergency alarm in less than 10 seconds for this sensor.

We observed that the other sensor transmits unencrypted packets that only contain data and a 1-byte sequence number. Therefore, in order to trigger the emergency response of

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMSCS.2017.2720660, IEEE Transactions on Multi-Scale Computing Systems

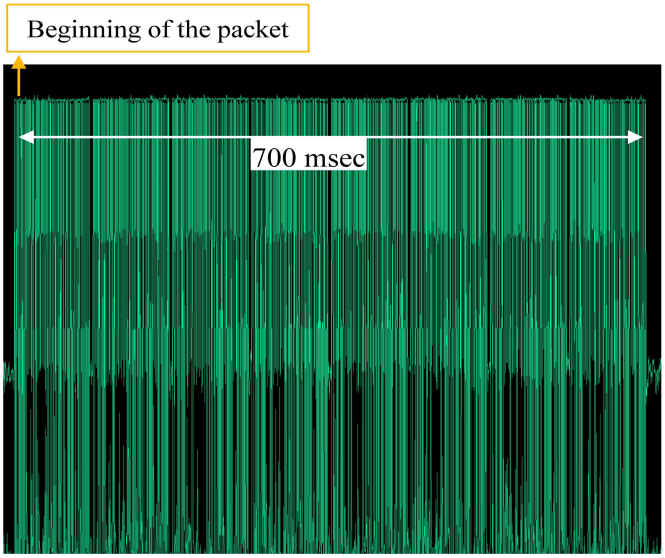

Fig. 6. The bitstream transmitted by one of the level sensors to the base station of the industrial CPS.

the CPS that uses this sensor, an attack can be implemented as follows. The attacker first captures a packet from the sensor and extracts the sequence number. Then he creates and transmits a packet in which the data field is set to its maximum value and the sequence number field is set to the sequence number extracted from the captured packet plus one. We implemented this attack and were able to successfully trigger the emergency responses of the system.

Table 5 specifies communication frequency, modulation type, and pin length for each level sensor.

TABLE 5
Communication frequency, modulation type, and pin length for each level sensor

Sensor	Communication frequency (MHz)	Modulation	Pin length (bits)	Sequence number
LS I	433	OOK	10	No
LS II	920	OOK	0 (no pin)	Yes (1 byte)

Moreover, we placed HackRF at different distances from the base stations of the two sensors to find the maximum transmission distance. Table 6 summarizes results of this experiment. In real-world industrial CPSs, where signal repeaters (i.e., electronic devices that receive signals and retransmit them at a higher power) are in widespread use to support long-range communications, sensors may be located several miles from the base station. Unfortunately, the attacker can also exploit these repeaters to extend the attack range up to tens of miles, e.g., the attacker can place a HackRF hundred meters away from a repeater that is located several miles away from the CPS base station under attack.

TABLE 6
Maximum transmission distance for each industrial level sensor examined in Experimental scenario 2

Sensor	Maximum transmission distance (m)
LS I	70
LS II	250

6 Suggested Countermeasures

In this section, we first suggest three approaches to mitigate the consequences of launching DISASTER, and for each approach, we briefly describe its limitations and disadvantages. Second, we discuss why preventing DISASTER may not always be feasible due to the existence of unpredictable situations.

6.1 Proactive countermeasures

Next, we describe three proactive approaches, which can be deployed in the design and verification phases of manufacturing, to prevent DISASTER.

6.1.1 Utilizing cryptographic mechanisms

As demonstrated in Section 5.2, using simple customized communication protocols to provide short-range communication between sensors and the centralized base station can lead to serious security issues and enable an attacker to reverse-engineer the protocol. The main weakness of the majority of non-standard communication protocols is that they do not offer any cryptographic mechanisms to ensure confidentiality and integrity of data transmitted by sensors. Utilizing standard communication protocols, which provide strong encryption mechanisms to ensure confidentiality and integrity, or adding encryption mechanisms to customized communication protocols can significantly limit the ability of an attacker to launch a security attack against the system. Bluetooth [19] and ZigBee [20] are instances of standardized communication protocols that offer lightweight encryption mechanisms (e.g., encrypted timestamp and data encryption), yet provide a long battery lifetime. Despite the advantages of encryption mechanisms, as described later in Section 8.2, utilizing strong encryption in the sensors used in CPS safety mechanisms may not be feasible due to several domain-specific limitations of safety system, e.g., latency requirements.

In addition to encryption, obfuscation (i.e., a procedure applied to data to intentionally make them hard to understand without knowing the procedure that was applied) can make reverse engineering of the protocol harder. However, obfuscation cannot truly secure the system since a skilled attacker may eventually be able to reverse-engineer the obfuscation procedure. In fact, obfuscation can only delay (not prevent) reverse engineering of the communication protocol.

6.1.2 Security/safety-oriented verification

As mentioned earlier, common design flaws and security weaknesses of safety mechanisms along with ignorance of common security-safety conflicts/trade-offs can endanger both the security and safety of the system. There is a great body of literature on different types of design verification approaches and several commercialized verification approaches that manufacturers can use to detect and address common design flaws before introducing their product into the market. Such verification approaches have traditionally

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMSCS.2017.2720660, IEEE Transactions on Multi-Scale Computing Systems

been used to ensure that a product, service, or system works correctly, meets specific requirements, and fulfills its intended purposes. Unfortunately, traditional verification mechanisms do not typically target a comprehensive set of security and safety requirements. Recently, a few verification approaches [21]–[24] have begun to take various security and safety considerations into account. Moreover, a few proposals [25] have offered theoretical approaches to detect different safety-security conflicts/trade-offs in CPSs.

Utilizing such newly-proposed approaches can enable designers and manufacturers to predict or detect design flaws, security weaknesses, and safety-security conflicts before releasing a product to the market. However, relying on such methods cannot completely prevent DISASTER due to: (i) the existence of serious disagreements between some security and safety requirements that forces the manufacturer to consider some requirements and ignore others, (ii) inefficiency and imperfections of verification mechanisms, and (iii) unpredictability of threats against CPSs.

6.1.3 Designing multimodal systems

In order to plan and execute an appropriate emergency response, a CPS must reason about its surroundings and obtain a precise description of the environment. Such a system can gather and process data from its surroundings using various types of sensors. Generally, a CPS that uses multiple types of sensors (referred to as a multimodal system) can obtain more information about the environment than a CPS that only relies on one type (referred to as a unimodal system). As a result, multimodal CPSs provide two fundamental advantages over unimodal systems. First, they offer a higher accuracy in detecting emergency situations due to the fact that they can obtain more information about the environment and utilize sensor fusion methods (that combine various sensory data to improve the resolution and accuracy of specific sensor data) to achieve an accurate description of the environment. Second, multimodal CPSs are typically more difficult to attack since the attacker has to simultaneously target several sensor types to launch DISASTER. For example, we have designed and implemented a simple multimodal residential CPS that offers theft and fire detection. Its safety system uses motion detection sensors in conjunction with door sensors for theft detection and smoke detectors and temperature sensor for fire detection. In fact, emergency responses will only be activated if the door sensor of a room indicates that the door is open and the motion detector of the same room detects the presence of an intruder. Similarly, the fire evacuation procedure is initiated only if the smoke is detected in a room and the temperature of the same room is quite elevated.

However, multimodal CPSs generally have two disadvantages over unimodal systems: (i) multimodal systems are more complex and expensive, (ii) since the multimodal systems need to process more information, they are slower in detecting emergency situations. Moreover, if all sensors utilized in a multimodal CPS have common design flaws and weaknesses, e.g., all sensors use the same communication protocol that support neither obfuscation nor encryption, launching DISASTER against the multimodal system may not be significantly harder than launching an attack against a unimodal system.

6.1.4 Domain-specific policies

Depending on the application domain, designers may be able to develop fine-grained policies and algorithms that enhance the security of the system with a negligible negative impact on the safety of the system. For example, for the fire evacuation system, designers can divide different physical spaces within a building into two categories: high-security and low-security, and set additional security policies for high-security environments. Upon the detection of a fire, the safety system only opens all doors of low-security spaces, however, it opens the door of the high-security spaces if it detects the presence of an occupant or authorized emergency responders, e.g., a firefighter carrying an approved badge.

6.2 Unpredictable situations

Although the above-mentioned approaches can significantly limit the ability of a potential attacker, providing a comprehensive solution for eliminating the security risk associated with DISASTER is hard for two reasons. First, DISASTER might be feasible as a result of the existence of a situation that is not predictable at design time. Second, modeling human errors, e.g., pitfalls in the installation procedure, which might make a system susceptible to the proposed attacks, is very difficult. For instance, suppose a company sends certified installers to install a residential automation system that provides a fire evacuation mechanism, which is able to unlock the doors upon the detection of a fire. A month later, another company installs an air conditioner unit that is accessible from outside the building. In this scenario, an attacker might be able to inject smoke into the residence using the routing paths of the air conditioner and trigger the fire evacuation mechanism even when there is no fire. In fact, the air conditioner provides the means to a potential attacker to have physical access to the home automation system and trigger its emergency responses.

7 Related Work

In this section, we briefly discuss related work. We first describe traditional CPSs, their architecture, and their main components. Then, we discuss previous attacks against CPSs, in general, and attacks against safety systems utilized in CPSs, in particular.

7.1 Traditional residential/industrial CPSs

Even before the introduction of state-of-the-art CPSs, several types of control systems were used in different residential and industrial settings for automatically operating equipment. Traditional systems had three main components: input devices, centralized control units, and output devices. Input devices include analog devices that provide an analog input, e.g., voltage and current, digital devices that provide on-off status input, e.g., two-state switches, data devices that provide data inputs like text and binary inputs, e.g., barcode readers, and operator interface devices that enables human operators to interact with the control system. Control units were commonly a proportional-integral-derivative (PID) controller, i.e., a controller that calculates an error value as the difference between a desired point and a

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMSCS.2017.2720660, IEEE Transactions on Multi-Scale Computing Systems

measured variable, and controls output devices to minimize the error, or a simple programmable logic controller (PLC), i.e., an industrial digital computer that has been optimized and adapted for the control of manufacturing processes. In early systems, these components were connected via communication protocols that were based on wire links, e.g., Fieldbus [26]. In such systems, emergency responses were usually very limited, e.g., turning on/off a machine or opening a valve, and were activated by human operators (using two-state red emergency buttons) or simple control systems.

Due to the absence of wireless technologies and network connectivity in early automation systems, accidents, inappropriate user activity, and disappointed employees accounted for most of the problems (see [27] for a survey). The emergence of CPSs and the incorporation of wireless connectivity, which offers numerous benefits to automation systems and enables a variety of new applications, has introduced several new challenges, e.g., reliability, availability, and security. In this paper, we focused on a new security concern associated with the use of insecure wireless sensors, which do not commonly tolerate the additional energy/delay overhead of cryptographic operations, in safety-related operations.

7.2 Previously-proposed attacks against CPSs

The incorporation of wireless modules and Internet-connected base stations in CPSs significantly boosted the market for residential and industrial CPSs since it enabled vendors to manufacture more cheaply and users to install more easily, with minimal environmental modification.

Early Internet-connected base stations transmitted data to back-end servers over unencrypted channels, e.g., a web-based industrial CPS that uses unencrypted channels [28]. Thus, they were vulnerable to a variety of remote security attacks, e.g., eavesdropping and integrity attacks. However, since base stations usually had sufficient computational and energy resources, several cryptographic mechanisms were proposed for preventing such attacks. In state-of-the-art residential CPSs, the communication link with the outside world, e.g., Cloud servers, is encrypted. Although encryption can prevent a variety of security attacks, it has been shown that even encrypted packets may leak private information about the operation of CPSs and their users [29], [30], e.g., if a residential CPS includes a security camera that only transmits a snapshot upon detection of a movement, the transmission of packet from the CPS to a specific server reveals the presence of occupants inside the house [29].

As we discovered in our experiments and by examining different documentations of in-market products, the link between battery-powered wireless sensors (in particular, sensors used in safety systems of CPSs) and the base station is commonly established based on customized non-standard protocols. Although it is a well-known fact that unencrypted channels offer a weak link to potential attackers, such channels were assumed to be secure enough for the majority of safety operations since the main objective of attackers was assumed to be disabling the emergency responses (as opposed to deliberately enabling them). In fact, real-world instances of attacks against the safety system utilized in CPSs have confirmed that attackers were interested in directly disabling the emergency responses, commonly by jamming communications. As a result, several previous studies [7], [31] on security of safety systems focused on how attackers can bypass the safety operation by jamming the communication channels and how designers can design low-cost anti-jamming circuitry for base stations to detect jamming attacks. In addition, it has been shown that malicious computer worms can be designed, e.g., Stuxnet [5], [6], to target PLCs and disable emergency shutdown in an emergency. Unlike previous studies, in this paper, we focused on attacks in which the attack intentionally activates the emergency responses in the absence of an emergency. We discussed why potential attackers may want to active emergency responses, how such attacks are possible in real-world scenarios, and why traditional solutions may fail to address these attacks.

8 Discussion

In this section, we discuss two items not yet explained in detail. First, we describe how pervasive use of the Internet-connected device paradigm may enable new IoT-enabled DISASTER from a remote distance. Second, we discuss the impact of widespread use of insecure sensors in state-of-the-art safety systems utilized in CPSs.

8.1 Feasibility of IoT-enabled DISASTER

The pervasive use of Internet-connected devices is one of the most spectacular phenomena of the last decade. The miniaturization of transceivers, along with reduced costs and sizes of on-device resources, has offered the opportunity to transform isolated devices into communicating things and led to the emergence of IoT in the last decade. Traditional CPSs are being transformed as more Internet-connected base stations and devices get added to such systems. As a side effect, the number of potential threats and possible attacks against the security of IoT-enabled CPSs is growing exponentially [32]. Although it has been suggested that the processing units of automation CPSs can be moved to Cloud services [33], [34], manufacturers have been skeptical about offloading data to Cloud servers for making mission-critical decisions due to the security, safety, and latency requirements of automation systems, in particular industrial systems. Furthermore, the potential unavailability of Internet connectivity has limited the use of Cloud servers in automation CPSs. As a result, the number of Cloud-based automation services is very limited in CPSs, in particular for safety operations. In fact, even in state-of-the-art residential/industrial CPSs, latency-sensitive safety-related operations are performed locally and safety-related components are not directly connected to the Internet. However, CPSs do commonly transfer monitoring-related data to Cloud servers. Thus, in the current state of the technology, attackers cannot easily take control of safety systems from a remote distance over the Internet.

To launch DISASTER, the effective methods are currently limited to attacks against the wireless communication of insecure sensors or infecting the control system via injection of a malware using a physical device, e.g., a USB stick [35].

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMSCS.2017.2720660, IEEE Transactions on Multi-Scale Computing Systems

As IoT expands and Internet becomes faster, we expect that more automation operations will be transferred to remote servers and rely on the Internet, providing new attack surfaces for potential attackers who aim to target safety-related operations of CPSs. Although Cloud-based safety operations have not been well-explored, relying on the Cloud for such operations may be useful in some scenarios, e.g., providing remote firmware update. With the emergence of real-world instances of Cloud-based safety operations, new security threats must be taken into account by designers and new methodologies for launching DISASTER over the Internet, e.g., uploading infected firmware to base stations or injecting false data into the Cloud to activate emergency responses or exploiting known vulnerabilities of safety systems remotely, must be studied.

8.2 Inefficiency of state-of-the-art cryptographic operations

While designing sensors for safety-related operations, manufacturers aim to maximize the sensor battery lifetime and, at the same time, minimize the detection latency, i.e., the time between the occurrence of an emergency and its detection by the safety system, and costs. Next, we discuss why even state-of-the-art encryption mechanisms may not be appropriate for safety-related sensors used in CPSs.

Long battery lifetime and low cost may be offered as a valuable feature to attract customers or may be required in some applications. In addition, low detection latency is needed for a majority of safety operations to maximize user safety, in particular in industrial settings. Smoke detectors that offer a ten-year battery lifetime are rapidly replacing older models that require annual battery replacement since a longer battery lifetime reduces the cost of maintenance and enhances user convenience. A typical tire pressure sensor used to detect tire blowouts is expected to work for six to ten years and expected to detect the blowout in few tens of milliseconds.

Although several cryptographic mechanisms (for encryption and authentication) [36]–[38] have been proposed in the literature to enhance the security of wireless sensors used in different CPSs, they impose a significant energy overhead on sensors, require hardware modifications, and/or increase the detection latency. For example, the lightweight secure communication protocol discussed in [38] for safety-related in-vehicle sensors increases both total energy consumption of sensors and detection latency by a factor of 1.5×. Energy overheads decrease the device battery lifetime and lead to increased maintenance costs and user inconvenience. Hardware modification, e.g., increasing memory capacity to support extra processing needed for the encryption, may increase development costs for manufacturers. The detection latency overhead imposed by cryptographic operations can endanger the safety of the users in emergency situations. Adding cryptographic operations (encryption, decryption, and authentication) to resource-limited sensors may significantly increase the time needed for both processing and transmitting the packets (for enabling such operations, extra bits must be added to packets and transmitted).

As mentioned in Section 3.2.1, in addition to the eight sensors (six sensors used in home CPSs and two industrial sensors) that we closely inspected and reverse-engineered, we examined the documentation of 70 sensors (in particular, motion, temperature, liquid level, gas pressure, acceleration, yaw rate, and glucose sensors) from 10 manufacturers. These sensors are widely used in safety-related operations of different CPSs (home automation, industrial, vehicles, medical automation systems). Our examinations confirmed that the most widely-used solution for increasing battery lifetime and decreasing detection latency costs is using customized low-energy wireless protocols with neither encryption nor secure authentication. We realized that a few (3 out of 10) made use of data obfuscation techniques, while others simply implemented communications over non-obfuscated unencrypted channels. Furthermore, we observed that two manufacturers used microcontrollers with a built-in hardware encryption module, however, they did not actually exploit it in their implementation due to encryption overheads.

9 Conclusion

In this paper, we introduced DISASTER, which exploits design flaws and security weaknesses of safety mechanisms deployed in CPSs along with safety-security conflicts to trigger the system's emergency responses even in the absence of a real emergency situation. This can lead to serious consequences, ranging from economic collateral damage to life-threatening conditions.

We examined several already-in-use sensors and listed common design flaws and security weaknesses of safety mechanisms. We discussed the various impacts of DISASTER and described potential consequences of such attacks. We also demonstrated the feasibility of launching DISASTER in realistic scenarios, e.g., residential and industrial automation/monitoring CPSs. We suggested several countermeasures against the proposed attacks and discussed how unpredictable situations may give rise to significant security problems in presumably secure CPSs. We discussed related work and briefly compared our attack to previous ones. Finally, we described the feasibility of IoT-enabled DISASTER and the inefficiency of state-of-the-art encryption.

References

[1] A. A. Cárdenas, S. Amin, B. Sinopoli, A. Giani, A. Perrig, and S. Sastry, "Challenges for securing cyber physical systems," in *Proc. Wkshp. Future Directions in Cyber Physical Systems Security*, 2009.

[2] R. R. Rajkumar, I. Lee, L. Sha, and J. Stankovic, "Cyber physical systems: The next computing revolution," in *Proc. ACM Design Automation Conference*, 2010, pp. 731–736.

[3] C. Neuman, "Challenges in security for cyber physical systems," in *Proc. Wkshp. Future Directions in Cyber Physical Systems Security*, vol. 7, 2009.

[4] N. Veerasamy, M. Grobler, and B. Von Solms, "Building an ontology for cyberterrorism," in *Proc. European Conf. Information Warfare and Security*, 2012, p. 286.

[5] N. Falliere, L. O. Murchu, and E. Chien, "W32. Stuxnet dossier," *White paper, Symantec Corp., Security Response*, vol. 5, 2011.

[6] A. A. Cárdenas, S. Amin, Z. S. Lin, Y. L. Huang, C. Y. Huang, and S. Sastry, "Attacks against process control systems: Risk assessment, detection, and response," in *Proc. ACM 6th Symp. Information, Computer and Communications Security*, 2011, pp. 355–366.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMSCS.2017.2720660, IEEE Transactions on Multi-Scale Computing Systems

[7] L. Lamb, "Home insecurity: No alarms, false alarms, and SIGINT," Oak Ridge National Laboratory, Tech. Rep., 2014.
[8] G. Sabaliauskaite and A. P. Mathur, "Aligning cyber-physical system safety and security," in *Proc. Complex Systems Design & Management Asia*. Springer, 2015, pp. 41–53.
[9] M. Li, W. Lou, and K. Ren, "Data security and privacy in wireless body area networks," *IEEE Wireless Communications*, vol. 17, no. 1, pp. 51–58, 2010.
[10] M. Sun, S. Mohan, L. Sha, and C. Gunter, "Addressing safety and security contradictions in cyber-physical systems," in *Proc. Wkshp Future Directions in Cyber-Physical Systems Security*, 2009.
[11] A. Parvulescu, "System for automatically unlocking an automotive child safety door lock," June 2000, US Patent 6,081,758.
[12] "HackRF One," https://greatscottgadgets.com/hackrf/, accessed: 2016-12-20.
[13] "Peeping into 73,000 unsecured security cameras," http://www.networkworld.com/article/2844283/microsoft-subnet/peeping-into-73-000-unsecured-security-cameras-thanks-to-default-passwords.html, accessed: 2016-12-20.
[14] "GNURadio," http://gnuradio.org/redmine/projects/gnuradio/wiki, accessed: 2016-12-20.
[15] "Federal Communication Commission," https://www.fcc.gov/general/fcc-id-search-page, accessed: 2016-12-20.
[16] P. J. Christensen, W. H. Graf, and T. W. Yeung, "Refinery power failures: Causes, costs and solutions," *Petroleum Technology Quarterly*, vol. 18, no. 4, 2013.
[17] "Sheffield forgemasters fined after worker's death," http://www.thestar.co.uk/news/local/breaking-sheffield-forgemasters-fined-after-worker-s-death-1-6326832, accessed: 2016-12-20.
[18] "If the fire doesn't kill you, the CO_2 might," http://www.analoxsensortechnology.com/blog/2015/08/if-the-fire-doesnt-kill-you-the-co2-might-carbon-dioxide-and-fire-suppression/, accessed: 2016-12-20.
[19] C. Gomez, J. Oller, and J. Paradells, "Overview and evaluation of Bluetooth Low Energy: An emerging low-power wireless technology," *Sensors*, vol. 12, no. 9, pp. 11 734–11 753, 2012.
[20] A. Dementyev, S. Hodges, S. Taylor, and J. Smith, "Power consumption analysis of Bluetooth Low Energy, ZigBee and ANT sensor nodes in a cyclic sleep scenario," in *Proc. IEEE Wireless Symposium*, 2013, pp. 1–4.
[21] R. Akella and B. M. McMillin, "Model-checking BNDC properties in cyber-physical systems," in *Proc. 33rd IEEE Int. Computer Software and Applications Conference*, vol. 1, 2009, pp. 660–663.
[22] R. Akella, H. Tang, and B. M. McMillin, "Analysis of information flow security in cyber-physical systems," *Int. J. Critical Infrastructure Protection*, vol. 3, no. 3, pp. 157–173, 2010.
[23] C. Li, A. Raghunathan, and N. K. Jha, "Improving the trustworthiness of medical device software with formal verification methods," *IEEE Embedded Systems Letters*, vol. 5, no. 3, pp. 50–53, 2013.
[24] Z. Jiang, M. Pajic, R. Alur, and R. Mangharam, "Closed-loop verification of medical devices with model abstraction and refinement," *Int. J. Software Tools for Technology Transfer*, vol. 16, no. 2, pp. 191–213, 2014.
[25] M. Sun, S. Mohan, L. Sha, and C. Gunter, "Addressing safety and security contradictions in cyber-physical systems," in *Proc. First Wkshp. Future Directions in Cyber-Physical Systems Security*, pp. 311–318.
[26] J.-P. Thomesse, "Fieldbus technology in industrial automation," *Proc. IEEE*, vol. 93, no. 6, pp. 1073–1101, 2005.
[27] E. Byres and J. Lowe, "The myths and facts behind cyber security risks for industrial control systems," in *Proc. Verband der Elektrotechnik Kongress*, vol. 116, 2004, pp. 213–218.
[28] D. Li, Y. Serizawa, and M. Kiuchi, "Concept design for a web-based supervisory control and data-acquisition system," in *Proc. IEEE Transmission and Distribution Conference*, vol. 1, 2002, pp. 32–36.
[29] N. Apthorpe, D. Reisman, and N. Feamster, "A smart home is no castle: Privacy vulnerabilities of encrypted IoT traffic," http://datworkshop.org/papers/dat16-final37.pdf, accessed: 2017-03-01.
[30] A. M. Nia, S. Sur-Kolay, A. Raghunathan, and N. K. Jha, "Physiological information leakage: A new frontier in health information security," *IEEE Trans. Emerging Topics in Computing*, vol. 4, no. 3, pp. 321–334, 2016.
[31] L. Zhang and H. Zhang, "A survey on security and privacy in emerging sensor networks: From viewpoint of close-loop," *Sensors*, vol. 16, no. 4, p. 443, 2016.
[32] A. Mosenia and N. K. Jha, "A comprehensive study of security of Internet of Things," *IEEE Trans. Emerging Topics in Computing*, DOI: 10.1109/TETC.2016.2606384, 7 Sept., 2016.
[33] H. Esen, M. Adachi, D. Bernardini, A. Bemporad, D. Rost, and J. Knodel, "Control as a service: Cloud-based software architecture for automotive control applications," in *Proc. Int. Wkshp. Swarm at the Edge of the Cloud*, 2015, pp. 13–18.
[34] O. Givehchi, J. Imtiaz, H. Trsek, and J. Jasperneite, "Control-as-a-service from the Cloud: A case study for using virtualized PLCs," in *Proc. IEEE Wkshp. Factory Communication Systems*, 2014, pp. 1–4.
[35] T. Chen, "Stuxnet, the real start of cyber warfare?" *IEEE Network*, vol. 24, no. 6, pp. 2–3, 2010.
[36] A. Perrig, R. Szewczyk, J. D. Tygar, V. Wen, and D. E. Culler, "SPINS: Security protocols for sensor networks," *Wireless networks*, vol. 8, no. 5, pp. 521–534, 2002.
[37] J. Zhen, J. Li, M. J. Lee, and M. Anshel, "A lightweight encryption and authentication scheme for wireless sensor networks," *Int. J. Security and Networks*, vol. 1, no. 3-4, pp. 138–146, 2006.
[38] M. Xu, W. Xu, J. Walker, and B. Moore, "Lightweight secure communication protocols for in-vehicle sensor networks," in *Proc. ACM Wkshp. Security, Privacy & Dependability for Cyber Vehicles*, 2013, pp. 19–30.

Arsalan Mosenia received his B.S. degree in Computer Engineering from Sharif University of Technology, Tehran, Iran, in 2012, and M.A. degree in Electrical Engineering from Princeton, NJ, in 2014. He is currently pursuing a Ph.D. degree in Electrical Engineering at Princeton University, NJ. His research interests include wireless sensor networks, Internet-of-Things, computer security, distributed computing, mobile computing, and machine learning.

Anand Raghunathan is a Professor and Chair of VLSI in the School of Electrical and Computer Engineering at Purdue University, where he leads the Integrated Systems Laboratory. His research explores domain-specific architecture, system-on-chip design, embedded systems, and heterogeneous parallel computing. Prof. Raghunathan has co-authored a book ("High-level Power Analysis and Optimization"), eight book chapters, 21 U.S patents, and over 200 refereed journal and conference papers. His publications have been recognized with eight best paper awards and four best paper nominations. He received the Patent of the Year Award (recognizing the invention with the highest impact), and two Technology Commercialization Awards from NEC. He was chosen by MIT's Technology Review to be among the TR35 (top 35 innovators under 35 years, across various disciplines of science and technology) in 2006, for his work on "making mobile secure". Prof. Raghunathan has chaired the ACM/IEEE International Symposium on Low Power Electronics and Design, the ACM/IEEE International Conference on Compilers, Architecture, and Synthesis for Embedded Systems, the IEEE VLSI Test Symposium, and the IEEE International Conference on VLSI Design. He has served as an Associate Editor of the IEEE Transactions on CAD, IEEE Transactions on VLSI Systems, ACM Transactions on Design Automation of Electronic Systems, IEEE Transactions on Mobile Computing, ACM Transactions on Embedded Computing Systems, IEEE Design & Test of Computers, and the Journal of Low Power Electronics. He was a recipient of the IEEE Meritorious Service Award (2001) and Outstanding Service Award (2004). He is a Fellow of the IEEE, and Golden Core Member of the IEEE Computer Society. Prof. Raghunathan received the B. Tech. degree in Electrical and Electronics Engineering from the Indian Institute of Technology, Madras, and the M.A. and Ph.D. degrees in Electrical Engineering from Princeton University.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TMSCS.2017.2720660, IEEE Transactions on Multi-Scale Computing Systems

Susmita Sur-Kolay (SM05) received the B.Tech. degree in electronics and electrical communication engineering from Indian Institute of Technology, Kharagpur, India, and the Ph.D. degree in Computer Science and Engineering from Jadavpur University, Kolkata, India. She was in the Laboratory for Computer Science, Massachusetts Institute of Technology, Cambridge, MA, USA, from 1980 to 1984. She was a Post-Doctoral Fellow in the University of Nebraska-Lincoln, Nebraska-Lincoln, NE, USA, in 1992, a Reader in Jadavpur University from 1993 to 1999, a Visiting Faculty Member with Intel Corporation, Santa Clara, CA, USA, in 2002, and a Visiting Researcher at Princeton University in 2012. She is a Professor in the Advanced Computing and Microelectronics Unit, Indian Statistical Institute, Kolkata. She has co-edited two books, authored a book chapter in the Handbook of Algorithms for VLSI Physical Design Automation, and co-authored about 100 technical articles. Her current research interests include electronic design automation, hardware security, quantum computing, and graph algorithms. Prof. Sur-Kolay was a Distinguished Visitor of the IEEE Computer Society, India. She has been an Associate Editor of the IEEE Transactions on Very Large Scale Integration Systems, and is currently an Associate Editor of ACM Transactions on Embedded Computing Systems. She has served on the technical program committees of several leading conferences, and as the Program Chair of the 2005 International Conference on VLSI Design, the 2007 International Symposium on VLSI Design and Test, and the 2011 IEEE Computer Society Annual Symposium on VLSI. Among other awards, she was a recipient of the President of India Gold Medal from IIT Kharagpur.

Niraj K. Jha (S'85-M'85-SM'93-F'98) received his B.Tech. degree in Electronics and Electrical Communication Engineering from Indian Institute of Technology, Kharagpur, India in 1981, M.S. degree in Electrical Engineering from S.U.N.Y. at Stony Brook, NY in 1982, and Ph.D. degree in Electrical Engineering from University of Illinois at Urbana-Champaign, IL in 1985. He is a Professor of Electrical Engineering at Princeton University.

He has served as the Editor-in-Chief of IEEE Transactions on VLSI Systems and an Associate Editor of IEEE Transactions on Circuits and Systems I and II, IEEE Transactions on VLSI Systems, IEEE Transactions on Computer-Aided Design, IEEE Transactions on Computers, Journal of Electronic Testing: Theory and Applications, and Journal of Nanotechnology. He is currently serving as an Associate Editor of IEEE Transactions on Multi-Scale Computing Systems and Journal of Low Power Electronics. He has served as the Program Chairman of the 1992 Workshop on Fault-Tolerant Parallel and Distributed Systems, the 2004 International Conference on Embedded and Ubiquitous Computing, and the 2010 International Conference on VLSI Design. He has served as the Director of the Center for Embedded System-on-a-chip Design funded by New Jersey Commission on Science and Technology and the Associate Director of the Andlinger Center for Energy and the Environment. He is the recipient of the AT&T Foundation Award and NEC Preceptorship Award for research excellence, NCR Award for teaching excellence, Princeton University Graduate Mentoring Award, and six Commendations for Outstanding Teaching from the School of Engineering and Applied Science. He is a Fellow of IEEE and ACM. He received the Distinguished Alumnus Award from I.I.T., Kharagpur in 2014.

He has co-authored or co-edited five books titled Testing and Reliable Design of CMOS Circuits (Kluwer, 1990), High-Level Power Analysis and Optimization (Kluwer, 1998), Testing of Digital Systems (Cambridge University Press, 2003), Switching and Finite Automata Theory, 3rd edition (Cambridge University Press, 2009), and Nanoelectronic Circuit Design (Springer, 2010). He has also authored 15 book chapters. He has authored or co-authored more than 430 technical papers. He has coauthored 14 papers, which have won various awards. These include the Best Paper Award at ICCD'93, FTCS'97, ICVLSID'98, DAC'99, PDCS'02, ICVLSID'03, CODES'06, ICCD'09, and CLOUD'10. A paper of his was selected for "The Best of ICCAD: A collection of the best IEEE International Conference on Computer-Aided Design papers of the past 20 years," two papers by IEEE Micro Magazine as one of the top picks from the 2005 and 2007 Computer Architecture conferences, and two others as being among the most influential papers of the last 10 years at IEEE Design Automation and Test in Europe Conference. He has co-authored another six papers that have been nominated for best paper awards. He has received 16 U.S. patents. He has served on the program committees of more than 150 conferences and workshops.

His research interests include embedded computing, FinFETs, low power hardware/software design, computer-aided design of integrated circuits and systems, machine learning, and secure computing. He has given several keynote speeches in the area of nanoelectronic design/test and embedded systems.