EAI Endorsed Transactions
on Security and Safety

Book review ICST.ORG

Reviewing the book "Principles of Cyber-physical Systems" from a security perspective

Minghui Zhu[1,*] and Peng Liu[2]

[1]Department of Electrical Engineering, Pennsylvania State University
[2]College of Information Science and Technology, Pennsylvania State University

Abstract

This is a review of the book "Principles of Cyber-physical Systems" authored by Rajeev Alur and published by the MIT Press at 2015.

Keywords: Cyber-physical systems, security.

Received on 01 October 2015, accepted on 01 October 2015, published on 05 October 2015

doi: 10.4108/eai.5-10-2015.150480

1. Introduction

Rapid advances in power, mobility and efficiency have led to pervasive usage of information and communication technologies. The Internet is among the greatest inventions of all time. Recently, information and communication technologies are increasingly integrated with the elements of the physical world; e.g., automobiles and medical devices. The new generation of systems are referred to as cyber-physical systems (CPS). The synergy of information and communication technologies and physical processes enables cyber-physical systems to achieve unprecedented intelligence. It is envisioned that the technologies of cyber-physical systems will transform the way people interact with engineered systems -- just as the Internet has transformed the way people interact with information [1]. Cyber-physical systems are ubiquitous, and cross a variety of sectors such as building automation, energy, healthcare, manufacturing, and transportation. A recent report estimates that the technical innovations of cyber-physical systems could find direct applications in sectors currently accounting for more than $32.3 trillion in economic activity, and with the potential to grow to $82 trillion of output by 2025 -- about one half of the global economy [2].

*Corresponding author. Email:muz16@psu.edu

2. This book focuses on four basic properties of cyber-physical systems

Last section pinpoints the vital significance of cyber-physical systems. However, there is still a lack of a systems science which can address the unique challenges of cyber-physical systems caused by the strong coupling of the cyber space and the physical world. This book is timely, and provides a comprehensive and coherent introduction to the principles of modelling, specification, analysis and design of cyber-physical systems.

Modelling. Cyber-physical systems contain a large number of heterogeneous components in the cyber space and the physical world. It presents a major barrier to model cyber-physical systems in such a way as mathematical abstractions can capture the system essences and meanwhile manage the complexity of system design. In this book, the author discusses a variety of efficient models of cyber-physical systems. In particular, Chapter 2 focuses on a discrete and synchronous model of reactive computation, where all components execute in lock-step in a sequence of rounds. In each round, a reactive component reads its inputs;

based on its current states and inputs, it computes outputs and updates the internal state. In Chapter 4, an asynchronous model of computation is introduced where different activities execute at independent speeds. Such models naturally capture multi-processor machines and distributed multi-agent systems. Then Chapter 7 presents the timed models of computation, where processes rely on the global physical time to achieve a loose form of synchronization. The timed model is able to express phenomena such as “measure the temperature every 5 seconds’’. Chapters 2, 4 and 7 focus on the discrete models of computation. In contrast, Chapter 6 discusses continuous-time models of dynamic systems which are well suited to model temporal evolution in the physical world. The models introduced include differential equations and state-space equations. Chapter 9 provides an exposition on hybrid systems where continuous-time evolution and discrete transitions co-exist. Such models serve as a unified framework to characterize the intricate interactions among computation, communication and control of the physical world.

Properties. Besides modelling, a substantial portion of the book is devoted to elaborating on important properties of cyber-physical systems, including safety, liveness, stability and real-time scheduling.

Safety. Safety is referred to that system states never enter a region which is harmful to the system. One example is that it is unsafe for two vehicles if they are too close to each other. Another example is that, in the power grid, the frequency deviations of power generators cannot exceed certain threshold. To verify system safety, Chapter 3 introduces the general techniques of inductive invariants and state-space exploration. Barrier certificates and symbolic search algorithms are studied for the verification of hybrid systems.

Liveness. Liveness is referred to that a system eventually achieves specified goals. For example, a missle eventually hits a given target. In Chapter 5, Linear Temporal Logic is used to formally express liveness requirements. Both enumeration and symbolic state-space exploration techniques are generalized to verify Linear Temporal Logic requirements of system models. This problem is referred to as model checking.

Stability. Stability refers to the ability of a system to return to desired equilibrium after experiencing small perturbations of initial states. This is a well-studied topic in control theory. Chapter 6 discusses Lyapunov stability and input-output stability where a particular emphasis is placed on linear time-invariant systems.

Real-time Scheduling. Many cyber-physical systems are real-time systems whose correctness depends on their temporal aspects as well as their functional aspects. A fundamental problem of real-time systems is the formalization of demands for processing times by different computational tasks and general policies for allocating processing times so that these demands are met. This subject is referred to as real-time scheduling. Chapter 8 focuses primarily on understanding two classic scheduling algorithms: Earliest Deadline First and Rate Monotonic.

3. Security is an emerging CPS property

Besides four important properties articulated in the book, security is an emerging one. In particular, information and communication technology systems are inherently vulnerable to cyber attacks. Such cyber vulnerabilities can be exploited by attackers to transcend cyber defenses and further compromise physical systems. The vulnerabilities of cyber-physical systems to cyber attacks are evidenced by a number of recent accidents. In 2006, two traffic engineers hacked Los Angeles traffic light control system, lengthened red lights and further snarled traffic. In 2010, the malware Stuxnet attacked industrial programmable logic controllers and reportedly destroyed almost one-fifth of Iran's nuclear centrifuges. In 2011, an American unmanned aerial vehicle was brought down by Irian cyber warfare unit. In 2013, the Industrial Control Systems Cyber Emergency Response Team received 181 vulnerability reports about industrial control systems throughout the year in the United States.

Generally speaking, security properties are associated with the confidentiality, integrity, and availability of cyber-physical systems. In Section 4, several concrete security properties will be presented through a case study. In the context of cyber-physical systems, security properties could not be clearly understood without clarifications on the relationships between security and the classic CPS properties.

Safety vs. Security. In cyber-physical systems, safety properties are usually defined in a way independent of the attacker's intent, objectives, and strategies. In contrast, such information typically determines security properties. Moreover, a critical question is whether security properties can be modelled as safety properties when security accidents are viewed as a *bad* thing. We believe that it is not the case for most of the time. Take security properties against zero-day attacks as an example. The fact that the corresponding vulnerability is unknown to the defender indicates that whether the zero-day attack happens or not is actually determined by the attacker, not cyber-physical systems. Without explicitly taking into account these uncertainties, it is unlikely to prove any security-centric safety properties. To the best of our knowledge, existing CPS safety research works have not yet systematically investigated the modelling of unknown vulnerabilities, the corresponding zero-day attacks, and the associated uncertainties.

Liveness vs. Security. One aspect of CPS security is availability, and in certain circumstances, availability properties could be specified and studied as liveness properties. For example, self-healing systems aim to ensure that compromised system can self-heal from an intrusion and eventually provide the requested services. However, availability is only one aspect of CPS security; it focuses on how to enable cyber-physical systems to continuously function under attacks. Several important security properties are actually outside the scope of availability. One example is information leakage centric properties. That is, instead of causing cyber-physical systems dysfunction, many attacks only want to steal sensitive data from authorized entities.

Stability vs. Security. Stability under uncertainties or disturbances has been extensively studied in the control literature. A number of control branches have been developed, including stochastic control, robust control and adaptive control. Each of them has certain limitations to deal with CPS security, especially zero-day attacks. In particular, stochastic control can address random failures. However, when and where a zero-day attack happens is actually determined by the attacker's subjective *attack plot* and schedule, not probabilities. Robust control can ensure stability or related as long as uncertainties are within pre-described sets. For zero-day attacks, uncertainty sets could be very difficult to obtain. Because the vulnerabilities are unknown before zero-day attacks happen. If their estimates are too conservative, system performance will be significantly degraded. Adaptive control also demands certain structural description of uncertainties, and its transient performance could be arbitrarily bad. As a result, new methodologies are of strong need to ensure stability under zero-day attacks.

Real-time Scheduling vs. Security. Zero-day attacks is an emerging threat to the real-time scheduling properties. More specifically, security defenses share the same set of computing and communication resources with other tasks of real-time systems; e.g., data processing and controller synthesis. If their runtime overheads are too high, security defenses would consume too many resources and other time-sensitive tasks may miss their deadlines. As a result, it would require the co-design of real-time scheduling policies and light-weight security defense schemes to ensure the security of real-time systems.

4. Case study on CPS security

Below, we use unmanned-vehicle-operator-networks, a representative class of cyber-physical systems, to show several concrete security properties.

System model. The unmanned-vehicle-operator network is depicted in Figure 1 and represents a cyber-physical system. In particular, the cyber components include wireless communication, embedded computation and

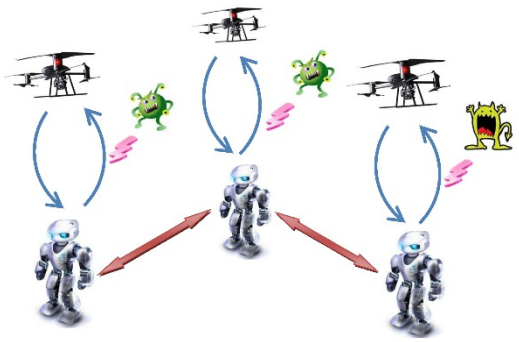

Figure 1: Unmanned-vehicle-operator-network

executable codes running inside the vehicles and handheld devices. Figure 2 shows some unmanned aerial and ground vehicles. The physical components include the sensing and maneuver devices of the vehicles as well as their surroundings; e.g., moving obstacles. Control or decision making synergizes heterogeneous functions of the cyber and physical components and enables them to act as a whole.

Figure 2: Unmanned-vehicle testbed at Penn State [3]

Each vehicle carries onboard sensors and so the vehicle can measure its own states; e.g., position and velocity, and estimate the states of nearby vehicles. In addition, each vehicle carries onboard communication devices which allow the vehicle and its assigned handheld device to exchange timely information; e.g., sensing data and control commands. Each vehicle is also equipped with embedded CPU, memory and executable codes such that the vehicle can process information and generate control commands on its own if necessary. On a vehicle that runs virtualization software, a hypervisor can be installed onboard and can overrule the vehicle at abnormal situations. Besides, the maneuver system enables the vehicle to move along specified paths and perform given

tasks. The handheld device is spatially distributed from its associated vehicle and responsible for remote monitoring and control. To achieve this, the handheld device needs to process data, communicate with its vehicle and serve as the human computer interface for its human operator. The handheld device generates control commands for its vehicle at normal situations. At abnormal situations, the handheld device passes the control authority to the operator and the operator makes decisions instead.

Threat model. In the unmanned-vehicle-operator network, the cyber and physical components are tightly coupled and their synergy allows the vehicles to accomplish complex tasks. However, the inherent vulnerability of information and communication technology systems to cyber attacks imposes significant security risks on the physical system. The adversaries can transcend cyber defense, intentionally modify control commands for the vehicles and eventually cause damage in the physical space. Here are three classes of cyber attacks: denial-of-service (DoS) attacks, replay attacks and malware. The information exchanges of vehicle-operator and inter-operator are carried out via wireless communication networks. The communication channels are subject to two classes of network attacks: DoS attacks and replay attacks. For DoS attacks, the adversary blocks the communications between the sender and receiver, and then no message can be transmitted. For replay attacks, the adversary intercepts the messages sent from the sender, stores the intercepted information in his/her memory and resends it to the receiver later. Resending the same message could happen for multiple times in a row. Each vehicle is equipped with embedded CPU, memory and executable codes which are subject to malware; e.g., Trojan. A Trojan could be a malicious driver provided by a third-party vendor. A Trojan could also be an infected executable file stored in the vehicle. By infected, we mean that the file is originally clean but modified by the attacker later on. Besides the Trojan, we assume that a program running on the vehicle could have vulnerabilities and bugs. Accordingly, the bugs could be exploited by onboard Trojan or remote malicious requests sent via wireless channels.

Security goal. To deal with network attacks, a main security goal is to enable multiple unmanned vehicles to accomplish given tasks; e.g., formation achieving, despite the threats of DoS attacks and replay attacks. Formation achieving can be interpreted as the stability of desired formation. So, distributed attack-resilient formation control can be viewed to ensure the basic stability property under the emerging network attacks.

To deal with malware attacks, a main security goal is to avoid damage and task failures caused by vehicle collisions. Once compromised, an unmanned vehicle could be largely controlled by malware for a period of time until a super-privileged module in the vehicle receives a remote “stop it” command. Viewing this security problem from the safety perspective, the corresponding safety properties cannot be specified until the following research questions are answered: How to model the malware and the unknown vulnerabilities exploited by the malware? How to model the effects of malware attacks on the states and inputs of the compromised vehicle? How to integrate intrusion detection systems in the cyber space and control theoretic estimation in the physical world to reveal zero-day vulnerabilities?

5. Conclusions

The development of cyber-physical systems is still at an early stage. This book provides a comprehensive and coherent introduction to the principles of modelling, specification, analysis and design of cyber-physical systems. These principles are drawn from a variety of sub-disciplines, including concurrency theory, distributed algorithms, control theory, formal methods, real-time systems and hybrid systems. The book is suitable for a semester-long graduate course in computer science, computer engineering, electrical engineering, mechanical engineering and aerospace engineering. In addition, the book provides valuable resource for researchers from both academia and industry who are interested in cyber-physical systems and emerging CPS security.

Acknowledgements.

M. Zhu and P. Liu are partially supported by the NSF CPS-Security project CNS-1505664.

References

[1] NSF Cyber-physical Systems (CPS) program, http://www.nsf.gov/funding/pgm_summ.jsp?pims_id=503286

[2] P. Evans and M. Annunziata, *Pushing the Boundaries of Minds and Machines*. General Electric, November 2012.

[3] http://nrsl.mne.psu.edu/