An Assurance Framework for Secure Cyber-Physical Systems

Eunsuk Kang, Massachusetts Institute of Technology (eskang@mit.edu)

Keywords: Security, systems engineering, architecture, modeling, formal methods

Computational devices are being deeply integrated into nearly every aspect of our lives, including energy, transportation, financial, and healthcare systems. These devices also bring an unprecedented level of complexity and risks into the existing infrastructures, introducing new behaviors that arise from complex interaction between physical and computational components. Due to their close proximity to us, a failure in one of these systems can lead to catastrophic outcomes, ranging from economic damages to loss of human lives.

Traditional techniques for securing a computer system---encryption, access control mechanisms, monitoring, and testing---are essential but not sufficient to establish the security of a cyber-physical system. Many infrastructural systems rely on legacy software that cannot be easily patched, and as a result, remains vulnerable to common exploits. Previous studies have shown that systems that use strong encryption are still susceptible to simple attacks, due to a design flaw or ill-trained human operator (Anderson, 1993). Despite recent advances in formal verification and static analysis, building a piece of complex software that is completely free of vulnerabilities is far from the reality. Above all, little is understood about security implications when components of two entirely different natures---computational and physical---are composed into a single system.

Instead of aiming for an impenetrable system, as is often the goal of traditional security, we believe that a more effective approach is to build *assurance* that the system will prevent the most catastrophic outcomes, *even* in the presence of security attacks. In this research project, we propose a new framework to assist the system engineer in evaluating the security of a cyber-physical system. The framework consists of the following elements, each of which we believe poses significant research challenges:

(1) **Heterogenous Modeling Tool**: Before performing a security analysis, the engineer must first construct an adequate model of the system. The main challenge in modeling a cyber-physical system is its *heterogenous* nature. A typical system is a mosaic of diverse parts---software products, hardware parts, physical plants, and users---at different levels of abstraction---program instructions, network messages, continuous dynamics, and human behavior. These components are not amenable to a traditional composition approach where they are joined at a common interface, since their behaviors are described using independent sets of alphabet. No single abstraction will

be sufficient to encompass the entire system. A new technique must be developed for combining multiple, independent abstractions into a global model that accurately captures the end-to-end behavior of the system.

(2) **Notation for Security Requirements**: In computer security, there are well-understood types of security requirements that users find desirable, such as confidentiality and integrity; for example, on a bank website, it is highly desirable that a user's account information be accessible to only him or her. In a cyber-physical system, some existing classes of security policies are no longer as applicable or critical. For example, in a smart grid system, leaking data about the energy usage of a household is undesirable, but itself might be of minor consequences; a far more critical requirement is that the household be provided with a basic level of heating at all times during a subzero temperature. New classes of *domain-specific* security requirements must be devised for cyber-physical systems, and the engineer must be provided with a notation for precisely expressing and *prioritizing* those requirements.

(3) **Scalable System-Level Analysis**: Given the scale and complexity of a typical cyber-physical system, guaranteeing a complete prevention of an attack is too difficult and expensive. A more realistic goal is to prevent catastrophic failures, assuming that an attack is likely, by minimizing the size of the *trusted base*---the parts of the system that are responsible for ensuring the most critical security requirements (Kang and Jackson, 2010). For example, it may be difficult to guarantee that a radiation therapy system functions correctly in all cases, but one may be able to show that it prevents patient overdose using multiple safety interlocks. An end-to-end analysis must be developed to assist the engineer in identifying the attack surface and evaluating the impact of a security attack on the trusted base.

A conventional approach that relies on testing, patches, and firewalls, with the goal of preventing all unauthorized access, is not only expensive, but unlikely to be effective against the scale of a cyber-physical system. We believe that the proposed framework, when implemented, can enable a more cost-effective security evaluation by focusing on the most critical security requirements; for example, ensuring that a radiation therapy system never overdoses a patient, or that a smart grid supplies households with minimal necessary electricity. This framework itself won't guarantee perfect security, but will be a first step towards providing assurance that we can rely on our critical infrastructures to protect us, even in the ever-growing presence of security threats.

References

Anderson, R. J. *Why Cryptosystems Fail*. ACM Conference on Computer and Communications Security (CCS), 1993.

Kang, E. and Jackson, D. *Dependability Arguments with Trusted Bases*. International Conference on Requirements Engineering (RE), 2010.