

Security Reasoning for Distributed Systems with Uncertainties

PI: André Platzer

Student: Erik P. Zawadzki

Team The project team includes André Platzer who is an assistant professor in the computer science department at Carnegie Mellon University. He is an expert in verification and analysis of hybrid, distributed, and stochastic dynamic systems, including cyber-physical systems. The team further includes Erik P. Zawadzki, who is a fourth year graduate student in the computer science department at Carnegie Mellon University and is developing reasoning techniques for first-order MILPs and fast propositional solvers for probabilistic model counting.

Objective Phenomena like Stuxnet make apparent to the public what experts knew long ago: security is not an isolated question of securing a single door against lockpicking or securing a single computer against a single hacker trying to gain access via a single network activity. Because the strength of a security system is determined by its weakest link, security is much more holistic and affects more and more elements of a system design. Most systems are not understood properly by simplistic finite-state abstractions like yes/no information about whether a node in a (sufficiently small) finite network has been compromised or not. Stuxnet, for example, is reported to have a sophisticated interaction of control affects, sensor modifications, and even exhibits hidden long-term effects on the physical world by changing the behavior of programmable logic controllers (PLCs). Moreover, security-relevant systems of today are more often than not characterized by distributed setups, both in the system and in the attack. The security analyst, furthermore, faces uncertainties that aggregate to paralytic “zero” knowledge—unless he takes a probabilistic view and quantitatively relates the relative likelihoods of symptoms and explanations via partial observations and incomplete prior knowledge. The scale and complexity of any affected system, however, makes the analysis hard. But, more crucially, it is becoming infeasible to scale between systems to craft and tune a new security analysis over and over again for each new particular application scenario.

Research We propose to address the scale problem in security analysis by developing representation and reasoning techniques that support higher level structure and that enable the security community to factor out common core reasoning principles from the particular elements, rules, and data facts that are specific to the application at hand. This principle, separation of reasoning engine and problem specification, has been pursued with great success in satisfiability modulo theories (SMT) solving. Probabilistic reasoning is powerful for many specific domains, but does not have any full-fledged extensions to the level of scalable higher-level representations that are truly first-order. Based on our preliminary results, we propose to develop first-order probabilistic programs and study how they can represent security analysis questions for systems with both distributed aspects and quantitative uncertainty. We propose to study instance-based methods for reasoning about first-order probabilistic programs. Instance-based methods enjoy a good trade-off between generality and efficiency. They can leverage classical advances in probabilistic reasoning for finite-dimensional representations and lift them to the full expressive and descriptive power of first-order representations. With this approach, we achieve inter-system scaling by decoupling the generics from the specifics and we hope to improve intra-system scaling by being able to combine more powerful reasoning techniques from classically disjoint domains.

Relevance This project is of relevance to the security community, since, if successful, it would provide more general and more flexible off-the-shelf solutions that can be used to address security analysis questions for particular systems. To simplify the design of problem-specific computer-aided security analysis procedures, this project addresses a separation of the problem description from the reasoning techniques about them. At the same time, it increases representational and computational power to scale to systems with uncertainty and distributed effects.

Impact This project has the potential to help solve security analysis questions that scale to distributed systems and the presence of uncertainty. Both aspects are central in security challenges like Stuxnet. Because each security analysis question is different, it is more economically feasible to assemble particular security analyses from suitable reasoning components. This project addresses one such component with a good trade-off between generality and efficiency.