Automated Verification of Privacy in Security Protocols: Back and Forth Between Theory & Practice

Lucca Hirschi

Context. Security protocols leverage *cryptographic primitives* in order to achieve various *security goals* such as secrecy, authentication, or *privacy*, which is becoming increasingly important. Unfortunately, designing secure protocols is extremely complex as witnessed by attacks regularly disclosed on protocols of utmost importance. In order to increase the confidence we can put in security protocols, *formal methods* based on the *symbolic model* have proved their usefulness by providing rigorous, mathematical frameworks to analyze them. For the case of *reachability properties* (*e.g.*, secrecy, authentication), this approach has lead to mature tools and industrial successes. Unfortunately, this is not yet the case for privacy, which can be explained by the recentness of this line of work and the more complex nature of privacy properties often modeled through *behavioral equivalences* instead of reachability properties, that are easier to check. To circumvent the undecidability of verifying behavioral equivalences, two main solutions have been developed: either recover decidability by only considering a bounded number of sessions, or only semi-decide the problem for the unbounded setting with pragmatic, efficient, procedures using well-chosen sound abstractions. Unfortunately, those approaches suffer from two different problems that significantly limit their practical impact, making analyses infeasible for large classes of real-life security protocols. First, (symbolically) exploring all possible executions, as done in the bounded setting, leads to the so-called *state space explosion problem* caused by the concurrency nature of security protocols. Second, the approximations that are made in the unbounded setting are too imprecise to meaningfully analyze some important privacy properties such as unlinkability.

Contributions. In my thesis, to overcome these severe limitations, I design **new algorithms** and **tools** to effectively and efficiently analyze privacy properties, that I apply on several **real-life case studies**, including protocols coming from the **e-passport application**. This involves developing **theoretical foundations** to formally define and justify the correctness of those methods. I also make the effort to confirm the **practical relevance** of proposed solutions by putting them into practice (creating new tools or modifying existing verifiers) and by analyzing real-world case studies (providing security guarantees or finding attacks).

In standard model-checking approaches for concurrent systems, state space explosion problems, from which the bounded setting dramatically suffers, are often handled using *Partial Order Reduction* (POR) techniques. However, it is known that standard POR techniques cannot be *directly* applied in the context of security protocol verification (reasons are *e.g.*, the adversarial setting, symbolic executions, equivalence more involved than reachability). In my thesis, I **devise and evaluate new POR techniques**, dedicated to the security setting, to mitigate the state space explosion. My techniques borrow ideas from concurrency theory, trace theory, and, perhaps more surprisingly, focusing from proof theory. Blending all these ingredients, and adapting them to the demanding framework of security, I have come up with the first POR techniques that can effectively be used in symbolic verification algorithms for equivalence properties of security protocols. I put those techniques into practice in the tool Apte: from the theoretical aspects of this integration (correctness result) to the implementation in the distributed code. I also present extensive benchmarks showing that those theoretical results do translate to **dramatic speedups** (*e.g.*, more than 5 order of magnitude on real-life case studies).

In the unbounded setting, the impractically strong form of equivalence is a serious limitation when considering the unlinkability property, that I address via an **innovative hybrid approach**: I am able to precisely characterize when unlinkability and anonymity are met via *sufficient conditions*, that can then be checked with precision with state-of-the-art verifiers. Those conditions, together with my soundness Theorem, yield a **new verification technique**, that I have implemented in a new tool (**UKano**), that enables the verification of a large class of protocols that were previously out of scope. I confirm the practical relevance by carrying out **first analyses of real-life case studies**, revealing new weaknesses along the way.

Impact. In addition to my implementations in the tools Apte and Spec, my POR techniques have also been integrated and implemented in two other state-of-the art verifiers by independent researchers. Now, all the verifiers in the bounded setting are equipped with my POR techniques and benefit from significant speed-ups. As a result, practical scenarios can now be analyzed in a reasonable time. My hybrid approach for the unbounded setting significantly expands the scope of formal methods for privacy: I was able to establish the **first proofs** of unlinkability for numerous protocols including, attribute-based, **e-passport** and other RFID protocols, to find a **new weakness** on the PACE protocol (that should replace BAC in e-passports) and a new **traceability attack** on the LAK protocol whereas it was previously claimed untraceable. I am pursuing those two lines of work by extending the underlying techniques to larger classes of protocols and other privacy properties. All these contributions have been published in international conferences and journals in the field of security, formal methods and concurrency theory: **S&P'16**, **CONCUR'15**, **LMCS**, **JLAMP** and **POST'14**.