Decision procedures for equivalence based properties (part II)

David Baelde[1], Stéphanie Delaune[1], and Steve Kremer[2]

[1] LSV, CNRS & ENS Cachan
[2] LORIA, Inria Nancy Grand Est

The results presented in this report are based on results that have been published in [5, 6, 16, 18, 17, 23, 2, 13]– works that have been supported by the VIP project and that are available on the website of the VIP project.

Abstract. This deliverable concerns the Task 3 of the VIP project: *Algorithmic and decidability issues.* This report is a follow-up of the report entitled "Decision procedures for equivalence based properties". Actually, extended versions of most of the results presented in the previous report are currently under submission to journals (TOCL and I&C). The purpose of the present report is to summarize new results that have been made in this area since the previous report.

1 Introduction

Security protocols are widely used today to secure transaction that rely on public communication channels such as the Internet. It is therefore essential to obtain as much confidence as possible in their correctness. Starting in the 80s, many works have been devoted to the use of formal methods to analyse the security of these protocols (*e.g.* [24, 30]). In the case of a bounded number of sessions, secrecy preservation is co-NP-complete [30], and for an unbounded number of sessions, several decidable classes have been identified (*e.g.* [24]). Many tools have also been developed to automatically verify cryptographic protocols (*e.g.* AVISPA [3], ProVerif [8]).

Until recently, most efforts and successes only concerned trace properties, *i.e.* security properties that can be checked on each individual sequence of messages corresponding to an execution of the protocol. Secrecy (at least weak forms of secrecy) and authentication are typical examples of trace properties. There are however several security properties, which cannot be defined as trace properties and require a notion of *behavioural equivalence*. We focus here on the notion of *trace equivalence* which is well-suited for the analysis of security protocols. Intuitively, two processes P and Q are trace equivalent, denoted $P \approx_t Q$, if any experiment performed by an attacker on both processes leads to the emission of two sequences of messages that are indistinguishable, *i.e.* the attacker cannot observe any difference between these two sequences.

Contributions. In the previous report, algorithms for analysing equivalence-based properties have been described. Since replication very quickly yields to undecidability, we mainly focussed on finite processes, and we only described very few results that have been obtained for general processes (*i.e.* processes with replication). Since then, we have obtained new results in several directions:

1. We proposed some partial order reduction techniques [5, 6] to mitigate the so-called state space explosion problem encountered when implementing procedures described in the previous report.
2. We improved existing results to deal with general processes [16–18]. In particular, we obtained the first decidability results for protocols with nonces and replications.
3. We also proposed some results to take into account some specificities of RFID and e-voting protocols [23, 2, 13]. In particular, we proposed a procedure to analyse protocols that rely on the exclusive-or operator.

2 Improving the efficiency of existing verification tools

A major challenge faced when building verification tools for security protocols arises when modelling the behaviour of the attacker, who can generate a potentially infinite number of messages. In order to cope with this prolific attacker problem and obtain decision procedures, approaches based on symbolic semantics and constraint resolution have been proposed [26, 30]. This has lead to tools for verifying reachability-based security properties such as confidentiality [26] or, more recently, equivalence-based properties such as privacy [31, 12, 11].

In both cases, the practical impact of most of these tools is limited by a typical state explosion problem caused by the exploration of the large number of interleavings in the protocol's behaviour. In standard model-checking approaches for concurrent systems, the interleaving problem is handled using partial order reduction techniques [29]. For instance, the order of execution of two independent (parallel) actions is typically irrelevant for checking reachability. Things become more complex when working with a symbolic semantics: the states obtained from the interleaving of parallel actions will differ, but the sets of concrete states that they represent will have a significant overlap. Earlier work has shown how to limit this overlap [27] in the context of reachability properties for security protocols, leading to high efficiency gains in the OFMC tool of the AVISPA platform [3].

In this work, we revisit the work of [27] to obtain a partial order reduction technique for the verification of equivalence properties. Specifically, we focus on trace equivalence, requiring that two processes have the same sets of observable traces and perform indistinguishable sequences of outputs. This notion is well-studied and several algorithms and tools support it [9, 14, 31, 12, 11]. Contrary to what happens for reachability-based properties, trace equivalence cannot be decided relying only on the reachable states. The sequence of actions that leads to this state plays a role. Hence, extra precautions have to be taken before discarding a particular interleaving: we have to ensure that this is done in both sides of the equivalence in a similar fashion.

We achieve our goal by refining the interleaving semantics in two steps, gradually eliminating redundant traces. The first refinement, called *compression*, uses the notion of polarity [1] to impose a simple strategy on traces. It does not rely on data analysis at all and can easily be used as a replacement for the usual semantics in verification algorithms. The second one, called *reduction*, takes data into account and achieves optimality in eliminating redundant traces. In practice, the reduction step can be implemented in an approximated fashion, through an extension of constraint resolution procedures. We have done so in the tool APTE, showing that our theoretical results do translate to significant practical optimisations. Furthermore, our study brings an improvement of the original technique [27] that would apply equally well for reachability checking.

These results have been obtained by David Baelde, Stéphanie Delaune and Lucca Hirschi and will be part of the PhD thesis of Lucca Hirschi. They have been published at POST'14 and CONCUR'15.

3 Dealing with an unbounded number of sessions

Given a security protocol, does it achieve its security goals? This question is actually undecidable for trace properties as well as equivalence properties, for an unbounded number of sessions [25]. Bounding the number of sessions suffices to retrieve decidability for standard primitives, both for trace properties [30] and equivalence properties [7]. However, analysing a protocol for a fixed (often low) number of sessions does not allow to *prove* security. Even if my favourite security protocol has no flaw when used three times, there is absolutely no guarantee that a flaw will not appear when used a fourth time.

How can we prove security without limiting the number of sessions? Some tools such as ProVerif [8] or Scyther [22] can actually handle an unbounded number of sessions although they are not guaranteed to terminate. Yet, in practice, these tools work well, at least for trace properties. So a remaining open problem for the last ten years is to characterize a decidable fragment of security protocols, that captures most real protocols.

Even in the context of reachability properties, most existing decidability results for an unbounded number of sessions focused at protocols without nonces (see *e.g.* [25, 19]). Actually, the first decidability for equivalence properties for an unbounded number of sessions (but a bounded number of nonces) has been obtained in the framework of the VIP project (see [15]).

Decidability results for type-compliant protocols without nonces. To go beyond the very limited class of protocols presented in [15], we propose a result to reduce the search space for attacks [16]. Specifically, we show that if there is an attack then there is one that is well-typed. Our result holds for a large class of typing systems and a large class of determinate security protocols. Assuming finitely many nonces and keys, we can derive from this result that trace equivalence is decidable for an unbounded number of sessions for the class of simple type-compliant protocols. These hypotheses are explained below:

– *Simple processes.* This notion has been introduced in [20] and used in subsequent works. Intuitively, we assume that each process communicates on a distinct channel. In practice, each machine has its own IP address and each session is characterized by some session identifier. We also assume that each process consists of a sequence of inputs and outputs (with some tests). This models very well standard security protocols (with no else branches).
– *Type compliant protocols.* Intuitively, we assume that ciphertexts cannot be confused. A similar notion has been formally introduced in [10] and was shown to ensure termination of ProVerif (without nonces). This condition is part of the good design practices and is easy to enforce by adding some identifier (a tag) in each ciphertext. Of course the same tags are re-used in all sessions.

As an intermediate result, we also provide a novel decision procedure in the case of a bounded number of sessions.

Getting rid of nonces. Analysing protocols without nonces may seem restrictive. While this abstraction is clearly sound in the context of secrecy properties (for protocols without else branches), this is no longer the case for equivalence properties. In [17], we study how to soundly get rid of nonces in the context of equivalence properties. We show that nonces can be replaced by constants provided that each nonce is associated to two constants (instead of typically one constant for secrecy properties). Our result holds for deterministic (simple) protocols and a large class of primitives that includes *e.g.* standard primitives, blind signatures, and zero-knowledge proofs.

The first decidability result in the general setting. We also built on top of our result presented in [16] to propose the first decidability result for trace equivalence, for an unbounded number of sessions and *with nonces*. Since even simple reachability properties are undecidable in this context, we make some assumptions. First, we consider simple and type-compliant protocols as defined in [16]. Second, we assume that the underlying protocols have an *acyclic dependency graph.* Indeed, considering constructions used in undecidability results, one can notice that the encodings rely on some form of cyclicity. Typically, the last message of the protocol is re-injected at the first step of the protocol, forming an infinite loop. We therefore introduce the notion of dependency graph taking into account *sequential dependency* (some action can only be taken after some other actions), and *data dependency* (some message can only be built once some information is learnt from another message). This graph can be computed automatically from the protocol's specification.

Our main contribution is to show that the equivalence between simple and type-compliant protocols with an acyclic dependency graph is decidable, for protocols using symmetric encryption, concatenation, and nonces. Our class encompasses most symmetric key protocols we considered, including Needham-Schroeder with symmetric key, Otway-Rees, Denning-Sacco, or Wide-Mouthed-Frog. For some of these protocols, we had to consider an explicitly tagged version.

These results have been obtained by Stéphanie Delaune and Rémy Chrétien (in collaboration with Véronique Cortier) and will be part of the PhD thesis of Rémy Chrétien. They have been published at CONCUR'14, CSF'15, and ESORICS'15.

4 Taking into account some specificities

We mainly focus on some specificities encountered in RFID and e-voting protocols.

4.1 A first result to deal with *e.g.* exclusive-or

In the context of a bounded number of sessions, it has been shown that the problem of deciding whether two protocols are in trace equivalence amounts to deciding the equivalence of two constraint systems, *i.e.*, whether they have the same set of solutions.

In this work, we continue the study of the problem of deciding the equivalence of constraint systems used to model security protocols. In particular we consider the case where cryptographic primitives are modelled using a *group theory*. Group theories are a special case of monoidal theories which have been extensively studied by F. Baader and W. Nutt [28, 4] who have provided a complete survey of unification in these theories. Group theories include theories for exclusive or and Abelian groups. These theories are useful to model many security protocols (see [21]), as well as for modeling low level properties of encryption schemes and chaining modes.

More precisely we provide several new decidability and complexity results for the equivalence of constraint systems. We consider exclusive or and Abelian Groups which may also contain a unary homomorphic symbol. Our results rely on an encoding of the problem in systems of equations on a ring associated to the equational theory under study. To the best of our knowledge these are the first results to decide equivalence of constraint systems for these theories.

These results have been obtained by Stéphanie Delaune and Steve Kremer (in collaboration with Daniel Pasaila). They have been published at IJCAR'12.

4.2 Length may break privacy

While some decision procedures have been proposed for automatically deciding trace equivalence, all existing approaches abstract away the information an attacker may get when observing the length of messages.

In this work, we study trace equivalence with length tests. We first show that, in the static case, almost all existing decidability results (for static equivalence) can be extended to cope with length tests. In the active case, we prove decidability of trace equivalence with length tests, for a bounded number of sessions and for standard primitives. Our result relies on a previous decidability result from

Cheval *et al.* [12] (without length tests). The procedure has been implemented and we have discovered a new flaw against privacy in the biometric passport protocol.

This result has been obtained by Vincent Cheval (in collaboration with Véronique Cortier), and has been published at CAV'13.

4.3 Analysing everlasting privacy

Will my vote remain secret in 20 years? This is a natural question in the context of electronic voting, where encrypted votes may be published on a bulletin board for verifiability purposes, but the strength of the encryption is eroded with the passage of time. The question has been addressed through a property referred to as *everlasting privacy*. Perfect everlasting privacy may be difficult or even impossible to achieve, in particular in remote electronic elections. We therefore propose a definition of *practical* everlasting privacy. The key idea is that in the future, an attacker will be more powerful in terms of computation (he may be able to break the cryptography) but less powerful in terms of the data he can operate on (transactions between a vote client and the vote server may not have been stored).

In [2], we propose a definition of everlasting privacy in the applied-pi calculus. We provide the means to characterize what an attacker can break in the future in several cases. In particular, we model this for perfectly hiding and computationally binding primitives (or the converse), such as Pedersen commitments, and for symmetric and asymmetric encryption primitives. We adapt existing tools, in order to allow us to automatically prove everlasting privacy. As an illustration, we show that several variants of Helios (including Helios with Pedersen commitments) and a protocol by Moran and Naor achieve practical everlasting privacy, using the ProVerif [8] and the AKiSs [11] tools.

These results have been obtained by Steve Kremer (in collaboration with Myrto Arapinis, Véronique Cortier, and Mark Ryan), and have been published at POST'13.

5 Conclusion

We now have several procedures to analyse privacy type properties. Most of them have been implemented (or are currently under implementation). Each procedure has its own specificities and limitations but they allow us to cover many interesting scenarios.

We still plan to further study the AKISS procedure to integrate some additional equational theories such as the exclusive or operator.

References

1. J.-M. Andreoli. Logic programming with focusing proofs in linear logic. *J. Log. Comput.*, 2(3), 1992.

2. M. Arapinis, V. Cortier, S. Kremer, and M. D. Ryan. Practical Everlasting Privacy. In D. Basin and J. Mitchell, editors, *Proceedings of the 2nd Conference on Principles of Security and Trust (POST'13)*, volume 7796 of *Lecture Notes in Computer Science*, pages 21–40, Rome, Italy, Mar. 2013. Springer.
3. A. Armando et al. The AVISPA Tool for the automated validation of internet security protocols and applications. In *Proc. 17th Int. Conference on Computer Aided Verification (CAV'05)*, volume 3576 of *LNCS*, pages 281–285. Springer, 2005.
4. F. Baader. Unification in commutative theories. *Journal of Symbolic Computation*, 8(5):479–497, 1989.
5. D. Baelde, S. Delaune, and L. Hirschi. A reduced semantics for deciding trace equivalence using constraint systems. In M. Abadi and S. Kremer, editors, *Proceedings of the 3rd International Conference on Principles of Security and Trust (POST'14)*, volume 8414 of *Lecture Notes in Computer Science*, pages 1–21, Grenoble, France, Apr. 2014. Springer.
6. D. Baelde, S. Delaune, and L. Hirschi. Partial order reduction for security protocols. In L. Aceto and D. de Frutos-Escrig, editors, *Proceedings of the 26th International Conference on Concurrency Theory (CONCUR'15)*, volume 42 of *Leibniz International Proceedings in Informatics*, pages 497–510, Madrid, Spain, Sept. 2015. Leibniz-Zentrum für Informatik.
7. M. Baudet. Deciding security of protocols against off-line guessing attacks. In *Proc. 12th Conference on Computer and Communications Security (CCS'05)*, pages 16–25. ACM Press, 2005.
8. B. Blanchet. An Efficient Cryptographic Protocol Verifier Based on Prolog Rules. In *Proc. 14th Computer Security Foundations Workshop (CSFW'01)*, pages 82–96. IEEE Comp. Soc. Press, 2001.
9. B. Blanchet, M. Abadi, and C. Fournet. Automated verification of selected equivalences for security protocols. *Journal of Logic and Algebraic Programming*, 75(1):3–51, 2008.
10. B. Blanchet and A. Podelski. Verification of cryptographic protocols: Tagging enforces termination. In *Foundations of Software Science and Computation Structures (FoSSaCS'03)*.
11. R. Chadha, Ş. Ciobâcă, and S. Kremer. Automated verification of equivalence properties of cryptographic protocols. In H. Seidl, editor, *Programming Languages and Systems —Proceedings of the 21th European Symposium on Programming (ESOP'12)*, volume 7211 of *Lecture Notes in Computer Science*, pages 108–127, Tallinn, Estonia, Mar. 2012. Springer.
12. V. Cheval, H. Comon-Lundh, and S. Delaune. Trace equivalence decision: Negative tests and non-determinism. In Y. Chen, G. Danezis, and V. Shmatikov, editors, *Proceedings of the 18th ACM Conference on Computer and Communications Security (CCS'11)*, pages 321–330, Chicago, Illinois, USA, Oct. 2011. ACM Press.
13. V. Cheval, V. Cortier, and A. Plet. Lengths may break privacy – or how to check for equivalences with length. In *Proceedings of the 25th International Conference on Computer Aided Verification (CAV'13)*, volume 8043 of *Lecture Notes in Computer Science*, pages 708–723, St Petersburg, Russia, July 2013. Springer.
14. Y. Chevalier and M. Rusinowitch. Decidability of symbolic equivalence of derivations. *Journal of Automated Reasoning*, 48(2), 2012.
15. R. Chrétien, V. Cortier, and S. Delaune. From security protocols to pushdown automata. In F. V. Fomin, R. Freivalds, M. Kwiatkowska, and D. Peleg, editors, *Proceedings of the 40th International Colloquium on Automata, Languages and Programming (ICALP'13) – Part II*, volume 7966 of *Lecture Notes in Computer Science*, pages 137–149, Riga, Latvia, July 2013. Springer.

16. R. Chrétien, V. Cortier, and S. Delaune. Typing messages for free in security protocols: the case of equivalence properties. In P. Baldan and D. Gorla, editors, *Proceedings of the 25th International Conference on Concurrency Theory (CONCUR'14)*, volume 8704 of *Lecture Notes in Computer Science*, pages 372–386, Rome, Italy, Sept. 2014. Springer.
17. R. Chrétien, V. Cortier, and S. Delaune. Checking trace equivalence: How to get rid of nonces? In P. Ryan and E. Weippl, editors, *Proceedings of the 20th European Symposium on Research in Computer Security (ESORICS'15)*, Lecture Notes in Computer Science, Vienna, Austria, Sept. 2015. Springer. To appear.
18. R. Chrétien, V. Cortier, and S. Delaune. Decidability of trace equivalence for protocols with nonces. In *Proceedings of the 28th IEEE Computer Security Foundations Symposium (CSF'15)*, Verona, Italy, July 2015. IEEE Computer Society Press.
19. H. Comon-Lundh and V. Cortier. New decidability results for fragments of first-order logic and application to cryptographic protocols. In *14th International Conference on Rewriting Techniques and Applications (RTA'2003)*, volume 2706 of *LNCS*. Springer, 2003.
20. H. Comon-Lundh and V. Cortier. Computational soundness of observational equivalence. In *15th ACM Conference on Computer and Communications Security (CCS'08)*. ACM Press, 2008.
21. V. Cortier, S. Delaune, and P. Lafourcade. A survey of algebraic properties used in cryptographic protocols. *Journal of Computer Security*, 14(1):1–43, 2006.
22. C. Cremers. The Scyther Tool: Verification, falsification, and analysis of security protocols. In *Computer Aided Verification, 20th International Conference, CAV 2008, Princeton, USA, Proc.*, volume 5123/2008 of *Lecture Notes in Computer Science*, pages 414–418. Springer, 2008.
23. S. Delaune, S. Kremer, and D. Pasailă. Security protocols, constraint systems, and group theories. In B. Gramlich, D. Miller, and U. Sattler, editors, *Proceedings of the 6th International Joint Conference on Automated Reasoning (IJCAR'12)*, volume 7364 of *Lecture Notes in Artificial Intelligence*, pages 164–178, Manchester, UK, June 2012. Springer-Verlag.
24. D. Dolev and A. C. Yao. On the security of public key protocols. In *Proc. 22nd Symposium on Foundations of Computer Science (FOCS'81)*, pages 350–357. IEEE Computer Society Press, 1981.
25. N. Durgin, P. Lincoln, J. Mitchell, and A. Scedrov. Undecidability of bounded security protocols. In *Workshop on Formal Methods and Security Protocols*, Trento, Italia, 1999.
26. J. Millen and V. Shmatikov. Constraint solving for bounded-process cryptographic protocol analysis. In *Proc. 8th ACM Conference on Computer and Communications Security (CCS'01)*. ACM Press, 2001.
27. S. Mödersheim, L. Viganò, and D. A. Basin. Constraint differentiation: Search-space reduction for the constraint-based analysis of security protocols. *Journal of Computer Security*, 18(4):575–618, 2010.
28. W. Nutt. Unification in monoidal theories. In *Proc. 10th International Conference on Automated Deduction, (CADE'90)*, volume 449, pages 618–632, Kaiserslautern (Germany), 1990. Springer.
29. D. Peled. Ten years of partial order reduction. In *Proc. 10th International Conference on Computer Aided Verification, CAV'98*, volume 1427 of *Lecture Notes in Computer Science*. Springer, 1998.

30. M. Rusinowitch and M. Turuani. Protocol insecurity with finite number of sessions is NP-complete. In *Proc. 14th Computer Security Foundations Workshop (CSFW'01)*, pages 174–190. IEEE Comp. Soc. Press, 2001.
31. A. Tiu and J. E. Dawson. Automating open bisimulation checking for the spi calculus. In *Proc. 23rd IEEE Computer Security Foundations Symposium (CSF'10)*, pages 307–321. IEEE Computer Society Press, 2010.