Results on the case studies

Stéphanie Delaune and Ivan Gazeau

LSV, CNRS & ENS Cachan, Université Paris Saclay, France

Abstract. The results presented in this report is a summary of all the experiments that have been performed during the project on the case studies presented in [15].

1 Introduction

The automated study of protocols in the symbolic model already has a great support for trace properties (reachability, secrecy,...). However, equivalence properties (unlinkability, anonymity,...) were hard to check at the start of the project. The most advanced tools were ProVerif and APTE. The ProVerif analyser was however limited in the kind of processes it can analyze. ProVerif analyses bi-processes only: two processes that differ only by values but that have the same control flow. On the other hand, the theory for messages, i.e. which primitives can be used, is limited to a finite set in APTE and ProVerif can handle a bigger set but not primitive like xor. As a consequence, all cases studies presented in the report was out of reach by all analysers either due to the primitives that they use or due to conditional branchings.

Concurrently of our work two other tools have been developed: Tamarin [16] and Maude-NPA [17]. These tools are based on rewriting techniques and rely on bi-processes. They can handle unbounded number of sessions but cannot grant termination.

2 Our tools in a nutshell

We give below a brief overview of the techniques that we have developed and implemented during the project. All these tools are available on the web page of the VIP project:

http://www.lsv.ens-cachan.fr/Projects/anr-vip/index.php

2.1 APTE

The APTE tool implements the decision procedure described in [11]. It is based on the constraint solving approach (as presented in [12]). It is suitable to analyze equivalence-based properties when considering a bounded number of sessions (*i.e.* processes without replication). It considers a fixed set of cryptographic primitives, namely symmetric/asymmetric encryptions, signature, pair,

and hash functions, and is able to reason about protocols with else branches, as well as private channels, and non-deterministic choice. This procedure explores all possible symbolic traces and compute all possible resulting symbolic constraint systems on both sides. This forward symbolic exploration of two processes is finite since all symbolic traces have a bounded length and the exploration is finitely branching since inputs are abstracted away by variables and constraints. The procedure then checks the symbolic equivalence of all the resulting pairs of sets of constraint systems.

The core of APTE has been developed before the start of the VIP project. During the VIP project, we significantly improve the scope of APTE and its performance allowing one to analyse several case studies that, until now, were out of reach. First, we have extended APTE to deal with some forms of *side-channel* attacks regarding the length of messages [13]. Second, to address the so-called state space explosion problem, we have developed dedicated partial order reduction techniques and therefore dramatically reduce the number of interleavings. These techniques, presented in [6, 5], have been implemented in APTE and brought significant speed-up.

The publications related to this tool and partially supported by the VIP project are [6, 5, 13].

2.2 AKISS

The AKISS tool check equivalences properties for protocols. The equivalences properties verified by AKISS are under and over approximations of trace equivalence which is the standard equivalence property for cryptographic protocols.

The AKISS tool implements the decision procedure described in [9]. It is suitable to analyse equivalence-based properties for a bounded number of sessions. Unlike APTE, this tool deals with rich user-defined term algebras provided that they can be defined using a convergent rewriting system enjoying the *finite variant property*. This especially includes all subterm convergent equational theories. In AKISS , protocols are modelled as sets of symbolic traces with equality tests. Further, first-order Horn clauses are used to model all possible instantiations of symbolic traces, and a saturation procedure allows one to put all clauses into *solved forms*. Finally, this finite description of all possible concrete executions is used to decide equivalence between the two processes under study.

This procedure is actually able to check an over approximation (called $\approx_{\mathsf{ct}}$) and an under approximation (called $\approx_{\mathsf{ft}}$) of trace equivalence, and it has been shown that $\approx_{\mathsf{ct}}$ actually coincides with trace equivalence for a large class of processes, namely the class of determinate processes.

A first version of the algorithm has been first presented in [9]. Since then, termination of the procedure has been established for subterm convergent theories [8]. Moreover, the technique has been recently extended to analyse protocols that use *exclusive or* [4] leading to the first tool able to effectively verify equivalence properties of protocols relying on various cryptographic primitives including xor.

The publications related to this tool and partially supported by the VIP project are [9, 8]. The paper regarding the exclusive-or extension is currently under submission.

2.3 ProVerif and UKANO

Regarding the verification of equivalence-based properties for an unbounded number of sessions, a possibility is to merge the protocols under study into a so-called bi-process, and to consider a strong form of equivalence. This method has first been presented in [7] and implemented in the ProVerif tool. The approach implemented in ProVerif is flexible enough to model for instance different flavors of encryptions (symmetric, asymmetric, randomized, ...), signature, and blind signature, but excludes exclusive-or, and more generally associate and commutative operators. ProVerif is quite efficient, and terminates on many examples. However, one of its main limitation is the fact that it is not able to analyse trace equivalence (but only diff-equivalence) which is strictly stronger than trace equivalence. Basically, the two processes have to be executed exactly in the same way notably for internal rules (*e.g.* evaluation of a conditional) whereas the attacker cannot observe these details. During the project, we proposed two solutions to overcome this limitation.

First, we developed an extension of the automatic protocol verifier ProVerif in order to prove more observational equivalences. This result is obtained by pushing away the evaluation of conditionals into terms, and this allow one in particular to automatically prove anonymity in the private authentication protocol. Second, we proposed a different approach: we designed two conditions on protocols which are sufficient to ensure anonymity and unlinkability, and which can then be effectively checked automatically using ProVerif. This approach has been implemented in UKANO. This theoretical result is general enough to apply to a wide class of 2-party protocols.

The publications related to the ProVerif *and* UKANO *tools and partially supported by the VIP project are [10, 14].*

3 Experimentation

In this section, we report on some case studies that have been done during the project. We mainly consider the case studies that have been described in [15] and that have been used as a guideline for our research agenda.

3.1 Some electronic voting protocols

Privacy-type properties are critical in e-voting protocols, and we conduct some experiments on the two e-voting protocols (namely FOO, and Okamoto) as mentioned in [15]. We consider *voter privacy*: the adversary should not be able to learn how each voter voted. Voter privacy is naturally modeled as an equivalence property: it is not possible to distinguish the situation where honest voter

A votes yes and honest B votes no from the situation that A votes no and B votes yes.

FOO protocol. The FOO protocol relies on blind signatures and a commitment function. We model these primitives through a rewrite system which is not subterm convergent, but it is optimally reducing (and therefore it falls into the class of primitives that can be handled by AKISS). Note that only two honest voters need to be modelled for showing anonymity. All remaining voters and election authorities are subsumed by the adversary. On a standard modern laptop, AKISS takes a few seconds to carry out the above verification.

Okamoto protocol. The Okamoto protocol is a variant of the FOO protocol which aims at achieving receipt-freeness. To avoid vote-selling, a voter should not be able to provide a receipt of how he voted to a potential coercer. In the FOO protocol this is possible by sending all private names to a coercer. The main tool to avoid this problem in the Okamoto protocol is the use of trapdoor commitment functions. These functions allow to change the value of committed vote using a secret value called the trapdoor. To model this new primitive, we rely on a rewrite system that is out of the scope of most tools, even in the simpler case of a passive adversary. However, this rewrite system is optimally reducing and can be handled by AKISS. The modeling of the Okamoto protocol requires private channels. As we do not have private channels in the calculus used by AKISS, we transformed the protocol so that every message sent by honest participants on a private channel is sent encrypted under a key not known to the adversary On a standard modern laptop, AKISS takes about 30 seconds to carry out the above verification.

To our knowledge, AKISS was the first tool able to handle these two e-voting protocols automatically. In addition, we have also proposed a way to model a stronger notion of privacy, namely *everlasting privacy*. We modified AKISS and ProVerif to verify several variants of the Helios e-voting protocol w.r.t. everlasting privacy. All the conclusions about this additional case studies are available in [3].

3.2 Some RFID protocols

Many protocols implemented on low-power devices, such as RFID tags, often rely on the bitwise exclusive or (xor) operator because of its computational efficiency [18], and unlinkability is an important issue for such an application. Again, unlinkability can be modeled through an equivalence property. Roughly, the idea is to see whether an attacker can observe a difference between an ideal situation where each tag is used only once from a situation where a tag can be used more than once.

To be able to deal with these case studies, we have developed a novel procedure (see [4]) that allows one to verify equivalence between finite processes, *i.e.*, processes without replication, for protocols that use *exclusive or* (xor). We have implemented our procedure in the recent tool AKISS, and we have used it

to check unlinkability on various RFID protocols. We obtain therefore the first tool that is able to effectively verify equivalence properties of protocols relying on various cryptographic primitives including xor.

In total we modelled 5 RFID protocols from [18]: the KCL, LD, LAK, OTYT and YPL protocols. On 4 of the 5 protocols we find (known) attacks which violate unlinkability. LAK was proved to be unlinkable in our setting. Note that for our tool finding attacks or proving their absence are mainly equally difficult tasks, as for both we need to complete the saturation of the traces. Regarding efficiency, while small examples take a few seconds, the largest examples take several hours.

3.3 E-passport application

Several protocols are used to protect the information stored inside the e-passport. As described in [15], the BAC protocol is a key establishment protocol that is supposed to be unlinkable. The purpose of this protocol is to establish a fresh key that is used in subsequent communications to encrypt personal data. We also consider two authentication protocols that are executed after the BAC protocol. Actually, we also performed some analyses considering these 3 protocols together.

All these protocols use quite standard primitives but require else branches to be modelled faithfully. Therefore the only suitable tools to perform a security analysis are APTE, ProVerif and UKANO. When considering anonymity, the strong notion of equivalence used in ProVerif is sufficient to conclude. However, diff-equivalence is too strong and does not allow one to conclude regarding the unlinkability property. The traceability attack (as described in [2]) on the French implementation of the passport can now be obtained automatically relying on the tool APTE. This traceability attack comes from the fact that the French implementation use error messages that are too precise at some places. Also, during experimentation, it has been noticed that the length of the messages that are outputted by different passports are not always the same, and this may also leak some information. Therefore, the extension devised during the project are particularly interesting for this case study. We have discovered a new flaw against privacy in the biometric passport protocol. We show that an attacker can break privacy by observing messages of different lengths depending on which passport is used, therefore discovering who between Alice or Bob is currently using her/his passport.

The technique implemented in UKANO also allows one to conclude on this case study considering an arbitrary number of sessions. This leads to the first formal security proof of protocol (the version that uses the same error message). This technique has also been successfully used to analyse the BAC protocol in combination with passive (resp. active) authentication protocol.

PACE protocol. The Password Authenticated Connection Establishment protocol (PACE) has been proposed by the German Federal Office for Information Security (BSI) to replace the BAC protocol. It has been studied in the literature but to the best of our knowledge, no formal proofs about privacy have been given. Similarly to BAC, the purpose of PACE is to establish a secure channel

based on an optically-scanned key k. Unlike BAC, it relies on Diffie-Hellman exchanges.

The technique implemented in UKANO (which used ProVerif as a back-end and therefore allows one to model modular exponentiation) is flexible enough to analyse several variant of the PACE protocol. Our security analysis allows one to highlight an imprecision in the official specification that may lead to practical attacks on unlinkability. Second, we report on an attack that that we discovered using our method on some modelling of PACE found in the literature. Third, we turn to PACE as properly understood from the official specification: when the latter test is present and the decryption may not fail. In that case, we report on a new attack. In practice, this flaw seems hard to exploit but it could be a real privacy concern: if a tag initiates multiple readers, an attacker may learn which ones it had initiated by forwarding messages from one to another. It does not seem to be realistic in the e-passport scenario, but could be harmful in other contexts.

3.4 3G mobile phones

Modelling the AKA protocol in a faithful way is challenging since it use a sequence number that is stored from one session to another, and therefore there is a need of mutable states. Instead, we consider a simplest version in which the sequence number (modeled as a nonce) is shared by the 2 parties. With such a modelling, a tool like ProVerif is able to establish anonymity and unlinkability. Therefore, we do not conduct specific experiments on this application relying on our own verification tools. However, this case study has been used to illustrate the results we have obtained regarding composition issues (see [1]).

3.5 Private authentication

This case study was not presented in [15]. However, it has served as a guideline to develop/improve some of procedures we have developed during the VIP project, and we obtain several interesting results on it. A description is available in [10].

APTE was the first tool to be able to analyse a privacy property on this protocol but only for a bounded number of sessions. Moreover, due to the high complexity of the algorithm, only a very few number of sessions could be analysed using APTE. The POR techniques presented in [6] and implemented in APTE greatly improved the situation allowing one to perform a security analysis considering 7 processes in parallel. Actually, thanks to the improvements presented in [10] and implemented in ProVerif, a security analysis for an unbounded number of sessions has been successfully performed using ProVerif.

4 Conclusion

This report describes the experiments that we have conducted during the project on the case studies that are presented in [15]. Most of the protocols that were

out of the scope of the existing verification tools can now be analysed (at least partially) using tools developed during the project. In particular, we obtain the first formal security proof (for an unbounded number of sessions) of the BAC protocol relying on our automatic tool UKANO. During our experimentations, we have also discovered several vulnerabilities of the PACE protocol (used in the e-passport application).

As future work, we are considering improving the implementation of AKISS with xor to make it more performant. This can be achieved by removing redundant clauses and by using a native unifier instead of Maude each time the unification belongs to a simple class. The techniques using partial-order reduction applied on APTE to reduce the number of equivalent trace are likely to be applicable on AKISS and should also speed-up the decision algorithm. On the other hand, the techniques used by UKANO, which have been implemented in top of ProVerif, does not rely on the ProVerif principle: it should be possible to adapt the tool such that it can be plugged on Tamarin. This would allow to prove protocol with modular exponentiation which are handled by Tamarin but not by ProVerif.

References

1. M. Arapinis, V. Cheval, and S. Delaune. Composing security protocols: from confidentiality to privacy. In R. Focardi and A. Myers, editors, *Proceedings of the 4th International Conference on Principles of Security and Trust (POST'15)*, volume 9036 of *Lecture Notes in Computer Science*, pages 324–343, London, UK, Apr. 2015. Springer.
2. M. Arapinis, T. Chothia, E. Ritter, and M. Ryan. Analysing unlinkability and anonymity using the applied pi calculus. In *Proc. of 23rd IEEE Computer Security Foundations Symposium (CSF'10)*, pages 107–121. IEEE Computer Society Press, 2010.
3. 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.
4. D. Baelde, S. Delaune, I. Gazeau, and S. Kremer. Verification of privacy-type properties for security protocols with XOR. Technical report, 2016.
5. D. Baelde, S. Delaune, and L. Hirschi. A reduced semantics for deciding trace equivalence using constraint systems. In *Proc. of the 3rd International Conference on Principles of Security and Trust (POST'14)*, volume 8414 of *LNCS*, pages 1–21, Grenoble, France, 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. B. Blanchet, M. Abadi, and C. Fournet. Automated Verification of Selected Equivalences for Security Protocols. In *Symposium on Logic in Computer Science*, pages 331–340, Chicago, IL, June 2005. IEEE Comp. Soc. Press.

8. R. Chadha, V. Cheval, Ş. Ciobâcă, and S. Kremer. Automated verification of equivalence properties of cryptographic protocol. *ACM Transactions on Computational Logic*, 2016. To appear.
9. 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.
10. V. Cheval and B. Blanchet. Proving more observational equivalences with proverif. In D. Basin and J. Mitchell, editors, *Proceedings of the 2nd International Conference on Principles of Security and Trust (POST'13*, Lecture Notes in Computer Science, pages 226–246, Roma, Italy, Mar. 2013. Springer.
11. 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.
12. V. Cheval, V. Cortier, and S. Delaune. Deciding equivalence-based properties using constraint solving. *Theoretical Computer Science*, 492:1–39, June 2013.
13. V. Cheval, V. Cortier, and A. Plet. Lengths may break privacy – or how to check for equivalences with length. In *Proc. of the 25th International Conference on Computer Aided Verification (CAV'13)*, volume 8044 of *LNCS*, pages 708–723. Springer, 2013.
14. L. Hirschi, D. Baelde, and S. Delaune. A method for verifying privacy-type properties: the unbounded case. In M. Locasto, V. Shmatikov, and Ú. Erlingsson, editors, *Proceedings of the 37th IEEE Symposium on Security and Privacy (S&P'16)*, San Jose, California, USA, May 2016. IEEECSP. To appear.
15. L. Hirschi and S. Delaune. Description of some case studies. Deliverable VIP 6.1, (ANR-11-JS02-0006), Sept. 2013. 14 pages.
16. S. Meier, B. Schmidt, C. Cremers, and D. Basin. *The TAMARIN Prover for the Symbolic Analysis of Security Protocols*, pages 696–701. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.
17. S. Santiago, S. Escobar, C. Meadows, and J. Meseguer. *A Formal Definition of Protocol Indistinguishability and Its Verification Using Maude-NPA*, pages 162–177. Springer International Publishing, Cham, 2014.
18. T. van Deursen and S. Radomirovic. Attacks on rfid protocols. *IACR Cryptology ePrint Archive*, 2008:310, 2008.