Math. Struct. in Comp. Science (2015), *vol.* 25, *pp.* 203–206. © Cambridge University Press 2014
doi:10.1017/S0960129513000583 First published online 10 November 2014

Preface to the special issue on quantitative information flow

MIGUEL E. ANDRÉS[†,1], CATUSCIA PALAMIDESSI[‡]
and GEOFFREY SMITH[§]

[†]*LIX, École Polytechnique, Palaiseau, France*
Email: mandres@lix.polytechnique.fr
[‡]*INRIA Saclay and LIX, Palaiseau, France*
Email: catuscia@lix.polytechnique.fr
[§]*School of Computing and Information Sciences, Florida International University, Miami, Florida, U.S.A.*
Email: smithg@cis.fiu.edu

A long-standing and fundamental issue in computer security is to control the *flow of information*, whether to prevent confidential information from being *leaked*, or to prevent trusted information from being *tainted*. While there have been many efforts aimed at preventing improper flows completely (see for example, the survey by Sabelfeld and Myers (2003)), it has long been recognized that perfection is often impossible in practice. A basic example is a login program – whenever it rejects an incorrect password, it unavoidably reveals that the secret password differs from the one that was entered. More subtly, systems may be vulnerable to *side channel* attacks, because observable characteristics like running time and power consumption may depend, at least partially, on sensitive information.

For these reasons, the possibility of *quantifying* information flow becomes attractive, as this could allow certain improper flows to be tolerated on the grounds that they are ‘small’. While there was early work on quantitative information flow by Denning (1983), Millen (1987), McLean (1990) and Gray (1991), the area received relatively little attention until the past decade, when it was revitalized starting with the efforts of Clark, Hunt, and Malacaria (2001).

In the past decade, there has been too much work for a comprehensive survey here, but we can briefly describe the main themes that have been explored.

From the perspective of *foundations*, there have been a variety of studies aimed at defining quantitative measures of information flow for a variety of system models, establishing the operational significance of the measures with respect to security, and establishing their mathematical properties, including relationships among the different measures that have been considered. Papers with a foundational focus include those of Clarkson *et al.* (2005), Köpf and Basin (2007), Malacaria (2007), Chatzikokolakis *et al.* (2008), Smith (2009), McIver *et al.* (2010), Barthe and Köpf (2011) and Alvim *et al.* (2012).

From the perspective of *verification techniques*, there have been studies of a variety of analysis methods. A type system for analysing information flow is presented by Clark

[1] Current affiliation: Google, Inc.

et al. (2007), while model checking techniques are considered by Backes *et al.* (2009), Newsome *et al.* (2009), Andrés *et al.* (2010) and Heusser and Malacaria (2010).

Statistical sampling is used by Chatzikokolakis *et al.* (2010) and Köpf and Rybalchenko (2010). Also, the computational complexity of quantitative information flow problems is studied by Yasuoka and Terauchi (2010).

Finally, *applications* are beginning to appear, in which quantitative leakage analyses have been done for real system vulnerabilities. For example, Köpf and Smith (2010) analyse timing attacks against blinded RSA cryptography, Heusser and Malacaria (2010) analyse leakage due to Linux kernel bugs, and Köpf *et al.* (2012) analyse cache attacks against AES cryptography.

After soliciting contributions to a special issue of *Mathematical Structures in Computer Science* on Quantitative Information Flow, we received submissions from leading researchers in the field. Following careful reviewing, we selected eight papers for inclusion in this special issue.

In *Quantification of Integrity*, Clarkson and Schneider consider how to quantify integrity, an important topic that has received much less attention than has the quantification of confidentiality. The authors argue that integrity has several facets, which they call ‘contamination’ and ‘suppression’. Contamination is concerned with the amount of untrusted information that flows into trusted outputs, while suppression is concerned with the amount of trusted information that fails to flow into trusted outputs. They model these concepts formally, establish some consequences, and develop applications to differential privacy and to belief-based information flow.

In *A Semiring-based Trace Semantics for Processes with Applications to Information Leakage Analysis*, Boreale, Clark and Gorla present a formalism to specify and calculate aspects of information leakage, inspired by ideas from language-based security, coalgebraic formalisms and process algebra. Processes are specified in a process algebra whose semantics is given as a formal power series over some generic semiring. They provide an equivalent compositional semantics, laying an abstract foundation for information leakage analysis.

In *Asymptotic Information Leakage under One-Try Attacks*, Boreale, Pampaloni and Paolini consider the leakage resulting from n independent repetitions of a probabilistic channel, using the same secret input in each repetition. Considering the leakage asymptotically, they show that (assuming a uniform prior distribution on the secret) the min-entropy leakage converges to the logarithm of the number of distinct rows in the channel matrix. Also, they use the information-theoretic method of types to prove bounds on the rate of convergence. Finally, they generalize to a hidden Markov model, in which each run of the channel produces an infinite trace of outputs.

In *Hidden-Markov Program Algebra with Iteration*, McIver, Meinicke and Morgan introduce a quantitative compositional semantics for programs with iteration, and a notion of refinement which compares two programs with respect to their information leakage. The authors also provide algebraic laws to help in reasoning about programs and their composition, supporting a stepwise refinement approach to secure system construction.

In *Quantifying Opacity*, Bérard, Mullins and Sassolas consider probabilistic generalizations of the *opacity*. In the purely nondeterministic setting, (symmetric) opacity holds

https://doi.org/10.1017/S0960129513000583 Published online by Cambridge University Press

of a system predicate ϕ if no observable output ever reveals whether or not ϕ holds. The authors generalize to probabilistic systems and consider several quantitative opacity notions based on the probability that an adversary will be able to deduce, or perhaps guess, whether or not ϕ holds. They also give algorithms to compute these quantities.

In *Algebraic Foundations for Quantitative Information Flow*, Malacaria explores the foundations of quantitative information flow in the special case of deterministic systems. A deterministic system induces a partition on the space of secret inputs, where each block corresponds to the set of secret inputs that map to a particular output. Under the lattice of information, partitions are partially ordered by the relation of *partition refinement*. The author explores the implications of this algebraic structure for the leakage of deterministic systems, showing for instance that the property that P_1's partition is refined by P_2's is *equivalent* to the property that the leakage of P_1 never exceeds that of P_2, no matter the prior; moreover, this strong leakage ordering is the same, whether leakage is measured by Shannon entropy, min-entropy, or guessing entropy.

In *An Analysis of Trust in Anonymity Networks in the Presence of Adaptive Attackers*, Sassone, Hamadou and Yang discuss anonymous communication systems and explore the idea of enhancing such systems with a notion of trust. Such a combination could, in principle, improve not only the reliability of the system (by avoiding nodes with bad reputation) but also the anonymity guarantees (by increasing the chances of communicating with honest nodes). The paper analyses the implications of trust on the crowds and onion routing protocols, showing benefits obtained by trust but also possible attacks on the trust mechanisms themselves.

Finally, in *Quantifying Information Flow in Cryptographic Systems*, Backes and Köpf introduce a new definition of quantitative information flow, *transmissible information*, that is able to capture both cryptographic systems (with computationally bounded adversaries) as well as information-theoretically secure systems. They show that transmissible information is preserved under universal composability, giving a way to lift quantitative bounds from idealized to actual cryptographic systems.

We are grateful to the authors for their excellent submissions to this special issue, and to the referees for their careful efforts to ensure the high quality of this special issue.

References

Alvim, M. S., Chatzikokolakis, K., Palamidessi, C. and Smith, G. (2012) Measuring information leakage using generalized gain functions. In: *Proceedings 25th IEEE Computer Security Foundations Symposium* 265–279.

Andrés, M. E., Palamidessi, C., van Rossum, P. and Smith, G. (2010) Computing the leakage of information-hiding systems. In: Esparza, J. and Majumdar, R. (eds.) Tools and Algorithms for the Construction and Analysis of Systems. *Lecture Notes in Computer Science* **6015** 373–389.

Backes, M., Köpf, B. and Rybalchenko, A. (2009) Automatic discovery and quantification of information leaks. In: *Proceedings 30th IEEE Symposium on Security and Privacy* 141–153.

Barthe, G. and Köpf, B. (2011) Information-theoretic bounds for differentially private mechanisms. In: *Proceedings 24th IEEE Computer Security Foundations Symposium* 191–204.

https://doi.org/10.1017/S0960129513000583 Published online by Cambridge University Press

Chatzikokolakis, K., Chothia, T. and Guha, A. (2010) Statistical measurement of information leakage. In: Esparza, J. and Majumdar, R. (eds.) Tools and Algorithms for the Construction and Analysis of Systems. *Lecture Notes in Computer Science* **6015** 390–404.

Clark, D., Hunt, S., and Malacaria, P. (2001) Quantitative analysis of the leakage of confidential data. In: Proceedings Workshop on Quantitative Aspects of Programming Languages. *Electronic Notes in Theoretical Computer Science* **59** (3) 238–251.

Clark, D., Hunt, S. and Malacaria, P. (2007) A static analysis for quantifying information flow in a simple imperative language. *Journal of Computer Security* **15** 321–371.

Clarkson, M., Myers, A. and Schneider, F. (2005) Belief in information flow. In: *Proceedings 18th IEEE Computer Security Foundations Workshop* 31–45.

Chatzikokolakis, K., Palamidessi, C. and Panangaden, P. (2008) On the Bayes risk in information-hiding protocols. *Journal of Computer Security* **16** (5) 531–571.

Denning, D. (1983) *Cryptography and Data Security*, Addison-Wesley.

Gray, J. W. III (1991) Toward a mathematical foundation for information flow security. In: *IEEE Symposium on Security and Privacy* 21–35.

Heusser, J. and Malacaria, P. (2010) Quantifying information leaks in software. In: *Proceedings of the Annual Computer Security Applications Conference* 261–269.

Köpf, B. and Basin, D. (2007) An information-theoretic model for adaptive side-channel attacks. In: *Proceedings of the 14th ACM Conference on Computer and Communications Security* 286–296.

Köpf, B., Mauborgne, L. and Ochoa, M. (2012) Automatic quantification of cache side-channels. In: *Proceedings of the 24th International Conference on Computer-Aided Verification* 564–580.

Köpf, B. and Rybalchenko, A. (2010) Approximation and randomization for quantitative information-flow analysis. In: *Proceedings of the 23nd IEEE Computer Security Foundations Symposium* 3–14.

Köpf, B. and Smith, G. (2010) Vulnerability bounds and leakage resilience of blinded cryptography under timing attacks. In: *Proceedings of the 23nd IEEE Computer Security Foundations Symposium* 44–56.

Malacaria, P. (2007) Assessing security threats of looping constructs. In: *Proceedings of the 34th Symposium on Principles of Programming Languages* 225–235.

McLean, J. (1990) Security models and information flow. In: *IEEE Symposium on Security and Privacy* 180–189.

Millen, J. K. (1987) Covert channel capacity. In: *IEEE Symposium on Security and Privacy* 60–66.

McIver, A., Meinicke, L. and Morgan, C. (2010) Compositional closure for Bayes risk in probabilistic noninterference. In: *Proceedings of the International Colloquium on Automata, Languages and Programming* 223–235.

Newsome, J., McCamant, S. and Song, D. (2009) Measuring channel capacity to distinguish undue influence. In: *Proceedings of the Fourth Workshop on Programming Languages and Analysis for Security* 73–85.

Sabelfeld, A. and Myers, A. C. (2003) Language-based information flow security. *IEEE Journal on Selected Areas in Communications* **21** (1) 5–19.

Smith, G. (2009) On the foundations of quantitative information flow. In: de Alfaro, L. (ed.) Proceedings of the 12th International Conference on Foundations of Software Science and Computational Structures. *Lecture Notes in Computer Science* **5504** 288–302.

Yasuoka, H. and Terauchi, T. (2010) Quantitative information flow—verification hardness and possibilities. In: *Proceedings of the 23nd IEEE Computer Security Foundations Symposium* 15–27.

https://doi.org/10.1017/S0960129513000583 Published online by Cambridge University Press