Lecture Notes in Computer Science **8902**

Commenced Publication in 1973
Founding and Former Series Editors:
Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

More information about this series at http://www.springer.com/series/7407

Matteo Maffei · Emilio Tuosto (Eds.)

Trustworthy Global Computing

9th International Symposium, TGC 2014
Rome, Italy, September 5–6, 2014
Revised Selected Papers

Springer

Editors
Matteo Maffei
Department of Computer Science
Saarland University
Saarbrücken
Germany

Emilio Tuosto
Computer Science
University of Leicester
Leicester
United Kingdom

ISSN 0302-9743
Lecture Notes in Computer Science
ISBN 978-3-662-45916-4
DOI 10.1007/978-3-662-45917-1

ISSN 1611-3349 (electronic)

ISBN 978-3-662-45917-1 (eBook)

Library of Congress Control Number: 2014959178

LNCS Sublibrary: SL1 – Theoretical Computer Science and General Issues

Springer Heidelberg New York Dordrecht London

Printed on acid-free paper

Springer-Verlag GmbH Berlin Heidelberg is part of Springer Science+Business Media
(www.springer.com)

Preface

This volume contains the proceedings of TGC 2014, the Ninth International Symposium on Trustworthy Global Computing. The symposium was held in Rome, Italy, during September 5–6, 2014. It was colocated with the CONCUR and IFIP-TCS. Informal pre-proceedings were made available in electronic form to the participants. The papers in this volume were further improved by the authors, in response to helpful feedback received at the symposium.

The Symposium on Trustworthy Global Computing is an international annual venue dedicated to safe and reliable computation in the so-called global computers, i.e., those computational abstractions emerging in large-scale infrastructures such as service-oriented architectures, autonomic systems, and cloud computing systems. It focuses on frameworks, tools, algorithms, and protocols for open-ended, large-scale systems and applications, and on rigorous reasoning about their behavior and properties. The underlying models of computation incorporate code and data mobility over distributed networks that connect heterogeneous devices, often with dynamically changing topologies. In this context, TGC 2014 focused on secure and reliable computation.

The first TGC event took place in Edinburgh in 2005, with the co-sponsorship of IFIP TC-2, as part of ETAPS 2005. TGC 2005 was the evolution of the previous Global Computing I workshops held in Rovereto in 2003 and 2004 (see LNCS vol. 2874) as well as of the workshops on Foundation of Global Computing held as satellite events of ICALP and CONCUR (see ENTCS vol. 85). Four editions of TGC were co-located with the reviews of the EU-funded projects AEOLUS, MOBIUS, and SENSORIA within the FP6 initiative. They were held in Lucca, Italy (TGC 2006, LNCS vol. 4661); in Sophia Antipolis, France (TGC 2007, LNCS vol. 4912); in Barcelona, Spain (TGC 2008, LNCS vol. 5474); and in Munich, Germany (TGC 2010, LNCS vol. 6084). Further editions of TGC were held in Aachen, Germany (TGC 2011, LNCS vol. 7173) and in Newcastle upon Tyne, UK (TGC 2012, LNCS vol. 8191). In 2013, TGC was held in Buenos Aires (LNCS vol. 8358).

TGC 2014 solicited contributions in all areas of global computing, including (but not limited to) theories, languages, models, and algorithms; language concepts and abstraction mechanisms; security, trust, privacy, and reliability; resource usage and information flow policies; software development and software principles; model checkers, theorem provers, and static analyzers.

The fruitful collaboration with CONCUR, initiated in 2013, was continued this year allowing concurrent submissions to CONCUR and TGC, with the reviewing schedule of TGC slightly delayed with respect to that of CONCUR and submissions accepted by CONCUR were automatically withdrawn from TGC. This year there were 5 papers concurrently submitted to TGC and CONCUR and 15 papers were submitted only to TGC (a paper initially submitted was later withdrawn). This time, the papers submitted at CONCUR and TGC were also reviewed by the Program Committee of TGC, praiseworthy for the effort put in the thorough revision process and the accurate discussion.

The Program Committee selected 12 papers to be included in this volume and to be presented at the symposium (2 of which were conditionally accepted in the first instance). The program was structured in sessions named "Theory" (chaired by Michele Loreti), "Session Types" (chaired by Massimo Bartoletti and Stephanie Delaune), "Security" (chaired by Peter Thiemann), "Cryptographic Protocol Analysis" (chaired by Matteo Maffei). Also, a session on "Brief Announcements" (chaired by Emilio Tuosto) gave the possibility to participants to present (on-going) work not included in this proceedings. Finally, TGC's program had invited lectures of Véronique Cortier (CNRS, France) and, jointly with CONCUR, of Catuscia Palamidessi (Inria Saclay and LIX, France).

We would like to thank the Steering Committee of TGC for inviting us to chair the conference; the members of the Program Committee and external referees for their detailed reports and the stimulating discussions during the review phase; the authors of submitted papers, the invited speakers, the session chairs, and the attendees for contributing to the success of the event. We are also grateful to Paolo Baldan and Daniele Gorla, the chairs of CONCUR 2014, with whom we had the opportunity to collaborate. Finally, we thank the providers of the EasyChair system, which was used to manage the submissions.

November 2014

Matteo Maffei
Emilio Tuosto

Organization

Program Committee

Stephen Chong	Harvard University, USA
Anupam Datta	CMU, USA
Stephanie Delaune	CNRS, LSV, France
Mariangiola Dezani-Ciancaglini	Università di Torino, Italy
Fabio Gadducci	Università di Pisa, Italy
Dan Ghica	University of Birmingham, UK
Andrew D. Gordon	Microsoft Research and University of Edinburgh, UK
Joshua Guttman	Worcester Polytechnic Institute, USA
Daniel Hirschkoff	ENS Lyon, France
Christos Kaklamanis	University of Patras and CTI, Greece
Boris Köpf	IMDEA Software Institute, Spain
Alberto Lluch Lafuente	IMT Institute for Advanced Studies Lucca, Italy
Michele Loreti	Università degli Studi di Firenze, Italy
Matteo Maffei	Informatik, Saarland University, Germany
Sergio Maffeis	Imperial College London, UK
Hernan Melgratti	Universidad de Buenos Aires, Argentina
António Ravara	Universidade Nova de Lisboa, Portugal
Alejandro Russo	Chalmers University of Technology, Sweden
Andrey Rybalchenko	TUM, Germany
Emilio Tuosto	University of Leicester, UK
Björn Victor	Uppsala University, Sweden
Roberto Zunino	Università degli Studi di Trento, Italy

Additional Reviewers

Barbanera, Franco
Capecchi, Sara
Castellani, Ilaria
Celestini, Alessandro
Coppo, Mario
Crafa, Silvia
Giunti, Marco
Grossi, Roberto
Knowles, Kenneth
Lozes, Etienne
Mostrous, Dimitris
Santini, Francesco
Tinacci, Marco
Vigliotti, Maria Grazia
Weber, Tjark

Abstracts

Generalized Bisimulation Metrics

Konstantinos Chatzikokolakis[1,2], Daniel Gebler[3], Catuscia Palamidessi[4,2], and Lili Xu[2,5]

[1] CNRS
[2] LIX, Ecole Polytechnique
[3] VU University Amsterdam
[4] INRIA
[5] Institute of Software, Chinese Academy of Science

Abstract

Originally proposed in the seminal works of van Breugel and Worrel [2,3] and of Desharnais et al. [6,7], the pseudometric based on the Kantorovich lifting has become very popular in the process algebra community. One reason for its success is that, when dealing with probabilistic processes, distances are more suitable than equivalences, since the latter are not robust wrt small variation of probabilities. Another important reason is that, thanks to the dual presentation of the Kantorovich lifting in terms of the mass transportation problem, the distance can be efficiently computed using linear programming algorithms [1,2]. Furthermore, this pseudometric is an extension of probabilistic bisimilarity, in the sense that two states have distance distance 0 if and only if they are bisimilar. Furthermore, it is defined as the greatest fixpoint of a transformation that has the same structure as the one used for bisimilarity. This allows to transfer some of the concepts and methods that have been extensively explored in process algebra, and to use lines of reasoning which the process algebra community is familiar with. Along the same lines, a nice property of this pseudometric is that the standard operators of process algebra are non-expansive wrt it. This generalizes the result that bisimulation is a congruence, and can be used in a similar way, for compositional reasoning and verification.

Last but not least, the Kantorovich bisimilarity metric provides a bound on the corresponding distance on probabilistic traces [5] (corresponding in the sense that the definition is based on the same Kantorovich lifting). This means that it can be used to verify certain probabilistic properties on traces. More specifically, it can be used to verify properties that are expressed in terms of difference between probabilities of sets of traces. These properties are linear, in the sense that the difference increases linearly wrt variations on the distributions.

Many properties, however, such as several privacy and security ones, are not linear. This is the case of the popular property of differential privacy [8], which is expressed in terms of ratios of probabilities. In fact, there are processes that have small Kantorovich distance, and which are not ϵ-differentially private for any finite ϵ. Another example are

This work has been partially supported by the project ANR-12-IS02-001 PACE, the project ANR-11-IS02-0002 LOCALI, the INRIA Equipe Associée PRINCESS, the INRIA Large Scale Initiative CAPPRIS, and the EU grant 295261 MEALS.

the properties used in quantitative information flow, which involve logarithmic functions on probabilities.

An interesting line of research, therefore, is to generalize the Kantorovich lifting to obtain a family of pseudometrics suitable for the verification of a wide class of properties, following the principles that:

i. the members of this family should depend on a parameter related to the class of properties (on traces) that we wish to verify,
ii. each member should provide a bound on the corresponding distance on trace distributions,
iii. the kernel of each member should correspond to probabilistic bisimilarity,
iv. the general construction should be coinductive,
v. the typical process-algebra operators should be non-expansive,
vi. each member should be feasible to compute.

In a recent work [4], we have achieved the first four desiderata. Regarding the last two, so far we have studied a particular case (hereafter called multiplicative variant of the Kantorovich lifting) based on the notion of distance used in the definition of differential privacy. We were able to find a dual form of the lifting, which allows to reduce the problem of its computation to a linear optimization problem solvable with standard algorithms. We have also proved that several typical process-algebra operators are non-expansive, and we have given explicitly the expression of the bound. As an example of application of our framework, we have shown how to instantiate our construction to obtain the multiplicative variant of the Kantorovich pseudometric, and how to use it to verify the property of differential privacy.

References

1. Bacci, G., Bacci, G., Larsen, K.G., Mardare, R.: On-the-fly exact computation of bisimilarity distances. In: TACAS. LNCS, vol. 7795, pp. 1–15. Springer (2013)
2. van Breugel, F., Worrell, J.: An algorithm for quantitative verification of probabilistic transition systems. In: Proc. of CONCUR'01. pp. 336–350. Springer (2001)
3. van Breugel, F., Worrell, J.: Towards quantitative verification of probabilistic transition systems. In: Proc. of ICALP. LNCS, vol. 2076, pp. 421–432. Springer (2001)
4. Chatzikokolakis, K., Gebler, D., Palamidessi, C., Xu, L.: Generalized bisimulation metrics. In: Proc. of CONCUR. LNCS, vol. 8704, pp. 32–46. Springer (2014)
5. Chen, D., van Breugel, F., Worrell, J.: On the complexity of computing probabilistic bisimilarity. In: FOSSACS. LNCS, vol. 7213, pp. 437–451. Springer (2012)
6. Desharnais, J., Gupta, V., Jagadeesan, R., Panangaden, P.: Metrics for labeled markov systems. In: Proc. of CONCUR. LNCS, vol. 1664, pp. 258–273. Springer (1999)
7. Desharnais, J., Jagadeesan, R., Gupta, V., Panangaden, P.: The metric analogue of weak bisimulation for probabilistic processes. In: Proc. of LICS. pp. 413–422. IEEE (2002)
8. Dwork, C.: Differential privacy. In: Proc. of ICALP. LNCS, vol. 4052, pp. 1–12. Springer (2006)

Electronic Voting: How to Ensure Privacy and Verifiability

Véronique Cortier

LORIA - CNRS, France

Electronic voting is now used in many countries such as Estonia, United States, Norway, Canada, or France. It provides a convenient way to tally the votes and enables voters to vote from their home. However, it also raises controversy about the security and transparency of the election. How to make sure each vote is counted as intended? How to protect vote privacy? Therefore, several countries like Germany, Netherlands, or the United Kingdom have stopped electronic voting, at least momentarily [5].

Electronic should offer at least the same guarantees than traditional paper-based voting systems. Namely, it should ensure both ballot privacy (no one should know my vote) and verifiability (the result of the election corresponds to the votes casted by voters). Even stronger than ballot privacy is receipt-freeness or coercion-resistance: each vote should remain private even when a voter is willing to tell how he voted. This is to protect against vote buying and coercion. These properties are antagonist: a voter should not be able to prove how he voted, yet he should be able to check that his vote has been counted. Designing a secure e-voting protocol therefore requires a fine tuning between these two properties.

The two main academic Internet voting protocols are Civitas [2] and Helios [1]. We refer the reader to [3] for a short survey on other electronic voting systems. Civitas offers the best security guarantees: coercion-resistance and full verifiability. It might however be difficult to use in practice. Helios is not coercion-resistant but offers ballot privacy and verifiability except for voter eligibility. The fact that Helios does not cover eligibility verifiability means that a compromised ballot box may add ballots. This is fixed by a variant of Helios, named Belenios [4].

Formal methods have been successful in the analysis of standard security protocols such as key exchange or authentication protocols. Electronic voting however challenges the current techniques. First, electronic voting systems often make use of more subtle cryptographic primitives that include blind signatures, homomorphic encryption, or zero-knowledge proofs. Most existing tools cannot yet handle these primitives, or at least not all of them. Moreover, most verification techniques for security protocols are dedicated to trace properties such as authentication or confidentiality. However, ballot privacy is typically expressed as a behavioral equivalence. Decision procedures for equivalence-based properties are yet under development. Another challenge in the context of electronic voting is to identify and model security properties. What is a good electronic voting system? This question is yet to be answered. Even the definition of a key property like ballot privacy has a lot of variants and is still under study. We refer

The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement no 258865, project ProSecure.

the reader to [4] for a more detailed survey on existing techniques for the analysis of e-voting systems.

References

1. Ben Adida, Olivier de Marneffe, Oliver Pereira, and Jean-Jacques Quisquater. Electing a university president using open-audit voting: Analysis of real-world use of Helios. In *Proceedings of the 2009 conference on Electronic voting technology/workshop on trustworthy elections*, 2009.
2. Michael R. Clarkson, Stephen Chong, and Andrew C. Myers. Civitas: Toward a secure voting system. In *Proc. IEEE Symposium on Security and Privacy*, pages 354–368, 2008.
3. Véronique Cortier. Electronic voting: How logic can help. In *Proceegings of the 12th International Joint Conference on Automated Reasoning (IJCAR 2014)*, volume 8562 of *LNAI*, pages 16–26, Vienna, Austria, 2014.
4. Véronique Cortier, David Galindo, Stéphane Glondu, and Malika Izabachene. Election verifiability for helios under weaker trust assumptions. In *Proceedings of the 19th European Symposium on Research in Computer Security (ESORICS'14)*, LNCS, Wroclaw, Poland, September 2014. Springer.
5. Jordi Barrat i Esteve, Ben Goldsmith, and John Turner. International experience with e-voting. Technical report, Norwegian E-Vote Project, 2012.

Contents

An Information Flow Monitor for a Core of DOM: Introducing References and Live Primitives 1
Ana Almeida-Matos, José Fragoso Santos, and Tamara Rezk

Finding a Forest in a Tree: The Matching Problem for Wide Reactive Systems 17
Giorgio Bacci, Marino Miculan, and Romeo Rizzi

Automata for Analysing Service Contracts 34
Davide Basile, Pierpaolo Degano, and Gian Luigi Ferrari

On Duality Relations for Session Types 51
Giovanni Bernardi, Ornela Dardha, Simon J. Gay, and Dimitrios Kouzapas

Characterising Testing Preorders for Broadcasting Distributed Systems 67
Andrea Cerone and Matthew Hennessy

Tests for Establishing Security Properties 82
Vincent Cheval, Stéphanie Delaune, and Mark Ryan

A Class of Automata for the Verification of Infinite, Resource-Allocating Behaviours 97
Vincenzo Ciancia and Matteo Sammartino

Multiparty Session Nets 112
Luca Fossati, Raymond Hu, and Nobuko Yoshida

Interaction and Causality in Digital Signature Exchange Protocols 128
Jonathan Hayman

Session Types with Gradual Typing 144
Peter Thiemann

Corecursion and Non-divergence in Session-Typed Processes 159
Bernardo Toninho, Luis Caires, and Frank Pfenning

Trust-Based Enforcement of Security Policies 176
Roberto Vigo, Alessandro Celestini, Francesco Tiezzi, Rocco De Nicola, Flemming Nielson, and Hanne Riis Nielson

Author Index 193