



DAGSTUHL REPORTS

Volume 1, Issue 12, December 2011

Secure Computing in the Cloud (Dagstuhl Seminar 11491) <i>Benny Pinkas, Ahmad-Reza Sadeghi, and Nigel P. Smart</i>	1
Secure Architectures in the Cloud (Dagstuhl Seminar 11492) <i>Sabrina De Capitani di Vimercati, Wolter Pieters, and Christian W. Probst</i>	11
Visualization and Processing of Tensors and Higher Order Descriptors for Multi-Valued Data (Dagstuhl Seminar 11501) <i>Carl-Fredrik Westin and Bernhard Burgeth</i>	27
Design of Reversible and Quantum Circuits (Dagstuhl Seminar 11502) <i>Kenichi Morita and Robert Wille</i>	47
Privacy and Security in Smart Energy Grids (Dagstuhl Seminar 11511) <i>Stefan Katzenbeisser, Klaus Kursawe, Bart Preneel, and Ahmad-Reza Sadeghi</i>	62

ISSN 2192-5283

Published online and open access by

Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, Dagstuhl Publishing, Saarbrücken/Wadern, Germany.

Online available at <http://www.dagstuhl.de/dagrep>

Publication date

March, 2012

Bibliographic information published by the Deutsche Nationalbibliothek

The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available in the Internet at <http://dnb.d-nb.de>.

License

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 Unported license: CC-BY-NC-ND.



In brief, this license authorizes each and everybody to share (to copy, distribute and transmit) the work under the following conditions, without impairing or restricting the authors' moral rights:

- Attribution: The work must be attributed to its authors.
- Noncommercial: The work may not be used for commercial purposes.
- No derivation: It is not allowed to alter or transform this work.

The copyright is retained by the corresponding authors.

Aims and Scope

The periodical *Dagstuhl Reports* documents the program and the results of Dagstuhl Seminars and Dagstuhl Perspectives Workshops.

In principal, for each Dagstuhl Seminar or Dagstuhl Perspectives Workshop a report is published that contains the following:

- an executive summary of the seminar program and the fundamental results,
 - an overview of the talks given during the seminar (summarized as talk abstracts), and
 - summaries from working groups (if applicable).
- This basic framework can be extended by suitable contributions that are related to the program of the seminar, e.g. summaries from panel discussions or open problem sessions.

Editorial Board

- Susanne Albers
- Bernd Becker
- Karsten Berns
- Stephan Diehl
- Hannes Hartenstein
- Frank Leymann
- Stephan Merz
- Bernhard Nebel
- Han La Poutré
- Bernt Schiele
- Nicole Schweikardt
- Raimund Seidel
- Gerhard Weikum
- Reinhard Wilhelm (*Editor-in-Chief*)

Editorial Office

Marc Herbstritt (*Managing Editor*)
Jutka Gasirowski (*Editorial Assistance*)
Thomas Schillo (*Technical Assistance*)

Contact

Schloss Dagstuhl – Leibniz-Zentrum für Informatik
Dagstuhl Reports, Editorial Office
Oktavie-Allee, 66687 Wadern, Germany
reports@dagstuhl.de

Digital Object Identifier: 10.4230/DagRep.1.12.i

www.dagstuhl.de/dagrep

Secure Computing in the Cloud

Edited by

Benny Pinkas², Ahmad-Reza Sadeghi¹, and Nigel P. Smart³

¹ TU Darmstadt, DE, ahmad.sadeghi@cased.de

² Bar Ilan University Ramat Gan, IL, benny@pinkas.net

³ University of Bristol, GB, nigel@cs.bris.ac.uk

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 11491 “Secure Computing in the Cloud”. Cloud Computing offers a lot of benefits for end customers: high-end machines, incredible amounts of storage, high availability and everything available at the touch of a button. In this seminar we concentrate on compute clouds: Clouds, that do not only offer storage but also computations that can be outsourced in form of virtual machines (VMs). Outsourcing computations as well as data to a third party, in this case the cloud provider, are accompanied by the qualms of confiding data to the cloud provider based on blindly trusted service level agreements. The participants of this seminar discuss the involved risks, create threat models as basic assumptions that describe the (un-)trusted entities and present solutions that augment trust in the cloud provider, the integrity and verifiability of computations and data processed in the cloud.

Seminar 04.–09. December, 2011 – www.dagstuhl.de/11491

1998 ACM Subject Classification K.6.5 Security and Protection

Keywords and phrases cloud computing, outsourced computation, verifiability, integrity, confidentiality, trust

Digital Object Identifier 10.4230/DagRep.1.12.1

Edited in cooperation with Stefan Nürnberger


1 Executive Summary

Stefan Nürnberger

Benny Pinka

Ahmad-Reza Sadeghi

Nigel P. Smart

License  Creative Commons BY-NC-ND 3.0 Unported license
© Stefan Nürnberger, Benny Pinkas, Ahmad-Reza Sadeghi, and Nigel P. Smart

Introduction

Cloud computing offers IT resources, including storage, networking, and computing platforms, on an on-demand and pay-as-you-go basis. The high usability of today’s cloud computing platforms makes this rapidly emerging paradigm very attractive for customers who want to instantly and easily provide web-services that are highly available and scalable to the current demands. In the most flexible and general cloud computing model (“*Infrastructure as-a Service*”, *IaaS*), customers are able to run entire Virtual Machines (VMs) inside the Cloud. VM images function as templates from which a virtually unlimited number of VM instances can be instantiated.



Except where otherwise noted, content of this report is licensed under a Creative Commons BY-NC-ND 3.0 Unported license
Secure Computing in the Cloud, *Dagstuhl Reports*, Vol. 1, Issue 12, pp. 1–10
Editors: Benny Pinkas, Ahmad-Reza Sadeghi, and Nigel P. Smart



Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Problem Description

Due to virtualisation, limited physical resources are made available for masses. The sharing of these resources and the complex configuration and maintenance of the needed infrastructure is accompanied by security threats [2, 10]. According to the Cloud Security Alliance (CSA), the major inhibitor of a widespread adaptation of cloud computing is the protection of data [4], as data is no longer under the physical control of the owner (in this case the cloud customer). The cloud provider has access to data stored on disks and data transferred through the cloud network. The fact, that the physical hardware of the cloud is shared with other customers, potentially with adversaries, further stresses the need to protect data in order to thwart the lack of physical control over the own data. Moreover, the outsourced computations must be entrusted to the cloud service provider and face the risk of

Sloppy/Lazy provider: A provider that makes mistakes or simplifies computations. The sloppy and lazy provider might compromise the integrity of the result of computations. Verification of results would be a countermeasure here, for example by executing the computations on multiple, independent clouds.

Greedy provider: A provider which reduces security in order to save money. Greedy providers are willing to violate policies for economic reasons, thereby exposing the data to insider or outsider threats.

Malicious Tenant: A cloud customer (tenant) who is deliberately exploits security vulnerabilities to gain access to data or intellectual insight of processes and computations.

The CSA recommends the use of encryption to protect data in transit and data at rest. However, cryptography in the cloud faces two problems:

1. cryptographic keys in a running VM instance are susceptible to run-time attacks like web server exploits, and
2. key provisioning to a VM is not feasible when we assume the cloud provider has access to data and VM images stored on disk.

Seminar Topics

The participants of this seminar were mainly concerned with the privacy of computation or data with respect to the cloud provider. From concrete examples like doctor-patient-confidentiality while processing genomic data at a third party [5, 6], to generic solutions that hide computations that are done at the cloud provider from the cloud provider itself [9]. Additionally, means to verify the result of an outsourced computation with significantly less computational effort than performing the calculation itself [1, 8, 7, 3]. And last, but not least, even outlooks to ad-hoc clouds that are formed by mobile devices on-demand.

References

- 1 M. Atallah, K. Pantazopoulos, J. Rice, and E. Spafford. Secure outsourcing of scientific computations. *Advances in Computers*, 54:216–272, 2001.
- 2 S. Bleikertz, M. Schunter, C.W. Probst, D. Pendarakis, and K. Eriksson. Security Audits of Multi-tier Virtual Infrastructures in Public Infrastructure Clouds. *Proceedings of the 2010 ACM workshop on Cloud computing security workshop*, 2010.
- 3 K. Chung, Y. Kalai, and S. Vadhan. Improved delegation of computation using fully homomorphic encryption. In *CRYPTO'10*, volume 6223, pages 483–501. 2010.
- 4 Cloud Security Alliance (CSA). Top threats to cloud computing, version 1.0. <http://www.cloudsecurityalliance.org/topthreats/csathreats.v1.0.pdf>, March 2010.

- 5 M. Franz, B. Deiseroth, K. Hamacher, S. Jha, S. Katzenbeisser, and H. Schröder. Secure computations on real-valued signals. In *IEEE Workshop on Information Forensics and Security (WIFS'10)*. IEEE Press, 2010.
- 6 M. Franz, B. Deiseroth, K. Hamacher, S. Jha, S. Katzenbeisser, and H. Schroder. Towards secure bioinformatics services. In *Financial Cryptography and Data Security: 15th International Conference, FC 2011, Gros Islet, St. Lucia, February 28-March 4, 2011, Revised Selected Papers*, volume 7035, page 276. Springer-Verlag New York Inc, 2012.
- 7 R. Gennaro, C. Gentry, and B. Parno. Non-interactive verifiable computing: outsourcing computation to untrusted workers. In *CRYPTO'10*, volume 6223, pages 465–482. 2010.
- 8 S. Hohenberger and A. Lysyanskaya. How to securely outsource cryptographic computations. In *TCC'05*, volume 3378, pages 264–282, 2005.
- 9 M.O. Rabin. How to exchange secrets by oblivious transfer. Technical report, Technical Report TR-81, Harvard Aiken Computation Laboratory, 1981.
- 10 T. Ristenpart, E. Tromer, H. Shacham, and S. Savage. Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds. In *CCS'09*, pages 199–212, 2009.

2 Table of Contents

Executive Summary

Stefan Nürnberg, Benny Pinkas, Ahmad-Reza Sadeghi, and Nigel P. Smart . . . 1

Overview of Talks

Introduction to cloud security issues
Benny Pinkas 5

Privacy for Genomic Computations
Stefan Katzenbeisser 5

Two New Models for Delegation of Computation
Ben Riva 5

Outsourcing Multi-Party Computation
Seny Kamara 6

Share Conversion and Private Information Retrieval
Yuval Ishai 6

Automatically Optimizing Secure Computation
Florian Kerschbaum 7

A New Approach to Practical Active-Secure Two-Party Computation
Claudio Orlandi 7

Fundamental Issues When Using Crypto in the Cloud
Stefan Nürnberg 8

On “device clouds”
N. Asokan 8

One Cloud for All – Virtual Revolution?
Marc Oliver Pahl 8

Panel Discussions 9

Participants 10

3 Overview of Talks

3.1 Introduction to cloud security issues

Benny Pinkas (Bar-Ilan University – Ramat-Gan, IL)

License © ⓘ ⊖ Creative Commons BY-NC-ND 3.0 Unported license
© Benny Pinkas

The talk serves as an introduction to cloud computing and to the features that make it so appealing. We then discuss several security issues that are new to cloud computing, which are relevant to either storage clouds or compute clouds.

3.2 Privacy for Genomic Computations

Stefan Katzenbeisser (TU Darmstadt, DE)

License © ⓘ ⊖ Creative Commons BY-NC-ND 3.0 Unported license
© Stefan Katzenbeisser

Joint work of Katzenbeisser, Stefan; Hamacher, Kay; Franz, Martin; Deiseroth, Björn; Jha, Somesh; Busch, Heike
Main reference M. Franz, B. Deiseroth, K. Hamacher, S. Jha, S. Katzenbeisser, H. Schroder, “Towards secure bioinformatics services,” in Proc. of 15th Int’l Conf. on Financial Cryptography and Data Security (FC’11), Revised Selected Papers, vol. 7035, LNCS, p. 276, Springer, 2012.

We show how privacy of genomic sequences can be protected while they are analyzed using Hidden Markov Models (HMM), which is commonly done in bioinformatics to detect certain non-beneficial patterns in the genome. Besides offering strong privacy guarantees, our solution also allows protecting the intellectual property of the parties involved, which makes the solution viable for implementation of secure bioinformatics services.

In particular, we show how two mutually mistrusting parties can obviously run the forward algorithm in a setup where one party knows a HMM and another party knows a genomic string; while the parties learn whether the model fits the genome, they neither have to disclose the parameterization of the model nor the sequence to each other.

Despite the huge number of arithmetic operations required to solve the problem, we experimentally show that HMMs with sizes of practical importance can obviously be evaluated using computational resources typically found in medical laboratories.

3.3 Two New Models for Delegation of Computation

Ben Riva (Tel Aviv University, IL)

License © ⓘ ⊖ Creative Commons BY-NC-ND 3.0 Unported license
© Ben Riva

Consider a weak client that wishes to delegate computation to an untrusted server and be able to succinctly verify the correctness of the result. We present two new natural relaxations to this model. Specifically:

1. We present a model where the weak client delegates the computation to two or more servers, and is guaranteed to output the correct answer as long as even a single server is honest.


In this model, we show: (1) a 1-round statistically sound protocol for any log-space uniform NC circuit; (2) a very efficient computationally sound protocol for any polynomial computation, with a logarithmic number of rounds.

- Next we present a model with a public offline stage. (That is, the offline stage involves no secret randomness and can be publicly verified separately.)

Here we show two computationally sound protocols for any circuit C , where the client runs in time $\text{poly}(\log(\text{size}(C)), \text{depth}(C))$. The first protocol requires only 1 round of interaction, and its soundness is guaranteed assuming the existence of poly-logarithmic PIR. The second protocol requires $\text{poly}(\log(\text{size}(C)), \text{depth}(C))$ rounds but is much more efficient.

3.4 Outsourcing Multi-Party Computation

Seny Kamara (Microsoft Research – Redmond, US)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Seny Kamara

Joint work of Kamra, Seny; Mohassel, Payman; Raykova, Mariana


We initiate the study of secure multi-party computation (MPC) in a server-aided setting, where the parties have access to a single server that (1) does not have any input to the computation; (2) does not receive any output from the computation; but (3) has a vast (but bounded) amount of computational resources. In this setting, we are concerned with designing protocols that minimize the computation of the parties at the expense of the server.

We develop new definitions of security for this server-aided setting, that generalize the standard simulation-based definitions for MPC, and allow us to formally capture the existence of dishonest but non-colluding participants. This requires us to introduce a formal characterization of non-colluding adversaries that may be of independent interest.

We then design general and special-purpose server-aided MPC protocols that are more efficient (in terms of computation and communication) for the parties than the alternative of running a standard MPC protocol (i.e., without the server). Our main general-purpose protocol provides security when there is at least one honest party with input. We also construct a new and efficient server-aided protocol for private set intersection and give a general transformation from any secure delegated computation scheme to a server-aided two-party protocol.

3.5 Share Conversion and Private Information Retrieval

Yuval Ishai (Technion – Haifa, IL)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Yuval Ishai

Joint work of Beimel, Amos; Ishai, Yuval; Kushilevitz, Eyal; Orlov, Ilan

We suggest a new framework for the construction of information-theoretic private information retrieval (PIR) protocols which relies on a generalized notion of "share conversion" in secret sharing schemes. Our framework unifies previous results in the area and gives rise to new protocols that improve the concrete complexity of PIR even for feasible real-life parameters.

In a nutshell, we use the following two-step approach:

- (1) apply share conversion to get a low-communication secure multiparty computation protocol P for a nontrivial class F of low-depth circuits;
- (2) use a lower bound on the VC dimension of F (a combinatorial measure of dimension) to get a good PIR protocol from P .

Our framework reduces the task of designing good PIR protocols to that of finding powerful forms of share conversion which support circuit classes of a high VC dimension.

3.6 Automatically Optimizing Secure Computation

Florian Kerschbaum (TU Dresden, DE)

License © © ⊖ Creative Commons BY-NC-ND 3.0 Unported license
© Florian Kerschbaum

Main reference F. Kerschbaum, “Automatically Optimizing Secure Computation,” in Proc. of the 18th ACM Conf. on Computer and Communications Security (CCS’11), pp. 703–714, ACM.

URL <http://dx.doi.org/10.1145/2046707.2046786>

On the one hand, compilers for secure computation protocols, such as FairPlay or FairPlayMP, have significantly simplified the development of such protocols. On the other hand, optimized protocols with high performance for special problems demand manual development and security verification.

The question considered in this paper is: Can we construct a compiler that produces optimized protocols? We present an optimization technique based on logic inference about what is known from input and output.

Using the example of median computation we can show that our program analysis and rewriting technique translates a FairPlay program into an equivalent – in functionality and security – program that corresponds to the protocol by Aggarwal et al. Nevertheless our technique is general and can be applied to optimize a wide variety of secure computation protocols.

3.7 A New Approach to Practical Active-Secure Two-Party Computation

Claudio Orlandi (Bar-Ilan University – Ramat-Gan, IL)

License © © ⊖ Creative Commons BY-NC-ND 3.0 Unported license
© Claudio Orlandi

Joint work of Nielsen, Jesper; Nordholt, Peter; Orlandi, Claudio; Sheshank, Sai

Main reference J. Nielsen, P. Nordholt, C. Orlandi, S. Sheshank, “A New Approach to Practical Active-Secure Two-Party Computation,” Cryptology ePrint Archive: Report 2011/091, 2011.

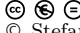
URL <http://eprint.iacr.org/2011/091>

We propose a new approach to practical two-party computation secure against an active adversary. All prior practical protocols were based on Yao’s protocol. We use an OT-based approach and get efficiency via OT extension.

To get a practical protocol we introduce a number of novel techniques for relating the outputs and inputs of OTs in a larger computation. We also report on an implementation of this approach, that shows that our protocol is more efficient than any previous one: As an example, evaluating a Boolean circuit of 34000 gates (oblivious AES encryption) takes less than 2 seconds using our protocol.

3.8 Fundamental Issues When Using Crypto in the Cloud

Stefan Nürnberger (TU Darmstadt, DE)


License  Creative Commons BY-NC-ND 3.0 Unported license
© Stefan Nürnberger

Traditionally, Smartcards or HSMs can be used in untrusted environments. However, a natural adaptation from cryptography used on physical machines to virtual machines used in the cloud is not possible due to a lack of securely storing a cryptographic key. Currently, there is no secure means to authenticate a running virtual machine (VM) or to put a key in the VM image before it gets started. Additionally, when a customer's VM is exposed to the Internet (e.g. to provide web services) it is susceptible to attacks and cryptographic keys might be compromised.

We propose an architecture that incorporates a component (CryptoProxy) that securely wraps high-value secret keys of the cloud customer and only exposes them as cryptographic primitives to the customer's VM or to transparently protect resources the VM uses. For instance, this can be authentication or encryption of data or even of a VM image. Keys can be provisioned from outside the cloud to the CryptoProxy over a trusted channel. Our architecture allows to authenticate running VM instances, protect cryptographic keys and acts as a trust anchor that can audit key usage.

3.9 On “device clouds”

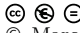
N. Asokan (NOKIA Research Center – Helsinki, FI)

License  Creative Commons BY-NC-ND 3.0 Unported license
© N. Asokan

Although local networking technologies like Bluetooth and WiFi have been common for over a decade, they have not led to local ad-hoc device-to-device networking. Several recent developments on energy-efficient and fast device-to-device connections may change this in the near future. This may make it possible to realize a different type of “cloud” consisting of devices (probably of other users) nearby. There are several interesting use cases that make the use of local device cloud interesting from efficiency or privacy points of view. But they also lead to several security and privacy concerns.

3.10 One Cloud for All – Virtual Revolution?

Marc Oliver Pahl (TU München, DE)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Marc Oliver Pahl

A Vision: Computation and storage are ambient resources that can be used by everyone at any place at any time.

With virtualisation every computing and storage resource can be used by everyone. Devices with wireless access are always connected to the cloud and can augment reality. The data of the users is spread everywhere. The dissemination is not controllable and this is also not wanted as the system is autonomous.

How can we have privacy in this vision?

Some thoughts on the example of automated functionality in and across buildings.

4 Panel Discussions

Two cloud security related seminars took place simultaneously at Schloss Dagstuhl from December 4 to December 9, 2011. This Seminar focused on the verifiability, auditability and confidentiality of computation and data, whereas seminar *11492 Secure Architectures in the Cloud* discussed architectures for verification of computations and configurations, as well as customisability of cloud security and privacy. The joint panel discussion featured the panellists Radu Sion, Martijn Warnier, and Marianne Winslett (11492), as well as Ari Juels, Ahmad Sadeghi and Nigel Smart (11491).

Topics of the panel discussion included, but were not limited to, the following. The panel discussed the “big question” whether small or medium-sized enterprises are more secure in the cloud or using their own systems. Naturally, no answer was found. Here as well, an estimation of the security of cloud providers compared to the security of local infrastructures is essential. For this purpose, self-regulatory or government-initiated penetration testing agencies were suggested in order to assess different cloud infrastructures in an objective fashion.

We discussed the security consequences of providing complete infrastructure-as-a-service (IaaS) images in an App-Store like fashion for clouds. This raises security concerns both for the users of such images (potentially malicious software pre-installed) and for the providers of the images (full erasure of sensitive, private data from the images). Consequently, automated checks are needed to address these problems – with some technical details still being challenging.

Other topics included the efficient verifiability of outsourced computation in the general setting that the cloud provider is not fully trusted. Moreover, the internet-of-things was also a topic. That includes car-to-X communication as well as device clouds. The latter allows the creation of ad-hoc clouds, e.g. for the purpose of sharing an internet connection with people who are travelling in order to save roaming fees.

Further, we discussed the possibility of buying insurance for the data stored in the cloud. This, however, requires precise definitions of (a) the coverage of the insurance (data loss, leakage or corruption) and (b) how to assess whether such an event has indeed occurred.

Participants

- N. Asokan
NOKIA Research Center –
Helsinki, FI
- Maxime Augier
EPFL – Lausanne, CH
- Amir Herzberg
Bar-Ilan Univ. – Ramat-Gan, IL
- Yuval Ishai
Technion – Haifa, IL
- Ari Juels
RSA Laboratories – Bedford, US
- Seny Kamara
Microsoft Res. – Redmond, US
- Stefan Katzenbeisser
TU Darmstadt, DE
- Florian Kerschbaum
TU Dresden, DE
- Thilo Mie
KIT – Karlsruhe Institute of
Technology, DE
- Stefan Nürnberger
TU Darmstadt, DE
- Claudio Orlandi
Bar-Ilan Univ. – Ramat-Gan, IL
- Marc Oliver Pahl
TU München, DE
- Alain Patey
Morpho, SAFRAN Group, FR
- Benny Pinkas
Bar-Ilan Univ. – Ramat-Gan, IL
- Ben Riva
Tel Aviv University, IL
- Ahmad-Reza Sadeghi
TU Darmstadt, DE
- Nigel P. Smart
University of Bristol, GB
- Francois-Xavier Standaert
Université Catholique de
Louvain, BE
- Eran Tromer
Tel Aviv University, IL



Secure Architectures in the Cloud

Edited by

Sabrina De Capitani di Vimercati¹, Wolter Pieters², and
Christian W. Probst³

1 Università degli Studi di Milano, IT, sabrina.decapitani@unimi.it

2 Delft University of Technology, NL, w.pieters@tudelft.nl*

3 Technical University of Denmark, DK, probst@imm.dtu.dk

Abstract

This report documents the outcomes of Dagstuhl Seminar 11492 “Secure Architectures in the Cloud”. In cloud computing, data storage and processing are offered as services, and data are managed by external providers that reside outside the control of the data owner. The use of such services reduces the burden of the owners in managing their data, and may provide significant cost savings. However, cloud computing introduces new security and privacy concerns. In fact, there is little consensus on how to guarantee the confidentiality, integrity, and availability of data in cloud computing scenarios. Also, it is unclear to what extent parties can be held accountable in case something goes wrong. In this seminar, we searched for architectures, modelling approaches, and mechanisms that can help in providing guarantees for cloud security. We proposed the concept of verification-as-a-service that can guide architectures for verification of cloud architectures and configurations, as well as results of computations. We also proposed architectures for organising customisability of security and privacy for cloud customers.

Seminar 4.–9. December, 2011 – www.dagstuhl.de/11492

1998 ACM Subject Classification K.6.5 Security and Protection

Keywords and phrases attestation, auditing, cloud computing, security architectures, security modelling, verification

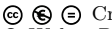
Digital Object Identifier 10.4230/DagRep.1.12.11

1 Executive Summary

Wolter Pieters

Christian W. Probst

Sabrina de Capitani di Vimercati

License  Creative Commons BY-NC-ND 3.0 Unported license
© Wolter Pieters, Christian W. Probst and Sabrina de Capitani di Vimercati

Introduction

In cloud computing, data storage and processing are offered as a service, and the data resides outside the control of the owner. It is often argued that clouds improve security, as the providers have more security expertise than their (smaller) customers. However, despite

* At the time of the seminar, the second editor was affiliated with the University of Twente. His research was supported by the research program Sentinels (<http://www.sentinel.nl>). Sentinels is being financed by Technology Foundation STW, the Netherlands Organization for Scientific Research (NWO), and the Dutch Ministry of Economic Affairs.



Except where otherwise noted, content of this report is licensed under a Creative Commons BY-NC-ND 3.0 Unported license

Secure Architectures in the Cloud, *Dagstuhl Reports*, Vol. 1, Issue 12, pp. 11–26

Editors: Sabrina De Capitani di Vimercati, Wolter Pieters, and Christian W. Probst



Dagstuhl Reports

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

theoretical breakthroughs in cryptography, there is little consensus on how we can provide architectural solutions guaranteeing that cloud data remains confidential, uncorrupted, and available. Also, it is unclear to what extent parties can be held accountable in case something goes wrong. In seminar *11492 Secure Architectures in the Cloud*, we searched for architectures, modelling approaches, and mechanisms that can help in providing guarantees for cloud security. The main question was which cloud-specific security architectures should and could be devised, and how they can be matched to security policies. The seminar was attended by researchers from different academic and industry communities, making it possible to propose integrated solutions and research directions that transcend disciplines. Four main topics have been the subject of this seminar (see also [14]):

1. *Data protection.* Data outside the data owner's control implies that privacy and even integrity can be put at risk. Guaranteeing the privacy and integrity of the data, whether stored in the system or communicated to external parties, becomes a primary requirement, and has raised the attention of both individuals and legislators. Cloud providers have to properly protect the privacy of (possible sensitive) information when storing, processing or sharing it with others [19], and have to adopt adequate access control solutions for enforcing selective access to the data. New approaches have emerged for identifying persons and roles and linking them to access privileges, such as identity-, attribute-, claims-, and data-based access control (e.g., [7, 15]). We discussed challenges of the cloud to the notions of privacy, accountability and user empowerment, their legal, ethical, and architectural implications, and possible solutions.
2. *Simulating physical constraints in the cloud.* In the cloud, we cannot easily enforce where data is stored and how long, and from where it is accessed. Location-based access control aims at limiting access to specific locations, thereby seemingly putting physical limitations back in place [23]. Measures proposed include use of GPS, trusted platform modules (TPMs), but also physically unclonable functions (PUFs) [21]. Also, data could be moved away from attacks [17]. With respect to time, mechanisms have been proposed to assure deletion of data in the cloud [9, 12, 22]. We assessed to which extent these approaches are sufficient to simulate physical constraints, and which architectural solutions are needed to make such forms of assurance possible in practice.
3. *Misuse detection.* Many methods have been proposed for intrusion detection, penetration testing and digital forensics. Are these sufficient for cloud environments? The seminar identified necessary adaptations to system and threat models as well as security metrics, to adequately indicate which attacks are possible and which are actually happening, and thereby reduce cybercrime.
4. *Splitting the clouds.* Public clouds, containing data from different parties, are not deemed suitable for particularly sensitive information. This means that decisions will have to be made about which data to put in the cloud and which data not, which security properties to outsource and which not, and how to make sure that the entire system conforms to the security requirements (cf. [4]). The seminar investigated suitable architectures for "splitting the clouds". For example, in "security-as-a-service", not only IT infrastructure is rented, but also the security that is added to it. For authentication this seems to work pretty well, but how far can this concept be stretched to other security properties such as confidentiality and integrity?

Processing encrypted data was discussed in the parallel seminar *11491 Secure Computing in the Cloud*. This report covers the results of the seminar on Secure Architectures in the Cloud, abstracts of presentations, and proceedings of the working groups. The topics have been restructured during the seminar, and we will refer back to the topics originally proposed where appropriate. Several follow-up initiatives have been assigned to the participants.

Main Findings

As a general observation, we concluded that clouds require a different kind of architectural decisions than traditional information systems. In complex systems such as clouds, we cannot do lots of things manually anymore. For example, there is usually no way to inspect a cloud for evidence manually after an incident. This means that the architecture needs to allow for automation of such tasks, by providing not only functional services, but also meta-services to perform automated maintenance, recovery, etc. Moreover, the processes that make use of such meta-services need suitable architectures themselves. In particular, the following meta-services are needed:

- Automated policy checking,
- Automated configuration verification,
- Automated incident management,
- Automated auditing, and
- Automated forensics.

These processes could be deployed again in (different) clouds, but then the same security concerns apply to them as well.

In this sense, the cloud paradigm begs the question whether we can do everything as a service (XaaS). The participants came up with many different XaaS concepts. In particular, we proposed the concept of verification-as-a-service, which can refer to both the verification of the results of computations, as well as the verification of the (security) architecture and configuration in place at the cloud provider. The former is well-known in the field of electronic voting systems (cf. [20]); the latter resonates with the practice of security auditing. Verification-as-a-service is the main focus in relation to original topics 3 and 4. Specific challenges relate to the verification of negative properties (something is not the case in the architecture) and verification of the results of randomised algorithms. Also, testing-as-a-service could be employed to test functional and non-functional properties of cloud services.

As an instance of simulating physical constraints (topic 2) in relation to verification-as-a-service, we discussed the possibilities for verifying the location of data in the cloud (cf. [23]). One proposal is to integrate GPS with trusted hardware (such as TPM) to prove locations.

Verification-as-a-service provides a paradigm to organise accountability in the cloud. This could be realised by different techniques, for example by:

- Transparency of architecture/configuration (inspection/attestation),
- Forensics (e.g., watermarking),
- Regulation (precaution) and enforcement,
- Incident response (logging), or
- Creating incentives.

Verifying the *integrity* of data seems to be more intuitive than verifying its *confidentiality*. With integrity, it is possible, for example, to compare two different copies. With confidentiality, one would have to prove that only certain parties possess a copy. It only seems to be possible to falsify this after the fact, when it is indeed discovered that data has been leaked. Even in that case, one would need some kind of watermark to prove who leaked the data, for it might have been the user as well as the provider. How to develop a service that provides such watermarking in relation to confidentiality-as-a-service has been identified as an open problem, in relation to topic 4 (splitting the clouds).

Especially on the user side, accountability can be further enhanced by modifiability, or customisability, which allows the user to adapt services to his or her own policies. This requires negotiation on policies, not only between the user and the initial provider, but

also between providers within the supply chain (cf. [24]). Customisability is the main topic discussed in relation to original topic 1 (data protection). Again, special services can be set up that allow the user to achieve this for multiple cloud services at the same time, which would amount to modifiability-as-a-service. Such services could be standardised to make sure that they really empower the user, by employing certain privacy policies themselves, and providing an understandable interface (cf. [6, 10]). We would then have achieved “standardised customisability”.

We formulated several attacker models that lie behind these proposals. Many standard attacker models are problematic in the cloud. An evil/malicious cloud service provider implies that we cannot solve anything without advanced encryption methods, which are costly or even infeasible in many scenarios. Assuming that computations are performed in the clear, we have to assume that the cloud service provider is *indifferent*, not curious. Thus, we trust the cloud provider on the issue of confidentiality, in the sense that we do not expect the provider to leak or misuse data intentionally. However, the provider may still be a:

- Sloppy provider (makes mistakes),
- Lazy provider (simplifies computations), or
- Greedy provider (reduces security to save money).

The sloppy and lazy provider might compromise the integrity of the result of computations. Verification of results would be a countermeasure here, for example by executing the computations on multiple, independent clouds.

Greedy providers are willing to violate policies for economic reasons, thereby exposing the data to insider or outsider threats. Although we do not assume malice on the side of the provider, we do assume malice on the side of other cloud users, who may or may not have specialised access (e.g., administrators). In relation to the greedy provider, one would want to have some means to verify the architecture in place.

Especially if services have been customised, one would want to have some kind of assurance that there is actually a change in configuration taking place based on the customisation. We proposed the development of a tool suite to support remote measurements of architectural variables, which would include existing proposals (cf. [2, 18, 25]). Care needs to be taken that acquiring such information does not violate customer privacy or company property rights [3]. Also, even if the architecture would be (partly) known, the user would then need meaningful support to choose among different providers (and thereby different architectures). This provides another incentive to develop quantitative models that can indeed calculate overall security risks from system architectures, based on existing qualitative approaches [1, 8, 11, 13, 16]. The user can then compare risks and costs to make decisions [5]. Such decisions could even be made in real-time based on information on the current security situation, leading to what has been called fluid information systems [17].

A remaining question is how to create incentives to invest in cloud security. If there is no immediate impact, investments may lag behind with respect to threat levels. Ironically, you can gain a competitive advantage by making your competitors invest in security. Do we really need big scandals to improve security? In any case, achieving more security by (self-)regulation, whether by law, seals, or otherwise, requires architectures such as proposed here, for it is impossible to impose constraints if they cannot be verified.

In conclusion, this seminar proposed architectures for verifying the results of cloud computations, verifying the configuration of cloud architectures, and supporting customisability of cloud services in terms of security. These were defined in relation to cloud-specific attacker models. Visual representations of the proposed architectures can be found under the results of the working groups. Open problems are defined at the end of this document.

References

- 1 S. Bleikertz, T. Groß, and S. Mödersheim. Automated verification of virtualized infrastructures. In *CCSW'10: Proceedings of the 2010 ACM Cloud Computing Security Workshop*, pp. 47–58. ACM, 2010.
- 2 S. Bleikertz, M. Schunter, C.W. Probst, D. Pendarakis, and K. Eriksson. Security audits of multi-tier virtual infrastructures in public infrastructure clouds. In *CCSW'10: Proceedings of the 2010 ACM Cloud Computing Security Workshop*, pp. 93–102. ACM, 2010.
- 3 T.D. Breaux and C.B. Lotrionte. Towards a privacy management framework for distributed cybersecurity in the new data ecology. In *HST'11: Proceedings of the IEEE International Conference on Technologies for Homeland Security*, pp. 6–12. IEEE, November 2011.
- 4 S. Bugiel, S. Nürnberger, A.-R. Sadeghi, and T. Schneider. Twin clouds: Secure cloud computing with low latency. In B. De Decker, J. Lapon, V. Naessens, and A. Uhl, editors, *Communications and Multimedia Security*, vol. 7025 of *Lecture Notes in Computer Science*, pp. 32–44. Springer Berlin / Heidelberg, 2011. DOI: 10.1007/978-3-642-24712-5_3.
- 5 Y. Chen and R. Sion. To cloud or not to cloud?: musings on costs and viability. In *Proceedings of the 2nd ACM Symposium on Cloud Computing*, SOCC'11, pp. 29:1–29:7, New York, NY, USA, 2011. ACM.
- 6 L. Coles-Kemp and E. Kani-Zabihi. On-line privacy and consent: a dialogue, not a monologue. In *Proceedings of the 2010 workshop on New security paradigms*, NSPW'10, pp. 95–106, New York, NY, USA, 2010. ACM.
- 7 S. De Capitani di Vimercati, S. Foresti, S. Jajodia, S. Paraboschi, G. Pelosi, and P. Samarati. Encryption-based policy enforcement for cloud storage. In *ICDCSW'10: IEEE 30th International Conf. on Distributed Computing Systems Workshops*, pp. 42–51. IEEE, 2010.
- 8 T. Dimkov, W. Pieters, and P.H. Hartel. Portunes: representing attack scenarios spanning through the physical, digital and social domain. In *Proceedings of the Joint Workshop on Automated Reasoning for Security Protocol Analysis and Issues in the Theory of Security (ARSPA-WITS'10). Revised Selected Papers, Paphos, Cyprus*, vol. 6186 of *LNCS*, pp. 112–129, Berlin, March 2010. Springer Verlag.
- 9 R. Geambasu, T. Kohno, A.A. Levy, and H.M. Levy. Vanish: increasing data privacy with self-destructing data. In *Proceedings of the 18th conference on USENIX security symposium*, SSYM'09, pp. 299–316, Berkeley, CA, USA, 2009. USENIX Association.
- 10 L. Jędrzejczyk, B.A. Price, A. Bandara, and B. Nuseibeh. “privacy-shake”: a haptic interface for managing privacy settings in mobile location sharing applications. In *MobileHCI'10: Proceedings of the 12th International Conference on Human Computer Interaction with Mobile Devices and Services*, pp. 411–412, New York, NY, USA, September 2010. ACM.
- 11 B. Kordy, S. Mauw, S. Radomirović, and P. Schweitzer. Foundations of attack–defense trees. In *Formal Aspects of Security and Trust, 7th International Workshop, FAST 2010*, vol. 6561 of *LNCS*, pp. 80–95. Springer, 2011.
- 12 R. Perlman. The ephemerizer: making data disappear. Technical Report SMLI TR-2005-140, Sun Microsystems, Inc., Mountain View, CA, USA, 2005.
- 13 W. Pieters. Representing humans in system security models: An actor-network approach. *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications*, 2(1):75–92, 2011.
- 14 W. Pieters. Security and privacy in the clouds: a bird’s eye view. In S. Gutwirth, Y. Pouillet, P. De Hert, and R. Leenes, editors, *Computers, Privacy and Data Protection: an Element of Choice*, pp. 445–457. Springer, Dordrecht, 2011.
- 15 W. Pieters and Q. Tang. Data is key: introducing the data-based access control paradigm. In E. Gudes and J. Vaidya, editors, *Data and Applications Security 2009, Montreal, Canada*, vol. 5645 of *Lecture Notes in Computer Science*, pp. 240–251. Springer, 2009.

- 16 C.W. Probst and R.R. Hansen. An extensible analysable system model. *Information security technical report*, 13(4):235–246, 2008.
- 17 C.W. Probst and R.R. Hansen. Fluid information systems. In *Proceedings of the 2009 workshop on New security paradigms workshop*, pp. 125–132. ACM, 2009.
- 18 C.W. Probst, M.A. Sasse, W. Pieters, T. Dimkov, E. Luysterborg, and M. Arnaud. Privacy penetration testing: How to establish trust in your cloud provider. In S. Gutwirth, R. Leenes, P. De Hert, and Y. Pouillet, editors, *European Data Protection: In Good Health?*. Springer, Dordrecht, 2012.
- 19 J. Ruiter and M. Warnier. Privacy regulations for cloud computing: Compliance and implementation in theory and practice. In S. Gutwirth, Y. Pouillet, P. De Hert, and R. Leenes, editors, *Computers, Privacy and Data Protection: an Element of Choice*, pp. 361–376. Springer Netherlands, 2011. DOI: 10.1007/978-94-007-0641-5_17.
- 20 P.Y.A. Ryan, D. Bismark, J. Heather, S. Schneider, and Z. Xia. Prêt à voter: a voter-verifiable voting system. *IEEE Transactions on Information Forensics and Security*, 4:662–673, 2009.
- 21 G.E. Suh and S. Devadas. Physical unclonable functions for device authentication and secret key generation. In *Proceedings of the 44th annual Design Automation Conference, DAC’07*, pp. 9–14, New York, NY, USA. ACM, 2007.
- 22 Q. Tang. Timed-ephemerizer: Make assured data appear and disappear. In F. Martinelli and B. Preneel, editors, *Public Key Infrastructures, Services and Applications*, vol. 6391 of *Lecture Notes in Computer Science*, pp. 195–208. Springer Berlin / Heidelberg, 2010. DOI: 10.1007/978-3-642-16441-5_13.
- 23 A. van Cleeff, W. Pieters, and R. J. Wieringa. Benefits of location-based access control: a literature study. In *Proceedings of the 3rd IEEE/ACM International Conference on Cyber, Physical and Social Computing (CPSCoM 2010), Hangzhou, China*, pp. 739–746, Los Alamitos, CA, November 2010. IEEE Computer Society.
- 24 M. Winslett, T. Yu, K.E. Seamons, A. Hess, J. Jacobson, R. Jarvis, B. Smith, and L. Yu. Negotiating trust in the web. *IEEE Internet Computing*, 6(6):30–37, 2002.
- 25 W. Xu, X. Zhang, H.n Hu, G.-J. Ahn, and J.-P. Seifert. Remote attestation with domain-based integrity model and policy analysis. *IEEE Transactions on Dependable and Secure Computing*, 99(Preliminary), 2011.

2 Table of Contents

Executive Summary

Wolter Pieters, Christian W. Probst and Sabrina de Capitani di Vimercati 11

Overview of Talks

Adaptive Information Security for Cloud Services: Relating Security Requirements to Design

Arosha Bandara 18

Security Assurance in Virtualized Infrastructures

Sören Bleikertz 18

Privacy and Security Requirements in the Cloud

Travis D. Breaux 19

Versatile Key Management for Secure Cloud Storage

Sebastian Graf 19

Energy Efficiency in Cloud and Related Security Issues

Toni Mastelic 20

Privacy and Integrity Issues in Data Mining Outsourcing

Anna Monreale 20

To Cloud or Not To. Simple Musings on Cloud Viability

Radu Sion 21

On Securing Untrusted Clouds with Cryptography

Radu Sion 21

BOTCLOUDS – The Future of Cloud-based Botnets?

Martijn Warnier 22

Working Groups

Privacy, Data Protection and User Empowerment 22

Verifying Configurations 22

Verifying Computations 23

Open Problems 24


Panel Discussions 25

Participants 26

3 Overview of Talks

3.1 Adaptive Information Security for Cloud Services: Relating Security Requirements to Design

Arosha Bandara (The Open University – Milton Keynes, GB)


License  Creative Commons BY-NC-ND 3.0 Unported license
© Arosha Bandara

Joint work of Bandara, Arosha; Yu, Yijun; Tun, Than Thien; Nuseibeh, Bashar

Information security involves protecting valuable information assets from possible harm. With the increasing use of cloud computing services, the technical and social contexts in which software applications are expected to operate become increasingly dynamic. As a result, the assets, their values, and attack scenarios can easily change. This increases the challenge of finding out what the information assets are, who their owners are, where in the system vulnerabilities lie, and the extent to which the security requirements need to be enforced. In such an environment, information security has to be highly context-sensitive: software applications must adapt to the changing contexts and respond quickly and appropriately to ensure that the requirements for information security are not violated. We call this notion Adaptive Information Security, and focus on three of its prerequisites in the context of cloud computing: (1) understanding user requirements for cloud applications; (2) traceability between requirements, design and implementation of cloud services; and (3) adaptive design for dynamic contexts.

3.2 Security Assurance in Virtualized Infrastructures

Sören Bleikertz (IBM Research – Zürich, CH)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Sören Bleikertz

Joint work of Bleikertz, Sören; Schunter, Matthias; Groß, Thomas; Eriksson, Konrad; Mödersheim, Sebastian;
Main reference S. Bleikertz, T. Groß, S. Mödersheim, “Automated Verification of Virtualized Infrastructures,” in Proc. of the 3rd ACM Workshop on Cloud Computing Security (CCSW’11) , pp. 47–58, 2011.

URL <http://dx.doi.org/10.1145/2046660.2046672>

Cloud computing and virtualized infrastructures are often accompanied by complex configurations and topologies. Dynamic scaling, rapid virtual machine deployment, and open multi-tenant architectures create an environment, in which local misconfiguration can create subtle security risks for the entire infrastructure. This situation calls for automated deployment as well as analysis mechanisms.

We present a platform that combines a static information flow analysis and a virtualization assurance language with state-of-the art verification methods. The system discovers the actual configurations of diverse virtualization environments and unifies them in a graph representation. Using graph traversal, it computes the transitive closure of information flow. The language integrates descriptions of virtualized infrastructures, their transformations, their desired security goals, and evaluation strategies. The different verification tools range from model checking to theorem proving; this allows us to exploit the complementary strengths of methods.

We demonstrate the feasibility of our approach by a real-world case study of a virtualized infrastructure of a global financial institution.

3.3 Privacy and Security Requirements in the Cloud

Travis D. Breaux (Carnegie Mellon University – Pittsburgh, US)

License © © ⊖ Creative Commons BY-NC-ND 3.0 Unported license
© Travis D. Breaux

Joint work of Gordon, David G.; Breaux, Travis D.

Main reference D.G. Gordon, T.D. Breaux, “Managing multi-jurisdictional requirements in the cloud: towards a computational legal landscape,” in Proc. 3rd ACM Workshop on Cloud Computing Security (CCSW’11), pp. 83–94, 2011.

URL <http://dx.doi.org/10.1145/2046660.2046678>

Cloud computing enables organizations to cheaply and quickly obtain computer resources on an as-needed basis, allowing them to more efficiently and effectively provide services to their consumers. Despite the cloud’s ubiquitous appearance, however, data provenance within the cloud presents a multi-jurisdictional challenge, as privacy laws and regulations that govern data may be applicable based upon the residence of the entity who owns the data, the type of organization that uses the data, and any intermediaries facilitating the handling of the data, such as cloud service providers. To this end, we are developing a modeling framework for determining jurisdictional applicability in which these entities are respectively designated as data subjects, data users, and data custodians. We foresee organizations using this framework during system design to determine and resolve the complex issues about where to provide services, store and transfer data.

3.4 Versatile Key Management for Secure Cloud Storage

Sebastian Graf (Universität Konstanz, DE)

License © © ⊖ Creative Commons BY-NC-ND 3.0 Unported license
© Sebastian Graf

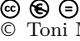
Storing data on cloud-based infrastructures facilitates infinite scalability and all-time availability. Putting data in the cloud additionally offers a convenient way to share any information with user-defined third-parties. However, storing data on the infrastructure of commercial third party providers, demands trust and confidence. Often simple approaches, like merely encrypting the data by providing encryption keys, which at most consists of a shared secret supporting rudimentary data sharing, do not support evolving sets of accessing clients to common data.

Based on well-established approaches regarding stream-encryption, we propose an adaption for enabling scalable and flexible key management within heterogeneous environments like cloud scenarios. Representing access-rights as a graph, we distinguish between the keys used for encrypting hierarchical data and the encrypted updates on the keys enabling flexible join-/leave-operations of clients. This distinction allows us to utilize the high availability of the cloud as updating mechanism without harming any confidentiality. Our graph-based key management results in a constant adaption of nodes related to the changed key. The updates on the keys generate a constant overhead related to the number of these updated nodes.

The proposed scalable approach utilizes cloud-based infrastructures for confidential data and key sharing in collaborative workflows supporting variable client-sets.

3.5 Energy Efficiency in Cloud and Related Security Issues

Toni Mastelic (TU Wien, AT)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Toni Mastelic


Joint work of Mastelic, Toni; Brandic, Ivona; Emeakaroha, Vincent; Maurer, Michael; Breskovic, Ivan
URL <http://www.infosys.tuwien.ac.at/linksites/FOSII/index.html>

Cloud computing is a promising approach for implementing scalable on- demand computing infrastructure. It includes business aspects like SLAs and customer-provider relationship, as well as organizational issues like scheduling, resource allocation, all the way to a technical details like VM monitoring and application deployment. While energy efficiency is mostly managed on an organizational level, it is realized by actions on the level of clusters, physical machines, VMs or even a single application.

By monitoring customer’s applications for a purpose of more efficient scheduling, provider reaches the privacy border. Also, by applying energy efficient measures like time-sharing VMs and running multiple VMs on a single physical machines, provider creates vulnerable environments for customer’s applications. Can customer trust provider’s measurements; how secure is his application; is customer’s privacy being threatened; these are all the questions which cannot be neglected for benefit of energy efficiency, but should certainly be considered.

3.6 Privacy and Integrity Issues in Data Mining Outsourcing

Anna Monreale (University of Pisa, IT)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Anna Monreale

Joint work of Monreale, Anna; Fosca, Giannotti; Dino, Pedreschi; Hui (Wendy), Wang; Laks, Lakshmanan V.S.;
Main reference F. Giannotti, L.V.S. Lakshmanan, A. Monreale, D. Pedreschi, H. (Wendy) Wang,
“Privacy-preserving data mining from outsourced databases,” in S. Gutwirth, Y. Pouillet, P. De Hert, R. Leenes (eds.): *Computers, Privacy and Data Protection: an Element of Choice*, pp. 411–426, Springer Netherlands, 2011.
URL http://dx.doi.org/10.1007/978-94-007-0641-5_19

Spurred by developments such as in cloud computing, there has been considerable recent interest in the paradigm of data mining-as-service. A company (data owner) lacking in expertise or computational resources can outsource its mining needs to a third-party service provider. In this paradigm two problems arise: i) both data and the knowledge extractable from the outsourced database are considered private property of the corporation (data owner), and so there arises serious privacy issues ii) a dishonest service provider may return inaccurate mining results to the data owner, so there arises serious integrity issue of the mining results.

To protect corporate privacy, the data owner has to transform its data and ship it to the server, send mining queries to the server, and recover the true mining results from the extracted knowledge received from the server. To detect security issues, the data owner has to apply an efficient and practical auditing approach that can verify the correctness and the completeness of mining results.

3.7 To Cloud or Not To. Simple Musings on Cloud Viability

Radu Sion (Stony Brook University, US)

License © © ⊖ Creative Commons BY-NC-ND 3.0 Unported license
© Radu Sion

Joint work of Sion, Radu; Chen, Yao;

Main reference Y. Chen, R. Sion, “To Cloud Or Not To Cloud? Musings On Costs and Viability,” in the ACM Symposium on Cloud Computing (SOCC 2011).

URL <http://www.cs.sunysb.edu/~sion/research/cloudc2010-draft.pdf>

In this talk we aim to understand the types of applications for which cloud computing is economically tenable, i.e., for which the cost savings associated with cloud placement outweigh any associated deployment costs.

We discover two scenarios. In an (i) “unified client” scenario, once cloud-hosted, applications are meant to be accessible only to a single cloud customer (or small set of associates). It then becomes important to ensure that the cost savings (mainly computation-related) can offset the often significant client-cloud distance (network costs etc).

In a (ii) “multi-client” setting on the other hand, outsourced applications serve numerous different third parties. We show that then clouds begin to act similarly in nature to content-distribution networks – by comparison, their better network integration is simply too good to pass on, when compared to locally hosting the applications (and incurring associated network costs).

Ultimately, we hope this work will constitute a first step in an objective evaluation of the technological side of costs of outsourcing and computing in general.

3.8 On Securing Untrusted Clouds with Cryptography

Radu Sion (Stony Brook University, US)

License © © ⊖ Creative Commons BY-NC-ND 3.0 Unported license
© Radu Sion

Joint work of Sion, Radu; Chen, Yao;


Main reference Y. Chen, R. Sion, “On Securing Untrusted Clouds with Cryptography,” in the ACM Workshop on Privacy in the Electronic Society (WPES 2010).

URL <http://www.cs.sunysb.edu/~sion/research/sion2010wpes-pcost.pdf>

In a recent interview, Whitfield Diffie argued that “the whole point of cloud computing is economy” and while it is possible in principle for “computation to be done on encrypted data, [...] current techniques would more than undo the economy gained by the outsourcing and show little sign of becoming practical”. In this talk we explore whether this is truly the case and quantify just how expensive it is to secure computing in untrusted, potentially curious clouds.

3.9 BOTCLOUDS – The Future of Cloud-based Botnets?

Martijn Warnier (TU Delft, NL)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Martijn Warnier

Joint work of Warnier, Martijn; Clark, Cassidy; Brazier, Frances

Main reference K. Clark, M. Warnier, F.M.T. Brazier, “BOTCLOUDS – The Future of Cloud-based Botnets?,” in Proc. of the 1st Int’l Conf. on Cloud Computing and Services Science (CLOSER 2011).

URL <http://homepage.tudelft.nl/68x7e/Papers/botclouds.pdf>

Many Cloud Service Providers (CSP) offer access to scalable, reliable computing resources following a pay-as-you-go model. Research into security of the Cloud focuses mainly on protecting legitimate users of Cloud services from attacks by external, malicious users. Little attention is given to prohibit malicious users from using the Cloud to launch attacks, such as those currently done by botnets. These attacks include launching a DDoS attack, sending spam and perpetrating click fraud. This paper discusses the threat of Cloud-based botnets, or botclouds and the need for new techniques to detect them. Two experiments show how simple and cheaply these attacks can be launched from botclouds.

4 Working Groups

4.1 Privacy, Data Protection and User Empowerment

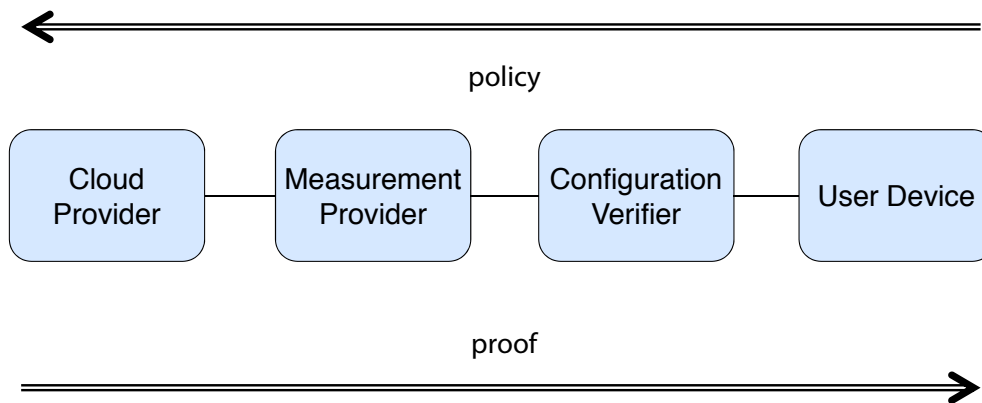
In this working group, the focus was on supporting the user in making meaningful choices with respect to security and privacy of data. To enable such choices, an interface is needed in which the user can obtain the relevant information, and make her choices known. This requires careful selection of the information and choices presented to the user, as well as the possibility to *actually* change decisions on the fly, that is, to move from one cloud service to another whenever the security requirements change. It also requires the propagation of information all the way up the supply chain from infrastructure to user interface, which requires cooperation of the different providers, and therefore standardisation. Important questions are:

- Modifiability: How to reach-back for security/ privacy customisation?
- Diversity / supply chains: How to manage these issues across multiple clouds?
- Scalability: How to make the approach work without overloading the user with information and choices?
- Mobility: How to swap in/out sub-clouds?
- User affordances: If clouds don’t enshrine real-world process complexity, what happens?
- Auditing: How to make auditing manageable if all security is customised?

4.2 Verifying Configurations

The objective of verifying configurations is to know that a certain architecture configuration is running on a cloud. Such an architecture consists of both hardware and software. An example application is when law requires an “adequate level of protection”, for example for privacy-sensitive data. The high-level research problems are:

- How does the cloud provider itself know what is running on its systems?
- How can one transfer such complicated knowledge to the cloud user?



■ **Figure 1** An architecture for verifying configurations

An additional challenge is that the cloud provider has right to its own privacy: one may not be able to study, or at least not publish, the underlying hardware because it's a business secret. This requires mechanisms to ensure the confidentiality of the architecture, while ascertaining its high-level security properties. Logging may help as a basis, but if the logs are maintained by the cloud provider, trust is still required.

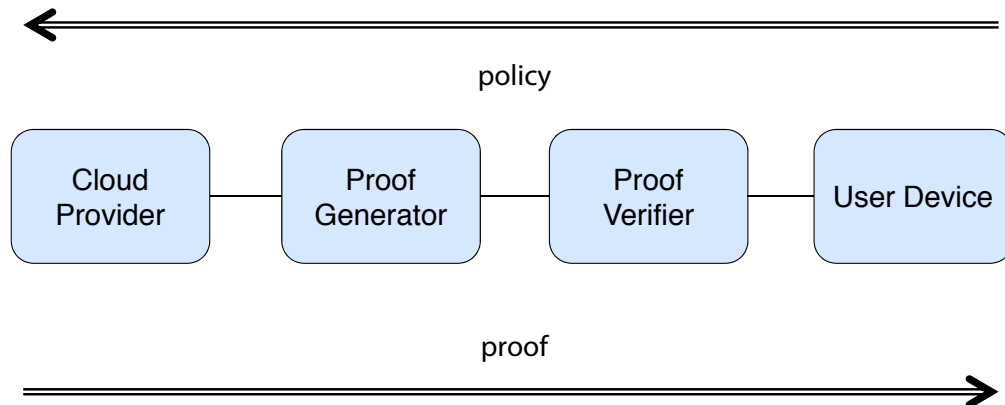
A basic architecture was proposed for verifying architecture configurations in the cloud (Figure 1). In this architecture, measurements performed by a Measurement Provider provide the basis for an assessment by a Configuration Verifier, which is then communicated to the user. Depending on confidentiality requirements on the architecture, different ways of communicating the information can be proposed.

Apart from verification of the configuration, the user would also be interested in verifying relevant data security properties:

- locality of data,
- integrity of data,
- confidentiality of data,
- availability of data,
- deletion of data, and
- non-repudiation of data leakage.

4.3 Verifying Computations

In the working group on verifying computations, the basic assumption was that, in order to assure integrity of results, it is not enough to check the systems. Following up on electronic voting research, there needs to be a way to check that the results of calculations are correct. One simple way to do this would be to perform the calculations in multiple independent clouds, and then use a voting algorithm to determine the correct result. This would provide statistical confidence about the correctness. However, it is not evident how to assure that clouds are independent, as for example to independent SaaS providers may rent infrastructures from the same IaaS provider. Also, if computations are not deterministic, one cannot expect the results to be the same for different instances.



■ **Figure 2** An architecture for verifying computations

A basic architecture was proposed for verifying the results of computations in the cloud (Figure 2). In this architecture, a proof of the result is generated by a Proof Generator service, which can be checked by a Proof Verifier service. The latter may or may not be localised on the user device. If it is not, for example because of the complexity of the task, another (simpler) mechanism is needed to assure the user of the correctness of the verification of the proof.

5 Open Problems

At the end of the seminar, we identified the following open problems and research directions:

- Architectures for trusted computing without trusted hardware,
- Security product lines for the cloud,
- Further refinement of cloud attacker models,
- Architectures to provide accountability for data leakage,
- Corporate social responsibility and transparency in relation to cloud security,
- Use cases for secure cloud architectures,
- Methods for security-related decision support for cloud consumers,
- Specification of the relation between technical measures and regulation,
- Further refinement of the Verification-as-a-Service concept,
- Development of a measurement/reconstruction suite for cloud architecture configurations, and
- Development of methods for cloud forensics.

All of these problems have been assigned to seminar participants as follow-up activities. For details, please contact the seminar organisers.

6 Panel Discussions

Two cloud security related seminars took place simultaneously at Schloss Dagstuhl from December 4 to December 9, 2011. Seminar *11491 Secure Computing in the Cloud* focused on the verifiability, auditability and confidentiality of computation and data, while seminar *11492 Secure Architectures in the Cloud* discussed architectures for verification of computations and configurations, as well as customisability of cloud security and privacy. The joint panel discussion featured the panellists Radu Sion, Martijn Warnier, and Marianne Winslett (11492), as well as Ari Juels, Ahmad Sadeghi and Nigel Smart (11491).

Topics of the panel discussion included, but were not limited to, the following. The panel discussed the “big question” whether small or medium-sized enterprises are more secure in the cloud or using their own systems. Naturally, no answer was found. Here as well, an estimation of the security of cloud providers compared to the security of local infrastructures is essential. For this purpose, self-regulatory or government-initiated penetration testing agencies were suggested in order to assess different cloud infrastructures in an objective fashion.

We discussed the security consequences of providing complete infrastructure-as-a-service (IaaS) images in an App-Store like fashion for clouds. This raises security concerns both for the users of such images (potentially malicious software pre-installed) and for the providers of the images (full erasure of sensitive, private data from the images). Consequently, automated checks are needed to address these problems – with some technical details still being challenging.

Other topics included the efficient verifiability of outsourced computation in the general setting that the cloud provider is not fully trusted. Moreover, the internet-of-things was also a topic. That includes car-to-X communication as well as device clouds. The latter allows the creation of ad-hoc clouds, e.g. for the purpose of sharing an internet connection with people who are travelling in order to save roaming fees.

Further, we discussed the possibility of buying insurance for the data stored in the cloud. This, however, requires precise definitions of (a) the coverage of the insurance (data loss, leakage or corruption) and (b) how to assess whether such an event has indeed occurred.

Participants

- Arosha Bandara
The Open University – Milton Keynes, GB
- Sören Bleikertz
IBM Research – Zürich, CH
- Travis D. Breaux
Carnegie Mellon University – Pittsburgh, US
- Julien Bringer
Morpho, SAFRAN Group, FR
- Sven Bugiel
TU Darmstadt, DE
- Lizzie Coles-Kemp
RHUL – London, GB
- Sabrina De Capitani di Vimercati
University of Milan, IT
- Trajce Dimkov
University of Twente, NL
- Sebastian Graf
Universität Konstanz, DE
- Fabio Massacci
University of Trento – Povo, IT
- Toni Mastelic
TU Wien, AT
- Sjouke Mauw
University of Luxembourg, LU
- Anna Monreale
University of Pisa, IT
- Sebastian Pape
TU Dortmund, DE
- Wolter Pieters
University of Twente, NL
- Christian W. Probst
Technical University of Denmark – Lyngby, DK
- Peter Y.A. Ryan
University of Luxembourg, LU
- Matthias Schunter
IBM Research – Zürich, CH
- Radu Sion
Stony Brook University, US
- André Van Cleeff
University of Twente, NL
- Marcel Waldvogel
Universität Konstanz, DE
- Martijn Warnier
TU Delft, NL
- Marianne Winslett
Univ. of Illinois – Urbana, US



Report from Dagstuhl Seminar 11501

Visualization and Processing of Tensors and Higher Order Descriptors for Multi-Valued Data

Edited by

Carl-Fredrik Westin¹ and Bernhard Burgeth²

1 Harvard Medical School – Boston, US, westin@bwh.harvard.edu

2 Universität des Saarlandes, DE, burgeth@math.uni-sb.de

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 11501 “Visualization and Processing of Tensors and Higher Order Descriptors for Multi-Valued Data”, taking place December 11–16, 2011. The seminar gathered 26 senior and younger researchers from various countries in the unique atmosphere offered by Schloss Dagstuhl. The focus of the seminar was to discuss modern and emerging methods for analysis and visualization of tensor and higher order descriptors from medical imaging and engineering applications. Abstracts of the talks are collected in this report.

Seminar 11.–16. December, 2011 – www.dagstuhl.de/11501

1998 ACM Subject Classification I.4 Image processing and computer vision

Keywords and phrases visualization, image processing, tensor fields, diffusion tensor imaging (DTI), diffusion MRI (dMRI), fiber tractography, higher-order tensors, partial differential equations (PDEs), structural mechanics, solid mechanics

Digital Object Identifier 10.4230/DagRep.1.12.27

Edited in cooperation with Lauren O’Donnell

1 Executive Summary

Carl-Fredrik Westin

Bernhard Burgeth

License  Creative Commons BY-NC-ND 3.0 Unported license
© Carl-Fredrik Westin and Bernhard Burgeth

Higher Order Descriptors in Medical Imaging and Engineering

This seminar is the 4th in a series of Dagstuhl Seminars devoted to the visualization and processing of tensor fields and higher order descriptors. Tensor fields play an important role in many different scientific disciplines and application domains such as medical imaging, image processing, fluid dynamics, and structural mechanics. Analysis and visualization of multi-valued data have gained significant importance in scientific visualization and image processing due to rapid advances in medical imaging and in the engineering industry.

In medical imaging, multi-valued data include diffusion weighted magnetic resonance imaging (dMRI), a medical imaging modality that allows the measurement of water diffusion in tissue *in vivo*. These measurements allow the description of diffusion in living fibrous tissue (e.g., white matter or muscle). The diffusion can be described by a diffusion tensor (i.e., a positive semidefinite 3×3 matrix). It is customary to acquire more complex data than can be described by the tensor model and recently the analysis has been extended to



Except where otherwise noted, content of this report is licensed under a Creative Commons BY-NC-ND 3.0 Unported license
Visualization and Processing of Tensors and Higher Order Descriptors for Multi-Valued Data, *Dagstuhl Reports*, Vol. 1, Issue 12, pp. 27–46

Editors: Carl-Fredrik Westin and Bernhard Burgeth



DAGSTUHL
REPORTS

Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

higher-order descriptors (i.e., higher-order tensors or spherical harmonics). There are several open questions how to best analyze and visualize such data.

In addition to tensor data from the medical field, a number of scientific and engineering applications produce tensor fields as a results of simulations of physical phenomena. The tensor concept is essential for the physical description of anisotropic behavior, especially in solid mechanics and civil engineering (e.g. stress-strain relationships, inertia tensors, permittivity tensor). The field of engineering faces many open problems in tensor field visualization and processing and novel technology is needed to address these problems.

Seminar Topics and Breakout Sessions

The emphasis of the seminar was on presenting the recent developments in the multidisciplinary field as well as identifying new challenges. We discussed a broad set of topics and challenges that cover both theoretical and practical issues related to analyzing and visualizing fields of tensors and higher order descriptors. During the workshop we discussed

- Higher-order models in dMRI beyond the diffusion tensor
- Higher-order tensors in image processing
- Computational analysis and visualization of airflow dynamics
- Novel differential geometric approaches to brain connectivity from dMRI
- Connectivity concepts in mathematical morphology for tensor fields
- Tensor concepts in structural mechanics and material science
- Visualization of uncertainty
- dMRI in brain studies for clinical applications

This year we scheduled time for breakout sessions that would foster focused discussions in smaller groups. During the first day of the meeting the group defined a list of important topics and open questions. Three of those were chosen and defined the breakout sessions:

- How do we define a suitable Finsler metric from diffusion data?
- How do we define biologically meaningful metrics from diffusion tensor and higher order model diffusion data?
- What are important questions in engineering that can be answered with visualization?

These breakout sessions turned out to be very successful and the groups scheduled extra time in the evenings for continued discussions. The format of the new breakout sessions fits very well in the Dagstuhl environment promoting discussions and interactions. If we get the chance to organize another meeting at Dagstuhl, the breakout sessions will definitely be a part of the schedule.

Outcomes

The participants all agreed that the meeting was successful and stimulating. Seminar participants are already collaborating on a Springer book summarizing the results of the meeting. The Springer book will have about twenty chapters authored by the meeting participants, and we expect the book to be published in early 2013. The participants

expressed interest in documenting the discussions in the breakout session in book chapters, in addition to the science described in their regular presentation.

The environment at Schloss Dagstuhl has generated several new scientific collaborations. The work in the engineering breakout session has resulted in a new project of four participants (Stommel, Burgeth, Scheuermann, Hotz) and a submission of a grant proposal to the Landesforschungsförderprogramm (LFFP) des Saarlandes. Meanwhile, the application for the grant has been approved. Three seminar participants who met at the meeting (O'Donnell, Hui Zhang, Schultz), from the USA, Great Britain, and Germany, are collaboratively organizing a workshop on computational diffusion MRI at the conference for Medical Image Computing and Computer-Assisted Intervention 2012.

It was voted that the group will apply for another meeting in this series, and that in addition to the current organizers (Carl-Fredrik Westin, Bernhard Burgeth, Anna Vilanova Bartoli), add Dr. Ingrid Hotz (ZIB – Berlin) as an organizer of the next event.

Acknowledgement

The organizers thank all the attendees for their contributions and extend special thanks to the team of Schloss Dagstuhl for helping to make this workshop a success. As always, we enjoyed the warm atmosphere of the Schloss, which supports both formal presentations as well as informal exchanges of ideas.

2 Table of Contents

Executive Summary

<i>Carl-Fredrik Westin and Bernhard Burgeth</i>	27
---	----

Overview of Talks

Operator-algebraic processing of matrix fields: potentials, shortcomings, and perspectives <i>Bernhard Burgeth</i>	32
Cycles of White Matter <i>Cagatay Demiralp</i>	32
Enhancement of Crossing Fiber-structures in DW-MRI via the Cartan Connection <i>Remco Duits</i>	32
Abstract: “Probing the Human Brain Connectome” <i>Luc M. J. Florack</i>	34
From diffusion tensor to Riemannian metric tensor <i>Andrea Fuster</i>	34
Estimation of 4th order tensors with positivity constraint in Diffusion MRI <i>Aurobrata Ghosh</i>	34
Metrics on Vector- and Tensorbundles <i>Hans Hagen</i>	35
Analyzing tensor fields to study flow fields <i>Mario Hlawitschka</i>	35
Finding Representative Subsets from 3d Second-Order Stress Tensor Fields <i>Ingrid Hotz</i>	36
Brain metrics breakout session <i>Derek K. Jones</i>	36
HOPE: Higher Order Phase Estimation <i>Hans Knutsson</i>	37
What Summarizes a Diffusion MRI Dataset? <i>David H. Laidlaw</i>	37
Nonnegative Definite Tensors and Neuroimaging <i>Lek-Heng Lim</i>	37
Interactive Exploration of Stress Tensors Used in Computational Turbulent Combustion <i>Georgeta Elisabeta Marai</i>	38
Tensor voting with vote clustering <i>Rodrigo Moreno</i>	38
Tractography in the clinic: What works and what is missing? <i>Lauren O’Donnell</i>	39
Estimation of Free-Water Corrected Diffusion Tensors <i>Ofer Pasternak</i>	39

Human cortical connectome reconstruction from diffusion weighted MRI: The effect of tractography algorithm <i>Alard Roebroeck</i>	39
Tensor lines: A good concept for solid mechanics applications? <i>Gerik Scheuermann</i>	40
Learning Higher Order Tensor Rank Estimates the Number of Fiber Compartments in Diffusion MRI <i>Thomas Schultz</i>	40
Application of Tensors to Model the Mechanical Properties of Short Fiber Reinforced Plastics <i>Markus Stommel</i>	40
Towards population studies with HARDI <i>Ragini Verma</i>	41
Rotationally invariant sampling and tensor Coulomb forces <i>Carl-Fredrik Westin</i>	41
Recent Developments in Visualization of Diffusion Tensor Data <i>Alexander Wiebel</i>	42
Tensor Field Analysis for Geometry Processing <i>Eugene Zhang</i>	42
In vivo imaging of brain microstructure using diffusion MRI <i>Gary Hui Zhang</i>	42
List of previous meetings in this workshop series	43
Schedule	43
Participants	46

3 Overview of Talks

3.1 Operator-algebraic processing of matrix fields: potentials, shortcomings, and perspectives

Bernhard Burgeth (Universität des Saarlandes, DE)

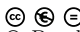
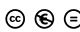
License  Creative Commons BY-NC-ND 3.0 Unported license
© Bernhard Burgeth

Image processing provides a variety of methods for the filtering and analysis of grayscale or color images. In this talk we report on an operator-algebraic framework that allows us to transfer concepts from PDE- or ordering-based image processing to the setting of matrix fields. We explain the fundamental concepts underlying this framework, such as a symmetric product of symmetric matrices or Loewner ordering, and discuss its virtues, difficulties, shortcomings, and perspectives.

3.2 Cycles of White Matter

Cagatay Demiralp (Brown University – Providence, US)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Cagatay Demiralp
URL <http://www.cs.brown.edu/~cad/>

Motivated by sheer curiosity in the question that if there exists a natural categorization of the brain based on white-matter topology, I will talk about our effort to characterize the topology of brain white matter using simplicial homology.

3.3 Enhancement of Crossing Fiber-structures in DW-MRI via the Cartan Connection

Remco Duits (TU Eindhoven, NL)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Remco Duits

Main reference (1) R. Duits, E. Creusen, A. Ghosh, T. Dela Haije, “Morphological and Linear Scale Spaces on $\mathbb{R}^3 \times S^2$ for Enhancement of Crossing Fibers in DW-MRI,” IJCV 2010 vol. 92, pp. 231–264, March 2011.

URL <http://www.bmia.bmt.tue.nl/People/RDuits/JMIVDuits2011final.pdf>

Main reference (2) R. Duits, E. Creusen, A. Ghosh, T. Dela Haije, “Diffusion, Convection and Erosion on $\mathbb{R}^3 \times S^2$ and their Application to the Enhancement of Crossing Fibers,” arXiv:1103.0656v4 [math.AP]

URL <http://arxiv.org/abs/1103.0656v4>

Diffusion-Weighted MRI (DW-MRI) measures local water diffusion in biological tissue, which reflects the underlying fiber structure. In order to enhance the fiber structure in the DW-MRI data we consider both (convection-)diffusions and Hamilton-Jacobi equations (erosions) on the coupled space $\mathbb{R}^3 \times S^2$ of 3D-positions and orientations, embedded as a quotient in the group $SE(3)$ of 3D-rigid body movements. These left-invariant evolutions are expressed in the frame of left-invariant vector fields on $SE(3)$, which serves as a moving frame of reference attached to fiber fragments. The linear (convection-) diffusions are solved by a convolution with the corresponding Green’s function, whereas the Hamilton-Jacobi equations are solved by a morphological convolution with the corresponding Green’s function.

The underlying differential geometry is induced by a Cartan connection on a principal fiber bundle within $SE(3)$: The evolutions locally take place along auto-parallel curves (exponential curves), which (due to torsion of the Cartan connection) do not coincide with the (sub-Riemannian) geodesics that we derived recently. All methods are tested on DTI-images of the brain. These experiments indicate that our techniques are useful to deal with both the problem of limited angular resolution of DTI and the problem of spurious, non-aligned crossings in HARDI. Finally, we propose new fiber tracking algorithms based on the evolved DW-MRI. The whole framework is a special case in our larger group theoretical framework with various imaging applications.

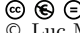
For more information, see [1, 2, 3, 4] and <http://arxiv.org/abs/1103.0656v4>. For other instances of our group-theoretical framework see [5, 6, 7] for evolutions on invertible orientation scores and see [8] and <http://arxiv.org/abs/1110.6087> for evolutions on Gabor transforms.

References

- 1 R. Duits, T. C. J. Dela Haije, A. Ghosh, E. J. Creusen, A. Vilanova, and B. ter Haar Romeny, “Enhancement of DW-MRI,” in *Scale Space and Variational Methods in Computer Vision (Lecture Notes in Computer Science)*, vol. 6667, (Heidelberg), pp. 1–13, Springer-Verlag, September 2011.
- 2 E. J. Creusen, R. Duits, and T. C. J. Dela Haije, “Numerical schemes for linear and non-linear enhancement of DW-MRI,” in *Scale Space and Variational Methods in Computer Vision (Lecture Notes in Computer Science)*, vol. 6667, (Heidelberg), pp. 14–25, Springer-Verlag, September 2011.
- 3 V. Prckovska, P. Rodrigues, R. Duits, A. Vilanova, and B. ter Haar Romeny, “Extrapolating fiber crossings from DTI data. can we infer similar fiber crossings as in HARDI ?,” in *CDMRI’10 MICCAI 2010 workshop on computational diffusion MRI*, vol. 1, (Beijing China), pp. 26–37, Springer, august 2010.
- 4 R. Duits and E. M. Franken, “Left-invariant diffusions on the space of positions and orientations and their application to crossing-preserving smoothing of HARDI images.,” *International Journal of Computer Vision, IJCV*, vol. 92, pp. 231–264, March 2011. Published digitally online in 2010 <http://www.springerlink.com/content/511j713042064t35/>.
- 5 R. Duits, M. Felsberg, G. Granlund, and B. M. ter Haar Romeny, “Image analysis and reconstruction using a wavelet transform constructed from a reducible representation of the Euclidean motion group,” *IJCV*, vol. 79, no. 1, pp. 79–102, 2007.
- 6 R. Duits and E. Franken, “Left invariant parabolic evolution equations on $SE(2)$ and contour enhancement via invertible orientation scores, part II: Nonlinear left-invariant diffusion equations on invertible orientation scores,” *Quarterly of Applied mathematics, AMS*, vol. 68, pp. 293–331, June 2010.
- 7 E. Franken and R. Duits, “Crossing-preserving coherence-enhancing diffusion on invertible orientation scores,” *International Journal of Computer Vision (IJCV)*, vol. 85, no. 3, pp. 253–278, 2009.
- 8 R. Duits, H. Fuehr, and B. Janssen, *Mathematical Methods for Signal and Image Analysis and Representation, Left Invariant Evolution Equations on Gabor Transforms*, ch. 8, pp. 151–172. Springer-Verlag, 2011. in press.

3.4 Abstract: “Probing the Human Brain Connectome”

Luc M. J. Florack (TU Eindhoven, NL)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Luc M. J. Florack


Joint work of Florack, L. M. J.; Fuster, A.

Human brain connectomics is the field of science that aims to integrate data and knowledge about structure and function of the human brain at all levels of scale. The comprehensive description that should result from this study, a.k.a. the connectome by analogy with the genome, is one of the grand challenges of the 21st century. As part of this endeavor we will develop new methods for tractography, the geometric delineation of neural fiber pathways, and for structural brain connectivity, at macroscopic scales (i.e. order of magnitude 1mm and above) that can be probed with state-of-the-art magnetic resonance imaging (MRI). Our methodology relies on the physics of anisotropic diffusion of water in brain white matter and its empirical manifestation under diffusion weighted magnetic resonance imaging.

We stipulate that local anisotropic diffusivities can be “geometrized away” similar to the geometrization of gravitational forces in general relativity. However, it turns out that a Riemannian framework, which has proven powerful in the case of mild anisotropies, is inappropriate for the description of the highly anisotropic diffusivity profiles observed in the brain. We propose to exploit a generalization of Riemannian geometry, viz. Riemann-Finsler geometry, to remove this limitation.

3.5 From diffusion tensor to Riemannian metric tensor


Andrea Fuster (TU Eindhoven, NL)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Andrea Fuster

Diffusion Tensor Imaging (DTI) can be investigated by using geometric methods. An active field of research in the past years has been the study of DTI in a Riemannian framework. The main idea is to associate a Riemannian metric to the diffusion tensor, by identifying the latter with the inverse metric. However, this might not always be the right choice of metric tensor, as already pointed out in the literature. We study this analogy from the point of view of the underlying diffusion equations, in different scenarios.

3.6 Estimation of 4th order tensors with positivity constraint in Diffusion MRI

Aurobrata Ghosh (INRIA Sophia Antipolis, FR)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Aurobrata Ghosh

Joint work of Ghosh, Aurobrata; Deriche Rachid

Main reference A. Ghosh, “High Order Models in Diffusion MRI and Applications,” PhD Thesis, April 2011.


URL <http://tel.archives-ouvertes.fr/tel-00645820/fr/>

Diffusion MRI, which is sensitive to the Brownian motion of molecules, has become today an excellent medical tool for probing the tissue micro-structure of cerebral white matter in vivo and non-invasively. It makes it possible to reconstruct fiber pathways and segment major fiber

bundles that reflect the structures in the brain which are not visible to other non-invasive imaging modalities. Since this is possible without operating on the subject, but by integrating partial information from Diffusion Weighted Images (DWIs) into a reconstructed complete image of diffusion, Diffusion MRI opens a whole new domain of image processing. Here we shall explore the role that Cartesian tensors play in the mathematical model. We shall begin with 2nd order tensors, since these are at the core of Diffusion Tensor Imaging. We shall then explore higher and even ordered symmetric tensors, that can take into account more complex micro-geometries of biological tissues such as axonal crossings in the white matter. We will emphasize the estimation of 4th order diffusion tensors with positivity constraint from DWIs.

3.7 Metrics on Vector- and Tensorbundles


Hans Hagen (TU Kaiserslautern, DE)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Hans Hagen

Geometry and Topology, especially Vector and Tensorfields are “the basic technologie” for Geometric Modelling and Scientific Visualization. A topological manifold M is locally connected, locally compact and a union of a countable collection of compact subsets. Such a topological space is metrizable! Vector- and Tensor-bundles do have such a manifold structure. We consider a vectorfield (tensorfield) being part of tangent-bundle (tensor-bundle) of a Riemannian manifold with a metric tensor. Special features of such a metric tensor are “used” to visualize features of the vectorfield (tensorfield).

3.8 Analyzing tensor fields to study flow fields


Mario Hlawitschka (Universität Leipzig, DE)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Mario Hlawitschka

Often, vector fields are “reduced” to scalar fields to study their properties. Opposing that path, I re-define methods based on the tensor field that is the derivative of the given vector field. In that space, we can rewrite known methods in a more-general framework. The methods are presented using the storyline of the analysis of compressible unstable flow around wind turbine fields where the focus lies in studying the effects blade-induced turbulences have on the energy generation as well as mechanical stability of the setup.

3.9 Finding Representative Subsets from 3d Second-Order Stress Tensor Fields

Ingrid Hotz (ZIB – Berlin, DE)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Ingrid Hotz

Tensor fields play an important role in many areas ranging from engineering to medicine. Compared to their significance, excluding diffusion tensor imaging, tensors are an underrepresented topic in visualization. This may be due to the fact that first there is no long tradition in tensor field analysis, second the terminology varies from application to application, and third the questions posed onto the data are very fuzzy. An additional challenge arises from the complexity of the tensor data itself. The goal of our work is to overcome some of these challenges creating schematic depictions of the data based on domain specific feature spaces.

The basis of these feature spaces is a decomposition of the tensor information in scalar and directional features. The scalar feature or shape space is a subset of the space spanned by the three eigenvalues, parametrized by invariants that are prevalent in a certain application. Often such invariants are already reflected by commonly used glyphs in the respective application. An example for invariants that are of importance for stress tensor fields in context with failure analysis is the maximum shear stress and the shape factor. Both entities can be nicely represented using Mohr's circles.

Our work can be summarized by the following steps: (1.) find appropriate feature spaces, (2.) structure these feature spaces into clusters, and (3.) find appropriate representatives for these subsets. Additional cluster analysis provides size, variance and directional distributions of the clusters. The resulting atlas like data representation can be used to intuitively interact with the data focusing onto trends or outliers respectively.

3.10 Brain metrics breakout session


Derek K. Jones (Cardiff University, GB)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Derek K. Jones

Obtaining good quality diffusion MRI and making sound and robust inferences from the data is not trivial, however, and involves a long chain of events from ensuring that the hardware is performing optimally, the pulse sequence is carefully designed, the acquisition is optimal, the data quality is maximized while artifacts are minimized, the appropriate post-processing is used, and, where appropriate, the appropriate statistical testing is used, and the data are interpreted correctly. In our breakout session, we discussed methods to compare brain metrics, what meaningful biological metrics do we get from HARDI data, and how can we get a formal probability of connection.

3.11 HOPE: Higher Order Phase Estimation

Hans Knutsson (Linköping University Hospital, SE)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Hans Knutsson

Joint work of Knutsson, Hans; Westin, Carl-Fredrik


Local phase is a powerful concept which has been successfully used in many image processing applications. For multidimensional signals the concept of phase is complex and there is no consensus on the precise meaning of phase. It is, however, accepted by all that a measure of phase implicitly carries a directional reference.

We present a novel representation of multidimensional phase that is shown to be equivalent to an extended Klein bottle. In contrast to previously suggested phase representations it is shown to be globally isometric for the simple signal class.

For 1-dimensional signals the new phase representation reduces to the original definition of amplitude and phase for analytic signals. Phase estimation using classical quadrature filter pairs is based on the analytic signal concept and requires a pre-defined filter direction. The new local phase representation removes this requirement by implicitly incorporating local orientation. The estimation approach uses spherically separable monomial monomial filter of orders 0, 1 and 2 which naturally extends to N dimensions.

3.12 What Summarizes a Diffusion MRI Dataset?


David H. Laidlaw (Brown University – Providence, US)

License  Creative Commons BY-NC-ND 3.0 Unported license
© David H. Laidlaw

Diffusion MRI datasets are challenging to comprehend; as a result, their potential has not yet been realized. Comprehension sometimes begins with visualization, but it usually needs to end with a summary – often a single number. What makes a good summary number? How do we know? I'm afraid I don't know the answers, so I hope that some of the participants will help to figure this out. I do have some examples that I hope will illustrate the problem, and I look forward to some stimulating discussion.

3.13 Nonnegative Definite Tensors and Neuroimaging

Lek-Heng Lim (University of Chicago, US)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Lek-Heng Lim


Joint work of Lim, Lek-Heng; Schultz, Thomas

One of the most important classes of matrices is the symmetric positive definite ones. They arise as covariance, density matrix, inner products, Laplacians, Mercer kernels, etc. So what is the equivalent of positive/nonnegative definiteness for higher order symmetric tensors? It turns out that there are two natural but different notions: One is that the homogeneous polynomial associated with the tensor be nonnegative valued while the other is that this polynomial be expressible as a sum of powers of linear forms. These two notions are in fact dual in an appropriate sense.

We show that both notions arise in diffusion MRI and lead to two methods for extracting nerve fibers crossing. We shall see that deciding nonnegative definiteness (either notions) of a higher-order tensor is an NP-hard problem but that due to a happy coincidence the cases relevant to these MRI applications yield readily computable convex problems.

3.14 Interactive Exploration of Stress Tensors Used in Computational Turbulent Combustion

Georgeta Elisabeta Marai (University of Pittsburgh, US)

License  Creative Commons BY-NC-ND 3.0 Unported license


© Georgeta Elisabeta Marai

Joint work of Marai, Georgeta Elisabeta; Yilmaz, Levent; Nik, Mehdi; Luciani, Timothy; Maries, Adrian; Haque, Abedul

Simulation and modeling of turbulent flow, and of turbulent reacting flow in particular, involves solving for and analyzing time-dependent and spatially dense tensor quantities, such as turbulent stress tensors. The interactive visual exploration of these tensor quantities can effectively steer the computational modeling of combustion systems. In this chapter, we discuss the challenges in dense symmetric tensor visualization applied to turbulent combustion calculation, and analyze the feasibility of using several established tensor visualization techniques in the context of exploring space-time relationships in computationally-simulated combustion tensor data. To tackle the pervasive problems of occlusion and clutter, we combine a detailed 3D inspection view based on volume rendering with glyph-based representations, used as 2D probes, while leveraging interactive filtering and flow salience cues to clarify the structure of the tensor datasets. Side-by-side views of multiple timesteps facilitate the analysis of time-space relationships. The result is a visual analysis tool to be utilized in debugging, benchmarking, and verification of models and solutions in turbulent combustion. We demonstrate this analysis tool on three example configurations and report feedback from combustion researchers.

3.15 Tensor voting with vote clustering

Rodrigo Moreno (Linköping University Hospital, SE)


License  Creative Commons BY-NC-ND 3.0 Unported license

© Rodrigo Moreno

Tensor voting is a robust technique to propagate and aggregate local information encoded through 2nd order tensors. Traditionally, tensor summation is used to aggregate the votes, which can be inappropriate in some applications. In this talk, I discuss the use of clustering-based aggregation for analyzing the votes cast by tensor voting. Two possible applications of this methodology are context-based tensor decomposition and extracting higher-order tensor fields from 2nd-order tensor fields.

3.16 Tractography in the clinic: What works and what is missing?


Lauren O'Donnell (Harvard Medical School – Boston, US)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Lauren O'Donnell

In this talk I review tractography methods used in clinical applications and discuss open issues. By measuring water diffusion in the brain, diffusion tensor MRI (DTI) gives information about the orientation and integrity of fiber tracts, the major neural connections in the white matter. DTI tractography follows directions of maximal water diffusion to estimate the trajectories of the fiber tracts. We summarize the current published state of the art of clinical tractography, focusing on correspondence with known anatomy, correspondence with electrical stimulation, and tractography's effect on measurable clinical endpoints. We focus on three tracts of interest for neurosurgery: the corticospinal tract, the arcuate fasciculus, and the optic radiation. We discuss the technology that is missing, including automated higher-level analyses of use to a doctor, such as answering the question “is this white matter normal?”

3.17 Estimation of Free-Water Corrected Diffusion Tensors


Ofer Pasternak (Harvard Medical School – Boston, US)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Ofer Pasternak

In Diffusion tensor Imaging, when partial volume of brain tissue and free-water occurs, the estimated tensor describes a mixture of both compartments. As a result, tensor indices such as FA and MD become less specific to tissue microstructure. We discuss tensor-regularization based approaches that provide estimations of free-water corrected tensors, and the implications of applying free- water correction in various case studies.

3.18 Human cortical connectome reconstruction from diffusion weighted MRI: The effect of tractography algorithm

Alard Roebroek (Maastricht University, NL)


License  Creative Commons BY-NC-ND 3.0 Unported license
© Alard Roebroek

Reconstructing the macroscopic human cortical connectome by Diffusion Weighted Imaging (DWI) is a challenging research topic that has recently gained a lot of attention. In the present work, we investigate the effects of intra-voxel fiber direction modeling and tractography algorithm on derived structural network indices (e.g. density, small-worldness and global efficiency). The investigation is centered on three semi-independent distinctions within the large set of available diffusion models and tractography methods: i) single fiber direction versus multiple directions in the intra-voxel diffusion model, ii) deterministic versus probabilistic tractography and iii) local versus global measure-of-fit of the reconstructed fiber trajectories. We discuss interactions in the combined effects of these methods, considerations of tractography sensitivity and specificity, and implications for future studies. It is concluded

that the choice of tractography algorithm along the three dimensions and thresholds (FA, angle, probabilistic) can affect structural network indices dramatically, which is crucial for the sensitivity of any human structural network study and for the validity of study comparisons.

3.19 Tensor lines: A good concept for solid mechanics applications?


Gerik Scheuermann (Universität Leipzig, DE)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Gerik Scheuermann

Tensorlines are among the very few known continuous techniques to visualize tensor fields. In addition, the whole area of tensor topology relies on this concept. Unfortunately, it can be observed that some engineers question their expression power for their application in solid mechanics. It is understood that tensorlines in solid mechanics (as already noted by Dickinson 1989 in the first paper on tensor lines) describe the principle stress directions. However, these directions may not give much information about material failure which is often a central concern of the engineer. The talk discusses the use of tensor lines in solid mechanics by looking at the history and comparing it to successful applications in neuroscience, liquid crystals and medical applications.

3.20 Learning Higher Order Tensor Rank Estimates the Number of Fiber Compartments in Diffusion MRI

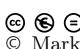
Thomas Schultz (MPI für Intelligente Systeme – Tübingen, DE)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Thomas Schultz

The need to decide on the number of fiber compartments is a fundamental limitation of multi-compartment models in high angular resolution diffusion imaging. This talk clarifies that when using higher-order tensor based spherical deconvolution to estimate multiple fiber directions, the number of compartments amounts to the numeric rank of the fODF tensor. Unfortunately, few practical results on finding numeric tensor rank are available. As a pragmatic alternative, we demonstrate that support vector machines, a standard tool from machine learning, can be used successfully to estimate tensor rank (and therefore fiber number) in this application.

3.21 Application of Tensors to Model the Mechanical Properties of Short Fiber Reinforced Plastics

Markus Stommel (Universität des Saarlandes, DE)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Markus Stommel


This contribution will focus on the application of tensors and tensors fields from an engineering perspective using the example of composite materials. Composite materials consist of at least two materials differing in physical properties and shape. Therefore, on the microscale

the material is inhomogeneous and consists of two or more phases like a matrix material that is reinforced by a second fibroid material. These materials are increasingly used in advanced technical applications which implicate the need for demanding simulation techniques.

The simulation of technical parts made of composites implies the use of continuum mechanics approaches. They usually require the homogenization of the composites microstructure toward a “smeared” continuum on a macroscale. This lecture deals with the past and ongoing effort in homogenization techniques that are used in engineering to simulate composites by the finite element method (FEA). It will demonstrate the basic procedure in performing FEA on short fiber reinforced plastics including the process and structural simulation concepts and their interaction. Afterwards some tensor related topics will be addressed that are open to further development of the homogenization methods.

3.22 Towards population studies with HARDI

Ragini Verma (University of Pennsylvania, US)


License  Creative Commons BY-NC-ND 3.0 Unported license
© Ragini Verma

Joint work of Verma, Ragini; Bloy, Luke; Ingalhalikar, Madhura; Smith, Alex

With the increase in HARDI studies, there has been a need for developing methods for analysis and processing of HARDI data. The talk will cover methods essential for HARDI analysis, motivated towards large population studies. The basic methods for doing population studies include voxel-based analysis of the brain (which needs the development of metrics), region-based analysis (which needs data-based clustering and atlas building), pattern classification and connectivity analysis. In addition to discussing these methods required to facilitate subsequent analysis, the talk will raise issues related to planning HARDI studies – do we need it, what do we need from it, how to compare it with DTI, and how to harness the unknown world of “connectivity”.

3.23 Rotationally invariant sampling and tensor Coulomb forces

Carl-Fredrik Westin (Harvard Medical School – Boston, US)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Carl-Fredrik Westin

Joint work of Westin, Carl-Fredrik; Knutsson, Hans

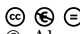
Minimizing the error propagation that a diffusion MRI (dMRI) gradient scheme introduces is an important task in the design of robust and un-biased experiments. Previous studies define the optimal single-shell scheme with respect to various parameters that include: 1) the angular distance between neighboring samples, minimized using an electrostatic optimization 2) the condition number, which estimates the effect of noise, and 3) rotational invariance, so the scheme produces rotationally unbiased estimates. Among the previously proposed schemes, the electrostatic optimization has been shown to produce the most balanced schemes.

It is common that acquisitions schemes are composed of a single b-value shell providing angular sampling of the diffusion profile. But the newer methods require, in addition, a radial sampling in order to observe phenomena such as restriction and hindrance. We propose two schemes for the construction of rotationally invariant multiple- shells. The first is a dual

frame method that optimizes the rotation invariance of any set of samples. The second uses a subset of the icosahedral set that can intuitively be used for nested rotationally invariant schemes with pre-defined number of samples.

3.24 Recent Developments in Visualization of Diffusion Tensor Data

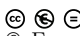
Alexander Wiebel (ZIB – Berlin, DE)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Alexander Wiebel

Many results of studies using diffusion tensor imaging can be summarized with some simple numbers or a validation of a hypothesis. For an examination of a single patient, however, medical practitioners and scientists often have to inspect the data itself or derived data like reconstructed fiber tracts and their bundles. This talk discusses two visualization techniques that can support such an inspection: one visualizes diffusion parameters of fiber bundles and one provides an illustrative rendering for probabilistic tractography data. Additionally, a short excursion to the use of diffusion MRI for monitoring inflammatory bowel diseases is given.

3.25 Tensor Field Analysis for Geometry Processing

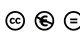
Eugene Zhang (Oregon State University, US)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Eugene Zhang
URL <http://web.engr.oregonstate.edu/~zhange>

Tensor field analysis has been a subject of interest in many scientific and engineering applications. In this talk I will review applications in geometry processing for which tensor fields play a prominent role. In addition, I will discuss how existing tensor field analysis can benefit these applications and what challenges remain in applying tensor field processing to geometry-related applications.

3.26 In vivo imaging of brain microstructure using diffusion MRI

Gary Hui Zhang (University College London, GB)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Gary Hui Zhang

Diffusion MRI is an imaging technique that provides unique insight into tissue microstructure. It makes measurements that are sensitive to the displacement pattern of water molecules undergoing diffusion. Because the tissue microstructure determines this displacement pattern, its properties can thus be deduced from the diffusion MRI measurements using appropriate computational techniques. This talk will review basics of diffusion MRI and discuss some of the advanced diffusion modeling techniques developed at UCL with a focus on neuroimaging applications.

4 List of previous meetings in this workshop series

- The 2004 Dagstuhl Perspective Workshop “Visualization and Processing of Tensor Fields” (Seminar 04172, April 2004, Organizers: Hans Hagen and Joachim Weickert) was the first international forum where leading experts on visualization and processing of tensor fields had the opportunity to meet, many for the first time. This workshop identified several key issues and triggered fruitful collaborations that have also led to the first book in this area. Springer book published in 2006: ISBN 978-3-540-25032-6.
- The 2007 Dagstuhl meeting “Visualization and Processing of Tensor Fields” (Seminar 07022, January 2007, Organizers: David Laidlaw and Joachim Weickert) was equally successful and the progress reported in a second book published with Springer published in 2009: ISBN 978-3-540-88377-7.
- The 2009 Dagstuhl meeting “New Developments in the Visualization and Processing of Tensor Fields” (Seminar 09302, July 2009, Organizers: Bernhard Burgeth and David Laidlaw) saw a shift in focus, and in addition to diffusion imaging, paid attention to engineering applications of tensors in fluid mechanics, material science, and elastography. Springer has also published a third book in this series: ISBN 978-3-642-27342-1 (*to appear*).

5 Schedule

Monday		Presenter
07:30-08:40	Breakfast	
09:15-09:40	Welcome and Introduction	C-F Westin, Bernhard Burgeth
09:40-10:20	Probing the Human Brain Connectome	Luc M. J. Florack, TU Eindhoven
10:20-10:40	Coffee break	
10:40-11:20	In vivo imaging of brain microstructure using diffusion MRI	Gary Hui Zhang, University College London
11:20-12:00	Human cortical connectome reconstruction from diffusion weighted MRI: The effect of tractography algorithm	Alard Roebroek, Maastricht University
12:15-13.30	Lunch Break	
13:30-14:10	Rotationally invariant sampling and tensor Coulomb forces	Carl-Fredrik Westin, Harvard Medical School – Boston
14:10-14.50	Tensor lines: A good concept for solid mechanics applications?	Gerik Scheuermann, Leipzig University
14:50-15.50	Coffee break	
15:50-16.30	Application of Tensors to Model the Mechanical Properties of Short Fiber Reinforced Plastics	Markus Stommel, Saarland University
16:30-17.30	Breakout sessions	
18:00	Dinner	

Tuesday		Presenter
07:30-08:40	Breakfast	
09:00-09:40	Tractography in the clinic: What works and what is missing?	Lauren O'Donnell, Harvard Medical School
09:40-10:20	Towards population studies with HARDI	Ragini Verma, University of Pennsylvania
10:20-10:40	Coffee break	
10:40-11:20	Estimation of Free-Water Corrected Diffusion Tensors	Ofer Pasternak, Harvard Medical School
11:20-12:00	Cycles of White Matter	Cagatay Demiralp, Brown University – Providence
12:15-13.30	Lunch Break	
13:30-14:10	HOPE: Higher order phase estimation	Hans Knutsson, Linköping University Hospital
14:10-14.50	Tensorbundles as a visualization tool?	Hans Hagen, TU Kaiserslautern
14:50-15.50	Coffee break	
15:50-16.30	Tensor Field Analysis and Processing for Geometry Processing	Eugene Zhang, Oregon State University
16:30-17.30	Breakout sessions	
18:00	Dinner	

Wednesday		Presenter
07:30-08:40	Breakfast	
09:00-09:40	What Summarizes a Diffusion MRI Dataset?	David H. Laidlaw, Brown University – Providence
09:40-10:20	What is a positive definite tensor and how it can be used in neuroimaging	Lek-Heng Lim, University of Chicago
10:20-10:40	Coffee break	
10:40-11:20	From diffusion tensor to Riemannian metric tensor	Andrea Fuster, TU Eindhoven
11:20-12:00	Breakout sessions	
12:15-13.30	Lunch Break	
13:30-14:00	Group photo	
14:00-15:00	Bus to Trier	
15:00-18:00	Visit Trier and its Christmas Market	
18:00-20:00	Dinner in Trier at nearby wine-cave	
20:00-21:00	Returning to Schloss Dagstuhl by buss	

Thursday		Presenter
07:30-08:40	Breakfast	
09:00-09:40	Estimation of 4th order tensors with positivity constraint in diffusion MRI	Aurobrata Ghosh, INRIA Sophia Antipolis
09:40-10:20	Enhancement of Crossing Fiber-structures in DW-MRI via the Cartan Connection	Remco Duits, TU Eindhoven
10:20-10:40	Coffee break	
10:40-11:20	Studying Tensors to Analyze Flow	Mario Hlawitschka, University of Leipzig
11:20-12:00	Learning Higher Order Tensor Rank Estimates the Number of Fiber Compartments in Diffusion MRI	Thomas Schultz , MPI Tübingen
12:15-13.30	Lunch Break	
13:30-14:10	Finding Representative Subsets from 3D Second-Order Stress Tensor Fields	Ingrid Hotz, ZIB – Berlin
14:10-14.50	Exploration of Stress Tensors Used in Computational Turbulent Combustion	G. Elisabeta Marai, University of Pittsburgh
14:50-15.50	Coffee break	
15:50-16.30	Operator-algebraic approach for matrix fields: potentials, shortcomings, and perspectives	Bernhard Burgeth, Saarland University
16:30-17.30	Results from breakout sessions	
18:00	Dinner	

Friday		Presenter
07:30-08:40	Breakfast	
09:00-09:40	Tensor Voting with Vote Clustering	Rodrigo Moreno, Linköping University Hospital
09:40-10:20	Recent Developments in Visualization of Diffusion Tensor Data	Alexander Wiebel, ZIB – Berlin
10:20-10:40	Coffee break	
10:40-12:00	Book, future seminar, and farewell	
12:15-13.30	Lunch and Departure	

Participants

- Bernhard Burgeth
Universität des Saarlandes, DE
- Cagatay Demiralp
Brown Univ. – Providence, US
- Remco Duits
TU Eindhoven, NL
- Luc M. J. Florack
TU Eindhoven, NL
- Andrea Fuster
TU Eindhoven, NL
- Aurobrata Ghosh
INRIA Sophia Antipolis, FR
- Hans Hagen
TU Kaiserslautern, DE
- Mario Hlawitschka
Universität Leipzig, DE
- Ingrid Hotz
ZIB – Berlin, DE
- Derek K. Jones
Cardiff University, GB
- Hans Knutsson
Linköping University, SE
- David H. Laidlaw
Brown Univ. – Providence, US
- Lek-Heng Lim
University of Chicago, US
- Georgeta Elisabeta Marai
University of Pittsburgh, US
- Rodrigo Moreno
Linköping Univ. Hospital, SE
- Lauren O'Donnell
Harvard Medical School –
Boston, US
- Ofer Pasternak
Harvard Medical School –
Boston, US
- Alard Roebroek
Maastricht University, NL
- Gerik Scheuermann
Universität Leipzig, DE
- Thomas Schultz
MPI für Intelligente Systeme –
Tübingen, DE
- Markus Stommel
Universität des Saarlandes, DE
- Ragini Verma
University of Pennsylvania, DE
- Carl-Fredrik Westin
Harvard Medical School –
Boston, US
- Alexander Wiebel
ZIB – Berlin, DE
- Eugene Zhang
Oregon State University, US
- Gary Hui Zhang
University College London, GB



Design of Reversible and Quantum Circuits

Edited by

Kenichi Morita¹ and Robert Wille²

1 Hiroshima University, JP, morita@iec.hiroshima-u.ac.jp

2 Universität Bremen, DE, rwille@informatik.uni-bremen.de

Abstract

It is a widely supported prediction that conventional computer hardware technologies are going to reach their limits in the near future. Consequently, researchers are working on alternatives. Reversible circuits and quantum circuits are one promising direction which allows to overcome fundamental barriers. However, no real design flow for this new kind of circuits exists so far. Physical implementations are in its infancy. Within this seminar, recent research questions of this emerging technology have been discussed.

Seminar 11.–14. December, 2011 – www.dagstuhl.de/11502

1998 ACM Subject Classification B.6 Logic Design, D.1 Programming Techniques, F.2 Analysis of Algorithms and Problem Complexity, I. Computing Methodologies


Keywords and phrases reversible computation, quantum computation, computer aided design, hardware and software design, physical implementation, applications

Digital Object Identifier 10.4230/DagRep.1.12.47

1 Executive Summary

Kenichi Morita

Robert Wille

License  Creative Commons BY-NC-ND 3.0 Unported license
© Kenichi Morita and Robert Wille

The development of computing machines found great success in the last decades. But the ongoing miniaturization of integrated circuits will reach its limits in the near future. Shrinking transistor sizes and power dissipation are the major barriers in the development of smaller and more powerful circuits. To further satisfy the needs for more computational power and further miniaturization, alternatives are needed that go beyond the scope of conventional technologies like CMOS. Reversible logic and quantum logic provide a promising alternative that may enhance or even replace conventional circuits in the future. More precisely:

■ Low Power Computation

While conventional circuits dissipate energy for each lost bit of information, reversible circuits are information lossless, i.e. theoretically they are not affected by this. Considering the increasing miniaturization, this makes reversible logic interesting for domains like low-power design. Besides this general paradigm, reversible circuits are particularly suited for complementary low-power solutions like adiabatic circuits or on-chip interconnect encoders.

■ Quantum Computation

Quantum Computation offers the promise of more efficient computing for problems that are of exponential difficulty for conventional computing paradigms. Considering that many of the established quantum algorithms include a significant Boolean component



Except where otherwise noted, content of this report is licensed under a Creative Commons BY-NC-ND 3.0 Unported license

Design of Reversible and Quantum Circuits, *Dagstuhl Reports*, Vol. 1, Issue 12, pp. 47–61

Editors: Kenichi Morita and Robert Wille



DAGSTUHL
REPORTS

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

(e.g. the oracle transformation in the Deutsch-Jozsa algorithm, the database in Grover's search algorithm, and the modulo exponentiation in Shor's algorithm), it is crucial to have efficient methods to synthesize quantum gate realisations of Boolean functions. Since any quantum operation inherently is reversible, reversible circuits can be exploited for this purpose.

However, no real design flow for these new kinds of circuits exists so far. Proposed approaches for synthesis, verification, and test are only applicable for very small circuits and systems. This is crucial since the design for reversible and quantum systems significantly differs from their conventional counterparts. Nearly all concepts and methods developed for conventional hardware design in the last decades have to be redeveloped in order to support the new technologies. Additionally, considering that today researchers are still faced with serious challenges for conventional technologies, it is worth working towards design solutions for reversible and quantum technologies already today.

The goal of the seminar was to bring together experts in order to present and to develop new ideas and concepts for the design of complex reversible and quantum circuits. In total, 17 presentations together with 1 tool demonstration (of the open source toolkit *RevKit*) and one panel sessions (on the different opinions on how reversible circuits can help to reduce power consumption during computation) have been conducted within the seminar. This has been accompanied by several working group and proposal preparation meetings. The most important topics which have been discussed were:

- **Design Methods**

How to (automatically) synthesize reversible and quantum circuits as well as check them for correctness?

Most of today's synthesis approaches for reversible and quantum circuits still rely on Boolean function descriptions like e.g. permutations, truth tables, or binary decision diagrams. In order to design complex circuits, higher levels of abstractions have to be considered. While for this purpose hardware description languages like VHDL, SystemC, or SystemVerilog have been established in conventional hardware design, high level synthesis of reversible and quantum circuits is just at the beginning. In order to advance this area, ideas about concepts, languages, and synthesis approaches for high level design have been presented and collected at the seminar.

- **Theoretical Consideration**

How can theoretical studies show us the way for realizing reversible/quantum computers?

In order to implement efficient reversible/quantum circuits and computers, we still need very basic and theoretical studies on them. This is because the paradigm of reversible/quantum computing has very different natures from that of conventional computing. Therefore, we shall still be able to find many novel and useful ideas for them through theoretical studies. In this seminar, theoretical consideration on various models in several levels have been presented and discussed. These models range from the element level to the software level, which include reversible logic elements and circuits, quantum automata, reversible Turing machines, and reversible programming languages.

- **Physical Realizations and Accuracy of Models**

How to physically implement the respective circuits?

How to close the gap between the theoretical models and the physical implementation?

In order to design reversible and quantum circuits, abstractions of the precise physical realizations are applied. These include the used gate library and the respective cost metrics, but also fault models for testing or abstractions for technology mapping. Due

to the progress in the development of physical realizations, these models constantly are subject to modifications which needed to be considered in the design phase. During the seminar, recent achievements in the development of physical realizations have been presented. This built the basis for discussions about updating and refining the applied models and abstractions.

■ **Applications**

How can reversible and quantum circuits be exploited in practically relevant application? How to measure and proof the benefits of these emerging technologies (e.g. how to substantiate improvements in the power consumption)?

So far, design methods have mostly been applied to “academic” examples only. However, the design of reversible circuits for precise applications is the next logical step. Possible directions (e.g. in the low-power domain) have been discussed at the seminar. This also included discussions of the requirements for such applications and how the benefits can be measured.

Results of the seminar are currently used in the preparation of upcoming scientific papers. As one example, a special issue on the results triggered by the ideas of this seminar is planned for 2013. Furthermore, the discussions encouraged the preparation of proposals for national and international research projects.

2 Table of Contents

Executive Summary

<i>Kenichi Morita and Robert Wille</i>	47
--	----


Overview of Talks

What do reversible Turing machines compute? <i>Holger Bock Axelsen</i>	51
Reversible CMOS computers <i>Alexis De Vos</i>	51
Reversible Logic Synthesis via Template Matching <i>Gerhard W. Dueck</i>	52
Quantum Finite automata on infinite words <i>Rusins Martins Freivalds</i>	52
Formal Verification of Quantum Systems <i>Simon Gay</i>	52
Principles of Reversible Computing <i>Robert Glück</i>	53
Quantum Automata Theory – A Review <i>Mika Hirvensalo</i>	54
Describing and Optimizing Reversible Logic using a Functional Language <i>Michael Kirkedal Thomsen</i>	54
Decision Diagram Techniques for Reversible and Quantum Circuit Equivalence Checking <i>D. Michael Miller</i>	56
Permutation Decision Diagrams (π DDs) and Analysis of Primitive Sorting Networks <i>Shin-ichi Minato</i>	56
Reversible logic elements with memory <i>Kenichi Morita</i>	57
Realization of Reversible Gates with New Quantum Gate Libraries <i>Zahra Sasanian</i>	57
RevKit: A Toolkit for Reversible Circuit Design <i>Mathias Soeken</i>	58
Testing and fault tolerance of reversible logic <i>Mehdi B. Tahoori</i>	58
Challenges in the Synthesis of Reversible Circuits: Today and Tomorrow <i>Robert Wille</i>	59
Logic level circuit optimization for topological quantum computation <i>Shigeru Yamashita</i>	60
A High-Level Reversible Programming Language <i>Tetsuo Yokoyama</i>	60
Participants	61

3 Overview of Talks

3.1 What do reversible Turing machines compute?

Holger Bock Axelsen (University of Copenhagen, DK)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Holger Bock Axelsen


We gave an overview of recent computability and complexity results for reversible Turing machines. We showed that garbage-free reversible Turing machines (without an extraneous input copy in its output) can compute injective functions only, but that all injective, computable functions are in range [1]. We gave a definition of universality for reversible programs, and briefly outlined a universal reversible Turing machine [2]. Moving on to computational complexity, we discussed some consequences of tape reduction [3], and raised the question of whether analogous results are obtainable for reversible circuits.

References

- 1 Axelsen, H.B., Glück, R.: What do reversible programs compute? In: Hofmann, M. (ed.) FOSSACS 2011. LNCS, vol. 6604, pp. 42–56. Springer-Verlag (2011)
- 2 Axelsen, H.B., Glück, R.: A simple and efficient universal reversible Turing machine. In: Dediu, A.H., Inenaga, S., Martn-Vide, C. (eds.) LATA 2011. LNCS, vol. 6638, pp. 117–128. Springer (2011)
- 3 Axelsen, H.B.: Time complexity of tape reduction for reversible Turing machines. In: De Vos, A., Wille, R. (eds.) RC 2011. LNCS, vol. 7165, pp. 1–13. Springer (2012)

3.2 Reversible CMOS computers

Alexis De Vos (Gent University, BE)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Alexis De Vos

Main reference A. De Vos, “Reversible computing,” Wiley-VCH, Weinheim (2010).
URL <http://users.elis.ugent.be/~aldevos/projects/computer.html>

CMOS (complementary metal-oxide-semiconductor) technology is the standard silicon technology for fabricating every-day electronic chips.

Applying the same transistor technology, but replacing AND, NAND, OR, NOR, and XOR gates by reversible gates, such as NOT gates, Feynman gates, Toffoli gates, and Fredkin gates, allows to build (classical) reversible circuits, i.e. chips that can compute either forwards or backwards, according to the applied driving force (i.e. applied voltages).

By using so-called adiabatic input signals (i.e. smooth voltage ramps), combined with dual-logic pass-transistor architecture, in principle, energy consumption per elementary computational step can be made arbitrarily small, and thus smaller than the Landauer limit. However, there is a fundamental trade-off with speed: the Landauer barrier can only be crossed if computing happens sufficiently slowly. A more practical obstacle is the threshold voltage of standard transistors. Only zero-threshold transistors can guarantee asymptotically zero energy consumption.


Standard threshold voltages lead to a more modest result: a reduction of the energy consumption by about a factor 10, compared to conventional digital circuit design.

References

- 1 A. De Vos: “Reversible computing”, Wiley-VCH, Weinheim (2010).

3.3 Reversible Logic Synthesis via Template Matching

Gerhard W. Dueck (University of New Brunswick at Fredericton, CA)


License  Creative Commons BY-NC-ND 3.0 Unported license
© Gerhard W. Dueck

Joint work of Dueck, Gerhard W.; Soeken, Mathias; Rahman, Mazder; Wille, Robert;

A review of reversible logic synthesis is presented. Heuristic synthesis of reversible circuits can be divided in two parts. First, find a reversible circuit (the circuit may be far from optimal). Second, optimize the circuit. One way to optimize the circuit is by applying rewriting rules, also known as templates. In this talk we prove two new results regarding template matching. First, the traditional definition of template does not guarantee minimality, since some templates are discarded. Secondly, we conjecture that given all templates with m lines and an exact template matching method an optimal circuit can be found for any circuit with m lines. A new solution for exact template matching is discussed. It should be feasible to find all templates with three lines. It remains to be seen how much reduction is possible for circuits with more than three lines. Some directions for further work are given.

3.4 Quantum Finite automata on infinite words


Rusins Martins Freivalds (University of Latvia, LV)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Rusins Martins Freivalds

The definition of language recognition by quantum finite automata in the case when the language consists of infinite words is somewhat non-trivial because there is no easy way how to consider a measurement after processing an infinite word. However a natural definition is possible based on measure theory developed by E.Borel and the standard definition of language recognition by deterministic and nondeterministic finite automata in the case when the language consists of infinite words developed by R.Büchi. We prove that Measure-Many finite quantum automata can recognize with probability 1 languages not recognizable by deterministic and nondeterministic finite automata. On the other hand, Measure-Once finite quantum automata can recognize with a bounded error only languages recognizable by deterministic finite automata.

3.5 Formal Verification of Quantum Systems

Simon Gay (University of Glasgow, GB)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Simon Gay

Joint work of Gay, Simon; Nagarajan, Rajagopal; Papanikolaou, Nikolaos;
Main reference S. J. Gay, N. Papanikolaou, R. Nagarajan, “QMC: a model checker for quantum systems,” in Proc. of the 20th Int’l Conf. on Computer Aided Verification (CAV). Springer LNCS 5123:543–547, 2008.
URL http://dx.doi.org/10.1007/978-3-540-70545-1_51

The field of formal methods is well established in classical computing. It defines formal languages in which to model systems and specify their desired properties, theories in which to express satisfaction of specifications by systems, and automated tools for verification. A range of formal methods techniques and tools have been applied to the analysis of classical


computing systems, especially concurrent and distributed systems, in diverse areas including networking, security, and safety-critical systems.

During the last several years, my collaborators and I have been developing formal methods for quantum systems and, more generally, systems that combine classical and quantum computation and communication. Our results include development of the theory of quantum process calculus, and a prototype model-checking tool for automatic verification of quantum systems (for example, communication protocols).

In my talk I motivated this research programme, argued that formal specification and verification should form part of the design process and tool chain for quantum systems, and presented a selection of results. The talk included a demonstration of the quantum model-checking tool QMC.

3.6 Principles of Reversible Computing

Robert Glück (University of Copenhagen, DK)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Robert Glück

We presented the main principles of reversible programming languages and discussed their relation to the design of reversible computer hardware. This includes a clear distinction between the mathematical specification and the operational properties of reversible programs which manifests itself in two main design steps: injectivization of the functional specification and the reversiblization of a program. The simple reversible programming language Janus was introduced as well as design principles for "good" reversible programming languages. We also outlined the main direction for what one may call 'reversible computing science'.

References

- 1 Abramov S. M., Glück R., Principles of inverse computation and the universal resolving algorithm. In: Mogensen T. A., Schmidt D. A., Sudborough I. H. (eds.), *The Essence of Computation: Complexity, Analysis, Transformation*. Lecture Notes in Computer Science, Vol. 2566, 269–295, Springer-Verlag 2002.
- 2 Axelsen H. B., Glück R., Yokoyama T., Reversible machine code and its abstract processor architecture. In: Diekert V., Volkov M.V., Voronkov A. (eds.), *Computer Science – Theory and Applications*. Proceedings. Lecture Notes in Computer Science, Vol. 4649, 56–69, Springer-Verlag 2007.
- 3 Axelsen H. B., Glück R., De Vos A., Thomsen M. K., MicroPower: towards low- power microprocessors with reversible computing. *ERCIM News*, Special Theme: Towards Green ICT, 79(1): 20–21, 2009.
- 4 Axelsen H. B., Glück R., What do reversible programs compute? In: Hofmann M. (ed.), *Foundations of Software Science and Computation Structures*. Proceedings. Lecture Notes in Computer Science, Vol. 6604, 42–56, Springer- Verlag 2011.
- 5 Glück R., Kawabe M., A method for automatic program inversion based on LR(0) parsing. *Fundamenta Informaticae*, 66(4): 367–395, 2005.
- 6 Yokoyama T., Glück R., A reversible programming language and its invertible self-interpreter. In: *Partial Evaluation and Program Manipulation*. Proceedings. 144–153, ACM Press 2007.
- 7 Yokoyama T., Axelsen H. B., Glück R., Principles of a reversible programming language. *Conference on Computing Frontiers*. Proceedings. 43–54, ACM 2008.

3.7 Quantum Automata Theory – A Review

Mika Hirvensalo (University of Turku, FI)

License © © © Creative Commons BY-NC-ND 3.0 Unported license
© Mika Hirvensalo

Main reference M. Hirvensalo, “Quantum Automata with Open Time Evolution,” *International Journal of Natural Computing Research* 1, pp. 70–85 (2010).

URL <http://dx.doi.org/10.4018/jncr.2010010104>

The main models of quantum finite automata are presented and their main properties are reviewed.

Measure-once -model by Moore and Crutchfield [5] is presented as a natural variant of probabilistic automata [6]. A further variant, Measure-many automata [4] were introduced by Kondacs and Watrous to enhance the language recognition capability. Ambainis & al. introduced Latvian automata [1], a model with more elegant closure properties and algebraic characterization. Hirvensalo [2], [3] introduced quantum automata with open time evolution, and this model can be seen as a true generalization, not only variant, of classical finite automata.

References

- 1 Andris Ambainis, Martin Beaudry, Marats Golovkins, Arnolds Kikusts, Mark Mercer and Denis Therien: Algebraic Results on Quantum Automata. *Theory of Computing Systems* 39: 1, 165–188 (2006).
- 2 Mika Hirvensalo: Various Aspects of Finite Quantum Automata. LNCS 5257 (Proceedings of DLT 2008), pp. 21–33 (2008)
- 3 Mika Hirvensalo: Quantum Automata with Open Time Evolution. *International Journal of Natural Computing Research* 1, pp. 70–85 (2010).
- 4 A. Kondacs and J. Watrous: On the power of quantum finite state automata. *Proceedings of the 38th Annual IEEE Symposium on Foundations of Computer Science*, pp. 66-75 (1997).
- 5 C. Moore, J. Crutchfield: Quantum automata and quantum grammars. *Theoretical Computer Science* 237, pp. 275–306 (2000).
- 6 M.O. Rabin: Probabilistic Automata. *Information and Control* 6, pp. 230–245 (1963).

3.8 Describing and Optimizing Reversible Logic using a Functional Language

Michael Kirkedal Thomsen (University of Copenhagen, DK)

License © © © Creative Commons BY-NC-ND 3.0 Unported license
© Michael Kirkedal Thomsen

The presentation showed the design of a language for the description and optimization of reversible logic circuits. The language is functional and based on combinators, where the most recognizable of these combinators is inversion, βf , that defines the inverse function of f using an efficient semantics.

It is important to ensure that all language constructs are reversible and for this language we will, furthermore, require that this is done by static analysis only.

For most reversible languages a run-time analysis is needed, but for circuit design a run-time check of the description is undesirable.


Future work will show that the combination of the functional language and the restricted reversible model will allow more possibilities in the term rewriting and, thus, we expect to do a better optimization than other optimization techniques for reversible circuits.

References

- 1 T. Altenkirch and J. Grattage. A functional quantum programming language. In 20th Annual IEEE Symposium on Logic in Computer Science, 2005. LICS 2005. Proceedings, pages 249–258. IEEE, June 2005.
- 2 H. B. Axelsen and R. Glück. What do reversible programs compute? In M. Hofmann, editor, FOSSACS, volume 6604 of LNCS, pages 42–56. Springer, 2011.
- 3 H. B. Axelsen, R. Glück, and T. Yokoyama. Reversible machine code and its abstract processor architecture. In V. Diekert, M. V. Volkov, and A. Voronkov, editors, CSR, volume 4649 of LNCS, pages 56–69. Springer, 2007.
- 4 J. Backus. Can programming be liberated from the von Neumann style? A functional style and its algebra of programs. In Communications of the ACM, volume 21, pages 613–641. ACM, 1978.
- 5 A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter. Elementary gates for quantum computation. Physical Review A, 52(5):3457–3467, 1995.
- 6 C. H. Bennett. Logical reversibility of computation. IBM Journal of Research and Development, 17:525–532, 1973.
- 7 E. Fredkin and T. Toffoli. Conservative logic. International Journal of Theoretical Physics, 21(3-4):219–253, 1982.
- 8 C. Lutz. Janus: A time-reversible language. A letter to R. Landauer. <http://tetsuo.jp/ref/janus.pdf>, 1986.
- 9 M. Sheeran. muFP, a language for VLSI design. In Proceedings of the 1984 ACM Symposium on LISP and functional programming, LFP 84, pages 104–112. ACM, 1984.
- 10 M. Sheeran. Hardware design and functional programming: a perfect match. Journal of Universal Computer Science, 11(7):1135–1158, jul 2005.
- 11 M. K. Thomsen, R. Glück, and H. B. Axelsen. Reversible arithmetic logic unit for quantum arithmetic. Journal of Physics A: Mathematical and Theoretical, 43(38):382002, 2010.
- 12 T. Toffoli. Reversible computing. In J. W. de Bakker and J. van Leeuwen, editors, ICALP, volume 85 of LNCS, pages 632–644. Springer, 1980.
- 13 V. Vedral, A. Barenco, and A. Ekert. Quantum networks for elementary arithmetic operations. Physical Review A, 54(1):147–153, July 1996.
- 14 R. Wille, S. Offermann, and R. Drechsler. SyReC: A programming language for synthesis of reversible circuits. In Proceedings of the Forum on Specification & Design Languages, pages 1–6, Southampton, UK, September 2010. IET.
- 15 T. Yokoyama, H. B. Axelsen, and R. Glück. Towards a reversible functional language. In A. De Vos and R. Wille, editors, RC 2011. Revised Selected Papers, LNCS. Springer, 2012, *to appear*.
- 16 T. Yokoyama and R. Glück. A reversible programming language and its invertible self-interpreter. In Partial Evaluation and Program Manipulation. Proceedings, pages 144–153. ACM Press, 2007.

3.9 Decision Diagram Techniques for Reversible and Quantum Circuit Equivalence Checking

D. Michael Miller (University of Victoria, CA)

License  Creative Commons BY-NC-ND 3.0 Unported license
© D. Michael Miller

Joint work of Miller, D. Michael; Wille, Robert; Grosse, Daniel; Drechsler, Rolf

Main reference R. Wille, D. Grosse, D.M. Miller, R. Drechsler, “Equivalence Checking of Reversible Circuits,” in Proc. ISMVL-2009, pp. 324–330, May, 2009.


URL <http://dx.doi.org/10.1109/ISMVL.2009.19>

The first part of this presentation is tutorial in nature. It reviews a selection of decision diagram techniques and how they have been applied to matrix representation for reversible and quantum circuits. The second part of the presentation shows how decision diagrams can be used to check the equivalence of two circuits regardless of the number of ancillary inputs and garbage outputs using novel techniques employing simple matrix operations.

The presentation addresses reversible circuits with binary inputs and outputs, but it is readily extended to the multiple-valued case, and to general quantum circuits. While this method has been developed using one particular type of decision diagram, QMDD, it should be straightforward to implement it using other decision diagram structures developed for quantum gates and circuits.

3.10 Permutation Decision Diagrams (π DDs) and Analysis of Primitive Sorting Networks

Shin-ichi Minato (Hokkaido University, JP)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Shin-ichi Minato

Main reference S. Minato, “ π DD: A New Decision Diagram for Efficient Problem Solving in Permutation Space,” in Proc. of 14th Int’l Conf. on Theory and Applications of Satisfiability Testing (SAT’11), pp. 90–104, 2011.

URL http://dx.doi.org/10.1007/978-3-642-21581-0_9

Recently, we proposed a new type of decision diagram named “ π DD,” for compact and canonical representation of a set of permutations. Similarly to an ordinary BDD or ZDD, π DD has efficient algebraic set operations such as union, intersection, etc. In addition, π DDs have a special Cartesian product operation which generates all possible composite permutations for two given sets of permutations. This is a beautiful and powerful property of π DDs.

In this talk, we present “permutation family algebra” based on π DDs, and how to describe and solve permutational problems using the algebra. We also show an application of π DDs for analyzing primitive sorting networks.

We succeeded in calculating the number of ways to construct minimum primitive sorting networks as 2,752,596,959,306,389,652 for $n = 13$, which has not been known in the past, where n is the width of the primitive sorting network.

3.11 Reversible logic elements with memory

Kenichi Morita (Hiroshima University, JP)

License © © ⊖ Creative Commons BY-NC-ND 3.0 Unported license
© Kenichi Morita

Main reference K. Morita, “Constructing a reversible Turing machine by a rotary element, a reversible logic element with memory,” Hiroshima University Institutional Repository, 2010.

URL <http://ir.lib.hiroshima-u.ac.jp/00029224>

We investigate a possibility of using reversible logic elements with memory (RLEM) as primitives for constructing reversible computing systems, from a theoretical standpoint. One of the characteristics of RLEMs is that there is no need to synchronize two or more input signals as in the case of reversible logic gates, because, in an RLEM, an incoming signal interacts only with its internal state, a stationary information. Another feature of an RLEM is that it can be modeled by an idealized physically reversible mechanical system easily, though its practical implementation in a reversible physical system is still very difficult as in the case of reversible gates. An important property of RLEMs is that some models of computing systems such as reversible Turing machines can be constructed by using a suitable RLEM very simply. We also discuss universality/non-universality of RLEMs. In the case of 2-state RLEMs, all but only four among the infinite kinds of RLEMs are universal in the sense that any reversible Turing machine can be built by each of them.

3.12 Realization of Reversible Gates with New Quantum Gate Libraries

Zahra Sasanian (University of Victoria, CA)

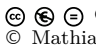
License © © ⊖ Creative Commons BY-NC-ND 3.0 Unported license
© Zahra Sasanian

Joint work of Sasanian, Zahra; Miller, D. Michael

The synthesized reversible circuits are usually evaluated based on quantum cost models. The most common quantum library for realizing a class of reversible gates called Multiple-Control Toffoli (MCT) gates is the NCV (NOT, CNOT, V, and V+) library. We presented two new quantum gate libraries as the target for technology mapping MCT cascades and showed their advantages over the well-known NCV library. The first new quantum library, NCVW, contains all the gates in the NCV library plus a gate implementing a fourth root of the NOT gate called the W gate and its inverse gate denoted by W+. The second new library includes NCV gates that are controlled by the quantum value, V1, instead of zero or one. A linear decomposition structure has also been proposed that utilizes this library to realize MCT gates with low linear quantum cost. The experimental results show that using these target libraries lead to less expensive circuits in terms of the number of gates compared to NCV circuits.

3.13 RevKit: A Toolkit for Reversible Circuit Design

Mathias Soeken (Universität Bremen, DE)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Mathias Soeken

Joint work of Soeken, Mathias; Frehse, Stefan; Wille, Robert; Drechsler, Rolf

Main reference M. Soeken, S. Frehse, R. Wille, R. Drechsler, “RevKit: An Open Source Toolkit for the Design of Reversible Circuits,” in *Reversible Computation 2011*, ser. Lecture Notes in Computer Science, vol. 7165, 2012, pp. 64–76, RevKit is available at www.revkit.org.

URL <http://www.revkit.org/>


RevKit is an open source toolkit for reversible circuit design. The motivation behind it is to make recent developments in the domain of reversible circuit design accessible to other researchers. Many approaches which are only available either independently or not at all can now be used under a common hood. This allows for the integration of different techniques in order to create new work flows. Furthermore, a C++ API and a Python interface make it possible to integrate own approaches. In this sense, RevKit addresses users who simply want to apply the framework and its tools as well as developers who actively want to develop further methods on top of the framework. RevKit is available at www.revkit.org.

References

- 1 M. Soeken, S. Frehse, R. Wille, and R. Drechsler, “RevKit: An Open Source Toolkit for the Design of Reversible Circuits,” in *Reversible Computation 2011*, ser. Lecture Notes in Computer Science, vol. 7165, 2012, pp. 64–76, RevKit is available at www.revkit.org.

3.14 Testing and fault tolerance of reversible logic

Mehdi B. Tahoori (KIT – Karlsruhe Institute of Technology, DE)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Mehdi B. Tahoori

Due to anticipated high failure rate of the emerging technologies to be used for realization of reversible logic, thorough testing and fault tolerance are crucial aspects in reversible circuits.

Fault masking techniques (to prevent error propagation) for reversible logic are presented and different implementations of reversible majority voters are shown. Using voter insertion techniques and taking advantage of available non-functional outputs, we are able to provide diagnosis information with adjustable resolution for reversible circuits. Such diagnosability has important applications in manufacturing testing as well as online repair. In contrast to previous work this voter is robust against single point of failure. We target missing and repeated gate fault models which are specific to reversible logic [1].

Also, an approach for online detection of faults in reversible circuits is presented. In this approach, reversible gates are modified in such a way that they can produce information on the number of cascaded gates. By using such information an appropriate reversible gate is added to detect missing and repeated gate faults [2].


By taking advantage of reversibility of these circuits, a new testing approach for reversible circuits is developed. In this method, the next test pattern is the response of the reversible circuit to the previous test pattern. This approach requires small amount of test information which makes it suitable for BIST implementation [3].

References

- 1 M. Zamani, and M.B. Tahoori. Online Missing/Repeated Gate Faults Detection in Reversible Circuits. In *Defect and Fault Tolerance of VLSI Systems*, pp. 435–442, oct. 2011.
- 2 M. Zamani, N. Farazmand, and M.B. Tahoori. Fault masking and diagnosis in reversible circuits. In *European Test Symposium*, pp. 69–74, may 2011.
- 3 M. Zamani, M. Tahoori, K. Chakrabarty, Ping-Pong Test: Compact Test Vector Generation for Reversible Circuits. In *VLSI Test Symposium*, 2012.

3.15 Challenges in the Synthesis of Reversible Circuits: Today and Tomorrow

Robert Wille (*Universität Bremen, DE*)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Robert Wille

Joint work of Drechsler, Rolf; Wille, Robert;

Main reference R. Drechsler, R. Wille, “From Truth Tables to Programming Languages: Progress in the Design of Reversible Circuits,” *ISMVL’11Tutorial*.

URL http://www.informatik.uni-bremen.de/agra/doc/konf/11_ismvl_reversible_circuit_design_tutorial.pdf


In the last decade, significant progress in the development of design methods for reversible circuits has been made. First synthesis approaches relied on function representations like truth tables or permutations. They enabled synthesis of reversible circuits with a minimal number of circuit lines, but were applicable to quite small functions only. As a result, researchers strived for more scalable synthesis approaches leading e.g. to ESOP-based synthesis, BDD-based synthesis or synthesis based on the preliminary hardware description language SyReC. While they enabled synthesis of more complex functions, they usually lead to circuits with too many circuit lines. Hence, after solving the “scalability”-problem, how to keep the number of circuit lines small remains as important research question.

References

- 1 R. Drechsler and R. Wille, “From Truth Tables to Programming Languages: Progress in the Design of Reversible Circuits,” in *Int’l Symp. on Multiple-Valued Logic*, 2011, pp. 78–85.
- 2 D. M. Miller, D. Maslov, and G. W. Dueck, “A transformation based algorithm for reversible logic synthesis,” in *Design Automation Conf.*, 2003, pp. 318–323.
- 3 D. Große, R. Wille, G. W. Dueck, and R. Drechsler, “Exact multiple control Toffoli network synthesis with SAT techniques,” *IEEE Trans. on CAD*, vol. 28, no. 5, pp. 703–715, 2009.
- 4 K. Fazel, M. Thornton, and J. Rice, “ESOP-based Toffoli gate cascade generation,” in *Communications, Computers and Signal Processing, 2007. PacRim 2007. IEEE Pacific Rim Conference on*, 2007, pp. 206–209.
- 5 R. Wille and R. Drechsler, “BDD-based synthesis of reversible logic for large functions,” in *Design Automation Conf.*, 2009, pp. 270–275.
- 6 R. Wille, S. Offermann, and R. Drechsler, “SyReC: A programming language for synthesis of reversible circuits,” in *Forum on Specification and Design Languages*, 2010, pp. 184–189.
- 7 R. Wille, O. Keszöcze, and R. Drechsler, “Determining the minimal number of lines for large reversible circuits,” in *Design, Automation and Test in Europe*, 2011, pp. 1204–1207.
- 8 R. Wille, M. Soeken, E. Schönborn, and R. Drechsler, “Circuit Line Minimization in the HDL-based Synthesis of Reversible Logic,” 2012.

3.16 Logic level circuit optimization for topological quantum computation

Shigeru Yamashita (Ritsumeikan University – Shiga, JP)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Shigeru Yamashita


Joint work of Yamashita, Shigeru; Devitt, Simon; Nemoto, Kae;

I formulated the logic level circuit optimization problem for topological quantum computation. Observing the properties of "braiding operations" in topological quantum computation, I formulated our problem as to find a good gate order and a good initial qubit order.

I then presented some heuristics for the problem.

3.17 A High-Level Reversible Programming Language

Tetsuo Yokoyama (Nanzan University, JP)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Tetsuo Yokoyama

We presented a summarizing talk that covered previous work on programming methodology in reversible computing including developing a better language for writing reversible programming, better ways to write reversible programming, and the fundamental concepts on reversible languages. First, we attempted to share brief overview of a high-level reversible programming language Janus.

Next, apart from particular computational models, the fundamental problems of irreversibility and approaches to their solutions were identified. The source of irreversibility stems from irreversible data update and backward nondeterministic control flow. The solutions for avoiding those irreversibility are mainly divided into unclean and clean approaches. The superiority of clean approach was discussed by using Janus, which is a clean reversible language.

Then, we described the details on programming techniques in a high-level reversible programming language. As other programming paradigms do, reversible programming gives us opportunity to access to its own way of modularity. After that, properties of reversible programming languages were briefly discussed. We posed a few of our current research questions. Our recent result on the domain-specific optimization of reversible simulation, twice as fast as Bennett method, is represented in the form of reversible logic gates, to show the potential usefulness of the reversible simulation in the domain-specific optimization of reversible logic gates.

Participants

- Holger Bock Axelsen
University of Copenhagen, DK
- Stéphane Burignat
Ghent University, BE
- Alexis De Vos
Ghent University, BE
- Rolf Drechsler
Universität Bremen, DE
- Gerhard W. Dueck
University of New Brunswick at
Fredericton, CA
- Rusins Martins Freivalds
University of Latvia, LV
- Simon Gay
University of Glasgow, GB
- Robert Glück
University of Copenhagen, DK
- Mika Hirvensalo
University of Turku, FI
- Pawel Kerntopf
Warsaw Univ. of Technology, PL
- Michael Kirkedal Thomsen
University of Copenhagen, DK
- D. Michael Miller
University of Victoria, CA
- Shin-ichi Minato
Hokkaido University, JP
- Kenichi Morita
Hiroshima University, JP
- Zahra Sasanian
University of Victoria, CA
- Julia Seiter
Universität Bremen, DE
- Mathias Soeken
Universität Bremen, DE
- Marek Szymowski
Warsaw Univ. of Technology, PL
- Mehdi B. Tahoori
KIT – Karlsruhe Institute of
Technology, DE
- Robert Wille
Universität Bremen, DE)
- Shigeru Yamashita
Ritsumeikan Univ. – Shiga, JP
- Tetsuo Yokoyama
Nanzan University, JP



Privacy and Security in Smart Energy Grids

Edited by

Stefan Katzenbeisser¹, Klaus Kursawe², Bart Preneel³, and
Ahmad-Reza Sadeghi⁴

1 TU Darmstadt, DE, katzenbeisser@seceng.informatik.tu-darmstadt.de

2 University of Nijmegen, NL, klaus.kursawe@gmail.com

3 K.U. Leuven, BE, Bart.Preneel@esat.kuleuven.be

4 TU Darmstadt, DE, ahmad.sadeghi@trust.cased.de

Abstract

The “smart energy grid” promises to improve the reliability and efficiency of the future energy grid by exchanging detailed usage information between the end consumers and the utilities. This application raises different questions with regard to privacy and security. For instance, detailed meter readings enable to infer detailed information on the private life of the consumers; furthermore, manipulations of meter readings open the possibility of fraud. The goal of the seminar was thus to raise awareness of the privacy and security problems associated with smart meters and bring together academic researchers as well as utility experts in order to start an open dialogue on smart grid privacy and security problems and potential solutions.

Seminar 18.–21. December, 2011 – www.dagstuhl.de/11511

1998 ACM Subject Classification K.4.1 Computers and Society, Public Policy Issues, Privacy

Keywords and phrases privacy, security, smart grid, digital metrology

Digital Object Identifier 10.4230/DagRep.1.12.62


1 Executive Summary

Klaus Kursawe

Stefan Katzenbeisser

Bart Preneel

Ahmad-Reza Sadeghi

License  Creative Commons BY-NC-ND 3.0 Unported license

© Klaus Kursawe, Stefan Katzenbeisser, Bart Preneel, Ahmad-Reza Sadeghi

The smart grid initiative is an attempt to improve reliability and efficiency of the electricity grid by adding communication and intelligence to its components all the way from end-user devices to the utilities. On the end user side, detailed usage information will be transferred to both home systems and the utilities; the utility can provide load- and pricing information to the meters and end-devices in real time. On the grid side, intelligent systems will allow for a more flexible energy distribution. Naturally, adding smartness to a critical and sizeable infrastructure system such as the electricity grid imposes extreme requirements on security and privacy, while facing numerous conflicting requirements from the different players. In addition, legislation is pushing hard to implement a large scale smart grid in a very short time: In Europe, the commission plans to achieve 80% smart grid coverage by 2020, with some countries starting to roll out meters at a large scale in 2012; in the US, the rollout has already started.



Except where otherwise noted, content of this report is licensed
under a Creative Commons BY-NC-ND 3.0 Unported license

Privacy and Security in Smart Energy Grids, *Dagstuhl Reports*, Vol. 1, Issue 12, pp. 62–68

Editors: Stefan Katzenbeisser, Klaus Kursawe, Bart Preneel, and Ahmad-Reza Sadeghi



Dagstuhl Reports

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

In such a setting, security and privacy are vital. A security breach of a smart energy grid can have severe consequences for power availability. With respect to privacy, the information gathered by the utility reveals a wealth of information about individual customers: examples are the day rhythm (power consumption data may reveal that a customer always comes home after the bars close, and has too little time between getting up and leaving the house to have breakfast), religious patterns (a devout Muslim may turn on the light for a morning prayer, or a catholic family may always leave home during the Sunday sermon), relationship patterns (energy usage may identify the days on which a group of people stayed in a house and the time when they went to bed), and even TV schedules (by combining electricity and water consumption measurements).

While it is not clear yet to which extent this data is going to be exploited, the potential privacy implications are substantial and have already been identified (after interoperability) as the second most important issue with the smart grid by NIST.

It is thus essential to build security and privacy protection into smart energy grids right from the start. The goal of this seminar was thus to raise awareness of this critical problem that may affect every European citizen within a couple of years and to bring together academic researchers as well as utility experts in order to start an open dialogue on smart grid privacy and security problems and potential solutions.

Topics covered during the seminar were:

- **Communication Security.** For the smart grid to work efficiently, end-user devices will need to communicate with the utility. The main challenge is that the end devices may be extremely limited in their capacity, and that commissioning—i.e., integration of a new device into a home- or office network—has to be simple and efficient. This will require new ways of secure communication between power consuming devices and smart meters as well as new ways to set up communication networks covering extremely small devices (such as light bulbs).
- **Privacy.** The amount of data collected about individual users in a smart grid setting is unprecedented, and leads to massive concerns about user's privacy. The setting is rather unique for privacy research: the data is not gathered for the profit of some company, but for the more noble cause of global energy savings, and the nature of the system makes it hard to temporarily opt out. Flexible Privacy-Enhancing Technologies are required to balance the conflicting requirements of privacy and data usage.
- **Implementation Security.** Already now, the first attacks on implementations of smart meters have been published. With a huge number of small embedded devices suddenly getting connected, implementation security becomes critical. Unfortunately, vendors of those devices are usually not experienced in protecting against network-based attacks, and resource constraints on such devices do not allow implementation of many standard security solutions designed to protect larger computer systems. Thus, new hardware security mechanisms are required.
- **Grid Architectures.** The smart grid combines architectural requirements that are inherently contradictory. On one side, control networks for critical systems should always put safety first, i.e., rather risk a data loss than a disruption in functionality. On the other side, this particular network deals with a huge amount of privacy related and security critical data, requiring adequate protection from data theft. New architectures need to be designed to accommodate both privacy and dependability at the same time.

2 Table of Contents

Executive Summary

Klaus Kursawe, Stefan Katzenbeisser, Bart Preneel, Ahmad-Reza Sadeghi 62

Overview of Talks

Metrology for the 21st Century: Security and Privacy
George Danezis 65

Smart Meter Security: Overview of European Initiatives and Member State Activities
Michael John 65

Some protective measures for privacy in Smart Grids
Florian Kerschbaum 65

Security in a changing DSO infrastructure
Erwin Kooi 66

Overview on Smart Grid Security and Privacy
Klaus Kursawe 66

Demand response of Smart Metering
Günter Müller 66

Can Security- and Privacy-Critical Applications be Cloudified? T-CLOUDS says YES!
Paulo Verissimo 67


What's going on in your neighborhood: Security and privacy analysis of utility meters
Wenyuan Xu 67

Participants 68

3 Overview of Talks

3.1 Metrology for the 21st Century: Security and Privacy


George Danezis (Microsoft Research UK – Cambridge, GB)

License  Creative Commons BY-NC-ND 3.0 Unported license
© George Danezis

Metrology as a field deals with measuring quantities, and legal metrology with devices that measure quantities relating to legal contracts. Modern meters are networked digital devices that are relied upon by multiple parties for their business, as in modern smart grid proposals. We argue that such meters should provide high integrity for their readings. Furthermore, through the use of modern signature and aggregation protocols we can require meters to support privacy: any computation can be performed on the readings privately by the data subject, without revealing those readings. These meters can be deployed in a variety of ways that are in line with current practices.

3.2 Smart Meter Security: Overview of European Initiatives and Member State Activities


Michael John (Elster GmbH – Mainz, DE)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Michael John

The European Commission initiated in 2009 the Task Force Smart Grids in order to facilitate the goals of the Third Energy Package. The mission of the Task Force was to advise the Commission's policy and regulation directions at EU level. This talk provided a summary of the activities of the Task Force's Expert Group on Smart Grid Security and Privacy as well as corresponding EU activities on standardization. Furthermore, an overview of the smart meter roll-outs in the different member states was given, outlining the diversity of the deployed solutions and their security levels, arguing for the need of EU-wide standards.

3.3 Some protective measures for privacy in Smart Grids


Florian Kerschbaum (TU Dresden, DE)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Florian Kerschbaum

Smart Metering collects time-granular consumption profiles. These profiles can be pseudonymized and then shared. We attempt a re-identification attack. First, we detect anomalies in the data, e.g. days of low or high consumption. Then we link all profiles based on similarity. We achieve 80%-90% accuracy in linking pseudonyms on our test data. The anomaly detection proves resistant against lower granularity of the collected data.

3.4 Security in a changing DSO infrastructure


Erwin Kooi (Alliander – Duiven, NL)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Erwin Kooi

The main drivers for smart grids from a DSO perspective are facilitating the energy transition and reducing time to repair of faults. The paradigm “supply follows demand” will change to “demand follows supply”, as supply of e.g. solar panels cannot be controlled. The grid will have to be able of transporting or managing the excess load or excess supply at lower levels in the grid. Having more insight in lower parts of the grid will help engineers troubleshoot faults and fault place locations. This insight will have to be done with respect for the privacy of our customers and should be done across all DSO’s.

3.5 Overview on Smart Grid Security and Privacy


Klaus Kursawe (Radboud University Nijmegen, NL)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Klaus Kursawe

Due to the perceived benefits and political pressure (e.g., the 20-20-20 rules in the 3rd energy package of the European Commission), the introduction of IT into the management of the electricity grid is rapidly progressing. This overview covers the motivation for this trend, the corresponding security and privacy issues, and the activities of the major regulatory bodies. On the privacy side, a protocol is described that allows the industry to operate the grid on aggregated data only, in a way that no personal data ever needs to exist outside the smart meter in unencrypted form.

3.6 Demand response of Smart Metering

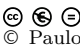
Günter Müller (Universität Freiburg, DE)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Günter Müller

Smart Metering is still a technical challenge. The objective is to collect enough information to manage the supply of energy. This generates a privacy issue, since the communication from smart meter to supplier is two ways. The smart meter reports not just the demand of energy but from the pattern of energy usage behavioral patterns can be deduced. To consider the privacy question as a technical question alone is not sufficient. The main questions are: Is energy in short supply? If yes, smart metering has a key role and privacy may be of secondary nature. If there is enough energy, just the supplier has an interest to know about behavioral patterns. Like in gas stations increase gasoline cost just short of vacation periods or for the weekend, power suppliers can adjust privacy according to demand. Privacy is to be designed to keep a balance of power between supplier and customer. This is beyond access control techniques as suggested today.

3.7 Can Security- and Privacy-Critical Applications be Cloudified? T-CLOUDS says YES!

Paulo Verissimo (University of Lisboa, PT)


License  Creative Commons BY-NC-ND 3.0 Unported license
© Paulo Verissimo

Joint work of Fernando André, Alysson Bessani, Miguel Correia, Pedro Costa, Marcelo Pasin, Bruno Quaresma, Paulo Sousa, Paulo Verissimo

As data and computation are moving to the cloud, worries about failures and disclosures increase. Whilst there is still some hesitation about moving critical applications onto the cloud, like e.g., medical records, financial data or smart energy grid control, time will come. This presentation discusses an approach towards a resilient cloud-of-clouds infrastructure, T-CLOUDS. A versatile architecture is introduced as well as some related algorithms and use cases, like dependable storage preserving integrity and confidentiality, or Byzantine fault-tolerant MapReduce.

3.8 What's going on in your neighborhood: Security and privacy analysis of utility meters

Wenyuan Xu (University of South Carolina, US)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Wenyuan Xu

Automatic meter reading (AMR) meters have been widely deployed in US and will be integrated into automatic meter infrastructure (AMI) in the near future. Thus, it is important to understand the security and privacy implication of the existing AMR meters. We have reverse-engineered a popular brand of AMR meters and shown that we are able to eavesdrop nearby AMR meters within 300 meters using a low noise amplifier and a 5 bBi antenna. Additionally, we can spoof AMR meters with arbitrary meter readings.

Participants

- Nikita Borisov
Univ. of Illinois – Urbana, US
- Binbin Chen
ADSC – Singapore, SG
- George Danezis
Microsoft Research UK –
Cambridge, GB
- Peter Ebinger
AGT Group (R&D) GmbH –
Darmstadt, DE
- Flavio D. Garcia
Radboud Univ. Nijmegen, NL
- Jorge Guajardo Merchan
Robert Bosch LLC –
Pittsburgh, US
- Matthias Hollick
TU Darmstadt, DE
- Bart Jacobs
Radboud Univ. Nijmegen, NL
- Michael John
Elster GmbH – Mainz, DE
- Stefan Katzenbeisser
TU Darmstadt, DE
- Florian Kerschbaum
TU Dresden, DE
- Erwin Kooi
Alliander – Duiven, NL
- Klaus Kursawe
Radboud Univ. Nijmegen, NL
- Leonardo Martucci
TU Darmstadt, DE
- Günter Müller
Universität Freiburg, DE
- Bart Preneel
K.U. Leuven, BE
- Carsten Rudolph
Fraunhofer SIT – Darmstadt, DE
- Ahmad-Reza Sadeghi
TU Darmstadt, DE
- Kazuo Sako
NEC – Kawasaki, JP
- Radu Sion
Stony Brook University, US
- Christian Stübke
Sirrix AG Bochum, DE
- Gene Tsudik
Univ. of California – Irvine, US
- Ingrid Verbauwhede
K.U. Leuven, BE
- Paulo Verissimo
University of Lisboa, PT
- Khan Ferdous Wahid
Fraunhofer SIT – Darmstadt, DE
- Jos Weyers
TenneT – Arnhem, NL
- Wenyan Xu
University of South Carolina, US

