DAGSTUHL
REPORTS

**Volume 4, Issue 1, January 2014**

**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*
June, 2014

*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`.



Digital Object Identifier: 10.4230/DagRep.4.1.i

*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.



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

www.dagstuhl.de/dagrep

Report from Dagstuhl Seminar 14021

# Symmetric Cryptography

**Edited by**

**Frederik Armknecht[1], Helena Handschuh[2], Tetsu Iwata[3], and Bart Preneel[4]**

1 **Universität Mannheim, DE,** armknecht@uni-mannheim.de
2 **Cryptography Research Inc. – San Francisco, US**
3 **Nagoya University, JP,** iwata@cse.nagoya-u.ac.jp
4 **K. U. Leuven, BE,** Bart.Preneel@esat.kuleuven.be

## Abstract

From 05.01.2014 to 10.01.2014, the Seminar 14021 in Symmetric Cryptography was held in Schloss Dagstuhl – Leibniz Center for Informatics. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

**Seminar** January 5–10, 2014 – http://www.dagstuhl.de/14021
**1998 ACM Subject Classification** E.3 Data Encryption, C.2 Computer-Communications Networks – General – Security and Protection, D.4.6 Security and Protection, H Information Systems – Security
**Keywords and phrases** Authenticity, Integrity, Privacy,Hash Functions, Block Ciphers, Provable Security, Cryptanalysis
**Digital Object Identifier** 10.4230/DagRep.4.1.1
**Edited in cooperation with** Qingju Wang

## 1 Executive Summary

*Frederik Armknecht*
*Helena Handschuh*
*Tetsu Iwata*
*Bart Preneel*



*Symmetric* cryptography deals with the case that both the sender and the receiver of a message are using the same key—the setting for symmetric encryption or authentication—as well as the case where there is no key at all—the setting for cryptographic hash functions. This differentiates symmetric cryptography from it *asymmetric* counterpart, where senders or verifiers use a "public key" and receivers or signers use a corresponding but different "private key." Although asymmetric cryptographic schemes provide in principle more flexibility, but are normally by orders of magnitude less efficient than symmetric cryptographic schemes. Thus, symmetric cryptosystems are the main workhorses of cryptography and highly relevant not only for academia, but also for industrial research, too.

The seminar was the fourth of its kind, the first one took place in 2007, the second in 2009, and the third in 2012. It concentrates on the design and analysis of

- *symmetric primitives* (block and stream ciphers, message authentication codes and hash functions), as well as
- *complex cryptosystems and cryptographic protocols* based on symmetric primitives.

One major topic was authenticated encryption. As already discussed at January 2012 Dagstuhl Seminar on Symmetric Cryptography, there is a demand for encryption schemes that ensure the confidentiality and integrity of data. This eventually led to an open cryptographic competition named CAESAR (Competition for Authenticated Encryption: Security, Applicability, and Robustness)[1] The goal of CAESAR is to identify a portfolio of authenticated ciphers that offer advantages over standard approaches like AES-GCM and (2) are suitable for widespread adoption. To this end cryptographic algorithm designers are invited to submit proposals of authenticated ciphers to CAESAR. All proposals will be made public for evaluation. As the deadline for first round submissions was in March 2014, i.e., only several weeks after the seminar, several groups were actively working on designing and analyzing new proposals for authenticated encryption schemes. Moreover, there was a discussion session that was mainly devoted to current CAESAR submissions. One result was a better understanding of necessary requirements and the current state of these schemes.

Another major topic was the analyis of Even-Mansour encryption schemes. Such schemes generalize common design approaches by reducing these to the composition of simple, idealized components like random permutations. Other topics focused during the discussion session include random number generation and provable security complex cryptosystems.

[1] See http://competitions.cr.yp.to/caesar.html.

## 2 Table of Contents

# 3 Overview of Talks

## 3.1 Solving LWE with BKW

*Martin R. Albrecht (Technical University of Denmark – Lyngby, DK)*


**Joint work of** Albrecht, Martin R.; Cid, Carlos; FaugÃre, Jean-Charles; Fitzpatrick, Robert; Perret, Ludovic
**Main reference** M. R. Albrecht, J.-C. Faugère, R. Fitzpatrick, L. Perret, "Lazy Modulus Switching for the BKW Algorithm on LWE," in Proc. of the 17th Int'l Conf. on Practice and Theory in Public-Key Cryptography (PKC'14), LNCS, Vol. 8383, pp. 429–445, Springer, 2014.
**URL** http://dx.doi.org/10.1007/978-3-642-54631-0_25

Some recent constructions based on LWE do not sample the secret uniformly at random but rather from some distribution which produces small entries. The most prominent of these is the binary-LWE problem where the secret vector is sampled from $\{0,1\}^*$ or $\{-1,0,1\}^*$. We present a variant of the BKW algorithm for binary-LWE and other small secret variants and show that this variant reduces the complexity for solving binary-LWE. We also give estimates for the cost of solving binary-LWE instances in this setting and demonstrate the advantage of this BKW variant over standard BKW and lattice reduction techniques applied to the SIS problem. Our variant can be seen as a combination of the BKW algorithm with a lazy variant of modulus switching which might be of independent interest.

## 3.2 On Increasing the Throughput of Stream Ciphers

*Frederik Armknecht (Universität Mannheim, DE)*


**Joint work of** Armknecht, Frederik; Mikhalev, Vasily
**Main reference** F. Armknecht, V. Mikhalev, "On Increasing the Throughput of Stream Ciphers," in Topics in Cryptology – Proc. of the Cryptographer's Track at the RSA Conf. 2014 (CT-RSA'14), LNCS, Vol. 8366, pp. 132–151, Springer, 2014.
**URL** http://dx.doi.org/10.1007/978-3-319-04852-9_7

Important practical characteristics of a stream cipher are its throughput and its hardware size. A common hardware implementation technique for improving the throughput is to parallelize computations but this usually requires to insert additional memory cells for storing the intermediate results, hence at the expense of an increased hardware size.

For stream ciphers with feedback shift registers (FSRs), we present an alternative approach for parallelizing operations with almost no grow of the hardware size by cleverly re-using existing structures. It is based on the fact that FSRs are usually specified in Fibonacci configuration, meaning that at each clock-cycle all but one state entries are simply shifted. The idea is to temporarily store values of the stream cipher outside of the FSR, e.g., intermediate results of the output function, directly into the FSRs.

We formally describe the transformation and its preconditions and prove its correctness. Moreover, we demonstrate our technique on Grain-128, one of the eSTREAM finalists with low hardware size. Our technique allows implementations, realized by the Cadence RTL Compiler considering UMC L180 GII technology, where the throughput is increased in the initialization mode by 18% and in the keystream generation mode by 24%, when the compiler was set to optimize the timing, and by 20 % in both modes when the compiler was set to optimize the area. As opposed to other solutions, no additional memory is required. In fact the hardware size even decreased from 1794 GE to 1748 GE in the time- optimized implementation and only slightly increased from 1627 GE to 1656 GE in the area-optimized implementation.

## 3.3 NORX

*Jean-Philippe Aumasson (Kudelski Security, CH)*


**Joint work of** Aumasson, Jean-Philippe; Jovanovic, Philipp; Neves, Samuel
**Main reference** NORX, submission to the CAESAR competition, submitted for publication.
**URL** https://norx.io

NORX is a parallel and scalable authenticated encryption algorithm with associated data (AEAD). The cipher is not patented and will be freely available for all applications. Likewise, its source code will be put under a public domain licence.

## 3.4 New Mobile Authentication and Key Agreement Algorithm

*Steve Babbage (Vodafone Group – Newbury, GB)*



Steve is the chair of ETSI SAGE (Security Algorithms Group of Experts), the standards group that specifies cryptographic algorithms for the 3GPP family of mobile telecoms standards (GSM, GPRS, UMTS, LTE) amongst other things. This talk is about a new authentication and key agreement algorithm that has recently been designed and ratified as a 3GPP standard. We explain the context in which the new algorithm operates, the motivation for standardising a new algorithm, and the rationale behind its design. The new algorithm (called TUAK) is built using the Keccak-f [1600] sponge function.

### References

1 3GPP. *Specification of the TUAK algorithm set: A second example algorithm set for the 3GPP authentication and key generation functions f1, f1*, f2, f3, f4, f5 and f5*; Document 1: Algorithm specification.* http://www.3gpp.org/DynaReport/35231.htm, 2013

## 3.5 Color Visual Cryptography with Scotch Tape and Polarizers

*Alex Biryukov (University of Luxembourg, LU)*



In this talk we have shown how to achieve color visual cryptography in practice without loss of resolution and contrast. Wide range of interference colors can be produced by placing variable thickness wave-plates between two crossed linear polarizers. We demonstrated simple arithmetic on this range of colors which can be used for symmetric encryption and secret sharing.

## 3.6 Complexity of Statistical Attacks

*Celine Blondeau (Aalto University, FI)*


**Joint work of** Blondeau, Céline; Nyberg, Kaisa
**Main reference** C. Blondeau, K. Nyberg, "Links Between Truncated Differential and Multidimensional Linear Properties of Block Ciphers and Underlying Attack Complexities," in Proc. of 33rd Annual Int'l Conf. on the Theory and Applications of Cryptographic Techniques (EUROCRYPT'14), to appear.

The various number of apparently different statistical attacks on block ciphers has raised the question about their relationships which would allow to classify them and determine those that give essentially complementary information about the security of block ciphers. While mathematical links between some of these attacks have been derived in the last couple of years, in this talk we present a relation between truncated differential and multidimensional linear attacks. By studying the data, time and memory complexities of a multidimensional linear key-recovery attack and its relation with the truncated differential one, we also show that in most cases a known-plaintext attack can be transformed into a less costly chosen-plaintext attack. In particular, we present a differential attack on 26 rounds of PRESENT in the chosen-plaintext model with less memory complexity than the previous attack.

Part of this presentation is also dedicated to the statistical saturation attack. In particular, we show that this attack is the same as a truncated differential attack, which allows us, for the first time, to provide a justifiable analysis of the complexity of the statistical saturation attack and discuss its validity on 24 rounds of the PRESENT block cipher. The link between the known plaintext multidimensional linear attack and the chosen plaintext truncated differential one can be generalized to the other statistical attacks and give further examples of attacks where the method used to sample the data required by the statistical test is more differentiating than the method used for finding the distinguishing property.

This is a joint work with Kaisa Nyberg accepted at Eurocrypt 2014.

## 3.7 About non-uniformity in anti-DPA threshold implementations

*Joan Daemen (STMicroelectronics – Diegem, BE)*


**Joint work of** Daemen, Joan; Bertoni, Guido; Nikov, Ventzislav; Nikova, Svetla; Peeters, Michaël; Van Assche, Gilles

We study threshold sharing schemes against DPA and investigate in what way the failure to meet the uniformity condition may jeopardize the immunity against first-order DPA.

For this we introduce a treatment of discrete distributions and vector Boolean mappings in the spectral domain using the Walsh-Hadamard transform that is of independent interest. We identify the characteristic properties of discrete distributions and mappings that are important in the macroscopic analysis: the total imbalance and imbalance contribution. We show that the total imbalance of the result of applying an iterated mapping to an input is the sum of the imbalance of that input plus the sum of the imbalances of the rounds of the iterated mappings. In the microscopic analysis we make use of (reduced) correlation matrices and imbalance vectors that are inherent in lossy mappings.

We apply this theory on non-uniform sharing and use the one for Keccak as a test bench for our techniques. It turns out that quantitatively the entropy loss in the 3-share Keccak architecture is not a problem, but that at the microscopic level some anomalies should be fixed.

## 3.8 Improved Linear Sieving Techniques with Applications to Step-Reduced LED-64

*Itai Dinur (ENS – Paris, FR)*


**Joint work of** Dinur, Itai; Dunkelman, Orr; Keller, Nathan; Shamir, Adi
**Main reference** I. Dinur, O. Dunkelman, N. Keller, A. Shamir, "Improved Linear Sieving Techniques with Applications to Step-Reduced LED-64," Cryptology ePrint Archive: Report 2013/634, 2013.
**URL** http://eprint.iacr.org/2013/634

In this paper, we present advanced meet-in-the-middle (MITM) attacks against the lightweight block cipher LED-64, improving the best known attacks on several step-reduced variants of the cipher in both single-key and related-key models. In particular, we present a known-plaintext attack on 2-step LED-64 with complexity of $2^{48}$ and a related-key attack on 3-step LED-64 with complexity of $2^{49}$. In both cases, the previously known attacks have complexity of $2^{60}$, i.e., only 16 times faster than exhaustive key search.

While our attacks are applied to the specific scheme of LED-64, they contain several general methodological contributions: First, we present the linear key sieve technique, which allows to exploit linear dependencies between key bits to obtain filtering conditions in MITM attacks on block ciphers. While similar ideas have been previously used in the domain of hash functions, this is the first time that such a technique is applied in block cipher cryptanalysis. As a second contribution, we demonstrate for the first time that a splice-and-cut attack (which so far seemed to be an inherently chosen-plaintext technique) can be used in the known-plaintext model, with data complexity which is significantly below the code-book size. Finally, we extend the differential MITM attack on AES-based designs, and apply it independently in two stages from both sides of the cipher, while using the linear key sieve and other enhancements.

## 3.9 Key Recovery Attack against HMAC/NMAC with Reduced Whirlpool

*Jian Guo (Nanyang TU – Singapore, SG)*


**Joint work of** Guo, Jian; Sasaki, Yu; Wang, Lei; Wang, Meiqin; Wen, Long; Wu, Shuang
**Main reference** J. Guo, Y. Sasaki, L. Wang, S. Wu, "Cryptanalysis of HMAC/NMAC-Whirlpool," in Proc. of the 19th Int'l Conf. on the Theory and Application of Cryptology and Information Security – Part II (ASIACRYPT'13), LNCS, Vol. 8270, pp. 21–40, Springer, 2013.
**URL** http://dx.doi.org/10.1007/978-3-642-42045-0_2

We presented the first key recovery attack against HMAC instantiated with Whirlpool hash function. Combining the generic state recovery attack developed recently, and progress in preimage attacks against Whirlpool, we are able to recover the original key with Whirlpool reduced to 6 rounds, and equivalent keys for 7 rounds. The later was based on recent progress in the MITM attacks against AES and AES-like block ciphers.

This talk is based on two pieces of work:

1. Jian Guo, Yu Sasaki, Lei Wang, Shuang Wu: Cryptanalysis of HMAC/NMAC-Whirlpool. ASIACRYPT 2013.
2. Jian Guo, Yu Sasaki, Lei Wang, Meiqin Wang, Long Wen: Equivalent Key Recovery Attacks against HMAC and NMAC with Whirlpool Reduced to 7 Rounds. FSE 2014.

## 3.10 Black-box, White-box, and Public-key ASASA schemes

*Dmitry Khovratovich (University of Luxembourg, LU)*


**Joint work of** Khovratovich, Dmitry; Biryukov, Alex; Bouillaguet, Charles

The informal notion of white-box cryptography was coined by Chow et al. 2002 as a method to protect cryptographic keys in a public implementation of encryption algorithms, which is fully accessed by an adversary. White-box implementations of the AES and DES ciphers were presented, but they were all badly broken. Subsequent attempts were no better. Whereas some theoretical foundations of white-box cryptography have been given recently in Wyseur's PhD thesis, so far they have not lead to any practical scheme.

I present an overview of the white-box cryptography concept along with the most common applications and proposed designs. I try to answer the question if the security of a white-box scheme can be relied on public scrutiny in contrast to the hardness assumptions behind RSA and other public-key schemes.

Alongside the theoretical results, I present some well-known attempts to construct a white-box cryptographic scheme from the AES and DES ciphers, and show their inherent weaknesses. Finally, I discuss some potential methods to construct a secure white-box cipher from scratch. Our first construction uses the results from finite fields theory and public-key cryptography and satisfies the strongest notion of white-box security. The second construction introduces the concept of a memory-hard cipher, that consumes a specified amount of memory and prohibits the adversary from recovering a compact representation of the scheme.

## 3.11 Tight Security Bounds for Triple Encryption

*Jooyoung Lee (Sejong University – Seoul, KR)*


**Main reference** J. Lee, "Tight Security Bounds for Triple Encryption," Cryptology ePrint Archive: Report 2014/015, 2014.
**URL** http://eprint.iacr.org/2014/015

In this talk, we revisit the old problem asking the exact provable security of triple encryption in the ideal cipher model. For a blockcipher with key length $k$ and block size $n$, triple encryption is known to be secure up to $2^{k+min\{k/2,n/2\}}$ queries, while the best attack requires $2^{k+min\{k,n/2\}}$ query complexity. So there is a gap between the upper and lower bounds for the security of triple encryption. We close this gap by proving the security up to $2^{k+min\{k,n/2\}}$ query complexity. With the DES parameters, triple encryption is secure up to $2^{82.5}$ queries, greater than the current bound of $2^{78.3}$ and comparable to $2^{83.5}$ for 2-XOR-cascade.

We also analyze the security of two-key triple encryption, where the first and the third keys are identical. We prove that two-key triple encryption is secure up to $2^{k+min\{k,n/2\}}$ queries to the underlying block cipher and $2^{min\{k,n/2\}}$ queries to the outer permutation. For the DES parameters, this result is interpreted as the security of two-key triple encryption up to $2^{32}$ plaintext- ciphertext pairs and $2^{81.7}$ block cipher encryptions.

**References**

1 Mihir Bellare and Phillip Rogaway, The Security of Triple Encryption and a Framework for Code-Based Game-Playing Proofs, EUROCRYPT'06, LNCS, Vol. 4004, pp. 409–426, DOI: 10.1007/11761679_25, Springer, 2006.

2 Peter Gazi and Ueli M. Maurer, Cascade Encryption Revisited, ASIACRYPT'09, LNCS, Vol. 5912, pp. 37–51, DOI: 0.1007/978-3-642-10366-7_3, Springer, 2009.

## 3.12 New Generic Attacks on Hash-based MACs

*Gaetan Leurent (University of Louvain, BE)*


**Joint work of** Leurent, Gaetan; Peyrin, Thomas; Wang Lei
**Main reference** New Generic Attacks on Hash-based MACs, Asiacrypt 2013, LNCS 8270, pages 1-20
**URL** http://dx.doi.org/10.1007/978-3-642-42045-0_1

In this paper we study the security of hash-based MAC algorithms(such as HMAC and NMAC) above the birthday bound. Up to the birthday bound, HMAC and NMAC are proven to be secure under reasonable assumptions on the hash function. On the other hand, if an n-bit MAC is built from a hash function with a l-bit state ($l < n$), there is a well-known existential forgery attack with complexity $2^{l/2}$. However, the remaining security after $2^{l/2}$ computations is not well understood. In particular it is widely assumed that if the underlying hash function is sound, then a generic universal forgery attack should still require $2^n$ computations and some distinguishing (e.g. distinguishing-H but not distinguishing-R) and state-recovery attacks should still require $2^l$ (or $2^k$ if $k < l$) computations.

In this work, we show that above the birthday bound, hash-based MACs offer significantly less security than previously believed. Our main result is a generic distinguishing-H and state-recovery attack against hash-based MACs with a complexity of only $\tilde{O}(2^{l/2})$. In addition, we show a key-recovery attack with complexity $\tilde{O}(2^{3l/4})$ against HMAC used with a hash functions with an internal checksum, such as GOST.

This surprising result shows that the use of checksum might actually weaken a hash function when used in a MAC. We stress that our attacks are generic, and they are in fact more efficient than some previous attacks proposed on concrete hash functions.

We use techniques similar to the cycle-detection technique proposed by Peyrin et al. at Asiacrypt 2012 to attack HMAC in the related-key model. However, our attack works in the single-key model for both HMAC/NMAC and without restriction on the key size.

## 3.13 Pipelineable On-Line Encryption (POE)

*Eik List (Bauhaus-Universität Weimar, DE)*



Correct authenticated decryption requires the receiver to buffer the decrypted message until the authenticity check has been performed. In high-speed networks, which must handle large message frames at low latency, this behavior becomes practically infeasible. This presentation proposes CCA-secure on-line ciphers as a practical alternative to AE schemes since the former

provide some defense against malicious message modifications. Unfortunately, all published on-line ciphers so far are either inherently sequential, or lack a CCA-security proof.

This paper introduces POE, a family of on-line ciphers that combines provable security against chosen-ciphertext attacks with pipelineability to support efficient implementations. POE combines a block cipher and an e-AXU family of hash functions. Different instantiations of POE are given, based on different universal hash functions and suitable for different platforms. Moreover, this presentation introduces POET, a provably secure on-line AE scheme, which inherits pipelineability and chosen-ciphertext-security from POE and provides additional resistance against nonce-misuse attacks.

## 3.14 Near-collisions in stream ciphers

*Willi Meier (FH Nordwestschweiz – Windisch, CH)*


**Joint work of** Zhang, Bin; Li Zhenqi; Meier, Willi
**Main reference** B. Zhang, Z. Li, "Near Collision Attacks on Grain v1m," in Proc. of the 20th Int'l Workshop on Fast Software Encryption (FSE'13), to appear.

Near-collision attacks on the stream cipher Grain v1 have been proposed by Zhang and Li at FSE 2013. The main idea is to identify near-collision internal states of a stream cipher at different time instants, and to retrieve the internal state. The original attacks on Grain v1 were based on assumptions that seem hard to verify. An improved framework for near-collision analysis of Grain v1 is developed that does not rely on assumptions and has assured success probability. The analysis is based on newly detected properties of Grain v1 and is a dedicated time-memory-data tradeoff attack that has lower data requirements than generic time-memory-data tradeoffs on stream ciphers.

## 3.15 APE: Authenticated Permutation-Based Encryption for Lightweight Cryptography

*Bart Mennink (KU Leuven, BE)*


**Joint work of** E. Andreeva, B. Bilgin, A. Bogdanov, A. Luykx, B. Mennink, N. Mouha, K. Yasuda, Kan, "APE: Authenticated Permutation-Based Encryption for Lightweight Cryptography," in Proc. of the 21st Int'l Workshop on Fast Software Encryption (FSE'14), to appear.

The domain of lightweight cryptography focuses on cryptographic algorithms for extremely constrained devices. It is very costly to avoid nonce reuse in such environments, because this requires either a hardware source of randomness, or non-volatile memory to store a counter. At the same time, a lot of cryptographic schemes actually require the nonce assumption for their security. In this paper, we propose APE as the first permutation-based authenticated encryption scheme that is resistant against nonce misuse. We formally prove that APE is secure, based on the security of the underlying permutation. To decrypt, APE processes the ciphertext blocks in reverse order, and uses inverse permutation calls. APE therefore requires a permutation that is both efficient for forward and inverse calls. We instantiate APE with the permutations of three recent lightweight hash function designs: Quark, Photon, and Spongent. For any of these permutations, an implementation that supports both encryption and decryption requires less than 1.9 kGE and 2.8 kGE for 80-bit and 128-bit security levels, respectively.

## 3.16 Triple and Quadruple Encryption: Bridging the Gaps

*Bart Mennink (KU Leuven, BE)*



Triple encryption is a cascade of three block cipher evaluations with independent keys, in order to enlarge its key size. This design is proven secure up to approximately $2^{\kappa+min\kappa/2,n/2}$ queries (by Bellare and Rogaway, EUROCRYPT 2006, and Gaži and Maurer, ASIACRYPT 2009), where kappa denotes the key size and n the block length of the underlying block cipher. On the other hand, the best known attack requires about $2^{\kappa+n/2}$ queries (by Lucks, FSE 1998, and Gaži, CRYPTO 2013). These bounds are non-tight for $\kappa \leq n$. In this work, we close this gap. By strengthening the best known attack as well as tightening the security bound, we prove that triple encryption is tightly secure up to $2^{\kappa+min\kappa,n/2}$ queries. Additionally, we prove that the same tight security bound holds for quadruple encryption (which consists of four sequentially evaluated block ciphers), and derive improved security and attack bounds for cascades consisting of five or more rounds. This work particularly solves the longstanding open problem of proving tight security of the well-known Triple-DES construction in the ideal model.

## 3.17 Towards Finding Optimal Differential Characteristics for ARX: Application to Salsa20

*Nicky Mouha (KU Leuven, BE)*


**Joint work of** Mouha, Nicky; Preneel, Bart
**Main reference** N. Mouha, B. Preneel, "Towards Finding Optimal Differential Characteristics for ARX: Application to Salsa20," Cryptology ePrint Archive: Report 2013/328,2013.
**URL** https://eprint.iacr.org/2013/328

An increasing number of cryptographic primitives are built using the ARX operations: addition modulo $2^n$, bit rotation and XOR. Because of their very fast performance in software, ARX ciphers are becoming increasingly common. However, there is currently no rigorous understanding of the security of ARX ciphers against one of the most common attacks in symmetric-key cryptography: differential cryptanalysis. In this paper, we introduce a tool to search for optimal differential characteristics for ARX ciphers. Our technique is very easy to use, as it only involves writing out simple equations for every addition, rotation and XOR operation in the cipher, and applying an off-the-shelf SAT solver. As is commonly done for ARX ciphers, our analysis assumes that the probability of a characteristic can be computed by multiplying the probabilities of each operation, and that the probability of the best characteristic is a good estimate for the probability of the corresponding differential. Using extensive experiments for Salsa20, we find that these assumptions are not always valid. To overcome these issues, we propose a method to accurately estimate the probability of ARX differentials.

## 3.18 LS-Designs: Bitslice Encryption for Efficient Masked Software Implementations

*Francois-Xavier Standaert (University of Louvain, BE)*


**Joint work of** Standaert, Francois-Xavier; Vincent Grosso, Gaëtan Leurent and Kerem Varici
**Main reference** V. Grosso, G. Leurent, F.-X. Standaert, K. Varici, "LS-Designs: Bitslice Encryption for Efficient Masked Software Implementations," in Proc. of the 21st Int'l Workshop on Fast Software Encryption (FSE'14), to appear; pre-print available from the author's webpage.
**URL** http://perso.uclouvain.be/fstandae/PUBLIS/142.pdf

Side-channel analysis is an important issue for the security of embedded cryptographic devices, and masking is one of the most investigated solutions to mitigate such attacks. In this context, efficient masking has recently been considered as a possible criteria for new block cipher designs. Previous proposals in this direction were applicable to different types of masking schemes (e.g. Boolean and polynomial). In this paper, we study possible optimizations when specializing the designs to Boolean masking. For this purpose, we first observe that bitslice ciphers have interesting properties for improving both the efficiency and the regularity of masked software implementations. Next we specify a family of block ciphers (denoted as LS-designs) that can systematically take advantage of bitslicing in a principled manner. Eventually, we evaluate both the security and performance of such designs and two of their instances, confirming excellent properties for physically secure applications.

## 3.19 Ketje and Keyak

*Gilles Van Assche (STMicroelectronics – Diegem, BE)*



In this presentation, we propose two variants of Keccak, called Ketje and Keyak, tailored for authenticated encryption. We focus on the design decisions behind these two proposals, which differ in the size of the state and consequently in the set of platforms they are best suited for: high-end platforms for Keyak and low-end platforms for Ketje.

Keyak uses the DuplexWrap construction, which is an improved version of SpongeWrap [SAC11]. Some of the Keyak instances are parallelizable, while others are serial. Cryptanalysis can be argued to be based on provable modes and a hermetic permutation.

Ketje uses a dedicated construction called MonkeyWrap. In contrast to Keyak, Ketje works at the round function level and cryptanalysis requires ad-hoc techniques for the function as a whole.

## 3.20 Uniformity on a diet

*Gilles Van Assche (STMicroelectronics – Diegem, BE)*


**Joint work of** Bilgin, Begül; Daemen, Joan; Nikov, Ventzislav; Nikova, Svetla; Rijmen, Vincent; Van Assche, Gilles
**Main reference** B. Bilgin, J. Daemen, V. Nikov, S. Nikova, V. Rijmen, G. Van Assche, "Efficient and First-Order DPA Resistant Implementations of Keccak," in Proc. of the 12th Smart Card Research and Advanced Application Conf. (CARDIS'13), to appear.

Besides hashing, Keccak can be used in many other modes, including ones operating on a secret value. Many applications of such modes require protection against side-channel

attacks, preferably at low cost. In this presentation, we present threshold implementations (TI) of Keccak with three and four shares that build further on unprotected parallel and serial architectures. We improve upon earlier TI implementations of Keccak in the sense that the latter did not achieve uniformity of shares. In our proposals we do achieve uniformity at the cost of an extra share in a four-share version or at the cost of injecting a small number of fresh random bits for each computed round. The proposed implementations are efficient and provably secure against first-order side-channel attacks.

## 3.21 Distance Bounding Protocols

*Serge Vaudenay (EPFL – Lausanne, CH)*


**Joint work of** Boureanu, Ioana; Mitrokotsa, Aikaterini; Vaudenay, Serge
**Main reference** Series of work presented at LATINCRYPT'12, FSE'13, LIGHTSEC'13, PROVSEC'13, ISC'13.
**URL** http://lasec.epfl.ch/infoscience/

Distance-bounding is the prominent solution to relay attacks against access control systems. In this talk, we review security models and existing protocols for distance-bounding. We identify two existing protocols with complete security results: SKI and FO. We compare them and identify open problems.

## 3.22 Automatic Search for Differential Trails in ARX Ciphers

*Vesselin Velichkov (University of Luxembourg, LU)*


**Joint work of** Biryukov, Alex; Velichkov, Vesselin
**Main reference** A. Biryukov, V. Velichkov, "Automatic Search for Differential Trails in ARX Ciphers", in Topics in Cryptology – Proc. of the Cryptographer's Track at the RSA Conf. 2014 (CT-RSA'14), LNCS, Vol. 8366, pp. 227–250, Springer, 2014; pre-print available at Cryptology ePrint Archive (Report 2013/328).
**URL** http://dx.doi.org/10.1007/978-3-319-04852-9_12
**URL** http://eprint.iacr.org/2013/853

In this talk we describe a tool for automatic search for differential trails in ARX (Addition, Rotation, XOR) ciphers. It is based on the well-known branch-and-bound technique, proposed by Mitsuru Matsui at EUROCRYPT'94 and used to find the best trails for up to 16 rounds of DES. Finding (near) optimal differential trails in symmetric-key primitives is critical for the accurate assessment of the security of the latter against one of the most powerful attacks – differential cryptanalysis. Being able to do this for ARX ciphers has always been a difficult problem. To the best of our knowledge this is the first application of Matsui's algorithm to ciphers that do not have S-boxes, such as ARX.

In the first part of the talk we briefly describe Matsui's algorithm and comment on the problems that arise when it is applied in its original form to ARX primitives. Next we present two solutions that address those problems: (1) partial difference distribution tables (pDDT) and (2) what we refer to as "The Highways and Country Roads Analogy". The latter two are combined into a modified version of Matsui's search technique called Threshold Search that is applicable to ARX.

Finally, we present results from the application of the threshold search tool to four ARX-based ciphers: TEA, XTEA, RAIDEN and SPECK. For TEA, the best trail found by our tool covers more than two times the number of rounds of the previous best (truncated) trail: 18 vs. 8. The best found trail for RAIDEN is 3-round iterative and covers all 32 rounds of the cipher. For three versions of the recently proposed cipher SPECK, namely SPECK-32/48/64, the best found trails cover 9, 10 and 13 rounds respectively. Those figures are comparable to a recent result by Abed et al. resp. 8, 10 and 13 rounds. The latter are however reported for differentials, while our results are for single trails. We also confirm the best known differential trail for XTEA that is on 14 rounds.

The source code of the tool is publicly available as part of a larger toolkit for the analysis of ARX at the following address: https://github.com/vesselinux/yaarx.

Extended version of the paper is available at IACR: http://eprint.iacr.org/2013/853.

## 3.23 Catena: A Memory-Consuming Password-Scrambling Framework

*Jakob Wenzel (Bauhaus-Universität Weimar, DE)*



It is a common wisdom that servers should better store the one-way hash of their clients' passwords, rather than storing the password in the clear. In this paper we introduce a set of functional properties a key-derivation function (password scrambler) should have. Unfortunately, none of the existing algorithms satisfies our requirements. Therefore, we introduce a novel and provably secure password-scrambling framework called Catena and derive an instantiation based, namely Catena-$\lambda$, which is based on a memory-consuming one-way function called $\lambda$ – bit-reversal graph ($\lambda$–BRG). It is characterized by its memory hardness, i.e., if one has only 1/c of memory available, one needs $c^{\lambda}$ processor units to gain the same time-memory trade-off. Thus, Catena-$\lambda$ excellently thwarts massively parallel attacks on cheap memory-constrained hardware, such as recent graphical processing units (GPUs). Additionally, we show that Catena-$\lambda$ is also a good key derivation function, since in the random oracle model it is indistinguishable from a random function. Furthermore, the memory access pattern of the $\lambda$–BRG is password-independent and therefore resistance against cache-timing attacks. Moreover, Catena supports (1) client-independent updates (the server can increase the security parameters and update the password hash without user interaction or knowing the password), (2) a server relief protocol (saving the server's resources at the cost of the client), and (3) a variant Catena-KG for secure key derivation (to securely generate many cryptographic keys of arbitrary lengths such that compromising some keys does not help to break others).

## Participants


- Martin R. Albrecht
Technical Univ. of Denmark – Lyngby, DK
- Elena Andreeva
KU Leuven, BE
- Frederik Armknecht
Universität Mannheim, DE
- Tomer Ashur
KU Leuven, BE
- Jean-Philippe Aumasson
Kudelski Security, CH
- Steve Babbage
Vodafone Group – Newbury, GB
- Daniel J. Bernstein
University of Chicago, US
- Eli Biham
Technion – Haifa, IL
- Alex Biryukov
University of Luxembourg, LU
- Céline Blondeau
Aalto University, FI
- Andrey Bogdanov
Technical Univ. of Denmark, DK
- Carlos Cid
Royal Holloway University of London, GB
- Joan Daemen
STMicroelectronics – Diegem, BE
- Itai Dinur
ENS – Paris, FR
- Orr Dunkelman
University of Haifa, IL
- Henri Gilbert
ANSSI – Paris, FR
- Jian Guo
Nanyang TU – Singapore, SG
- Tetsu Iwata
Nagoya University, JP
- Pascal Junod
HEIG-VD – Yverdon-les-Bains, CH
- Dmitry Khovratovich
University of Luxembourg, LU
- Matthias Krause
Universität Mannheim, DE
- Tanja Lange
TU Eindhoven, NL
- Nils Gregor Leander
Ruhr-Universität Bochum, DE
- Jooyoung Lee
Sejong University – Seoul, KR
- Gaetan Leurent
University of Louvain, BE
- Eik List
Bauhaus-Universität Weimar, DE
- Stefan Lucks
Bauhaus-Universität Weimar, DE
- Willi Meier
FH Nordwestschweiz – Windisch, CH
- Florian Mendel
TU Graz, AT
- Bart Mennink
KU Leuven, BE
- Nicky Mouha
KU Leuven, BE
- Kaisa Nyberg
Aalto University, FI
- Kenneth G. Paterson
Royal Holloway University of London, GB
- Thomas Peyrin
Nanyang TU – Singapore, SG
- Bart Preneel
KU Leuven, BE
- Christian Rechberger
Technical Univ. of Denmark – Lyngby, DK
- Greg Rose
Qualcomm Inc. – San Diego, US
- Yu Sasaki
NTT Labs – Tokyo, JP
- Francois-Xavier Standaert
University of Louvain, BE
- John Steinberger
Tsinghua Univ. – Beijing, CN
- Gilles Van Assche
STMicroelectronics – Diegem, BE
- Serge Vaudenay
EPFL – Lausanne, CH
- Vesselin Velichkov
University of Luxembourg, LU
- Qingju Wang
KU Leuven, BE
- Jakob Wenzel
Bauhaus-Universität Weimar, DE
- Kan Yasuda
NTT Labs. – Tokyo, JP


Report from Dagstuhl Perspectives Workshop 14022

# Connecting Performance Analysis and Visualization to Advance Extreme Scale Computing

**Edited by**

**Peer-Timo Bremer[1,3], Bernd Mohr[2], Valerio Pascucci[3], and Martin Schulz[4]**

1 **Lawrence Livermore National Laboratory*, US,** bremer5@llnl.gov
2 **Jülich Supercomputing Centre, DE,** B.Mohr@fz-juelich.de
3 **University of Utah – Salt Lake City, US,** pascucci@sci.utah.edu
4 **Lawrence Livermore National Laboratory*, US,** schulz6@llnl.gov

## Abstract

In the first week of January 2014 Dagstuhl hosted a Perspectives Workshop on "Connecting Performance Analysis and Visualization to Advance Extreme Scale Computing". The event brought together two previously separate communities – from Visualization and HPC Performance Analysis – to discuss a long term joined research agenda. The goal was to identify and address the challenges in using visual representations to understand and optimize the performance of extreme-scale applications running on today's most powerful computing systems like climate modeling, combustion, material science or astro-physics simulations.

**Perspectives Workshop** January 5–10, 2014 – http://www.dagstuhl.de/14022
**1998 ACM Subject Classification** C.4 Performance of Systems, D.1.3 Concurrent Programming (Parallel programming), D.4.8 Performance, I.3 Computer Graphics
**Keywords and phrases** Large scale data presentation and analysis, Exascale class machine optimization, Performance data analysis and root cause detection, High dimensional data representation
**Digital Object Identifier** 10.4230/DagRep.4.1.17

## 1 Executive Summary

*Peer-Timo Bremer*
*Bernd Mohr*
*Valerio Pascucci*
*Martin Schulz*



Over the last decades an incredible amount of resources has been devoted to building ever more powerful supercomputers. However, exploiting the full capabilities of these machines is becoming exponentially more difficult with each new generation of hardware. In the systems coming online at this moment, application developers must deal with millions of cores, complex memory hierarchies, heterogeneous system architectures, high-dimensional network topologies as well as a host of other hardware details that may affect the performance of a

* Part of this work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344 (LLNL-TR-653252).


Connecting Performance Analysis and Visualization to Advance Extreme Scale Computing, *Dagstuhl Reports*, Vol. 4, Issue 1, pp. 17–35
Editors: Peer-Timo Bremer, Bernd Mohr, Valerio Pascucci, and Martin Schulz

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

code. To help understand and optimize the behavior of massively parallel simulations a new subfield of computer science has grown devoted to developing tools and techniques to collect and analyze performance relevant data, such as execution time, operation counts, and memory or network traffic to help application developers pinpoint and ultimately fix performance problems. There now exist a number of standardized tools and APIs to collect a wide range of performance data at the largest scale. However, this success has created a new challenge, as the resulting data is far too large and too complex to be analyzed in a straightforward manner. While there exist some tools for performance analysis and visualization, these are predominately restricted to simple plots of the raw data and rely virtually exclusively on the users to infer connections between measurements and the observed behavior and to draw conclusions. Unfortunately, as the number of cores increases, this approach does not scale. The raw data is typically rather abstract, low-level, and unintuitive and it is difficult to understand within the context of the highly complex interaction of an application with the middle- and system software and the underlying hardware. For this reason, new automatic and more scalable analysis approaches must be developed to allow application developers to intuitively understand the multiple, interdependent effects that their algorithmic choices have on the resulting performance.

Following classical visualization mantra, the natural first step towards automatic analysis is to display an overview of the collected data to provide some insight into general trends. This helps both application developers and performance experts to form new hypotheses on potential causes of and solutions to performance problems. Furthermore, intuitive visualizations are highly effective in conveying the results of any analysis and thus are a valuable tool throughout the entire process. Unfortunately, visualizing performance data has proven challenging as the information is highly abstract, non-spatial, and often categorical. While some early attempts at including more advanced visualizations in performance tools have been proposed, these are rudimentary at best and have not found widespread adoption.

At the same time there exists a vibrant community in the area of information visualization and lately visual analytics that is exclusively aimed at developing techniques to visualize, illustrate, and analyze complex, non-spatial data. In particular, there exists a large body of work on general design principles of visualization tools, color spaces, and user interfaces as well as a wide array of common techniques that tackle a broad range of applications. The Dagstuhl Perspectives Workshop, for the first time, gathered leading experts from both the fields of visualization and performance analysis for joint discussions on existing solutions, open problems, and the potential opportunities for future collaborations.

The week started with a number of keynote sessions from well-known authorities in each area to introduce the necessary background and form a common baseline for later discussions. It soon became apparent that there exists a significant overlap in the common tasks and challenges in performance analysis and the abstract problem definitions and concepts common in visualization research. Subsequently, the workshop continued with short talks focusing on various more specific aspects of either existing challenges or potential solutions interspersed with increasingly longer group discussions. Theses extensive, inclusive, and in-depth exchanges ultimately shaped the second half of the workshop and in this form were only made possible through Dagstuhl's unique collaborative and discussion stimulating environment.

Ultimately, the workshop has started a number of collaborations and research projects between previously disparate fields with the potential of significant impact in both areas. Furthermore, the participants distilled the open challenges into three high-level recommendations: First, joined funding for the various open research questions. Second, support to

build and foster a new community on the border of visualization and performance analysis. And Third, the need to better integrate the anticipated results into the entire lifecycle of a massively parallel application from design to optimization and production.

## 2 Table of Contents

# 3 Keynotes and Topic Introduction

## 3.1 Really Quick Introduction to Parallel Performance Analysis

*Bernd Mohr (Jülich Supercomputing Centre, DE)*



This presentation provides a really quick introduction to parallel performance analysis. First, the common parallel programming paradigms MPI (Message Passing Interface) and OpenMP (Open specification for Multi-Processing) are introduced. Then the basic terminology and methodology for parallel program instrumentation, measurement, analysis and visualization is introduced. Finally, the two main sorts of performance data, summary profiles and detailed event trace, and their typical visualizations are explained in more detail.

## 3.2 Information Visualization and Visual Analytics

*Daniel A. Keim (Universität Konstanz, DE)*



There is a wide range of information visualization techniques that may be useful in analyzing the performance of HPC applications. The talk provides a brief overview of interesting visualization techniques and elaborates on the role of analysis and visualization. In putting visual analysis to work on big data, it is not obvious what can be done by automated analysis and what should be done by interactive visual methods. In dealing with massive data, the use of automated methods is mandatory – and for some problems it may be sufficient to only usefully automated analysis methods, but there is also a wide range of HPC problems where the use of interactive visual methods is necessary. The talk demonstrates a number of successful visual analytics applications from IP network monitoring, financial analysis, time series and high-dimensional data analysis.

## 3.3 Visualizing Performance Profiles with CUBE

*Felix Wolf (German Research School for Simulation Sciences, DE)*


**Joint work of** Wolf, Felix; Abraham, Erika; Bhatia, Nikhil; Böhme, David; Dennis, John; Geimer, Markus; Hatem, Youssef; Lücke, Monika; Pulatova, Farzona; Saviankou, Pavel; Schumacher, Peter; Song, Fengguang; Szebenyi, Zoltán; Visser, Anke; Voigtländer, Felix; Wylie, Brian
**URL** http://www.scalasca.org/software/cube-4.x/download.html

In this talk, we presented CUBE, an interactive browser that supports the visual exploration of performance profiles from parallel applications. CUBE's design emphasizes simplicity by combining a small number of orthogonal features with a limited set of user actions. The talk first described the underlying performance data model and the general usage model. It then focused on the visualization of process topologies, including Cartesian meshes with three and more dimensions and with potentially irregular subdomains. As further visualization targets, the talk covered time-dependent performance behavior and irregular domains. Finally, the talk discussed how CUBE interoperates with other visualization tools such as Vampir.

## 3.4 Vampir: Process Visualization of High Performance Software

*Holger Brunst (TU Dresden, DE)*


**Joint work of** Brunst, Holger; Knüpfer, Andreas
**Main reference** H. Brunst, A. Knüpfer, "Vampir," in David A. Padua (ed.), Encyclopedia of Parallel Computing, pp. 2125–2129, Springer, 2011.
**URL** http://dx.doi.org/10.1007/978-0-387-09766-4_60

Performance optimization is a key issue for the development of efficient parallel software applications. Vampir is a framework for performance analysis, which enables developers to quickly study program behavior at a fine level of detail. Performance data obtained from a parallel program run can be analyzed with a collection of specialized performance views. Intuitive navigation and zooming are the key features of the tool, which help to quickly identify inefficient or faulty parts of a program code.An important and unique feature of Vampir is its intuitive and interactive graphical representation of detailed performance event recordings over time (timelines) and as aggregated profiles. Extensive searching and filtering capabilities allow to quickly identify critical bottlenecks. In contrast to traditional profiling the details that caused a problem remain close at hand. The performance charts include rich sets of performance information and can be customized to the needs of both beginners and experts.New high performance computing systems are currently designed with a constantly increasing number of processing entities (cores), which is motivated by Moore's law and physical limitations. Performance tool visualizations have to follow this trend in order to be helpful in the future. Some of the resulting scalability issues in performance data visualization are presented in this talk in the scope of the Vampir framework.

## 3.5 Gathering and Interpreting Performance Data Visualizations – Initial Scaffolding

*Martin Schulz (Lawrence Livermore National Laboratory, US)*



Performance data can be gathered from a variety of different sources. This talk provides a straw-man for classifying these into four domains: hardware, tasks, code, and application data. Visualizing this data can in some cases be done within the same domain the data was collected in. An example is the visualization of hardware properties such as network links in the hardware domain. In other cases, though, the data is more helpful after it has been projected into a different domain and visualized there. An example is performance data of a CFD code mapped into the application domain to understand performance artifacts caused by the physics being simulated. The talk concludes with a proposal for such mappings and open research questions in establishing such a domain/mapping model, which decouples the data collection and visualization domains.

## 3.6 A(n Opinionated) Tour of InfoVis

*Carlos E. Scheidegger (AT&T Labs Research – New York, US)*



In this talk I present the following polemic: the entirety of information visualization design can be summarized as identifying mathematical structures in the input data space, and making sure they match in the visual encoding. In this sense, effective data visualizations are effective not because our eyes have more bandwidth, but rather because this visual representation of the data is actually closer its true nature. In addition, formulating our visualizations explicitly to elicit different structures in the data naturally leads to a classification of a large portion of dimensionality reduction techniques and graph drawing. Crucially, we don't require a rigid taxonomy to achieve this classification, so new data modalities can be comfortably incorporated by thinking about the properties of the input space that can or should be preserved.

## 3.7 Visual Analytics for High-Dimensional Data

*Klaus Mueller (Stony Brook University and SUNY Korea, US)*



Extreme scale computing can generate vast amounts of profile and trace data with many variables. In addition, the underlying machine has processors that are interconnected in a multi-dimensional (3-5) torus topology. The challenge is to visualize these data in such a way that multivariate relationships can be revealed, and optionally in the context of the multi-dimensional topology of the inter-connection network. I presented the underlying principles of visual analytics which is formed by a synergistic triad of visualization, data analytics, and interaction. The latter puts the analysts directly into the loop of the analytics data exploration process. I discussed various high-dimensional visualization techniques that come to live by the element of user interaction. All of these have data analytics and even machine learning algorithms under the hood to assist the analyst in the online reasoning process [5, 6]. Specifically, I discussed the method of (1) parallel coordinates as a ways to visualize the raw data along parallel dimension axes [1, 2], (2) correlation maps that allow users to see dependencies of patterns in the data [3], and (3) dynamic scatterplots that exploit motion parallax for an interactive high-dimensional viewing experience [4]. I also discussed our recent software framework, called "The ND-Scope" [8], which incorporates various facilities for high-dimensional data exploration and reasoning with high-dimensional data, also in the context of the earth's geography [7].

### References

1 A. Inselberg, B. Dimsdale. *Parallel Coordinates: A Tool for Visualizing Multi-dimensional Geometry*. IEEE Visualization, pp. 361–378, 1990.
2 K. McDonnell and K. Mueller. *Illustrative Parallel Coordinates*. Computer Graphics Forum, 27(3):1031–1038, 2008.
3 Z. Zhang, K. McDonnell, K. Mueller. *A Network-Based Interface for the Exploration of High-Dimensional Data Spaces*. IEEE Pacific Vis, Songdo, Korea, pp. 17–24, March, 2012.

4 J. Nam, K. Mueller. *TripAdvisorN-D: A Tourism-Inspired High-Dimensional Space Exploration Framework with Overview and Detail*. EEE Transactions on Visualization and Computer Graphics, 19(2): 291–305, 2013.
5 S. Garg, IV Ramakrishnan, K. Mueller. *A Visual Analytics Approach to Model Learning*. IEEE Conference on Visual Analytics Science and Technology (VAST'10), pp. 67–74, Salt Lake City, October, 2010.
6 S. Garg, E. Nam, IV. Ramakrishnan, K. Mueller. *Model-Driven Visual Analytics*. IEEE Symposium on Visual Analytics Science and Technology (VAST'08), pp. 19–26, 2008.
7 Z. Zhang, X. Tong, K. McDonnell, A. Zelenyuk, D. Imre, K. Mueller. *An Interactive Visual Analytics Framework for Multi-Field Data in a Geo-Spatial Context*. Tsinghua Science and Technology on Visualization and Computer Graphics, 18(2), April, 2013.
8 http://nd-scope.net/

## 3.8 Information Visualization for Performance Analysis – Initial Scaffolding

*Peer-Timo Bremer (Lawrence Livermore National Laboratory, US)*



Building on previous talks on the foundations of performance analysis this talk will provide an initial grouping and classification of existing visualization and analysis approaches. The goal is to provide the audience an overview of general approaches in visualization that may be applicable to performance analysis. In particular, with respect to the definition of different data domains we will highlight existing approaches in various domains and group them by different visualization paradigms such as hierarchical / focus+ context views or the type of interactions they allow. Finally, we provide some examples of gaps in the current landscape of techniques to start the discussion on how these may ultimately be addressed.

## 3.9 Some Remarks on the "User's" Perspective

*Hans-Joachim Bungartz (TU München, DE)*



The performance analysis community develops tools, and the workshop aims at identifying advanced visualization techniques, esp. related to information visualization and to visual analytics, that can enhance performance analysis tools – in terms of functionality, but in particular in terms of usability for those who are expected to use performance analysis. So what is the perspective of these "users" on risks and potentials, on automation and interaction, and on complexity and possible gains with respect to the ease of working with performance analysis tools? This perspective is the mandate of that presentation.The talk takes this point of view by clarifying the background, research practice, interest, and behavioral patterns of different "user specimens"; by clarifying the notion of performance from an algorithmic point of view; by pointing out the extreme importance of (numerical) algorithms, concerning both discretization and solution, as a link between "code" and "application" that must not be forgotten; by showing some examples of how code development,evaluation, and optimization in a CSE/HPC context works, when performance analysis enters the stage typically, and what is out of reach for performance analysis (currently); by briefly depicting typical users'

problems with performance analysis tools, their output data, and their current visualization capabilities.The talk ends with several somewhat provocative questions and statements, also taking into account the overview-type presentations on performance analysis and visualization from the workshop's first day.

# 4 Participant Contributions

## 4.1 Looking at Time-stamped Data

*Judit Gimenez (Barcelona Supercomputing Center, ES)*



Humans are visual creatures. Our brain is an expert on correlating, matching and extracting insight from visual information. In order to benefit from this capability the BSC Performance Tools (www.bsc.es/performance_tools) target to understand the applications behavior trough the analysis of detailed performance views. Our tools are based on traces to expose to the performance analyst all the details on the variability and distribution of the measured metrics. Paraver has only two types of views but a lot of flexibility to define and correlate them. Timelines provide the evolution along time and tables (histograms,profiles) a measurement of the metrics distribution. With Paraver, the analyst drives the full process having the power to decide how to explore the performance data generating and validating his/her hypothesis about the application behavior. The initial steps would be similar for all the analysis and are provided as basic methodology, the results of these steps would allow to decide the path to follow. To facilitate extracting insight from detailed performance data, during the last years we have been exploring the potential of applying data analytics techniques. Clustering, tracking and folding allow the performance analyst to identify the program structure, study its evolution and look at the internal structure of the computation phases. Paraver is a very flexible tool but it can also provide a very quick analysis of the application efficiency. A Paraver demo was given showing two complementary strengths of the tool: to demonstrate how using very few views and mouse clicks you can get a short diagnosis of your code as well as the ability of the tool to identify details and navigate through the collected data.

## 4.2 Extracting Logical Structure from MPI Traces

*Todd Gamblin (Lawrence Livermore National Laboratory, US)*


**Joint work of** Isaacs, Katherine; Gamblin, Todd; Bremer, Timo; Schulz, Martin; Bhatele, Abhinav; Hamann, Bernd
**Main reference** K. E. Isaacs, T. Gamblin, A. Bhatele, P.-T. Bremer, M. Schulz, B. Hamann, "Extracting logical structure and identifying stragglers in parallel execution traces," in Proc. of the 19th ACM SIGPLAN Symp. on Principles and Practices of Parallel Programming (PPoPP'14), pp. 397–398, ACM, 2014.
**URL** http://dx.doi.org/10.1145/2555243.2555288

Event traces are a valuable tool for understanding the performance of parallel programs, but analyzing large traces is difficult. Most current visualization techniques adopt a Gantt-chart like paradigm to show events over time. Unfortunately, laying out computation and

communication events in the order of their wall clock time introduces a large amount of clutter, which obscures important information in the visualization. In this paper, we describe a method to extract logical structure from a trace, which maintains order in terms of Lamport happened-before relations. We use this structure to define a notion of lateness among peers in a logical step. We show that laying out processes in terms of logical time and coloring them by lateness removes clutter, highlights parallel dependency chains, and uncovers the sources of parallel delays.

## 4.3 Prospective Visual Analysis Tools for Pattern Detection in Computing Performance Data

*Tobias Schreck (Universität Konstanz, DE)*



The analysis of computing performance data is an important problem to help understand computing resource consumption and to improve computing processes. It is also a difficult problem due to the size and complexity of data arising.Measurements may include e.g., large amounts of time- or network-oriented data.The visualization of the raw measurement data may not be most effective for the analysis, but techniques based on data aggregation and interactive filtering maybe needed. We present a selection of approaches from various application domains, in which data reduction, data visualization and interactive search are combined, and which may be a starting point for developing computing performance data analysis tools. Specifically, we survey techniques for exploration of clusters of network graphs and time-dependent scatter plot charts; for clustering of large numbers of heat map displays; and example-based search inline chart and scatter plot data. The aim of this talk is to survey a number of tools as to stimulate discussion on their application and extension for the analysis of computing performance data.

## 4.4 Mathematical Modeling of HPC Performance Analysis

*Hans Hagen (TU Kaiserslautern, DE)*



Analysis and simulation of complex physical phenomena, for example by CFD calculations, involves high performance computing on massively parallel machines. Often data is divided into subregions, transferred to a cluster or supercomputer and on to computing nodes. Overall performance of the computation is often measured in terms of efficiency, which is defined by the ratio of time used to actually solve the scientific problem by total run time. Performance heavily depend on various aspects like, for example, software implementation, system set-up, system resource occupancy during run time, communication behavior, and algorithmic parameter settings. In order to analyse HPC performance, various performance data, such as traces or energy use, are logged during run time. Traces capture the course of events during an HPC session and provide "insight" on what is happening where and under which condition.

Naturally, there are many challenges to the visualization of HPC performance data. For example, the mere data size brings numerous algorithmic and technical challenges. Nevertheless, at a first glance, this is looks like a typical visual analytics problem. But larger problems emerge at second glance. What is the mathematical characterization of the involved objects? How is the space or topology defined? In this talk, we formulate open problems from a visualization perspective that can serve as a basis for an in-depth discussion of this community.

## 4.5 A Potential Approach to Address Scalability Problems of Classical Performance Analysis by Combining Techniques from Visualization and Performance Analysis

*Matthias S. Mueller (RWTH Aachen, DE)*


**Joint work of** Mueller, Matthias S.; Henschel, Bernd

The rapid advancement of HPC has lead to huge scalability requirements for performance analysis tools. Huge progress has been made in the last 15 years (my personal time period with close contact to visualization and performance tools).However, facing limited financial resources all tools had to make their individual compromise between ease of use and scalability. As a consequence all tools share some problems: (a) they tend to solve problems of the past and fail at the highest scale or latest architecture/software and (b) the required abstraction level to cope with the scalability is difficult to understand even for advanced users.After all, I believe that this is an inherent problem that cannot be solved. But I also believe that an approach that combines different technologies and allows the users to switch between different abstraction levels and perspectives/views will be of great value at extreme scale. I will show three examples where current tools lack the flexibility to support that and how techniques from visualization can help to achieve this:

- Interactivity (Brushing , ....) and linked displays/views [2]
- Uncertainties
- Graphs

Especially the linked display/views technology combined with brushing us a key technology to maintain the understanding of the data that is analyzed and achieve the necessary insight. Because after all, performance analysis is not about tiny performance improvements, but it wants to achieve a general understanding of complex applications with their algorithms on complicated hardware architectures.

### References

1 B. Shneidermann *The eyes have it: a task by datatype taxonomy for information visualizations*. Proceedings of IEEE Symposium onVisual Languages, 1996.
2 Laramee, Robert S., et al. *Visual Analysis and Exploration of Fluid Flow in a Cooling Jacket*. Visualization, 2005. VIS 05. IEEE. IEEE, 2005.2005.

## 4.6 Scalable Representations for Performance Data

*Christopher Muelder (University of California – Davis, US)*


**Joint work of** Muelder, Christopher; Sigovan, Carmen; Ma, Kwan-Liu; Ross, Robert; Gygi, Francois; Cope, Jason; Lang, Sam; Iskra, Kamil; Beckman, Pete

Many performance visualization approaches rely on representations of per-node or per-process level data over time to instantiate an initial view of the data, but such representations generally can only depict minimally small examples, and rarely scale to the size of modern machines. We have been developing a number of representations of various parallel performance data classes aimed at being capable of scaling up to handle current and forthcoming scales, such as rank agnostic views, animated views, or system level metric behavioral similarity views. While these representations may not be as intuitive as classical representations, they have enabled views of data at scales unheard of before. Integration of these and similar techniques with more complex analytics and more intuitive detailed representations would lead to much more scalable visual analytic approaches for handling and exploring large scale performance data.

### References

1 Carmen Sigovan, Chris Muelder, and Kwan-Liu Ma *Visualizing Large-scale Parallel Communication Traces Using a Particle Animation Technique*. In Computer Graphics Forum 32(3), 2013, pages 141–150 (EuroVis 2013)
2 Carmen Sigovan, Chris Muelder, Kwan-Liu Ma, Jason Cope, Kamil Iskra, and Robert Ross. *A Visual Network Analysis Method for Large Scale Parallel I/O Systems*. In proceedings of 27th IEEE International Parallel and Distributed Processing Symposium (IPDPS), 2013, pages 308–319
3 Chris Muelder, Carmen Sigovan, Kwan-Liu Ma, Jason Cope, Sam Lang, Kamil Iskra, Pete Beckman, and Robert Ross *Visual Analysis of I/O System Behavior for High-End Computing*. In proceedings of 3rd Workshop on Large-Scale System and Application Performance (LSAP), 2011, pages 19–26 (Best Paper Award)
4 Chris Muelder, Francois Gygi, and Kwan-Liu Ma *Visual Analysis of Inter-Process Communication for Large-Scale Parallel Computing*. In IEEE Transactions on Visualization and Computer Graphics 15(6), 2009, pages 1129–1136 (InfoVis 2009)

## 4.7 Debugging and Hacking the User in Visual Analytics

*Remco Chang (Tufts University, US)*



The design of visualization systems always involves two primary considerations: the visualization system itself, and the user who uses the visualization. In this talk, I emphasize the consideration of the user, in particular in regards to how we could learn and leverage from the user's interactions so that the visualization system can "better understand" the user. The three examples I present are: (a) crowd sourcing past user histories using Scented Widgets [1]; (b) machine-learn the parameters of a distance metric from a user's interactions [2]; and (c) predict the user's analysis profile and behavior analyzing their interactions in real-time [3]. The three examples together demonstrate that there exist a tremendous

amount of information in the user's interactions that reveal their domain knowledge, past experiences, and individual differences. For the purpose of evaluating and creating better more adaptive visualization systems, incorporating these techniques would result in more precise evaluations that could in turn lead to the design of more personalized and intuitive interfaces and systems.

### References

1 W. Willett, J. Heer, and M. Agrawala. *Scented Widgets: Improving Navigation Cues with Embedded Visualizations*. EEE Transactions on Visualization and Computer Graphics, 13(6), pp. 1129–1136, 2007.
2 E. T. Brown, J. Liu, C. Brodley, R. Chang. *Dis-Function: Learning Distance Functions Interactively*. IEEE Visual Analytics Science and Technology (VAST), 83-92, 2012.
3 J. H. Zhao, Q. Lin, A. Ottley, and R. Chang. *Modeling User Interactions For Complex Visual Search Tasks*. Poster, IEEE Conference on Visual Analytics (VAST), 2013.

## 4.8 Re-constructing Events from Scalable and Streaming Unstructured Data

*Derek Xiaoyu Wang (University of North Carolina at Charlotte, US)*


**URL** http://cs.uncc.edu/~xwang25

Events are key elements for decision makers to form, discuss and react to social phenomena (e.g., movements, protests or campaigns) and natural phenomena (e.g. hurricane or coastal infrastructure changes). Yet, the increasing scale of digital data related to these phenomena presents challenges to sift through noisy information in order to discover meaningful events, which are characterized by changes in topics/features, geolocation, people and time. In addition, fully automated event analysis without human investigation could only provide limited evidence to decision makers. Therefore, identifying meaningful events is a fundamental challenge to visual analytics, machine learning and natural language processing research: it requires scalable architecture that balances the off-line data analytics and the interactive visual analysis with human assessment of the significance of the extracted event patterns. In this talk, I will discuss our research regarding this challenge through two novel visual analytics architectures for analyzing textual data and spatial data.First, I will introduce our scalable architecture that seamlessly integrates "big- data" text analytics (e.g. latent topic extraction and named entity recognition) with interactive visual exploration and pattern discovery. Specifically, the analytical environment presents a narrative visual interface regarding the investigative four W's: who, what, when and where for each event. It further allows users to interactively examine meaningful events using the four W's to develop an understanding of how and why. Designed to utilize Parallel Computing Clusters and other data- intensive platforms, our architecture enables stakeholders to make timely and more informed decisions with respect to on- going events, through knowing what the event is about, who is involved, where and when the event first occurred.Second, we have also applied our event-structuring concept to terrain analysis, where emergency responders must analyze temporal patterns and track changes in terrain features in response to possible natural disasters. With our narrative feature-extraction visualization platform, end-users can examine terrain features and ingest the related event structures through direct manipulations. Our platform uses search-by-example methods to automatically suggest similar terrain features and construct

animations to preview the changes. The end result is not simply looking at the terrain features disjointly, but a scripted-feature analysis that can be applied to LiDAR and SONAR data at scale.Collectively, my research focuses on advancing visual analytics research to reconstruct events from unstructured data, and identify a new visual analytics design framework for improving the scale and efficiency of the problem-solving process for analyzing textual and spatial data.

## 4.9 Some Inspiration From Scientific Visualization ...

*Hank Childs (University of Oregon, US)*


**URL** http://ix.cs.uoregon.edu/~hank/childs_dagstuhl.pdf

Although performance visualization will draw most heavily on techniques from information visualization, lessons learned from scientific visualization are likely useful. Specifically, the scientific visualization community has placed significant effort into flexible infrastructure that allows for dynamic composition of large numbers of algorithms. In this talk, I will discussing some of my own performance visualization endeavors to motivate why performance visualization could benefit from such infrastructure, and then present a high-level overview of the systems scientific visualization infrastructures use to achieve this user environment.

## 4.10 Towards an Integrated Performance Oriented Co-Design Methodology

*Ulrich Rüde (Universität Erlangen-Nürnberg, DE)*



Scalable algorithms are not necessarily fast, and highly optimized parallel software is not necessarily efficient. From the algorithmic perspective alone, the problems of an efficient implementation may be ignored, but just parallelizing and optimizing a given algorithm often leads also to poor results. Thus, in current practice, expensive supercomputer resources are often used inefficiently. A holistic approach will be necessary to better exploit present and future extreme scale parallel systems for computational science applications. A true co-design should incorporate the full computational modeling pipeline: It must begin where the physical model is developed, it must include the construction of the mathematical model, the discretization, the solver and of course the parallelization on all levels. Truly efficient software must equally exploit instruction level, node level, and system level parallelism. But the world of computational science has become complex and interdependent: An unfavorable choice of the discretization may inhibit vectorization, a shortsighted choice in developing the physical model may restrict parallel scalability, and an unwisely chosen solver library may lead to poor node level parallelism.

The presentation will illustrate our work towards developing a performance oriented co-design methodology for computational science and engineering together with a wish list what tool support will be necessary. The goal must be that the performance impact of any design decision along the whole pipeline should be critically quantified. This is primarily a question of a novel design methodology, but analysis tools and visualization methods are essential to make it feasible.

## 4.11 Scalable Visualization of Highly Distributed Computing Resources!?

*Wolfgang E. Nagel (TU Dresden, DE)*



ExaScale systems are expected to arrive between 2018 and 2020. However, the efficient usage of such huge and powerful computing systems is by no means trivial. The hardware gets increasingly complex, involving memory and cache hierarchies, accelerators, and hierarchical interconnects, leading to heterogeneous computing architectures. Complex parallel software will run on such hardware, composed of many complex components, including hybrid parallel paradigms, facing dynamic work loads, difficult load balancing, and more. The talk presents actual developments on the scalability of performance tools. Especially difficult gets the visualization of highest-scale performance data. The key problem is the graphical representation of a rapidly increasing number of concurrent processing elements on a display of limited size. This applies to both profile and trace data if the data has been recorded for individual processing elements. Recent research addresses the issue by means of data clustering, data hierarchies, and interactive 3D visualization.

# 5 Tool Demonstrations

## 5.1 Boxfish: a Tool for Projecting Performance Data

*Todd Gamblin (Lawrence Livermore National Laboratory, US)*


**Joint work of** Isaacs, Katherine; Landge, Aaditya; Gamblin, Todd; Bremer, Peer-Timo; Pascucci, Valerio; Hamann, Bernd
**Main reference** K. E. Isaacs, A. G. Landge, T. Gamblin, P.-T. Bremer, V. Pascucci, B. Hamann, "Exploring performance data with boxfish," in Companion of High Performance Computing, Networking, Storage and Analysis (SCC'12), pp. 1380–1381, IEEE, 2012.
**URL** http://dx.doi.org/10.1109/SC.Companion.2012.202

Understanding parallel performance data is difficult because the data is large, high-dimensional, and the relationships among measurements are not always clear. Measured data is associated with a wide variety of domains, such as network topology, multi-core node topology, logical processes and tasks, communication patterns, and the simulated physical domain. A performance problem measured in one domain may have a root cause in another domain. Boxfish is a performance data visualization tool that can project data from one domain to another according to mapping functions defined on relations of values in data domains. We demonstrate a prototype implementation of Boxfish and the power of its projection capabilities for understanding performance data.

## 5.2 Projections: Scalable Performance Analysis and Visualization

*Abhinav Bhatele (Lawrence Livermore National Laboratory, US)*



Projections is a performance visualization tool co-developed with the Charm++ parallel programming system. Performance of a message-driven, load-balanced system such as Charm++ becomes complicated to analyze. However, the same message-driven runtime can be leveraged to automatically instrument important events at very low cost. Based on such instrumentation, Projections provides tools for extensive post-mortem analysis. A distinguishing feature of Projections is the comprehensive set of views for analyzing performance that leads to insights about factors impacting performance. These include a rich timeline view and an outlier analysis tool. It also supports a novel, highly-scalable, live visualization tool for analyzing performance characteristics of running parallel programs.

## 5.3 Advanced Cube Features

*Markus Geimer (Jülich Supercomputing Centre, DE)*


**Joint work of** Lorenz, Daniel; Böhme, David; Mohr, Bernd; Strube, Alexandre; Szebenyi, Zoltán
**Main reference** D. Lorenz, D. Böhme, B. Mohr, A. Strube, Z. Szebenyi, "Extending Scalasca's Analysis Features," in A. Cheptsov, S. Brinkmann, J. Gracia, M. Resch, W. E. Nagel, (eds.), Tools for High-Performance Computing 2012, pp. 115–126, Springer, 2013.
**URL** http://dx.doi.org/10.1007/978-3-642-37349-7_8

In this presentation, we will demonstrate some of the more advanced features of the Cube profile browser. The main focus will be on the different views of the system dimension of profile experiments, such as the box plot and topology views showing the statistical or topological distribution of metric values across the system. In particular, the approaches implemented to display Cartesian topologies with more than three dimensions will be shown.

# 6 Working Groups Results

## 6.1 A Unified Data Model for Parallel Performance Data

*Peer-Timo Bremer (Lawrence Livermore National Laboratory, US)*
*Bernd Mohr (Jülich Supercomputing Centre, DE)*
*Valerio Pascucci (University of Utah – Salt Lake City, US)*
*Martin Schulz (Lawrence Livermore National Laboratory, US)*


**Joint work of** All participants of the Dagstuhl Perspectives Workshop 14022

One of the key insights for both areas – Performance Analysis and Visualization – has been that to precisely specify the problems, categorize potential methodologies, and to communicate across the area boundaries, the communities needed to establish a common unambiguous language and develop unifying models for data representation and processing. As will be discussed in more detail in the corresponding Dagstuhl manifesto, the workshop

participants developed a simple yet complete framework covering the entire pipeline from data collection, aggregation, and processing to the analysis and visualization. Furthermore, for each stage of the process the new methodology allows a comprehensive classification of existing techniques and well-posed descriptions of open challenges.

The developed data model consists of a set of **spaces** with unique indices, which describe the location of a sample, a set of **metrics**, representing the values or attributes of a sample, and a number of **maps** between spaces, which induce maps between values. A space describes where in the hardware (e.g., node or core id), software (e.g., function name or call path), parallel runtime (e.g., MPI rank or thread id), application space (e.g., on which mesh elements or simulated particles), or time (e.g., wallclock timestamp or phase id) a sample data point is taken. By introducing an experiment id as another space dimension, the data model can be used to describe the data of single performance measurement experiments as well as whole experiment series. Typical metrics for HPC performance analysis are inclusive and exclusive execution time, operations counts (e.g., loads, stores, integer or floating-point arithmetic) or network packet counts, but can also be more complex and higher dimensional data points like particle coordinates, physical machine locations, or references into complex data structures.

A **measurement** is a function that for a concrete set of space indices stores a set of metric values. Given different measurement, maps between their respective spaces allow one to correlate the measurements. A typical example for a map is the mapping between MPI ranks and the node or core id on which the MPI rank executed. The map between these space – typically called the *node mapping* – enables us to understand per-core measurements, e.g., FLOP counts, with per-rank information, e.g., number of bytes send or received. Another use case for map is to model the very common technique of profile measurements where instead of keeping every single data sample, the metric values get aggregated for the whole execution of an application. This can be modeled as a mapping from the original experiment space that includes the wallclock time to a reduced space without the time dimension, where the corresponding metric values get aggregated into the sum of the values. Finally, the model incorporates semantic **contexts** to certain space, e.g., the physical layout of the hardware interconnect, which enables corresponding visual encodings, e.g., a graph layout of the interconnect. Using these basic building blocks provides a comprehensive model in which to express existing measurement approaches, processing, and post-process analysis and visualization.

Using this model not only allowed the participants to more precisely and more easily explain issues and problems to each other, but also enabled them to very rapidly identify a number of open challenges as well as cross-cutting concerns all of which will require long term in-depth collaborations between both fields. Some of the missing components are new maps between certain spaces, which are crucial to allow correlating different measurements something widely considered indispensable for root-cause analysis. Other open problems are the need for scalable data representations and processing, and new attribution techniques. On the analysis and visualization side, intuitive visual metaphors for many semantic contexts are still missing, the handling of different notions of time, e.g., wallclock vs. logical time, remains rudimentary as does the treatment of ensembles of data. Furthermore, there exist a number of general concerns such as scalability, the need for multi-resolution and adaptive techniques, and a better understanding of the concept of time in the analysis and visualization.

## Participants

- Abhinav Bhatele
LLNL – Livermore, US
- Peer-Timo Bremer
LLNL – Livermore, US
- Holger Brunst
TU Dresden, DE
- Hans-Joachim Bungartz
TU München, DE
- Remco Chang
Tufts University, US
- Hank Childs
University of Oregon, US
- Todd Gamblin
LLNL – Livermore, US
- Markus Geimer
Jülich Supercomputing Centre, DE
- Judit Gimenez
Barcelona Supercomputing Center, ES
- Hans Hagen
TU Kaiserslautern, DE
- Daniel A. Keim
Universität Konstanz, DE
- Joshua A. Levine
Clemson University, US
- Naoya Maruyama
RIKEN – Kobe, JP
- Bernd Mohr
Jülich Supercomputing Centre, DE
- Christopher Muelder
Univ. of California – Davis, US
- Klaus Mueller
Stony Brook University, US
- Matthias S. Müller
RWTH Aachen, DE
- Wolfgang E. Nagel
TU Dresden, DE
- Valerio Pascucci
University of Utah, US
- Ulrich Rüde
Univ. Erlangen-Nürnberg, DE
- Carlos E. Scheidegger
AT&T Labs Research – New York, US
- Tobias Schreck
Universität Konstanz, DE
- Martin Schulz
LLNL – Livermore, US
- Derek Xiaoyu Wang
University of North Carolina at Charlotte, US
- Felix Wolf
German Resarch School for Simulation Sciences, DE

Report from Dagstuhl Seminar 14031

# Randomized Timed and Hybrid Models for Critical Infrastructures

**Edited by**

**Erika Ábrahám[1], Alberto Avritzer[2], Anne Remke[3], and William H. Sanders[4]**

1 **RWTH Aachen University, DE,** abraham@cs.rwth-aachen.de
2 **Siemens – Princeton, US**
3 **University of Twente, NL,** anne@cs.utwente.nl
4 **University of Illinois – Urbana, US,** whs@illinois.edu

## Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 14031 "Randomized Timed and Hybrid Models for Critical Infrastructures".

*Critical Infrastructures*, such as power grid and water and gas distribution networks, are essential for the functioning of our society and economy. *Randomized Timed and Hybrid Models* appear as a natural choice for their modeling, and come with existing algorithms and tool support for their analysis. However, on the one hand, the Critical Infrastructures community does not yet make full use of recent advances for Randomized Timed and Hybrid Models. On the other hand, existing algorithms are not yet readily applicable to the special kind of problems arising in Critical Infrastructures.

This seminar brought together researchers from these fields to communicate with each other and to exchange knowledge, experiences and needs.

**Seminar** January 12–17, 2014 – http://www.dagstuhl.de/14031
**1998 ACM Subject Classification** D.2.4 Software/Program Verification, D.4.5 Reliability, D.4.7 Organization and Design, F.1.2 Modes of Computation, G.3 Probability and Statistics
**Keywords and phrases** Critical Infrastructures, Smart Grids, Modeling, Randomized Timed and Hybrid Models, Analysis
**Digital Object Identifier** 10.4230/DagRep.4.1.36

## 1 Executive Summary

*Erika Ábrahám*
*Alberto Avritzer*
*Anne Remke*
*William H. Sanders*



### Seminar Description

More and more, our society and economy rely on the well-operation of, often hidden, Information and Communication Technology Infrastructures. These infrastructures play an ever-increasing role in other *Critical Infrastructures*, such as the power grid and water and gas distribution networks. Such systems are highly dynamic and include assets that are essential for the functioning of our society and economy. Users need to be able to place a high level of




trust in the operation of such systems, however, uncertainty in the environment, security and physical attacks, and errors in physical devices pose a serious threat to their reliable operation. Hence, it is very important that Critical Infrastructures survive catastrophic events.

Hence, modeling Critical Infrastructures and developing methods to analyze their *safety* and *dependability*, in the presence of failures and disasters is of utmost importance. It is of special interest to analyze, how quickly systems recover to acceptable levels of service after the occurrence of disasters, the so-called *survivability*. However, both failure and repair processes are random and a probability distribution is needed to describe how they evolve over time.

*Randomized Timed Models* are able to take the dependency of such processes on time into account and powerful techniques exist for their analysis. However, for Critical Infrastructures a modeling formalism is needed that allows describing both discrete and continuous quantities. Examples of discrete quantities are the number of spare parts and the state of sensors, actuators and Information and Communication Technology components, whereas the physical quantities, like the amount of produced energy or the quality of the treated water in terms of temperature and pressure naturally constitute continuous quantities.

*Randomized Hybrid Models* have been successfully applied to model safety-critical applications. Due to the flexible combination of discrete and continuous state components, Randomized Hybrid Models appear as a natural choice to accurately model Critical Infrastructures. Some formalisms were proposed for the analysis of Randomized Hybrid Models, and an increasing interest and activity can be observed in this field. Still, the industrial application that we are considering is far too large for state-of-the-art approaches; either they are applicable to specific applications only or they do not scale.

Up till now, most modeling in Critical Infrastructures is still fairly "classical" using reliability block diagrams, fault-trees or simplistic stochastic Petri nets. While researchers from the Critical Infrastructures community could benefit from recent advances for Randomized Hybrid Models and their formal analysis, existing algorithms are not yet readily applicable to the special kind of problems arising in Critical Infrastructures.

This clearly shows the need for bringing together experts in the areas of Randomized Timed Models and Randomized Hybrid Models with those from Critical Infrastructures. In the following we describe interesting advances in all three fields and comment on how they can help to bridge the current gap between the fields.

## Critical Infrastructures

Critical Infrastructures are in general controlled by SCADA (supervisory control and data analysis) systems, which are potentially vulnerable to attacks and misuse. SCADA systems consist of sensors, actuators, controllers and a human-machine interface through which human operators control the physical process. It is important to correctly capture interdependencies that arise between the SCADA network and the physical network, but also interdependencies between different Critical Infrastructures.

The complex nature of Critical Infrastructures requires a flexible and scalable compositional modeling framework that is able to accommodate different levels of abstraction. At design time, usually not all parameters and not all usage patterns are known exactly. Also the specific details of vulnerabilities and failures might be unknown, such as the mean time to failure and the impact of a given vulnerability. In such cases it is appropriate to make

*stochastic* assumptions about the system and the disaster behavior.

The heterogeneity of typical Critical Infrastructures may require a *combination* of different formalisms and techniques to describe the various components of a system and their dependencies. For example, the combination of continuous and discrete phenomena may need to be captured in the modeling framework, e.g, to model the process automation and the production process which is the essential part of several Critical Infrastructures.

Interactions and dependencies between subsystems of different nature inside a Critical Infrastructure or among cooperating Critical Infrastructures require advanced methods to reconcile different aspects under a common development and assessment framework. *Compositional* modeling can simplify the modeling process and can lead to intuitive formalisms. Furthermore, it enables compositional analysis techniques, which might reduce the complexity of verification and build a challenging topic that requires additional research.

In the seminar we discussed questions like the following:

- Which modeling methods are suitable for which types of Critical Infrastructures?
- Which are the crucial system issues that must be considered when accurately modeling Critical Infrastructures?
- How to distinguish the crucial parameters, thereby keeping the state space of the models as small as possible?

## Randomized Timed Models

Randomized Timed Models have been widely used for the modeling and evaluation of, e.g., computer and communication systems. They are in general well understood, suited to model complex systems, and efficient methods and tools exist for their analysis and simulation. Different modeling formalisms differ, e.g., in the model of time (discrete or continuous), in the existence or absence of nondeterminism, or the support of rewards.

Discrete-Time Markov Chains (DTMCs) belong to the most basic probabilistic models, offering a *discretized* model of time in the absence of nondeterminism. Continuous-Time Markov Chains (CTMCs) extend DTMCs by a *continuous* model of time. Several temporal logics were extended to specify relevant properties of Randomized Timed Models, and model checking algorithms were developed to check their validity for the above models. For example, Probabilistic CTL (PCTL) properties for DTMCs can be checked efficiently by solving systems of linear equations. Furthermore, efficient computation algorithms have been developed for model checking Continuous Stochastic Logic (CSL) properties of CTMCs (Baier, Haverkort, Hermanns, Katoen, 2003).

High-level formalisms like General Stochastic Petri Nets (GSPNs) and Stochastic Activity Networks allow to describe complex systems in a more compact way. Their evaluation can be lead back to methods for Markov chains.

Failure and repair processes of Critical Infrastructures often exhibit *nondeterminism*. Markov Decision Processes (MDPs) and Continuous-Time Markov Decision Processes (CTMDPs) extend DTMCs respectively CTMCs with the notion of nondeterminism. These powerful models can be analyzed by determining an optimal scheduler that removes the nondeterminism from the system and allows to apply the model checking approaches for DTMCs and CTMCs. Algorithms exist that compute such optimal schedulers based on solving the underlying optimization problems.

The non-functioning of Critical Infrastructures easily results in huge economic losses. To model the costs of failure and repair, a notion of *reward* can be added to the above models,

resulting in so-called Markov Reward Models (MRMs). To specify properties related to rewards, CSL has been extended to Continuous Stochastic Reward Logic (CSRL). Adding rewards to Randomized Timed Models makes the model checking problem very challenging. However, numerical algorithms exist for, e.g., model checking CSRL properties with arbitrary time and reward intervals for CTMCs with rewards. This is extremely useful for Critical Infrastructures, since these algorithms provide a direct and precise method for model checking survivability properties (Cloth, Haverkort, 2005).

There is quite a number of *tools* available for the analysis of the above model types. The most prominent ones are PRISM, MRMC, Möbius, Smart, CADP, or LiQuor. Besides formal verification, there are also simulation-based tools (e.g., APMC, VESTA). Most of these tools were successfully applied to different industrial case studies. However, these formalisms and tools are only partially suited for the model checking of Critical Infrastructures, mainly due to the lack of scalability and modeling power.

Model checking for the above models suffers from the well-known state explosion problem when applied to highly complex and large models of Critical Infrastructures. This problem could be tackled by compositional modeling and verification. However, though the models themselves support compositionality, there are no methods and tools readily available for compositional verification. Moreover, all the above models lack the power to model continuous physical processes, which is an essential part of Critical Infrastructures. Hence, the following section focuses on Randomized Hybrid Models.

In the seminar we discussed questions like the following:

- What are the (dis)advantages of the different modeling formalisms available?
- Which properties of Critical Infrastructures can already be efficiently analyzed with existing techniques?
- What are the requirements for compositional modeling and verification?

## Randomized Hybrid Models

When adding continuous behavior to discrete systems, the *hybrid* models become very powerful and in general undecidable. The most popular modeling formalism for hybrid systems are Hybrid Automata. Several analysis techniques were proposed for their reachability analysis, based on, e.g., approximation, hybridization, linearization, the usage of theorem provers, and interval-arithmetic.

Different approaches exist to extend hybrid models with *randomized* behavior. The most important difference between the extensions is *where* randomness is introduced. Timed Automata and Hybrid Automata were extended with *probabilistic discrete jumps* (in the style of DTMCs and MDPs) to Probabilistic Timed Automata respectively Probabilistic Hybrid Automata. In contrast to probabilistic discrete jumps, other formalisms, e.g., Piecewise Deterministic Markov Processes (Davis, 1993), allow *initialized jumps* to take place at *random times* (in the style of CTMCs and CTMDPs).

An orthogonal extension lies in introducing *stochastic differential equations* for modeling perturbations in the dynamic time behavior. When combined with probabilistic discrete jumps, this yields the model of Stochastic Hybrid Systems (Hu, Lygeros, Sastry, 2000). Another possibility considers the combination with CTMC-style stochastic jumps resulting in Switching Diffusion Processes (Gosh, Araposthatis, Marcus, 1997).

Only some simple classes of these models are decidable; their analysis can be lead back to the analysis of corresponding decidable classes of Hybrid Automata (Sproston, 2000).

Despite the undecidability of the above general classes, there are incomplete approaches available for their analysis, based on, e.g., Markov Chain approximation (Prandini, Hu, 2006) or discrete approximation (Koutsoukos, Riley, 2008). Latest work considers CEGAR-style abstraction that allows the application of model checking methods for Hybrid Automata (Zhang, She, Ratschan, Hermanns, Hahn, 2010).

Also the high-level Petri Net models can be extended with hybrid and randomized behavior. Including a notion of time, as in Timed Automata, results in Timed Petri Nets. Hybrid Petri Nets (David, Alla, 2001) are a high-level formalism for general Hybrid Automata. Colored Petri Nets correspond to Piecewise Deterministic Markov Processes (Everdij, Blom, 2009), supporting initialized stochastic jumps. Fluid Stochastic Petri Nets can be seen as a generalization[1] of Piecewise Deterministic Markov Processes, allowing for jumps to take place after a negative exponentially distributed amount of time. Besides the stochastic jumps, these models resolve nondeterminism by introducing discrete probability distributions for concurrently enabled transitions. This way, these models support both a probabilistic choice of jumps and a stochastic randomization of the time point of jumps, making the models extremely expressive and hard to formally analyze. Fluid Stochastic Petri nets can be solved analytically for up to three fluid places. For more general classes, simulation has to be used.

This variety illustrates the emerging interest of the research community in Stochastic and Probabilistic Hybrid Models. Traditionally, academic research focuses stronger on decidable subclasses than on efficient algorithms applicable to more expressive models. However, especially for Critical Infrastructures, models are needed that are able to specify complex continuous dynamics, e.g, in order to study recoverability processes.

For more expressive hybrid models, available analysis methods apply techniques like simulation, dynamic programming, and approximation. The Critical Infrastructures community would strongly benefit from the developments of modern model checking algorithms for models combining randomized and hybrid behavior.

In the seminar we discussed questions like the following:

- What particular hybrid model classes are suitable for Critical Infrastructures?
- How can initialized models be evaluated?
- How can efficient analysis (especially model checking) techniques be adapted for Randomized Hybrid Models?

## Achievements of the Research Seminar

This seminar offered a platform to bring together researchers, both from academia and industry, working on *Randomized Timed Models*, *Randomized Hybrid Models* and *Critical Infrastructures*. The program of the seminar was a balanced combination of (i) tutorials and presentations from all three fields to motivate collaboration and to develop a common ground for discussions and (ii) time for collaboration, where actual progress is expected to be made on increasing the efficiency, applicability and application of formal modeling and analysis techniques for Critical Infrastructures.

More specifically, we feel that this seminar helped to improve the development in the given area in the following points:

[1] by skipping the requirement of initialized jumps

1. The seminar *increased the interest* for both the academic development and the industrial application of formal methods to Critical Infrastructures and draw attention to open issues. We discussed industrially relevant case studies and specific requirements on modeling formalisms and evaluation techniques in this context.
2. While most of the existing work on Critical Infrastructures focuses on simulation, this seminar aimed at a thorough discussion of the requirements for appropriate formal analysis techniques. We provided an *overview* of the modeling and analysis methods already available in *Randomized Timed and Hybrid Models*, including a thorough discussion of their *suitability* for Critical Infrastructures.
3. We initiated *discussions and cooperations* that advance the state-of-the-art in Critical Infrastructures, regarding both the *development* and the *application* of suitable modeling formalisms and analysis techniques for Critical Infrastructures. We offered a platform to join expertise from different fields, to exchange knowledge about existing methods and applications, to push forward the communication of needs and interests, and to draw attention to challenging research fields and promising applications in the area of Critical Infrastructures.

## 2 Table of Contents

# 3 Overview of Talks

The talks in this section are listed in the order in which they were given.

## 3.1 Optimization Strategies for the Future Electricity Infrastructure – Smart Grid Research and Current Market Opportunities

*Albert Molderink (University of Twente, NL)*



Emerging technologies and a growing awareness of the drawbacks of our conventional energy supply increase the stress on the electricity infrastructure. In this presentation we briefly addressed these trends and the effects they have. Next, algorithms and strategies developed at the University of Twente and proposed in literature to deal with these effects were described. Finally, a few already introduced and emerging market opportunities were mentioned.

## 3.2 Engineering Cyber-Physical Systems/Critical Infrastructure Systems: A Craftsman Approach

*Peter Langendörfer (IHP GmbH – Frankfurt/Oder, DE)*


**Joint work of** Peter Langendörfer, Oliver Stecklina, Krzysztof Piotrowski and Steffen Peter

In this talk we shortly reported on CPS/CIS we built in the last years, e.g. in the project WSAN4CIP (http://www.wsan4cip.eu), to provide a practical view on what current problems are and how we tried to solve them. On the one hand we did the whole selection of soft- and hardware components manually, on the other hand we started to develop tools [1] that assist the developer in selecting appropriate components, getting an idea of potential deployment settings etc. Even though our tools provide some benefit compared to fully manual design there are still a lot of open questions. Our tools focus on functional aspects, of individual components, a thorough assessment compiled system is still missing. Timing aspects are currently also neglected, which is a serious problem given the real time requirements of CPS/CIS.

### References

1 K. Piotrowski and S. Peter. *Sens4U: Wireless sensor network applications for environment monitoring made easy*. In Proc. of the 4th Int. Workshop on Software Engineering for Sensor Network Applications (SESENA'13), in conjunction with ACM/IEEE International Conference on Software Engineering (ICSE'13), 2013.

## 3.3 Design of Distribution Automation Networks using Survivability Modeling and Power Flow Equation

*Daniel Sadoc Menasche (University of Rio de Janeiro, BR)*


**Joint work of** Daniel Sadok Menasche, Alberto Avritzer, Sindhu Suresh, Rosa M. Leão, Edmundo Souza e Silva, Morganna Diniz, Kishor Trivedi, Lucia Happe, and Anne Koziolek

Smart grids are fostering a paradigm shift in the realm of power distribution systems. Whereas traditionally different components of the power distribution system have been provided and analyzed by different teams, smart grids require a unified and holistic approach taking into consideration the interplay of distributed generation, distribution automation topology, intelligent features, and others. In this talk, we presented how we use transient survivability metrics to create better distribution automation network designs. Our approach combines survivability analysis and power flow analysis to assess the survivability of the distribution power grid network. Additionally, we presented an initial approach to automatically optimize available investment decisions with respect to survivability and investment costs. We have evaluated the feasibility of this approach by applying it to the design of a real distribution automation circuit. Our empirical results indicate that the combination of survivability analysis and power flow can provide meaningful investment decision support for power systems engineers.

### References

1 D. S. Menasché, A. Avritzer, S. Suresh, R. M. M. Leão, E. Souza e Silva, M. Diniz, K. Trivedi, L. Happe and A. Koziolek. *Assessing survivability of smart grid distribution network designs accounting for multiple failures*. Concurrency and Computation: Practice and Experience, Wiley Online Library, 2014.

2 A. Avritzer, S. Suresh, D. S. Menasché, R. M. M. Leão, E. de Souza e Silva, M. C. Diniz, K. Trivedi, L. Happe and A. Koziolek. *Survivability models for the assessment of smart grid distribution automation network designs*. In Proc. of the ACM/SPEC Int. Conf. on Performance Engineering, pp. 241–252, ACM, 2013.

3 A. Koziolek, A. Avritzer, S. Suresh, D. S. Menasche, K. Trivedi and L. Happe. *Design of distribution automation networks using survivability modeling and power flow equations*. In Proc. of the 24th IEEE Int. Symposium on Software Reliability Engineering (ISSRE'13), pp. 41–50, IEEE, 2013.

## 3.4 A Common Analysis Framework for Smart Distribution Networks Applied to Security and Survivability Analysis

*Lucia Happe and Anne Koziolek (Karlsruhe Institute of Technology, DE)*



Existing analysis approaches for power networks focus on analyzing the power network components separately. For example, communication simulation provides failure data for communication links, while power analysis makes predictions about the stability of the traditional power grid. However, these insights are not combined to provide a basis for design decisions for future smart distribution networks.

In this talk, we described an envisioned common model-driven analysis framework for smart distribution networks based on the Common Information Model (CIM [3]). This framework shall provide scalable analysis of large smart distribution networks by supporting analysis on different levels of abstraction. We plan to apply the framework to security analysis. Furthermore, we have applied our framework to holistic survivability analysis: We mapped the CIM on a survivability model [2] to enable assessing design options with respect to the achieved survivability improvement [1].

### References

1 A. Koziolek, L. Happe, A. Avritzer and S. Suresh. *A common analysis framework for smart distribution networks applied to survivability analysis of distribution automation*. In Proc. of the 1st Int. Workshop on Software Engineering Challenges for the Smart Grid (SE-SmartGrids'12), pp. 23–29, IEEE, 2012.

2 A. Avritzer, S. Suresh, D. Sadoc Menasché, R. M. Meri Leão, E. de Souza e Silva, M. Carmem Diniz, K. Trivedi, L. Happe and A. Koziolek. *Survivability models for the assessment of smart grid distribution automation network designs*. In Proc. of the 4th ACM/SPEC Int. Conf. on Performance Engineering (ICPE 2013), pp. 241–252, ACM, New York, NY, USA, 2013.

3 *IEC 61970 energy management system application program interface (EMS-API) – Part 301 Common Information Model (CIM) Base*. Edition 3.0, IEC, Aug 2011.

## 3.5 Tutorial: Formal Methods for Hybrid Systems

*Erika Ábrahám (RWTH Aachen University, DE)*


**Joint work of** Erika Ábrahám, Xin Chen, and Sriram Sankaranarayanan

Critical infrastructures often exhibit both dynamic and discrete behavior. Typically, the dynamic behavior stems from the continuous evolution of the physical system state, whereas the discrete behavior stems from the execution steps of discrete controllers. In this sense, critical infrastructures can be seen as *hybrid* systems. Models for hybrid systems can be formalized in different languages. Tools like for example Simulink are popular, because they offer rich libraries of model components and come with additional useful functionalities like, e.g., simulation. Unfortunately, such powerful modeling languages often lack a formal semantics. As an alternative, *hybrid automata* [2], extending discrete automata with continuous dynamics, can be used. Once a hybrid system is modeled in a formal language, formal analysis techniques can be applied to it. The perhaps most basic question one could be interested in is whether certain model states can be reached. This *reachability problem* formulation is simple, its solution is hard (undecidable for all but some very simple subclasses of hybrid automata). Nevertheless, there are different techniques to solve the reachability problem in an incomplete manner. Besides abstraction and model transformation techniques, just to mention some popular ones, a natural approach is to compute an *approximation* of the reachable states in an appropriate *representation*. For both over- and under-approximative computations, we first need to fix a data type to *represent* sets of states. State-of-the-art methods use different *geometric objects* like polytopes, zonotopes, ellipsoids etc. The choice of the geometry has a crucial effect on the practicability of the reachability computation. Once the state set representation is fixed, one way to determine the set of reachable states is to apply a forward fixedpoint-based search: Starting from the set of initial states, we

iteratively compute successor sets until a fixedpoint is reached, i.e., until we computed the whole set of reachable states. This method needs to determine successor sets under both discrete jumps and continuous evolution. The latter is often done by *flowpipe* construction, paving the whole flow by a set of geometric objects of the chosen type.

During the seminar we discussed different possibilities to apply such reachability analysis techniques to critical infrastructures. We especially focused on possible applications of our tool Flow* [1] in this context. Also interesting is the question how suitable are hybrid automata to model critical infrastructures, and what are the problems and the alternatives.

### References

1 X. Chen, E. Ábrahám and S. Sankaranarayanan. *Flow*: An analyzer for non-linear hybrid systems*. In Proc. of the 25th Int. Conf. on Computer Aided Verification (CAV'13), pages 258–263, volume 8044 of LNCS, Springer-Verlag, 2013.

2 R. Alur, C. Courcoubetis, N. Halbwachs, T. A. Henzinger, P.-H. Ho, X. Nicollin, A. Olivero, J. Sifakis and S. Yovine. *The algorithmic analysis of hybrid systems*. Theoretical Computer Science, 138(1):3–34, 1995.

## 3.6 Modeling Stochastic Hybrid Systems in Modelica: Some Results Obtained in the MODRIO Project

*Marc Bouissou (EDF, F)*



**Main reference** M. Bouissou, H. Elmqvist, M. Otter, A. Benveniste, "Efficient Monte Carlo simulation of stochastic hybrid systems," in Proc. of the Modelica Conference 2014, Linköping Electronic Conference Proceedings, Issue 96, Article No. 075, pp. 715–725, Linköping University Electronic Press, Linköpings universitet, 2014.
**URL** http://dx.doi.org/10.3384/ECP14096715

Usually, Modelica models are deterministic; they are built to simulate the nominal behavior of the systems they represent. In order to challenge the functioning of these systems in diverse situations, or in the presence of a varying environment, a degree of randomness is sometimes added to the system inputs. But the kind of models we want to be able to build in the MODRIO project are quite different: here, the random behavior can be due to the system itself, mainly because of failures (and repairs) of components. The purpose of reliability, and more generally, of dependability studies is to calculate probabilities of undesirable events such as the failure of the mission of a system, or to estimate the probability distribution of some performances of the system: total production on a given time interval, maintenance cost, number of repairs etc. The presentation showed extensions of the Modelica language that were proposed in order to facilitate the construction of such models. Some intermediary implementations of these extensions were demonstrated. The presentation was based on a joint work with other partners of the MODRIO project, which led to a remarkable result: a particularly efficient procedure to run Monte Carlo simulations of stochastic hybrid systems. This result is detailed in the reference above.

## 3.7 Tutorial: Probabilistic Model Checking

*Christel Baier (Technische Universität Dresden, DE)*



This talk gave an introduction to popular discrete-time probabilistic models and state-of-the-art model checking procedures for them. Discrete-time Markov chains (DTMCs) are purely probabilistic models, which can be extended by allowing non-determinism to discrete-time Markov decision processes (MDPs). Besides techniques for model checking $\omega$-regular properties of MDPs, further related topics like abstraction techniques and the computation of conditional probabilities were discussed in the talk.

## 3.8 Time-Dependent Analysis of Attacks

*Holger Hermanns (Universität des Saarlandes, DE)*


**Joint work of** Holger Hermanns, Florian Arnold, Reza Pulungan, and Mariëlle Stoelinga
**Main reference** F. Arnold, H. Hermanns, R. Pulungan, M. Stoelinga, "Time-dependent analysis of attacks," in Proc. of the 3rd Int'l Conf. on Principles of Security and Trust (POST'14), LNCS, Vol. 8414, pp. 285–305, Springer, 2014.
**URL** http://dx.doi.org/10.1007/978-3-642-54792-8_16

The success of a security attack crucially depends on time: the more time available to the attacker, the higher the probability of a successful attack; when given enough time, any system can be compromised. Insight in time-dependent behaviors of attacks and the evolution of the attacker's success as time progresses is therefore a key for effective countermeasures in securing systems. This paper presents an efficient technique to analyze attack times for an extension of the prominent formalism of attack trees. If each basic attack step, i.e., each leaf in an attack tree, is annotated with a probability distribution of the time needed for this step to be successful, we show how this information can be propagated to an analysis of the entire tree. In this way, we obtain the probability distribution for the entire system to be attacked successfully as time progresses. For our approach to be effective, we take great care to always work with the best possible compression of the representations of the probability distributions arising. This is achieved by an elegant calculus of acyclic phase type distributions, together with an effective compositional compression technique. We demonstrate the effectiveness of this approach on three case studies, exhibiting orders of magnitude of compression.

## 3.9 Parameter Identification and Synthesis from Qualitative Data and Behavioural Constraints

*Luca Bortolussi (University of Trieste, IT)*


**Joint work of** Luca Bortolussi, Guido Sanguinetti, Ezio Bartocci, and Laura Nenzi

In many applications, it is not always feasible to obtain quantitative measures of the process, but it is generally easier to capture qualitative properties of the dynamics. These properties

can be formalised in a suitable temporal logic, and their observations can be used to estimate parameter values, combining statistical model checking and machine learning tools in a Bayesian framework. A similar approach can be used to find a parametrisation forcing a system to satisfy robustly qualitative properties expressed in temporal logic.

### References

1 L. Bortolussi and G. Sanguinetti. *Learning and designing stochastic processes from logical constraints*. In Proc. of the 10th Int. Conf. on Quantitative Evaluation of SysTems (QEST'13), 8054:89–105, 2013.
2 E. Bartocci, L. Bortolussi, L. Nenzi and G. Sanguinetti. *On the robustness of temporal properties for stochastic models*. In Proc. of the 2nd Int. Workshop on Hybrid Systems and Biology, 125:3–19, 2013.

## 3.10 Randomized Methods for Design in the Presence of Uncertainty

*Maria Prandini (Technical University of Milan, IT)*


**Joint work of** Maria Prandini, Marco Campi, and Simone Garatti
**Main reference** M. C. Campi, S. Garatti, M. Prandini, "The scenario approach for systems and control design," Annual Reviews in Control, 33(2):149–157, 2009.
**URL** http://dx.doi.org/10.1016/j.arcontrol.2009.07.001

In this presentation, we described randomized methods to solve optimization problems in presence of uncertainty, focusing on the scenario approach to robust and chance-constrained optimization. The effectiveness and versatility of the scenario approach have been pointed out through some examples in systems and control.

## 3.11 Proving Safety of Complex Control Software: Three "Test Tube" Applications in Robotics

*Armando Tacchella (University of Genova, IT)*


**Joint work of** Armando Tacchella, Shashank Pathak, Luca Pulina, and Giorgio Metta
**Main reference** S. Pathak, L. Pulina, G. Metta, A. Tacchella, "Ensuring safety of policies learned by reinforcement: Reaching objects in the presence of obstacles with the iCub," in Proc. of the 2013 IEEE/RJS Int'l Conf. on Intelligent Robots and Systems (IROS'13), pp. 170–175, IEEE, 2013.
**URL** http://dx.doi.org/10.1109/IROS.2013.6696349

A great deal of current research is focused on making robots accomplish complex tasks in unstructured environments with increasing degrees of autonomy. Witnessing this trend, some recent contributions in the literature include perspectives on autonomy in exploration rovers [1], challenges for robot companions [2], and the impressive results obtained in the DARPA robotics challenge [3]. From a designer's point of view, autonomy can be seen as the robot's capability of adapting to unforeseen circumstances by evaluating the effects of its actions, and then taking appropriate strategic decisions. Operational scenarios where autonomy is required for robots to be effective, require rich and complex control architectures, which are usually organized in several levels, from those closest to hardware, e.g., motor control loops, to those farthest from it, e.g., object recognition, manipulation, locomotion, speech and combinations thereof. Since robots must be trustworthy, layered control architectures must be dependable. However, ensuring dependability in any complex architecture is difficult, and

it becomes an open challenge when autonomy clashes with basic requirements, e.g., safety. In this talk, we present three computer-augmented software engineering approaches to improve dependability of control architectures in autonomous robots. These approaches are targeted to different levels and different kinds of components inside the control architecture, and they rely on different formal models and techniques. However, they share the fundamental vision that formal models can be automatically (*i*) extracted, (*ii*) analyzed and (*iii*) exploited to obtain additional confidence in the properties of the control architecture. Our basic philosophy is to keep the amount of additional developer's knowledge as small as possible, while at the same time ensuring a precise analysis about whether the architecture matches its requirements. The final goal is to obtain a development environment wherein critical components in control architectures can be analyzed in a "push-button" fashion using state-of-the-art verification techniques.

### Verification of Embedded Control Software

In modern robots, powerful embedded controllers are commonly adopted to enable the implementation of sophisticated planning and control strategies – see, e.g., [4] for a discussion about this topic. The growing complexity of control strategies entails a growing complexity of embedded software which, in turn, may increase the occurrence of programming bugs that can disrupt the correct behavior of the controller. To detect these bugs before they can produce unwanted effects, we would like to apply software model checking – see [5] for a recent survey – in order to ensure that control programs cannot drive the robot to unwanted states. However, this is made challenging by the fact that the correctness of the control software relies on implicit assumptions about the system it controls, and properties are expressed in terms of the behavior of the controlled systems, not in terms of the behavior of the software itself. In [6], a methodology to enable embedded software model checking is introduced. The main idea is to apply system identification techniques to obtain a computational model of the physical system which can be checked together with the control software. We present an experience report along the lines of [6], where we consider the verification of an embedded control program in a two-wheeled self-balancing robot. The goal of the report is to highlight the current limitations of this methodology, and to propose further research to improve its feasibility and applicability.

### Middleware identification

Insofar a component of a control architecture is assigned a precise semantics, formal correctness verification is made possible. However, developing a formal model can be difficult for a "black-box" component, i.e., an overly-complex, poorly-documented, or closed-source piece of software. This can be critical when such component is located in middleware APIs used, e.g., to orchestrate uniform access to hardware resources. A viable solution to this problem is to adopt automata-based identification techniques - see, e.g., [7] for a comprehensive list of references. The key idea is that the internal structure of a black-box component can be inferred by analyzing its interactions with an embedding context. Learning algorithms supply the component with suitable input test patterns to populate a "conjecture" automaton by observing the corresponding outputs; then, they check whether the conjecture is behaviorally equivalent to the actual component. When such an abstract model of the original black-box component is obtained, it can be used as a stub to test and/or verify software components relying on it. Practical identification of different kinds of abstract models of middleware is enabled by our tool AIDE (Automata IDentification Engine) [8], an open-source software

written in C#. We sketch the design and the implementation of AIDE, we show the results of an experiment about the identification of a YARP [9] component, and we provide an example to demonstrate how the identified models can support bug-finding in control software relying on YARP.

### Safe Reinforcement Learning

Reinforcement Learning (RL) is one of the most widely adopted paradigms to obtain intelligent behavior from interactive robots – see, e.g., [10]. RL can be seen as a way to synthesize control programs when knowledge about the external environment is limited. RL methods have shown robust and efficient learning on a variety of robot-control problems – see, e.g. [11]. However, as mentioned in [12], the asymptotic nature of guarantees about the performance of RL makes it difficult to bound the probability of damaging the controlled robot and/or the environment. An interesting research question is thus how to guarantee that, given a control policy synthesized by RL, such policy will have a very low probability of yielding undesirable behaviors, e.g., damaging the robot or the environment wherein it operates. In particular, we consider Probabilistic Model Checking techniques – see, e.g., [13]. We describe the interactions between the robot and the environment using Markov chains, and the related safety properties using probabilistic logic. Both the encoding of the interaction models and their verification are fully automated, and only the properties have to be manually specified considering the project requirements. Our research goes even beyond automating verification, to consider the problem of automating repair, i.e., if the policy is found unsatisfactory, it is fixed with no manual inspection.

### References


1. M. Bajracharya, M. Maimone and D. Helmick. *Autonomy for mars rovers: Past, present, and future*. Computer, 41(12):44–50, 2008.
2. M. Beetz, U. Klank, I. Kresse, A. Maldonado, L. Mosenlechner, D. Pangercic, T. Ruhr and M. Tenorth. *Robotic roommates making pancakes*. In Proc. of the 11th IEEE-RAS Int. Conf. on Humanoid Robots (Humanoids'11), pp. 529–536, IEEE, 2011.
3. G. Pratt and J. Manzo. *The DARPA robotics challenge [competitions]*. Robotics & Automation Magazine, 20(2):10–12, IEEE, 2013.
4. C. Belta, A. Bicchi, M. Egerstedt, E. Frazzoli, E. Klavins and G. J. Pappas. *Symbolic planning and control of robot motion [grand challenges of robotics]*. Robotics & Automation Magazine, 14(1):61–70, IEEE, 2007.
5. R. Jhala and R. Majumdar. *Software model checking*. ACM Computing Surveys (CSUR), 41(4):21, 2009.
6. S. Scherer, F. Lerda and E. M. Clarke. *Model checking of robotic control systems*. In Proc. of ISAIRAS'05, pp. 5–8, 2005.
7. M. Shahbaz. *Reverse engineering enhanced state models of black box software components to support integration testing*. PhD thesis, Institut Polytechnique de Grenoble, Grenoble, France, 2008.
8. A. Khalili and A. Tacchella. *AIDE: Automata-identification engine*. http://aide.codeplex.com.
9. P. Fitzpatrick, G. Metta, and L. Natale. *Towards long-lived robot genes*. Robotics and Autonomous systems, 56(1):29–45, 2008.
10. R. S. Sutton and A. G. Barto. *Reinforcement Learning – An Introduction*. MIT Press, 1998.
11. J. A. Bagnell and S. Schaal. *Special issue on Machine Learning in Robotics (Editorial)*. The International Journal of Robotics Research, 27(2):155–156, 2008.

12 J. H. Gillula and C. J. Tomlin. *Guaranteed safe online learning via reachability: Tracking a ground target using a quadrotor*. In Proc. of ICRA'12, pp. 2723–2730, 2012.

13 M. Kwiatkowska, G. Norman, and D. Parker. *Stochastic model checking*. Formal methods for performance evaluation, pp. 220–270, 2007.

14 R. E. Kalman et al. *Contributions to the theory of optimal control*. Bol. Soc. Mat. Mexicana, 5(2):102–119, 1960.

15 P. Lancaster and L. Rodman. *Algebraic Riccati equations*. Oxford University Press, 1995.

16 *MATLAB version 8.1.0 (R2013a)*. The MathWorks Inc., Natick, Massachusetts, 2013.

17 L. Cordeiro, B. Fischer, and J. Marques-Silva. *SMT-Based bounded model checking for embedded ANSI-C software*. In Proc. of the Int. Conf. on Automated Software Engineering, pp. 137–148, 2009.

18 N. Mohamed, J. Al-Jaroodi, and I. Jawhar. *Middleware for robotics: A survey*. In 2008 IEEE Conf. on Robotics, Automation and Mechatronics, pp. 736–742, IEEE, 2008.

19 G. Metta, L. Natale, F. Nori, G. Sandini, D. Vernon, L. Fadiga, C. von Hofsten, K. Rosander, M. Lopes, J. Santos-Victor et al. *The iCub humanoid robot: An open-systems platform for research in cognitive development*. Neural Networks: The Official Journal of the International Neural Network Society, 2010.

20 M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler and A. Y. Ng. *ROS: An open-source robot operating system*. In Proc. of the ICRA workshop on Open Source Software, volume 3, 2009.

21 D. Angluin. *Learning regular sets from queries and counterexamples*. Information and computation, 75(2):87–106, 1987.

22 A. Gargantini. *Conformance testing*. Model-Based Testing of Reactive Systems, pp. 87–111, 2005.

23 O. Niese. *An integrated approach to testing complex systems*. PhD thesis, Technische Universität Dortmund, Dortmund, Germany, December 2003.

24 F. Aarts and F. Vaandrager. *Learning I/O automata*. Proc. of CONCUR'10, pp. 71–85, 2010.

25 A. Khalili and A. Tacchella. *Learning nondeterministic Mealy machines*. Technical report, University of Genoa, 2013.

26 D. C. Bentivegna, C. G. Atkeson, A. Ude and G. Cheng. *Learning to act from observation and practice*. International Journal of Humanoid Robotics, 1(4), December 2004.

27 G. Metta, L. Natale, S. Pathak, L. Pulina and A. Tacchella. *Safe and effective learning: A case study*. In Proc. of ICRA'10, pp. 4809–4814, 2010.

28 S. Pathak, L. Pulina, G. Metta and A. Tacchella. *Ensuring safety of policies learned by reinforcement: Reaching objects in the presence of obstacles with the iCub*. In Proc. of IROS'13, pp. 170–175, 2013.

29 E. Ábrahám, N. Jansen, R. Wimmer, J.-P. Katoen and B. Becker. *DTMC model checking by SCC reduction*. In Proc. of the 7th Int. Conf. on the Quantitative Evaluation of Systems (QEST'10), pp. 37–46. IEEE, 2010.

30 J.-P. Katoen, I. S. Zapreev, E. M. Hahn, H. Hermanns and D. N. Jansen. *The ins and outs of the probabilistic model checker MRMC*. Performance Evaluation, 68(2):90–104, 2011.

31 M. Kwiatkowska, G. Norman and D. Parker. *Prism: Probabilistic symbolic model checker*. Computer Performance Evaluation: Modelling Techniques and Tools, pp. 113–140, 2002.

32 L. Pulina and A. Tacchella. *An abstraction-refinement approach to verification of artificial neural networks*. In Proc. of the 22nd Int. Conf. on Computer Aided Verification (CAV'10), volume 6174 of LNCS, pp. 243–257, Springer-Verlag, 2010.

33 X. C. Ding, S. L. Smith, C. Belta and D. Rus. *MDP optimal control under temporal logic constraints*. In Proc. of the 50th IEEE Conf. on Decision and Control and European Control Conference (CDC-ECC), pp. 532—538, IEEE, 2011.

## 3.12 The Theory of Stochastic State Classes: Applications

*Laura Carnevali (University of Firenze, IT)*



Tools play a crucial role in supporting theoretical developments and in making them applicable. Oris implements the method of stochastic state classes, allowing formal design and quantitative analysis of models that include multiple non-Markovian timers with possibly bounded domain. These features fit a general class of safety-critical systems, providing support for their development and assessment. Applications of stochastic modeling and analysis through the Oris Tool are discussed referring to the evaluation of availability measures for maintenance procedures and gas distribution networks.

### References

1 L. Carnevali, M. Paolieri, F. Tarani and E. Vicario. *Quantitative evaluation of availability measures of gas distribution networks*. In Proc. of the Int. Conf. on Performance Evaluation Methodologies and Tools, IEEE CS, 2013.
2 L. Carnevali, M. Paolieri, K. Tadano and E. Vicario. *Towards the quantitative evaluation of phased maintenance procedures using non-Markovian regenerative analysis*. In Proc. of the European Workshop on Performance Engineering (EPEW'13), LNCS, pp. 176–190, Springer-Verlag, 2013.
3 S. Ballerini, L. Carnevali, M. Paolieri, K. Tadano and F. Machida. *Software rejuvenation impacts on a phased-mission system for Mars exploration*. In Proc. of the Int. Workshop on Software Aging and Rejuvenation (WoSAR'13), 2013.

## 3.13 Analysis of a Sewage Treatment Facility using Hybrid Petri Nets

*Anne Remke (University of Twente, NL)*


**Joint work of** Anne Remke, Hamed Ghasemieh, Boudewijn Haverkort, and Marco Gribaudo
**Main reference** H. Ghasemieh, A. Remke, B. R. Haverkort, "Analysis of a sewage treatment facility using hybrid Petri nets," in Proc. of the 7th Int'l Conf. on Performance Evaluation Methodologies and Tools (VALUETOOLS'13), ACM, to appear; available as pre-print.
**URL** http://eprints.eemcs.utwente.nl/24179/

Waste water treatment facilities clean sewage water from households and industry in several cleaning steps. Such facilities are dimensioned to accommodate a maximum intake. However, in the case of very bad weather conditions or failures of system components the system might not suffice to accommodate all waste water. In this talk we described the modeling of a real waste water treatment facility, situated in the city of Enschede, The Netherlands, as Hybrid Petri net with a single general one-shot transition (HPnGs) and analyses under which circumstances the existing infrastructure will overflow. This required extending the HPnG formalism with *guard arcs* and *dynamic continuous transitions* to model dependencies both on continuous places and on the rate of continuous transitions. Using recent algorithms for model checking STL properties on HPnGs, we compute survivability measures that can be expressed using the path-based until operator. After computing measures for a wide range of parameters, we provide recommendations as to where the system can be improved to reduce the probability of overflow.

### References

1 M. Gribaudo and A. Remke. *Hybrid Petri nets with general one-shot transitions for dependability evaluation of fluid critical infrastructures*. In Proc. of the IEEE 12th Int. Symposium on High Assurance Systems Engineering, IEEE CS Press, 2010, http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5634312.

2 H. Ghasemieh, A. Remke, B. R. Haverkort and M. Gribaudo. *Region-based analysis of hybrid Petri nets with a single general one-shot transition*. Formal Modeling and Analysis of Timed Systems, Springer-Verlag, 2012, http://dx.doi.org/10.1007/978-3-642-33365-1_11.

3 H. Ghasemieh, A. Remke and B. R. Haverkort. *Survivability evaluation of fluid critical infrastructures using hybrid Petri nets*. In Proc. of the 19th IEEE Pacific Rim International Symposium on Dependable Computing, 2013, http://eprints.eemcs.utwente.nl/24178/.

4 H. Ghasemieh, A. Remke and B. R. Haverkort. *Analysis of a sewage treatment facility using hybrid Petri nets*. In Proc. of the 7th Int. Conf. on Performance Evaluation Methodologies and Tools, 2013, http://eprints.eemcs.utwente.nl/241.


## 3.14 Resilience of Data Networking and Future Power Networks

*Hermann de Meer (Universität Passau, DE)*


**Main reference** http://resumenet.eu/

The intelligent power grid ("Smart Grid") will replace our current rigid and hierarchical power grid in the near future. The Smart Grid is realized by a strong entanglement of the power grid and modern communication infrastructures. The arising challenges in this field cover two opposing directions, namely the energy efficiency as well as the security and safety of the Smart Grid infrastructure.

The ResumeNet and HyRiM projects investigate ways to protect both the network part as well as the utility network infrastructures. To achieve this, system-wide approaches are developed that take into account the increased complexity of the Smart Grid as well as the diverse origins of possible failures, such as random or intentional faults or human errors at the operational as well as strategic corporate level.


### References

1 A. Berl, A. Fischer and H. de Meer. *Virtualisierung im Future Internet – Virtualisierungmethoden und Anwendungen*. Informatik-Spektrum, 33(2):186–194, 2010.

2 A. Fischer, A. Fessi, G. Carle and H. de Meer. *Wide-Area virtual machine migration as resilience mechanism*. In Proc. of the Int. Workshop on Network Resilience: From Research to Practice (WNR'11), pp. 72–77, IEEE, 2011.

3 A. Fischer, J. F. Botero, M. Duelli, D. Schlosser, X. Hesselbach and H. de Meer. *ALEVIN – A framework to develop, compare, and analyze virtual network embedding algorithms*. Electronic Communications of the EASST, 37:1–12, 2011.

4 J. P. G. Sterbenz, D. Hutchison, E. G. Cetinkaya, A. Jabbar, J. P. Rohrer, M. Schöller and P. Smith. *Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines*. Computer Networks, Special Issue on Resilient and Survivable Networks, 54(8):1245–1265, 2010.

5 P. Smith, D. Hutchison, J. P. G. Sterbenz, M. Schöller, A. Fessi, M. Karaliopoulos, C. Lac and B. Plattner. *Network resilience: A systematic approach*. IEEE Communications Magazine, 49(7):88–97, 2011.

## 3.15 Issues in Modelling Smart Grid Infrastructures to Assess Resilience-Related Indicators

*Felicita Di Giandomenico (CNR – Pisa, IT)*


**Joint work of** Felicita Di Giandomenico, Silvano Chiaradonna, and Nadir Murru
**Main reference** S. Chiaradonna, F. Di Giandomenico, N. Murru, "On a modeling approach to analyze resilience of a smart grid infrastructure," in Proc. of the 10th European Dependable Computing Conference (EDCC'14), to appear.

The evolution of electrical grids, both in terms of enhanced ICT functionalities to improve efficiency, reliability and economics, as well as the increasing penetration of renewable distributed energy resources to favor sustainability of the production and distribution of electricity, results in a more sophisticated electrical infrastructure which poses new challenges from several perspectives, including resilience and quality of service analysis. In addition, the presence of interdependencies, which more and more characterize critical infrastructures (including the power sector), exacerbates the need for advanced analysis approaches, to be possibly employed since the early phases of the system design, to identify vulnerabilities and appropriate countermeasures. In this presentation, we outline an approach to model and analyze smart grids and discuss the major challenges to be addressed in stochastic model-based analysis to account for the peculiarities of the involved system elements. Representation of dynamic and flexible behavior of generators and loads, as well as representation of the complex ICT control functions required to preserve and/or re-establish electrical equilibrium in presence of changes (both nominal ones, such as variable production by a photovoltaic energy source, and failures/disruptions both at electrical and ICT level) need to be faced to assess suitable indicators of the resilience and quality of service of the smart grid.

### References

1 S. Chiaradonna, F. Di Giandomenico and P. Lollini. *Definition, implementation and application of a model-based framework for analyzing interdependencies in electric power systems*. Int. Journal of Critical Infrastructure Protection, 4(1):24–40, 2011.
2 S. Chiaradonna, F. Di Giandomenico and N. Nostro. *Modeling and analysis of the impact of failures in electric power systems organized in interconnected regions*. In Proc. of the 41st Int. Conf. on Dependable Systems & Networks (DSN'11), pp. 442–453, IEEE Computer Society Press.
3 S. Chiaradonna, F. Di Giandomenico and N. Nostro. *Analysis of electric power systems accounting for interdependencies in heterogeneous scenarios*. In Proc. of EDCC'12, pp. 84–93, 2012.

## 3.16 Energy-Autonomous Smart Micro-Grids

*Gerard Smit (University of Twente, NL)*


**Joint work of** Gerard Smit and Johann Hurink; also supported by Alliander (Bram Reinders) and RWE (Stefan Nykamp)

When enough (renewable) generation like PV panels, biomass installations and wind-turbines in combination with storage assets are installed, it may be possible to create a self-supplying (autonomous) neighbourhood in a so-called energy autonomous smart micro-grid. The main objective of our work is: to develop methods and techniques to support the development of

energy-autonomous smart micro-grids. This broad main objective can be decomposed in a number of detailed research questions:

- In an energy-autonomous smart micro-grid demand/supply matching (DSM) has to be done on a local level. How to find local balance of demand/supply/storage. A related research question is: how (and for how long) can a micro-grid continue autonomously without a connection to the main electricity grid?
- What distributed energy management systems can be used for a local micro-grid and a cluster of micro-grids (systems of systems) attached to the smart grid.
- Find and use the flexibility of appliances in a micro-grid e.g. storage, charging time of EV, starting time of dishwashers.
- What kind of (wireless) communication networks will support reliable, real-time and efficient communication in a micro-grid?

### References

1 S. Nykamp, M. G. C. Bosman, A. Molderink, J. L. Hurink and G. J. M. Smit. *Value of storage in distribution grids-competition or cooperation of stakeholders?* IEEE Transactions on Smart Grid, 4 (3). pp. 1361–1370, 2013.

2 S. Nykamp, A. Molderink, J. L. Hurink and G. J. M. Smit. *Statistics for PV, wind and biomass generators and their impact on distribution grid planning*. Energy, 45(1):924–932, 2013.

## 3.17 Cyber-Security of SCADA Systems: A Case Study on Automatic Generation Control

*John Lygeros (ETH Zürich, CH)*



**Main reference** P. Mohajerin Esfahani, M. Vrakopoulou, G. Andersson, J. Lygeros, "A tractable nonlinear fault detection and isolation technique with application to the cyber-physical security of power systems," in Proc. of the 2012 IEEE 51st Annual Conf. on Decision and Control, pp. 3433–3438, IEEE, 2012.

**URL** http://dx.doi.org/10.1109/CDC.2012.6426269

Cyber-security issues in SCADA systems have concentrated considerable attention, due in part to highly publicized security threats such as the STUXNET computer worm. The research presented in this talk is motivated by security issues for SCADA systems used to monitor and control the power transmission grid. We specifically concentrate on the implications and possible countermeasures of attacks on the Automatic Generation Control (AGC) system, one of the few control loops closed over such SCADA systems without the intervention of human operators. We show how an attacker who gains access to the AGC signal of the SCADA system in one control area can robustly destabilize the transmission system. We then proceed to design countermeasures against such attacks. To this end, we develop a novel fault detection/isolation filter applicable to high dimensional nonlinear systems, based on randomized optimization methods.

### References

1 P. Mohajerin Esfahani, M. Vrakopoulou, J. Lygeros and G. Andersson. *Intrusion detection in electric power networks*. Patent Application EP 2690511 (January 29, 2014), WO 2014/015970 (January 30, 2014).

2 E. Tiniou, P. Mohajerin Esfahani and J. Lygeros. *Fault detection with discrete-time measurements: An application for the cyber security of power networks*. In Proc. of the IEEE Conf. on Decision and Control, 2013.
3 G. Andersson, P. Mohajerin Esfahani, M. Vrakopoulou, K. Margellos, J. Lygeros, A. Teixeira, G. Dan, H. Sandberg and K. Johansson. *Cyber-secirity of SCADA systems*. In Proc. of Innovative Smart Grid Technologies (ISGT IEEE PES), 2012.
4 P. Mohajerin Esfahani, M. Vrakopoulou, G. Andersson and J. Lygeros. *A tractable nonlinear fault detection and isolation technique with application to the cyber-physical security of power systems*. In Proc. of the IEEE Conf. on Decision and Control, 2012.
5 P. Mohajerin Esfahani, M. Vrakopoulou, K. Margellos, J. Lygeros and G. Andersson. *A robust policy for automatic generation control cyber attack in two area power network*. In Proc. of the IEEE Conference on Decision and Control, 2010.
6 P. Mohajerin Esfahani, M. Vrakopoulou, K. Margellos, J. Lygeros and G. Andersson. *Cyber attack in a two-area power system: Impact identification using reachability*. In Proc. of the American Control Conference, 2010.

## 3.18 Towards Quantitative Modeling of Reliability for Critical Infrastructure Systems

*Sahra Sedigh Sarvestani (University of Missouri – Rolla, US)*


**Joint work of** Sahra Sedigh Sarvestani, Ali Hurson, Bruce McMillin, and Ann Miller

Critical infrastructure systems are increasingly reliant on cyber infrastructure that enables intelligent real-time control of physical components. This cyber infrastructure utilizes environmental and operational data to provide decision support intended to increase the efficacy and reliability of the system and facilitate mitigation of failure. Realistic imperfections, such as corrupt sensor data, software errors, or failed communication links can cause failure in a functional physical infrastructure, defying the purpose of intelligent control. As such, justifiable reliance on cyber-physical critical infrastructure is contingent on rigorous investigation of the effect of intelligent control, including modeling and simulation of failure propagation within the cyber-physical infrastructure. We present and invite discussion on challenges in and recent advances towards development of quantitative models and accurate simulation methods for cyber-physical critical infrastructure systems, with focus on smart grids and intelligent water distribution networks.

### References

1 J. Lin and S. Sedigh. *Reliability modeling for intelligent water distribution networks*. Int. Journal of Performability Engineering, Special issue on Performance and Dependability Modeling of Dynamic Systems, 7(5):467–478, 2011.
2 J. Lin, S. Sedigh and A. R. Hurson. *Ontologies and decision support for failure mitigation in intelligent water distribution networks*. In Proc. of the 45th Hawaii Int. Conf. on System Sciences (HICSS-45), 2012.
3 J. Lin, S. Sedigh and A. R. Hurson. *Knowledge management for fault-tolerant water distribution*. In Proc. of Large Scale Network-Centric Computing Systems, John Wiley & Sons, 2012.
4 A. Faza, S. Sedigh and B. McMillin. *Reliability analysis for the advanced electric power grid: From cyber control and communication to physical manifestations of failure*. In Proc. of

the 28th Int. Conf. on Computer Safety, Reliability and Security (SAFECOMP'09), winner of best paper award, 2009.

5 A. Faza, S. Sedigh and B. McMillin *Integrated cyber-physical fault injection for reliability analysis of the smart grid*. In Proc. of the 29th Int. Conf. on Computer Safety, Reliability and Security (SAFECOMP'10), pp. 277–290, 2010.

6 K. Marashi, M. Woodard, S. Sedigh and A. Hurson. *Quantitative reliability analysis for intelligent water distribution networks*. In Proc. of the Embedded Topical Meeting on Risk Management for Complex Socio-Technical Systems (RM4CSS), Annual Meeting of the American Nuclear Society, 2013.

7 M. Woodard and S. Sedigh. *Modeling of autonomous vehicle operation in intelligent transportation systems*. In Software Engineering for Resilient Systems, volume 8166 of LNCS, pp. 133–140, Springer-Verlag, 2013.

## 3.19 Design Challenges for Systems of Systems

*Boudewijn Haverkort (University of Twente, NL)*



**Main reference** B. R. Haverkort, "The dependable systems-of-systems design challenge," IEEE Security & Privacy, 11(5):62–65, 2013.
**URL** http://dx.doi.org/10.1109/MSP.2013.124

Over the last few years there has been an increased interest in so-called systems-of-systems. In the control and management of infrastructural systems, systems-of-systems are widespread. However, the size of these systems and their management challenges make it a formidable task to really design them such that performance and dependability properties can be guaranteed. In this talk I addressed the background of systems-of-systems, and discussed the challenges associated with their design, especially in light of model-driven design approaches.

## 3.20 Multiformalism to Support Software Rejuvenation Modeling

*Marco Gribaudo (Politecnico di Milano, IT)*



**Joint work of** Marco Gribaudo, Mauro Iacono, and Enrico Barbierato

The study of software aging and rejuvenation is based on models that conjugate the complexity of architectural models with the problem of time dependence of parameters. Exploiting the metaphors of common performance-oriented modeling formalisms (such as Petri nets or queuing networks) with the support of proper solution techniques can help mod- elers in approaching the analysis of complex software-based systems. In this talk we showed how SIMTHESys (a multiformalism modeling framework) can be used to approach the modeling problem by implementing a new user-defined modeling formalisms and the related fluid-based solution engine.

## 3.21 Quantitative Evaluation of Smart Grid Control Traffic

*Katinka Wolter (FU Berlin, DE)*


**Joint work of** Katinka Wolter, Tilman Krauss, and Manfred Hartmann

The expected decentralised nature of the Smart Grid on the producer as well as on the consumer side requires a large amount of control in order to match supply and demand in an optimal way. Very likely the smart grid control traffic will not use dedicated communication lines but it will be transmitted using various communication channels, such as wireless or cellular networks or the public Internet. In consequence, Smart Grid control traffic will suffer from all kinds of disturbances and reliable transmission must be guaranteed using different kinds of redundancy mechanisms.

In this talk I presented stochastic models for traffic flow that were developed in collaboration with Bell Labs Berlin and show the insights we gained from varying the network topology, configuration parameters as well as the background load.

## 3.22 Zero-Defect Cyber-Physical Systems in Space: A True Mission

*Joost-Pieter Katoen (RWTH Aachen, DE)*


**Joint work of** Joost-Pieter Katoen, Marco Bozzano, Alessandro Cimatti, Marie-Claude Esteve, Viet Yen Nguyen, Thomas Noll, Marco Roveri, and Yuri Yushstein
**Main reference** M. Bozzano, A. Cimatti, J.-P. Katoen, V. Y. Nguyen, T. Noll, M. Roveri, "Safety, Dependability and Performance Analysis of Extended AADL Models," The Computer Journal, 54(5):754–775, 2011.
**URL** http://dx.doi.org/10.1093/comjnl/bxq024
**URL** http://compass.informatik.rwth-aachen.de

Building modern aerospace systems is highly demanding. They should be extremely dependable. They must offer service without interruption (i.e., without failure) for a very long time – typically years or decades. Whereas "five nines" dependability, i.e., a 99.999 % availability, is satisfactory for most safety-critical systems, for on-board systems it is not. Faults are costly and may severly damage reputations. Dramatic examples are known. Fatal defects in the control software of the Ariane-5 rocket and the Mars Pathfinder have led to headlines in newspapers all over the world. Rigorous design support and analysis techniques are called for. Bugs must be found as early as possible in the design process while performance and reliability guarantees need to be checked whenever possible. The effect of fault diagnosis, isolation and recovery must be quantifiable. Tailored effective techniques exist for specific system-level aspects. Peer reviewing and extensive testing find most of the software bugs, performance is checked using queueing networks or simulation, and hardware safety levels are analysed using a profiled Failure Modes and Effects Analysis (FMEA) approach. Fine. But how is the consistency between the analysis results ensured? What is the relevance of a zero-bug confirmation if its analysis is based on a system view that ignores critical performance bottlenecks? There is a clear need for an integrated, coherent approach! This is easier said than done: the inherent heterogeneous character of on-board systems involving software, sensors, actuators, hydraulics, electrical components, etc., each with its own specific development approach, severely complicates this. About three years ago we took up this grand challenge. Within the ESA- funded COMPASS (COrrectness, Modeling

and Performance of Aerospace SyStems) project, an overarching model-based approach has been developed. The key is to model on-board systems at an adequate level of abstraction using a general-purpose modeling and specification formalism based on AADL (Architecture Analysis & Design Language) as standardised by SAE International. This enables engineers to use an industry-standard, textual and graphical notation with precise semantics to model system designs, including both hardware as well as software components. Ambiguities about the meaning of designs are abandoned. System aspects that can be modeled are, amongst others,

- (timed) hardware operations, specified on the level of processors, buses, etc.,
- software operations, supporting concepts such as processes and threads,
- hybrid aspects, i.e., continuous, real-valued variables with (linear) time- dependent dynamics, and
- faults with probabilistic failure rates and their propagation between components.

A complete system specification describes three parts: (1) nominal behavior, (2) error behavior, and (3) a fault injection—how does the error behavior influence the system's nominal behavior? Systems are described in a component-based manner such that the structure of system models strongly resembles the real system's structure. This coherent and multi-disciplinary modeling approach is complemented by a rich palette of analysis techniques. The richness of the AADL dialect gives the power to specify and generate a single system model that can be analysed for multiple qualities: reliability, availability, safety, performance, and their mixture. All analysis outcomes are related to the same system's perspective, thus ensuring compatibility. First and foremost, mathematical techniques are used to enable an early integration of bug hunting in the design process. This reduces the time that is typically spent on a posteriori testing – in on-board systems, more time and effort is spent on verification than on construction! – and allows for early adaptations of the design. The true power of the applied techniques is their almost full automation: once a model and a property (e.g., can a system ever reach a state in which the system cannot progress?) are given, running the analysis is push-button technology. In case the property is violated, diagnostic feedback is provided in terms of a counterexample which is helpful to find the cause of the property refutation. These model-checking techniques are based on a full state space exploration, and detect all kinds of bugs, in particular also those that are due to the intricacies of concurrency: multiple threads acting on shared data structures. This type of bugs are becoming increasingly frequent, as multi-threading grows at a staggering rate. Analysing system safety and dependability is supported by key techniques such as (dynamic) fault tree analysis (FTA), (dynamic) Failure Modes and Effects Analysis (FMEA), fault tolerance evaluation, and criticality analysis. System models can include a formal description of both the fault detection and isolation subsystems, and the recovery actions to be taken. Based on these models, tool facilities are provided to analyze the operational effectiveness of the FDIR (Fault Detection, Isolation and Recovery) measures, and to assess whether the observability of system parameters is sufficient to make failure situations diagnosable. All techniques and the full modeling approach are supported by the COMPASS toolset, developed in close cooperation with the Italian research institute Fondazione Bruno Kessler in Trento, and is freely downloadable for all ESA countries from the website compass.informatik.rwth-aachen.de. The tool is graphical, runs under Linux, and has an easy-to-use GUI. Industrial case studies, carried out by key players in the aerospace industry, have shown the maturity of the approach and tool-set. An in-house case study at the ESA of modelling and analysing a modern satellite has been published at ICSE 2012 and comprises of the analysis of state spaces of hundreds of millions of states. Current research

focuses on compositional verification – how can we exploit the component-based structure of AADL models effectively in the verification process – and on applying the techniques to launchers. One of the main issues in that application domain is the wide range of timing granularity that is needed.

## 3.23 Smart Railroad Maintenance Engineering with Stochastic Model Checking

*Dennis Guck (University of Twente, NL)*


**Joint work of** Dennis Guck, Joost-Pieter Katoen, Mariëlle Stoelinga, Ted Luiten, and Judi Romijn
**Main reference** D. Guck, J.-P. Katoen, M. I. A. Stoelinga, T. Luiten, J. Romijn, "Smart railroad maintenance engineering with stochastic model checking," in Proc. of the 2nd Int'l Conf. on Railway Technology: Research, Development and Maintenance (Railways'14), Saxe-Coburg Publications, to appear.

RAMS (Reliability, Availability, Maintenance, Safety) requirements are utmost important for safety-critical systems like railroad infrastructure and signalling systems, and often imposed by law or other government regulations. Fault tree analysis (FTA, for short) is a widely applied industry standard for RAMS analysis, and is often one of the techniques preferred by railways organisations. FTA yields system availability and reliability, and can be used for critical path analysis. It can however not yet deal with a pressing aspect of railroad engineering: maintenance. While railroad infrastructure providers are focusing more and more on managing cost/performance ratios, RAMS can be considered as the performance specification, and maintenance the main cost driver. Methods facilitating the management of this ratio are still very uncommon. Therefore we present a flexible and transparent technique to incorporate maintenance aspects in fault tree analysis, based on stochastic model checking.

## 3.24 Cascading Events in Probabilistic Dynamical Networks

*Alessandro Abate (University of Oxford, GB)*


**Joint work of** Alessandro Abate, Ilya Tkachev and Pepijn Cox

The assessment of cascading events over probabilistic dynamical networks can be of interest in applications dealing with energy grids, computer networks, and banking systems. Small, abrupt events may lead to global cascades over such networks: the objective of this ongoing work is to propose a framework to characterise, assess, and possibly control such propagating events.

In this talk, the occurrence of contagious bankruptcies over an interconnected banking system was studied by means of randomised approaches. We also investigated the related sensitivity of networks dynamics and topologies.

## 3.25 Analysis of Complex Socio-Cyber-Physical Systems

*Martin Fränzle (Universität Oldenburg, DE)*


**Joint work of** Martin Fränzle, Stefan Puch and Bertram Wortelen

Cyber-physical systems are all about interaction; hence, getting interaction straight – at all aggregation levels and over a diverse range of time frames – is the real challenge. Interaction in cyber-physical systems inherently is heterogeneous, involving local or networked control loops, service compositions, cooperation protocols, but also humans in the loop. This forces us to accept and seamlessly integrate a diversity of models during system design and analysis. Some of these models are well-established in engineering and computer science, others have to be imported from other disciplines. The former include automata, ODE, Markovian stochastic processes of various flavors, as well as their various combinations into forms of hybrid systems. We have thus made quite some mileage on our way to the necessary model integration, but the selection and seamless integration of suitable models of human behavior still remains largely unexplored. Candidate models are supplied by other disciplines, especially cognitive psychology, but wait to be integrated with engineering models of the environment to faithfully reflect human behavior in feedback loops. Reasoning about heterogeneous models incorporating components modelling humans provides a challenge, in particular given the inherent epistemological limits to their validity, but also the extreme sparsity of fatal events in human-controlled systems. In the seminar talk at Dagstuhl, we have demonstrated this socio-technical perspective on cyber-physical systems by means of the example of advanced driver assistance systems. Setting up a model-based design and analysis chain for such systems hinges, first, on selecting models for human behaviour which would seamlessly integrate with hybrid models of the environment and, second, on devising appropriate analysis tools. For addressing the first issue, we exposed a cognitive architecture that interfaces nicely to engineering models in the style of hybrid systems by the fact that it internally is heterogeneous too, with the interfacing layers of perception and action being stochastic hybrid models, while internal layers of cognitive and associative capabilities are linked to these through a control-dominated, autonomous layer. For the other challenge, namely safety analysis, we argued that exhaustive methods in the vein of model checking are currently out of scope due to the extreme heterogeneity of the models, rendering co-simulation the only reasonable analysis technique available at the moment. Unfortunately, the safety targets are so high and the fault-masking capabilities of humans – be it real ones or adequate cognitive models – so thorough that statistical model checking by straightforward randomized co- simulation is bound to fall short when trying to substantiate the safety case, as even a single hazardous situation is extremely unlikely to show up in weeks of simulation time. Importance sampling is no cure either, unless an adequate proposal distribution is uncovered automatically, as the randomized decisions in the cognitive model are so fine-granular (they tend to decide between conflicting goals like keeping an eye on traffic for the next 20ms or initiating moving attention to checking the odometer) that manually devising a proposal distribution which is likely to yield a statistically relevant number of hazardous situations is infeasible. We showed that a criticality-driven variant of reinforcement learning can nevertheless be used as a guiding mechanism able to adjust the individual probabilities of the plethora of small randomized decisions along a reasonably long (1.2km) driving scenario, thus automatically uncovering a useful proposal distribution.

## 3.26 Optimal Counterexamples for Markov Models

*Ralf Wimmer (Universität Freiburg, DE)*


**Joint work of** Ralf Wimmer, Nils Jansen, Erika Ábrahám, Joost-Pieter Katoen, and Bernd Becker

Discrete-time Markov chains and Markov decision processes are not only commonly used for modeling discrete-time systems, but also as abstractions, e.g., of probabilistic hybrid systems after discretization. Counterexamples for violated system properties in general are not only helpful for the reproduction of errors during debugging, but can also be used for automatic refinement of abstractions of large systems. Counterexamples for Markov models can be defined at different levels: (a) on the level of system executions, i.e., a counterexample is a set of paths through the system whose joint probability mass exceeds a given upper bound, (b) the the level of the state space; here, a counterexample is a minimal subset of the states such that the probability to reach, e.g., a safety-critical state just visiting the chosen states is beyond the given bound, and (c) at the level of the modeling language. Then a counterexample is a minimal set of commands which together already induce an erroneous system. In this talk I gave an overview on these different kinds of counterexamples and present methods for their computation.

## 3.27 Quantitative Multi-Objective Verification for Probabilistic Systems

*Gethin Norman (University of Glasgow, GB)*


**Joint work of** Gethin Norman, Vojtech Forejt, Marta Kwiatkowska, David Parker, and Hongyang Qu
**Main reference** M. Kwiatkowska, G. Norman, D. Parker, H. Qu, "Compositional probabilistic verification through multi-objective model checking," Information and Computation, 232:38–65, 2013.
**URL** http://dx.doi.org/10.1016/j.ic.2013.10.001

In the first half of the talk I presented a method for analysing multiple quantitative objectives of systems that exhibit both nondeterministic and stochastic behaviour. These systems are modelled as Markov decision processes, enriched with reward structures that capture, for example, energy usage or performance metrics. The quantitative properties considered incorporate probabilistic safety and liveness properties and expected total rewards. In the second half of the talk, I showed how this approach can be applied to controller synthesis and its relevance to this seminar.

# 4 Panel Discussions

## 4.1 Open Problems

In order to prepare the panel discussion, our panelists, both from industry and from academia, had the opportunity to share their view on open problems in smart grids and critical infrastructures with the seminar group. The following open issues and needs were identified by the panelists:

### Peter Langendörfer

Information about the individual parts of the system is required, however, this requires a scalable modeling and analysis approach. Hence, how holistic can such a model be? A model would be needed that tells us about security and redundancy. This is currently not available. Moreover, interaction between different metrics would be good.

### Albert Molderink

On-line decisionmaking is highly relevant for smart grids. Real-time analysis in a dynamic model and model-based run-time decision making would be very beneficial to balance and control the power grid.

### Gerard Smit

On-line optimazation is very important for smart grids. The hierarchy that is present in smart grids can be used to divide and conquer, when modeling and analysing smart grids. Next to the analysis also the synthesis of smart grid models is very important.

### Daniel Sadoc Menasche

Models for smart grids and critical infrastructures must be hybrid and holistic. At the same time they should be easy and quick to evaluate. Also model validation would be important.

### Bil Sanders

Industry does not care how complex your approach is, you should try and challenge problems with various tools. I see three challenges:

- impact of failures and attacks on physical side of the system
- relationship between cyber and physical
- model compositional techniques

### Lucia Happe

I see challenges for metrics and tools, as well as for compositional models that ensure scalability. Mechanisms are required that compose meta models, i.e., language that is used to describe the models. Furthermore, the following items are required:

- a flexible notion of abstraction,
- decision making during design-time,
- immediate feedback at design-time.

## 4.2 From Research to Application: Open Problems, Needs and Wishes

***Panel discussion lead by Boudewijn Haverkort***

The goal of this panel discussion, which took place on Monday afternoon, was to identify interesting topics for discussion and exchange for the forthcoming days, and for possible future research and cooperations. Some of the identified areas were intensively discussed in smaller groups in the break-out sessions of the next days.

First, representatives from both academia and industry were interviewed. The questions focussed on the current research interests and directions on the one side, and the needs and problems for applications in the industry on the other side.

After the interviews a common discussion took place to identify interests and promising topics for discussions during the rest of the week. The seminar participants first identified larger topic areas together. After that, each participant could suggest special topics of interests, attached to one of the larger areas. We collected all the ideas on the blackboard.

This process resulted in a long list of interesting and relevant topic suggestions. Though we did not have the time to discuss them all, we give here a complete list. For most of the ideas we also received some additional explanations, which we list below.

We thank all participants for the fruitful discussion!

1. Understanding cyber ~ physical
   a. *Dennis: Modeling failure and security behavior with dynamic fault trees and attack trees*
      Fault trees (FT) are a wide-spred and preferred model in industry for reliability, availability, maintenance and security (RAMS) analysis. With dynamic FTs, new dynamic behaviour is added and further, it is possible to add attack scenarios into the tree. This model gives a (visual) description of the failure and security behaviour of the real system and can be transformed e.g. into a Markov automata to be analysed.
   b. *Link attacks and failures to physical processes*
   c. *Zbigniew: security metrics derived based on combination of heterogeneous evidence on system security throughout the system lifetime*
   d. *Martin: integration of cognitive models into hybrid system models*
      Interaction in cyber-physical systems is inherently heterogeneous, involving local or networked control loops, service compositions, cooperation protocols, but also humans in the loop. This forces us to accept and seamlessly integrate a diversity of models during system design and analysis. Some of these models are well-established in engineering and computer science, others have to be imported from other disciplines. The former include automata, ODE, Markovian stochastic processes of various flavors, as well as their various combinations into forms of hybrid systems. We have thus made quite some mileage on our way to the necessary model integration, but the selection and seamless integration of suitable models of human behavior still remains largely unexplored, as is the investigation of the inherent epistemological limits to the validity of the resulting models.
   e. *Martin: co-simulation of cognitive models technical systems (as of now: cars and driver assistance systems)*
      Cyber-physical systems are increasingly socio-technical: The heart of the CPS vision is having remote physical processes within the sphere of control of hand-held, wearable, or even in-body devices, which changes the way we interact with the physical environment. Reasoning about heterogeneous models incorporating both large-scale, geographically distributed, etc. technical systems and humans in the loop provides a challenge.

Co-simulation probably is the most direct line of attack towards model-based analysis of socio-technical CPS.

f. *Sarah: holistic models that help in prediting the fact of cyber-failures on tangible physical operations*
   This may be best illustrated through an example. I would like to know the likelihood of the occurrence of a cascading failure in a power grid, assuming that a given software error occurs in the software that calculates the power flow distribution.
g. *Bill: security argument graph technology*
h. *Bill: ADVISE cyber-human-physical modeling language*
i. *Hermann: resilience model and methods architecture*
j. *Felicita: analysis of the impact of failures affecting ICT-control of grid infrastructure and vice versa in presence of simultaneous failures affecting both contro land electrical grid*
   The issue here is that interdependencies existing between the ICT control infrastructure ad the controlled infrastructure (in the studies we addressed, this last is the Electrical Infrastructure) become formidable vehicles through which failures that may affect either of the two subsystems propagate to the other, increasing the entity of the resulting damage. It is therefore very important to understand and analyze the presence of such interdependencies and assess the impact of failures in presence of such interdependencies, in terms of indicators of interest to final users, as well as distribution system operators (such as black-out related indicators). This has been raised as a hot topic in the protection of critical infrastructures, and several projects and initiatives (local to specific countries but also as International efforts) have been originated to tackle it in the last decade, also triggered by major blackouts that have been experienced in Europe, US and Asia, which have affected a large part of the population (several tenths of millions of people).
k. *Enrico: diagnosis of the current state based on partial observations in a distributed/stochastic (and possibly non-deterministic) environment*
l. *Gerard: UNITY tool for modeling and simuation of discrete-event + simulation of continuous time*
m. *Rom: modeling physical entities environment and cyber entities using networks of timed automata*
   In modeling cyber physical systems, it would be nice to have some generic guidelines and support for modeling both physical and cyber networks, together with their intewractions, in the framework of timed automata.
n. *Jeremy: capture human user behavior and impact on large-scale systems with inhomogeneous stochastic models*

2. Model composition, scalability, hierarchy
   a. *Ralf: handling large discrete state spaces*
      At the university of Freiburg, Germany, we have done research on symbolic methods for discrete-time Markov models (i.e., DTMCs and MDPs): Using (MT)BDD-based methods we developed algorithms for state space minimization, counterexample generation, and computation of long-run expected rewards. It turned out that dedicated symbolic algorithms can often handle systems that are far out of reach for explicit representations.
   b. *Lucia: model join approach*
   c. *Markus: methods and tools for constructing structural Markovian models, stochastic process algebra, exponential delays+immediate actions, BDD-based scales quite well*

Methods and tools are available for constructing and analyzing structured Markovian models, based on formalisms such as stochastic process algebra, stochastic Petri nets, etc.. If a model uses only exponential and immediate delays, numerical analysis techniques are available. In case of general distributions, statistical model checking (discrete event simulation) is an option. Some tools (such as PRISM, SMART, CASPA) exploit the power of decision diagrams for compactly encoding the underlying transition systems and state sets, thereby making it possible to analyze also highly scaled model instances, which is of great importance when modelling critical infrastructures.

d. *Sarah: system-level performance and performability analysis based on component-level data*
   I would like to know how to compose a model of a large cyber-physical system from information about its components. We almost never have independence, so we cannot do what we usually do and multiply probabilities.
e. *Jeremy: large-scale composition of multiple component classes*
f. *Bill: Rep-Join model composition*
g. *Mauro: multi-formalism modeling*
   Since the elements that have to be considered and evaluated in the field are diverse and provide peculiar contribution, based onto the effects of parts of the system that substand phenomena of different nature, it is not straightforward nor probably profitable to try and force all the elements in a single representation language. Multi-formalism modeling allows the coexistence and coordination, in a single model, of different modeling formalisms, each of which can be the most natural (and familir to the domain experts) for a given subsystem, provided that proper inter-formalism semantics is defined. This can lower the learning curve, keep the efficiency of domain modelers, ensure the absence of mismatches or lack of synchronization between the representations of submodels and pave the way for custom high-level representation, where needed, to abstract the general supervision layer from the details of submodels.
h. *Mauro: model composition*
   The possibility of having proper mechanisms that enable the use of submodels in a model allows the separation of responsibilities between modelers that are experts of different subsystems/domains, the partial/parametrical/complete reuse of existing submodels, a structured management of large models and, in some cases, support for more efficient solutions.
i. *Gethin: computational and abstraction techniques for probabilistic and real-time systems*
   The corresponds to work I have done concerning both compositional verification and abstraction refinement techniniques for quantitative models. Both of which will be very useful tools for the analysis of large complex systems.

3. Tools/management of meta-model composition
   a. *Armando: engineering formal methods for applications to cyber-physical systems*
      In the past two decades, the research communities in Formal Methods and Computer Aided Verification have been extending the traditional deterministic, discrete-state and discrete-time techniques to systems exhibiting random and hybrid behavior. Thanks to powerful computation hardware and effective algorithms, the trade-off between model expressiveness and computational costs is shifting towards making more complex systems amenable to formal analysis. However, also CPS complexity kept increasing at a steady pace, so that there is still a gap to be filled in order to make theoretical contributions usable in the practice of CPS engineering. This may involve the development of domain-specific heuristics and abstractions, as well as improving the

usability of formal techniques through automation of domain and property encoding.

b. *Boudewijn: you have to adhere to industry tooling to have impact*
Scientist can make various tools, but only if they are integrated in tool-chains that are in actual use in industry, they will be used in practice. Academics should not expect industrial people to just use their tools. Moreover, industrial parties will only adhere to tools that have full 24x7 support. Scientific tools do usually not provide that!

c. *Mauro: meta-modeling exploitation*
The use of meta-modeling allows the definition of a substanding general framework, that is not dependent on the specific modeling formalism of a given submodel, on which inter-submodel or inter-formalism interactions can be founded. By using metamodeling, it is possible to manipulate the description of modeling formalisms, besides submodels, and to easily implement general model transformations, reductions, translations that are not bound to a certain model instance and use general rules, that can be extended in the future to not-yet-existing modeling formalisms as well.

4. Flexible abstraction
   a. *Daniel: TANGRAM-II modeling tool*
   b. *Gerard: how to find flexibility in SG*
   c. *Rom: flexible abstraction in timed automata modeling*
   A timed automata model of a cyber physical system may be too large to be analysed. Therefore it is necessary to use abstraction techniques, that should be easily adapted to the type of analyses that one is interested in.
5. Immediate feedback at design time
   a. *Jeremy: rapid analysis using fluid techniques*
   b. *Marco: advanced graphical user interfaces*
   c. *Laura: model-driven development*
   Formal methods can definitely contribute to increase the quality of software components by supporting multiple activities along the development life cycle. Specifically, formal modeling provides a well-defined semantics, which enables rigorous analysis through comprehensive exploration of system behaviors and supports derivation of a proof of correctness of software design. As a relevant point, early assessment of requirements allows immediate feedback at design time, which may have an impact on the quality and the cost of the final product.
   d. *Enrico: timing analysis of models beyond the limits of the Markov assumption*
   e. *Boudewijn: adaptive systems → design never ends*
   In the near future (partially now already) systems will not be delivered once. Over time, systems will improve (updates, etc) and extend their functionality over time. Hence, systems will continuously be redesigned. This will also require design methods to be able to operate 'online', without interfering with the system itself.
6. Cross-metric / property modeling (incl. cost)
   a. *Dennis: cost/rewards on a continuous probabilistic models / Markov automata*
   We extend Markov automata with state and impulse rewards. This leads to a richer set of properties, like the probability to reach a state until time T with costs lower than C. The most problematic part are the impulse rewards for time bounded properties.
   b. *Anne: model checking for survivability*
   c. *Felicita: cross metrics*
   More and more, requirements at the basis of systems employed in critical applications span several properties in the domain of resilience, security and, in general, quality of service. Given the need to satisfy a variety of such properties, which may also show

contrasting effects, a trade-off is usually required. Therefore, from the point of view of quantitatively assessing the level reached by such trade-off, metrics going across several properties would be desirable. In the past, the performability measure has been successfully proposed to trade between performance and reliability. Going in this direction, new metrics need to be explored, e.g., to trade between security and reliability, securety and safety, etc.

d. *Bill: Moebius modeling tool*
e. *Hermann: cross metric modeling*
f. *Albert: multi-objective optimization*
   The energy system is a complex system with multiple stakeholders and different parameters determining the costs and quality of service. Therefore, it is impossible to determine one single objective; it is an optimization over multiple-interrelated parameters. Moreover, it is even an optimization over multiple commodities: it is sometimes possible to interchange electrical consumption with gas consumption.
g. *Gethin: game models producers vs. consumers + analyzing trade-offs between metrics*
   The idea here is to model the system as a multi-player game as different parts of the system have different goals/metrics. For example a producer wants to optimise load-balance and profit while consumer would want to minimize cost whilst achieving some level of quality of service. Using game-models one can then look into tradeoffs between these metrics.

7. Runtime decision making
   a. *Erika: model-based predictive control*
      This technique is very popular and successful in engineering. The basic idea is to use a (sufficiently fine but not too complicated) model of a plant or system to predict how the system would behave under a certain control (e.g., via simulation). With this prediction, different techniques can be used to search for optimal control sequences. Similar techniques could be probably also applied to critical infrastructures.
   b. *Enrico: quantitative evaluation of models of operation procedures*
   c. *Jeremy: fluid analysis for rapid decision making*
   d. *Maria: approximate linear programming*
      When addressing optimal control of a Markov decision process through dynamic programming, the problem of determining the optimal value function can be rephrased as a Linear Programming (LP) program where a cost function is minimized subject to an infinite number of linear constraints, one for each control and state pair. This LP program is quite challenging to solve, since the optimization variable is infinite-dimensional (indeed, it is a function) and the number of constraints is infinite. Function approximation and relaxation of the resulting semi-infinite optimization problem can be exploited to compute an approximate linear programming solution.
   e. *Enrico: scheduling of activities with probabilistic durations*
   f. *John: real-time receding horizon optimization*
      This is really the same as model predictive control, I thought people would not know what this is so I decided to write something long winded and descriptive, then realized several other people just said MPC! The idea is to on-line run an optimization algorithm to select optimal future actions for the system (usually based on a model to predict the future under different choices of actions). One then executes the first of these optimal actions, throw the rest away, measure where the system ended up, and repeat the process. The periodic measurement introduces feedback, which makes the process

robust against uncertainties, most notably mismatch between the model used in the predictions and reality.

g. *Albert: model-predictive control*
A complicating aspect of energy optimization is that choices made influence the future status of the system too; switching on the washing machine at this moment leads to energy consumption for the next hour. Therefore, it is useful to take some future into account. A technology for this is Model Predictive Control.

h. *Bill: recovery and response engine, runtime for resiliency*

i. *Felicita: investigations on refining/adapting models to deal with initial inaccuracy evolutions of the system under analysis*
Modern software applications are increasingly pervasive, dynamic and heterogeneous. More and more they are conceived as dynamically adaptable and evolvable sets of components that must be able to modify their behaviour at run-time to tackle the continuous changes happening in the unpredictable open-world settings. The need for research advancement in the assessment of evolving, ubiquitous systems is recognized by the dependability/resilience community, since the involved aspects make traditional methods largely inadequate. Therefore, new approaches to tackle the involved challenges are under investigation. One direction to cope with the issues raised in the addressed context resorts to integrate pre-deployment stochastic model-based analysis with run-time monitoring, to achieve adaptive dependability assessment through re-calibration and enhancement of the dependability and performance prediction along time.

8. Deployment support
9. Synthesis
   a. *Rom: using model checking timed automata for control synthesis*
   There is an ample body of experience in deriving controllers using model checking, where the desired control is generated as a counterexample to the property "this system cannot be controlled in the right way". We would like to try to adapt this approach to timed automata models for cyber physical systems.
   b. *Erika: bounded-model-checking-based controller synthesis*
   Bounded model checking encodes system paths of a certain length satisfying certain properties as formulas. Checking these formulas for satisfiability answers the question for the existence of such paths. If a model of a critical infrastructure is available, why not to use bounded model checking for controller synthesis, i.e., for getting control sequences satisfying certain safety properties.
   c. *Gethin: synthesis for control strategies, optimum load balance subject to constraints for performance reliability etc.*
   This is related to the multi-objective model checking work which I gave a talk on. Using this approach one can find optimal policies/strategies for some metric/goal subject to meeting a number of constructs. A simple example is a power manager where one would want to optimise power consumption while providing a sufficient level of service.
   d. *Erika: CEGAR-based controller synthesis*
   Counterexample-Guided Abstraction Refinement (CEGAR) can be used for the safety analysis of complex systems. Starting from a course (over-approximating) abstraction of the system model, either the abstraction can be proven to be safe (in which case the concrete model is also safe), or the abstraction is unsafe which leads to an abstract counterexample. If this counterexample has a concrete counterpart, the system is

unsafe. Other we say that the counterexample is spurious. Spurious counterexamples can be used to direct the refinement of the model.
Similarly, if we want to lead the system to reach certain safe goal states, we could use CEGAR also for controler synthesis for criticaal infrastructures.

10. Hybrid nature
    a. *Anne: hybrid Petri net approach*
    b. *Enrico: timing analysis of models beyond the limits of Markov assumption*
    c. *Marco: spatial models*
    d. *Rom: using timed automata for hybrid and uncertainty modeling*
    Timed automata clocks have a very simple dynamics, but tricks can be used to effetively model and analyze continuous and stochastic behaviour. Such tricks would be important for cyber physical systems.
    e. *Erika: safety analysis for hybrid systems*
    In the last years great improvement can be observed in the development of tools for the automatic reachability analysis of hybrid systems (e.g., SpaceEx or Flow*). Can such tools be used for the safety analysis of critical infrastructures?
    f. *Boudewijn: there is no single correct model → do cooperating models*
    some academics/scientist claim that 'their model' is the true model for all to come. I do not believe in this. Various modelling approaches have different strengths. By smartly combining models, along clearly defined interfaces, I think more can be achieved than by adhering to just one 'model that does all'.
11. Model robustness validation (data driven)
    a. *Daniel: get data from communities that enforce "open data" policy*
    b. *Maria: quality assessment based on simulation*
    Suppose that one is interested in the probabilistic verification of a finite-horizon property for a given stochastic system that depends on the evolution of some output signal. According to the notion of approximate stochastic bi-simulation, the quality of a model as an approximate abstraction of the system can be quantified through the maximal distance between the system and the model outputs over all possible input realizations except for a set of them of probability epsilon. The evaluation of such distance, however, is a difficult task, computationally demanding in general. A possibility is then to assess the quality of the approximation by resorting to a randomized solution which prescribes to simulate the system and the model over a finite number N of realizations of the stochastic input only and then compute the maximal distance between the corresponding output signals. The finiteness of the considered realizations makes the problem computationally affordable. Probabilistic guarantees on the obtained solution can also be provided.
    c. *Daniel: insensitivity analysis of parameters*
    d. *Rom: validation of timed automata models*
    Timed automata validation of cyber physical models provides a challenge, given the size and complexity of such systems. How to make use of characteristics of the system to fight the state space explosion?
    e. *Armando: CEGAR-based model repair*
    Given a model and a property to be assessed, verification algorithms can give counter-examples when the property does not hold. Usually, counterexamples are used by the developers to fix the system manually, by tracing back to the causes of the anomalous behavior and then removing them. Model repair aims to automatize this process by calculating the fixes to the system, e.g., in terms of parameter tuning or structural

alterations. This technique can be useful in all the cases in which the manual fix does not make sense, e.g., in the case of control policies synthesized by means of real time dynamic programming or reinforcement learning.

12. Easy & quick & cheap (DSL)
    a. *Lucia: Vitruvius project*
    b. *Markus: language and tool set for modeling reliable systems*
    Recently, the LARES language (LAnguage for REconfigurable Systems) and an associated toolset have been developed. LARES focuses on dependability, fault-tolerance and reconfigurabiliy and is therefore particularly suited to the modelling of critical infrastructures. The modelling language supports the concepts of modularity and hierarchy. Different types of model validation are implemented in the LARES Integrated Development Environment (IDE). Model transformations from LARES to different target formalisms have already been realized, and more transformations, as well as extensions of the language, could be easily added in this open source project.
    c. *Usage of Modelica*
    d. *Mauro: user-orientation of modeling languages and solution descriptions, holistic representation*
    The use of a framework that enables the fast definition (and interpretation) of user defined modeling languages, such as domain specific languages, helps in encouraging users to adopt the framework, as they can naturally interact with it as they are used with other tools, and allows different representations of the same model at different complexity levels, offering each category of user the right perspective on the model. In this way even a complex model can be viewed as one, even if its submodels are very heterogeneous and are based on different premises.
13. Uncertainty
    a. *Erika: formal methods for probabilistic hybrid models*
    A lot of work was done on model checking discrete- and continuous-time Markov models. Can these methods be applied also to probabilistic models for critical infrastructures?
    b. *Dennis: non-deterministic behavior in continuous and probabilistic models*
    Markov automata are a model incorporating continuous stochastic timing, nondeterministic choices and discrete probability distributions. They provide a well-defined semantics for generalised stochastic Petri nets. Algorithms for timed reachability probabilities and expected durations until a certain event are already available.
    c. *Jeremy: stochastic process reward models*
    d. *Martin: automatic analysis of stochastic hybrid models*
    While some first prototype tools for the automatic analysis of stochastic hybrid models are available, all of them are severely limited as none of them scales well, none covers a comprehensive range of different stochastic phenomena (e.g., component failures, measurement errors in sensors, uncertain continuous dynamics, response time distributions, classification errors in signal processing and interpretation, ...), to name just a few shortcomings. We do thus need coordinated research concerning tool development, including novel notions of (automated, adaptive, etc.) abstractions for state space reduction.
    e. *Maria: randomized methods based on scenarios*
    The 'scenario approach' is an innovative technology that has been introduced to solve convex optimization problems with an infinite number of constraints, a class of problems which often occurs when dealing with uncertainty. This approach relies on random sampling of constraints, and provides a powerful means for solving a variety

of design problems, and, in particular, problems in systems and control such as model reduction, prediction, and constrained control design.

f. *Ralf: analysis of probabilistic systems, in particular generation of counterexamples*
The generation of counterexamples for violated system properties is an important part of the debugging process and also applied to refine system abstractions which are too coarse. For digital circuits, for instance, counterexamples are often obtained with little additional effort from the model checking process. For refuting LTL properties, a counterexample consists of a single system run that exhibits unwanted behavior. For probabilistic systems, the situation is different: model checking yields the mere probabilities, but no debugging information. Therefore dedicated counterexample generation algorithms are required. A counterexample has to certify that the probability of unwanted behavior is beyond a given limit. Therefore potentially large sets of runs are necessary whose joint probability mass exceeds the limit.
We have developed methods not only to compute counterexamples for very large systems, but also methods which compute smallest possible counterexamples both for DTMCs and MDPs on different levels of the system representation: traces at the lowest level, critical parts of the state space, and critical parts of the model description at the highest level.

g. *Sarah: investigation of parameter of uncertainty in overall model robustness*
I would like to know the extent of inaccuracy that will result in my model for say, likelihood of cascading failure in a power grid if I have over- or under-estimated the value of a parameter such as line capacity.

14. Real data & workload
    a. *Marco: workload generator and simulator*
15. Multi-aspect anomalie detection & response
    a. *John: model-based fault detection & isolation*
    Again model based, use a model of the nominal process to filter any data collected about the system. If the data is incompatible with the nominal process model, the filter will show a large residual. This serves as an indication that the process is non-nominal, suggesting an error or attack has occurred. In this case the system raises an alarm.
16. Workflow-driven security assessment
    a. *Bill: HiTop modeling language*
    b. *Zbigniew: Evaluation assessment based on security use cases → analogous to safety cases*
17. Feature interaction "unheard properties"
    a. *Armando: verification of emergent behaviors*
    When considering large distributed systems whose behavior results from the composition of several "locally consistent" control laws, it is possible that some global behavior emerges during runtime as a result of non-trivial interactions between the components. The behavior is emergent in the sense that no global control is enforced to produce it, yet it arises and it self-maintains consistently. Emergent behaviors can be desirable, e.g., fast routing in a large network using only local control policies, or undesirable, e.g., cascading failures. In both cases, modeling of the system and automated verification of properties entailing emergent behaviors can be useful tools in the analysis of complex CPS.

# 5 Working Groups

## 5.1 Preface

*Erika Ábrahám, Anne Remke, William H. Sanders, and Alberto Avritzer*



We have chosen a facultative approach towards forming working groups, since we believe in the power of self-management. Just before the first working sessions we invited people upfront to present their ideas for working groups, to also give people from other communities the possibility to join working groups with researchers, they were unfamiliar with. We repeated this process before the second break-out session, to allow people to join other working groups and to check whether new groups have formed during the first day. On Friday morning, we reserved time for short presentations of the working groups, which was very well received by the seminar participants.

## 5.2 From the Application Point of View

*Zbigniew Kalbarczyk*



Use of IEDs (Intelligent Electronic Device) in substations to monitor the power grid and communicate between the control centers and substations makes this infrastructure susceptible to transient errors and malicious attacks. We discuss experimental study of the impact of errors on the micro-processor based power grid equipment. Two case studies are presented:

1. Characterization of error resiliency of substation devices using fault/error injection
2. PMUs (phasor measurement units) and bad data detection algorithms (GPS spoofing attack).

## 5.3 Two Issues in Modeling Critical Infrastructures

*Rom Langerak*


Joint work of Rom Langerak, Felicita di Giandomenico and Zbigniew Kalbarczyk

As a newcomer to the field of critical infrastructures, I would like to raise two issues:

- What are the characteristics of critical infrastructures?
- How to avoid model bias in modeling critical infrastructures?

### Characteristics

I have some experience in modeling and analyzing several kind of networks (e.g. communication networks, biological networks, wireless sensor networks, switching networks). Which part of that experience could still be valid in the context of critical infrastructures, and what are new problems that need creativity to solve them? In order to answer this question, I would like to have a better idea of the specific characteristics of critical infrastructures (and I hope this is useful for other people as well :-) ).

### Model Bias

We all have our favorite models: timed automata, hybrid automata, markov chains, etc. etc. Now the model you choose has a big influence on the analysis methods and the kind of questions you are going to use. Model checking often concentrates on reachability, markov chains on steady state properties, control theorists focus on stability and robustness properties, and so on. Choosing a model in an early stage of tackling a problem may put such a bias on what you are going to study, that it may lead to a distortion of the actual problem, to answering the wrong questions, and in general to waisting a lot of time and to missing what is important for the domain experts. Therefore it is important to get a good understanding of a problem area, before you have formalized it! This means it would be helpful to have a some good "informal" concepts in order to understand and communicate about your understanding.

### A Tentative: Networks of Cyberphisical Nodes

What we came up with as a first tentative for an "informal" modeling framework is networks of nodes, where each node consists of two parts:

- a cyber part, with an associated cyber state (think of e.g. discrete variables)
- a physical part, with an associated physical state (think of e.g. a vector in $R^n$) and three types of interactions:
- interactions between the cyber and physical part in a node
- cyber interactions with cyber parts in other nodes
- physical interactions with physical parts in other nodes, where the interactions may lead to changes in the cyber or physical state of the corresponding part of the node. In addition, there may be global constraints on the network (e.g. physical laws or topology constraints or invariants of some nature).

We would like to point out that we do not prescribe any formal description or level of abstraction for these aspects (e.g. the interaction "camera sees a vehicle" could be described in many different ways, from abstract to physically concrete).

The idea would now be to form first a general idea of some critical infrastructure problem using this informal framework, and discuss issues like components, interactions, hierarchies, scenarios, metrics, goals, etc. etc. first on this informal model, before going to the phase of formal modeling (and getting the "real work" done :-) ).

### Questions

The above proposed framework seems quite general and natural. Its main feature is the distinction between a cyber and a physical part of a node. This does not seem to be very deep or shocking; still it might be quite useful as a way of trying to characterize the specific flavor of critical infrastructures, as a way of communicating with domain experts, and as a way of avoiding modeling bias.

What we might do together, is to take a look at several papers that we contributed together (and that are included in the attachment), and try to answer the following questions for a paper:

1. is the informal framework useful for understanding and/or describing the specific features of the application in the paper?
2. what is missing, what do we need to describe other characteristics?
3. is it possible to understand how the informal framework could be mapped to the modeling formalism in the paper?

## 5.4 Assessement of Strom Impacts

*Laura Carnevali*


**Joint work of** Laura Carnevali, Enrico Vicario, Anne Koziolek, Daniel Sadoc Menasche and Lucia Happe

In Dagstuhl, we have discussed how to model and analyze the impacts of large hurricanes on a power distribution network. In particular, we have considered smart grids equipped with reclosers and tie switches, and we have focused on the evaluation of survivability related metrics. The group discussion pointed out the opportunity to relate the survivability assessment with the hurricane characterization as well as the necessity to have a scalable survivability model to address large critical infrastructures. After the Dagstuhl seminar, we have carried on the study, developing a formal approach to the evaluation of different alternatives for storm hardening. We have recently submitted a conference paper and plan to go on with the collaboration, especially to take into account cascading failures and to evaluate different investment strategies with respect to customer affecting metrics.

## 5.5 Smart City Survivability

*Anne Remke*


**Joint work of** Anne Remke, Boudewijn Haverkort, Hamed Ghasemieh, Laura Carnevali, Enrico Vicario, Sahra Sedighsedigh, Alberto Avritzer, Daniel Sadoc Menasche, Lucia Happe, Anne Koziolek

More and more aspects of our daily life depend heavily on large-scale infrastructural systems, think of rail and road networks, but also about telecommunication networks (internet, wired and wireless telephony). More recently, also the networks that provide gas, water and electricity have become much more "ICT-based", implying that their well-operation is becoming dependent on the correct operation of the supporting ICT. And although the embedded ICT does provide more functionality, it is also often a source of failures, or the victim of attacks. Nevertheless, it is essential for all these critical infrastructural systems to survive catastrophic events. In this paper we address approaches towards so-called "survivability evaluation" of infrastructural systems; our focus thereby lies on water, gas and electricity infrastructures, infrastructures that used to be run by municipalities, but now are mostly run by large internationally operating companies.

We note here that the concept of survivability is not restricted to just this class of infrastructural systems. It is also known for military devices, for example, aircraft combat survivability, and even in agriculture [1].

The literature is abundant with different definitions of survivability. For an overview see for example [2, 3]. Distinct definitions stress different aspects of survivability, be it the detection of faults, the defence against attacks or the recovery from various types of disasters. We will focuss on the behaviour of a system after a disaster has occurred. Note that we do not introduce a new definition of survivability but state a slightly generalised version of the one in [4]; it reflects an intuitively appealing view on survivability of systems but is therefore also quite informal:

> *Survivability is the ability of a system to* ***recover*** *predefined* ***service*** *levels in a* ***timely manner*** *after the occurrence of* ***disasters****.*

A disaster might be any kind of severe disturbance of the infrastructural system, for example, a power breakdown, a complete or partial cut of communication lines, a flood, heavy rain or a thunderstorm. The possible causes are manifold and include purposeful attacks as well as natural disasters like earthquakes or thunderstorms.

A system is survivable if it includes mechanisms to return to normal service within an acceptable time even though a disaster occurred. What kind of mechanisms are used and how they are implemented is not part of the survivability definition. One possible mechanism to achieve survivability is fault tolerance or any other form of redundancy [5].

The above definition of survivability does not give at all a precise recipe how to decide whether a system is survivable or not. To overcome this, many approaches have been followed in the literature for the quantitative determination of survivability [6, 7, 3, 8, 9]. Most of them are model-based and suggest some measure on the system (model) behaviour and study its evolution after the occurrence of a disaster. It, thus, is the deliberate decision of the person performing the survivability evaluation to choose an appropriate measure.

What is typical for the approaches presented in this overview paper, is that the application field requires some form of hybrid model, taking into account discrete state components, continuous state components (for the physical issues playing a role), in combination with both deterministic and stochastic behaviour. This combination makes analytical approaches very challenging, however, there is a clear need for these, as purely simulation-based approaches are very costly, overly costly, to use in practice.

The rest of this paper is organised as follows. In the three sections that follow, we give a brief introduction into recent approaches on survivability evaluation of three infrastructures, being, smart gas, water and electricity networks.

## References


1 S. Ling, Z. Zesheng and G. Hengshen. *A GIS-based agricultural disaster evaluation system.* In Proc. of the ESRI International User Conference, 1998.

2 J. C. Knight, E. A. Strunk, and K. J. Sullivan. *Towards a rigorous definition of information system survivability*. In Proc. of the 3rd DARPA Information Survivability Conference and Exposition (DISCEX'03), pp. 78–89, IEEE Press, 2003.

3 Y. Liu and K. S. Trivedi. *A general framework for network survivability quantification*. In Proc. of the 12th GI/ITG Conference on Measuring, Modelling and Evaluation of Computer and Communication Systems together with 3rd Polish-German Teletraffic Symposium (MMB & PGTS 2004), pp. 369–378, VDE Verlag, 2004.

4 B. Ellison, D. A. Fisher, R. C. Linger, H. F. Lipson, T. Longstaff and N. R. Mead. *Survivable network systems: An emerging discipline*. Carnegie Mellon Software Engineering Institute, Tech. Rep. CMU/SEI-97-TR-013, 1997.

5 D. Pradhan, Ed. *Fault-tolerant and dependable computer system design*. 2nd edition, Prentice Hall, 2003.

6 S. C. Liew and K. W. Lu. *A framework for characterizing disaster-based network survivability*. IEEE J. Selected Areas in Communications, 12(1):52–58, 1994.

7 Y. Liu, V. B. Mediratta and K. S. Trivedi. *Survivability analysis of telephone access network*. In Proc. of the 5th IEEE International Symposium on Software Engineering (ISSRE'04), pp. 367–378, 2004.

8 D. Medhi. *A unified approach to network survivability for teletraffic networks: Models, algorithms and analysis*. IEEE Trans. Comm., 42(2-4):534–548, 1994.

9 A. Zolfaghari and F. J. Kaudel. *Framework for network survivability performance*. IEEE J. Selected Areas in Communications, 12(1):46–51, 1994.

## 5.6 Modeling Smart Grids

*Anne Remke*


**Joint work of** Anne Remke, Marijn Jongerden, Albert Molderink, Gethin Norman, Maria Prandini, Gerard Smit

In our discussion group we had experts from the areas of Smart grids, stochastic models and hybrid models present. We discussed several modeling challenges that are present in Smart grids and what could be suitable approaches for modeling and analysis.

We learned that one of the key issues in smart grids is the balancing of demand and production. This applies on all scales, for example within one household, but also within a neighborhood. For each house but also for each neighborhood an energy profile can be constructed, which predicts the overhead (positive and negative) within the next 24 hours. A prediction of prizes for energy is normally available 24 hours in advance, together with a cost function, that specifies how much a household would be willing to pay for extra energy, it is possible to exactly specify what will happen during the next 24 hours.

Moreover, we clarified many questions and assumptions regarding the functioning of Smart grids. For example, we discussed the following questions:

- Do you only load your storage battery from renewable sources, or would it be possible to load it from the grid, in case prizes are really cheap?
- Do you always use the capacity that might be left in the battery, or would you rather use the grid if prizes are cheap?
- Is it possible to charge the battery while using it to power devices?

We discussed the possibility to model different appliences, like a thermostat, microCHP, heat pumps, dish washer , freezer, fridge etc. as timed automata or as hybrid automata in order to control and balance their energy usage. Several optimization criteria are possible, while networking companies strive to keep the maximum peaks as low as possible, an individual would strive to minimize the amount of money spend for energy.

## 6 Seminar Program

| Monday (January 13, 2014) | |
|---|---|
| 8:45–9:00 | Welcome – Anne Remke |
| 9:00–9:30 | *William H. Sanders* |
| | **Challenges and opportunities in modeling the power grid cyber-physical infrastructure** |
| 9:30–10:00 | *Albert Molderink* |
| | **Optimization strategies for the future electricity infrastructure – Smart Grid research and current market opportunities** |
| 10:00–10:30 | *Peter Langendörfer* |
| | **Engineering cyber-physical systems/critical infrastructure systems: A craftsman approach** |
| 10:30–11:00 | **Coffee break** |
| 11:00–11:30 | *Daniel Sadoc Menasche* |
| | **Design of distribution automation networks using survivability modeling and power flow equations** |
| 11:30–12:00 | *Lucia Happe and Anne Koziolek* |
| | A common analysis framework for smart distribution networks applied to security and survivability analysis |
| 12:00–14:00 | **Lunch** |
| 14:00–14:30 | *Erika Ábrahám* |
| | **Tutorial: Formal methods for hybrid systems** |
| 14:30–15:00 | *Marc Bouissou* |
| | **Modeling stochastic hybrid systems in Modelica: Some results obtained in the MODRIO project** |
| 15:00–15:30 | **Coffee break** |
| 15:30–17:00 | **From research to application: Open problems, needs and wishes. Panel discussion lead by Boudewijn Haverkort** |
| | *Peter Langendörfer, Albert Molderink, William H. Sanders, Gerard Smit, N.N.* |
| 18:00–19:00 | **Dinner** |
| 19:30 | Opening of the art exhibit **Neun Minuten vor Vegas** by the German artist *Fabian Treiber* |

| Tuesday (January 14, 2014) | |
|---|---|
| 9:00–10:00 | *Christel Baier* |
| | **Tutorial: Probabilistic Model Checking** |
| 10:00–10:30 | **Coffee break** |
| 10:30–11:00 | *Holger Hermanns* |
| | **Time-dependent analysis of attacks** |
| 11:00-11:30 | *Luca Bortolussi* |
| | **Parameter identification and synthesis from qualitative data and behavioural constraints** |
| 11:30–12:00 | *Maria Prandini* |
| | **Randomized methods for design in the presence of uncertainty** |
| 12:00–14:00 | **Lunch** |
| 14:00–14:30 | *Enrico Vicario* |
| | **Quantitative evaluation of non-Markovian models through the method of stochastic state classes and the Oris tool** |
| 14:30–15:00 | *Armando Tacchella* |
| | **Proving safety of complex control software: A review of three "test tube" applications in robotics** |
| 15:00–15:30 | **Coffee break** |
| 15:30–18:00 | **Break out session (coffee available)** |
| 18:00–19:00 | **Dinner** |

| Wednesday (January 15, 2014) | |
|---|---|
| 9:00–9:30 | *Laura Carnevali* |
| | **The theory of stochastic state classes: Tool support and applications** |
| 9:30–10:00 | *Anne Remke* |
| | **Analysis of a sewage treatment facility using hybrid Petri nets** |
| 10:00–10:30 | **Coffee break** |
| 10:30–11:00 | *Hermann de Meer* |
| | **Resilience of data networking and future power networks** |
| 11:00–11:30 | *Felicita Di Giandomenico* |
| | **Issues in modelling smart grid infrastructures to assess resilience-related indicators** |
| 11:30–12:00 | *Gerard Smit* |
| | **Energy-autonomous smart micro-grids** |
| 12:00–14:00 | **Lunch** |
| 14:00–14:30 | *John Lygeros* |
| | **Cyber-security of SCADA systems: A case study on automatic generation control** |
| 14:30–15:00 | *Sahra Sedighsarvestani* |
| | **Towards quantitative modeling of reliability for critical infrastructure systems: advances and challenges** |
| 15:00–15:30 | **Coffee break** |
| 15:30–18:00 | **Break out session (coffee available)** |
| 18:00–19:00 | **Dinner** |

| Thursday (January 16, 2014) | |
|---|---|
| 9:00–9:30 | *Boudewijn Haverkort* |
| | **Systems of systems design challenges** |
| 9:30–10:00 | *Aad van Moorsel* |
| | **Data collection strategies for model-based analysis** |
| 10:00–10:30 | **Coffee break** |
| 10:30–11:00 | *Marco Gribaudo* |
| | **Multiformalism to support software rejuvenation modeling** |
| 11:00–11:30 | *Jeremy T. Bradley* |
| | **Rapid evaluation of time-critical service level objectives** |
| 11:30–12:00 | *Katinka Wolter* |
| | **Quantitative evaluation of smart grid control traffic** |
| 12:00–14:00 | **Lunch** |
| 14:00–14:30 | *Joost-Pieter Katoen* |
| | **A rigorous approach towards reliable and dependable train and space systems** |
| 14:30–15:00 | *Dennis Guck* |
| | **Smart railroad maintenance engineering with stochastic model checking** |
| 15:00–15:30 | **Coffee break** |
| 15:30–16:00 | *Alessandro Abate* |
| | **Cascading events in probabilistic dynamical networks** |
| 16:00–16:30 | *Martin Fränzle* |
| | **Symbolic analysis of complex systems** |
| 16:30–18:00 | **Break out session (coffee available)** |
| 18:00–19:00 | **Dinner** |

| Friday (January 17, 2014) | |
|---|---|
| 9:00–9:30 | *Ralf Wimmer* |
| | **Optimal counterexamples for Markov models** |
| 9:30–10:00 | *Gethin Norman* |
| | **Verification of probabilistic timed automata** |
| 10:00–10:30 | **Coffee break** |
| 10:30–12:00 | *Discussion of results* |
| 12:00–14:00 | **Lunch** |

## Participants


- Alessandro Abate
University of Oxford, GB
- Erika Ábrahám
RWTH Aachen, DE
- Christel Baier
TU Dresden, DE
- Bernd Becker
Universität Freiburg, DE
- Luca Bortolussi
University of Trieste, IT
- Marc Bouissou
Ecole Centrale Paris, FR
- Jeremy T. Bradley
Imperial College London, GB
- Laura Carnevali
University of Firenze, IT
- Hermann de Meer
Universität Passau, DE
- Felicita Di Giandomenico
CNR – Pisa, IT
- Martin Fränzle
Universität Oldenburg, DE
- Hamed Ghasemieh
University of Twente, NL
- Marco Gribaudo
Technical University of Milan, IT
- Dennis Guck
University of Twente, NL
- Lucia Happe
KIT – Karlsruhe Institute of Technology, DE
- Boudewijn Haverkort
University of Twente, NL
- Holger Hermanns
Universität des Saarlandes, DE
- Mauro Iacono
The Second Univ. of Naples, IT
- Marijn R. Jongerden
University of Twente, NL
- Zbigniew Kalbarczyk
University of Illinois – Urbana Champaign, US
- Joost-Pieter Katoen
RWTH Aachen, DE
- Anne Koziolek
KIT – Karlsruhe Institute of Technology, DE
- Peter Langendörfer
IHP GmbH – Frankfurt/Oder, DE
- Rom Langerak
University of Twente, NL
- John Lygeros
ETH Zürich, CH
- Daniel Sadoc Menasche
University of Rio de Janeiro, BR
- Albert Molderink
University of Twente, NL
- Gethin Norman
University of Glasgow, GB
- Maria Prandini
Technical University of Milan, IT
- Anne Remke
University of Twente, NL
- William H. Sanders
University of Illinois – Urbana Champaign, US
- Sahra Sedigh Sarvestani
University of Missouri – Rolla, US
- Markus Siegle
Universität der Bundeswehr – München, DE
- Gerard J. M. Smit
University of Twente, NL
- Oliver Stecklina
IHP GmbH, DE
- Armando Tacchella
University of Genova, IT
- Aad van Moorsel
Newcastle University, GB
- Enrico Vicario
University of Florence, IT
- Ralf Wimmer
Universität Freiburg, DE
- Verena Wolf
Universität des Saarlandes, DE
- Katinka Wolter
FU Berlin, DE


Report from Dagstuhl Seminar 14032

# Planning with Epistemic Goals

**Edited by**

**Thomas Ågotnes[1], Gerhard Lakemeyer[2], Benedikt Löwe[3], and Bernhard Nebel[4]**

1 **Universitetet i Bergen, NO,** thomas.agotnes@infomedia.uib.no
2 **RWTH Aachen, DE,** gerhard@kbsg.rwth-aachen.de
3 **Universiteit van Amsterdam, NL and Universität Hamburg, DE,** b.loewe@uva.nl
4 **Universität Freiburg, DE,** nebel@uni-freiburg.de

**Abstract**

This report documents the outcomes of Dagstuhl Seminar 14032 "Planning with epistemic goals". It brought together the communities of so far relatively separate research areas related to artificial intelligence and logic: automated planning on the one hand, and dynamic logics of interaction on the other. Significant overlap in motivation, theory and methods was discovered, and a good potential for cross fertilization became apparent.

**Seminar** January 12–15, 2014 – http://www.dagstuhl.de/14032
**1998 ACM Subject Classification** F.4.1 Mathematical Logic
**Keywords and phrases** planning, epistemic logic, modal logic
**Digital Object Identifier** 10.4230/DagRep.4.1.83
**Edited in cooperation with** Christian Becker-Asano

## 1 Executive Summary

*Thomas Ågotnes*
*Gerhard Lakemeyer*
*Benedikt Löwe*
*Bernhard Nebel*



Automatic planning is a subarea of Artificial Intelligence that was initiated in the 70s. The main idea was to develop efficient methods to generate action plans, for example for robot missions. The initial attempts were based on first order logic. However, most approaches quickly adapted simpler logics and focused on search techniques. The recent years have brought a huge advance on scalability by employing smart search techniques such as heuristic search, SAT, BDDs, and other techniques. Currently, planning researchers explore widening the scope of planning tasks and to connect back to logic oriented approaches of describing dynamics such as GOLOG. At the same time, planning researchers strive to capture planning settings that are more challenging than the classical setting. For instance, planning under uncertainty and planning taking into account beliefs are current research topics.

The research area of dynamic logics of interaction is part of the larger field of applied and interactive logic: the use of logical methods in order to formalize procedures in social and communication contexts. The systems are typically based on the semantics of modal logic, and often focus on information (ex)change and the dynamics of knowledge and beliefs.




Paradigmatic examples are public announcement logic and dynamic epistemic logic. One of the main technical features is the incorporation of agency and events into the modal framework as encapsulated by the notion of product update. Recently, some authors have proposed to use the ideas (or, more generally, the methodology) of dynamic approaches to logic for planning.

Epistemic goals, or more generally, goals that have to be expressed in some intensional language (epistemic, doxastic, deontic, others) have been discussed in several papers in the logic community, but are mostly absent from automatic planning. The development of a research community dealing with these goals in planning will require a close interaction between the two involved communities. The main goal of this workshop was to bring the two communities together and develop a vision of the mid-term goals of such a collaboration. In order to facilitate this, the organizers decided to arrange the workshop around work in four groups: after four tutorial lectures by Gerhard Lakemeyer, Hans van Ditmarsch, Thomas Bolander, and Hector Geffner on Monday, the participants were split up into four groups labelled APPL, BENCH, COMP, and LANG. Tuesday was largely reserved for work in the groups and for preparing the group reports included in this report. Tuesday evening also saw a concert in the *Weisser Saal* with François Schwarzentruber playing the piano and Hans van Ditmarsch playing the cello. The final day had some short presentations and a closing discussion.

For the four themes, the organizers had provided some guiding questions, but left the discussion open for the group participants:

**APPL** *Applying epistemic planning in the real world.* Theme coordinator: Ron Petrick; group participants: Maduka Attamah, Christian Becker-Asano, Martin Holm Jensen, Benedikt Löwe, Sheila McIlraith, Leora Morgenstern, and François Schwarzentruber. Guiding questions: What are promising applications that will convince the outside world to use epistemic planning? Which areas outside of academia could be interested in epistemic planning? How do we get other academic disciplines (such as roboticists) interested in epistemic planning? Can we come up with a concrete research plan for such an application within the next three years?

**BENCH** *Establishing benchmarks and concrete goals for epistemic planning.* Theme coordinator: Bernhard Nebel; group participants: Carmel Domshlak, Hector Geffner, Malte Helmert, Andreas Herzig, Jörg Hoffmann, Jérôme Lang, and Hans van Ditmarsch. Guiding questions: Can we come up with standardized problems to measure and compare systems for epistemic planning? Which standardized problems could help to calibrate the expressive power of epistemic planning formalism? What currently unsolved problems will serve as milestones and success criteria for the next three to five years?

**COMP** *Taming the complexity of epistemic planning.* Theme coordinator: Thomas Bolander; group participants: Gerhard Lakemeyer, Yongmei Liu, Robert Mattmüller, Sunil Simon, Jan van Eijck, and Yanjing Wang. Guiding questions: Which aspects of epistemic planning are responsible for the increase of computational complexity? Are there fragments of epistemic planning that allow for an efficient implementation? Can we devise sufficiently expressive planning formalisms that still have acceptable complexity?

**LANG** *Finding adequate languages for epistemic planning.* Theme coordinator: Thomas Ågotnes; group participants: Guillaume Aucher, Mikkel Birkegaard Andersen, Jens Claßen, Tiago de Lima, Valentin Goranko, and Gabriele Röger. Guiding questions: Which formalisms are adequate to represent epistemic planning problems? Can we devise languages for epistemic planning that are intuitive to understand and use? Can we extend

existing plan definition languages with epistemic features? Are the epistemic logics we have sufficiently expressive to serve as a basis for such planning formalisms?

In the final discussion, the participants discussed the immediate future of the interaction between the two fields. One idea was to edit a special issue of the journal *Annals of Mathematics and Artificial Intelligence*, and the seminar organisers are currently in negotiation with the journal editors about that. Thomas Bolander, Hans van Ditmarsch, Jan van Eijck, and R. Ramanujam are planning a follow-up meeting at the Lorentz Center in Leiden in the spring of 2015, and we hope to reconvene with many of the Dagstuhl participants at that meeting.

## 2 Table of Contents

# 3 Overview of Talks

## 3.1 The Situation Calculus and Golog

*Gerhard Lakemeyer (RWTH Aachen, DE)*



In this tutorial I present the basics of the situation calculus and the action programming language Golog. The situation calculus was invented by John McCarthy as a rich logical language to represent dynamical domains and to reason about action and change. Here I consider the variant introduced by Ray Reiter, as it has the very desirable feature that it comes equipped with a solution to the frame problem in the form of so-called successor state axioms, which determine precisely how the values of fluents change from one situation to another as a result of performing an action. An important benefit of successor state axioms is that they allow us to solve the projection problem, that is, reasoning about what is true after a sequence of actions, by performing regression, which reduces a query about the future to a query about the initial situation. After going over these and other properties of the situation calculus, I then move on to Golog, which has many features from imperative programming such as while-loops, but also constructs allowing for non-deterministic choice. The latter enables a user to encode planning problems as part of a Golog program. In the tutorial I briefly go over the semantics of the various constructs, which are all defined within the situation calculus. I also briefly sketch how knowledge and sensing can be handled in Golog and the situation calculus.

## 3.2 Epistemic Planning: The DEL approach

*Thomas Bolander (Technical University of Denmark – Lyngby, DK)*



In my talk I will present one of the possible approaches to planning with epistemic goals: planning based on Dynamic Epistemic Logic (DEL). I will show how planning based on DEL, henceforth called "epistemic planning", generalises classical propositional planning in the most obvious way. Epistemic planning is essentially obtained by generalising both the states and actions of propositional planning to multi-sets of such states and actions – and then add indistinguishability relations for each agent. In my talk I will introduce the general framework of epistemic planning, relate it to classical planning, and discuss complexity results for the plan existence problem. Multi-agent epistemic planning is undecidable, already with 2 agents and no ontic actions. Some relevant fragments have been shown to be decidable, but the quest for fragments of epistemic planning with decent complexities still goes on.

## 3.3 Merging DEL and ETL for epistemic planning

*Yanjing Wang (Peking University – Beijing, CN)*


**Joint work of** Wang, Yanjing; Yanjun Li
**Main reference** Y. Wang, Y. Li, "Not All Those Who Wander Are Lost: Dynamic Epistemic Reasoning in Navigation," in Proc. of Advances in Modal Logic 2012 (AiML'12), pp. 559-580, College Publications, 2012.
**URL** http://www.aiml.net/volumes/volume9/Wang-Li.pdf

In this talk, we start with a general approach of axiomatizing dynamic epistemic logics using epistemic temporal axioms, as developed in [1] and [2]. This method does not rely on whether the dynamic epistemic logic in question is reducible to a fragment without the dynamic operators. We identify four important (and meaningful) axioms underlying the standard public announcement logic and action model based dynamic epistemic logic: the invariance of propositional valuation, the definition of executability of actions, no miracles and perfect recall. It turns out that the first two are not necessary in order to obtain a complete logic using our method, and this gives us the technical preparation of a new dynamic epistemic logic on transition systems with a uncertainty set. The crucial idea is to include temporal information in the model but still handle the epistemic updates, in spirit, as in the semantics of DEL.

This approach is carried out initially in [3], and its model is essentially the non-probabilistic model of contingent planning in AI. We completely axiomatize the logic, demonstrate its normal form and show the decidability of the logic. The plan verification can be then turned into a model checking problem of this logic, and the plan existence problem can be turned into a model checking problem of a PDL-like extension of this logic. The advantages of using this modal logic approach to continent planning include the following: more general (epistemic) goals are handled (without increasing complexity), (conditional) plans are specified formally thus relationship (abstraction, refinement, equivalence) between plans can be studied precisely, and this gives us a basic common platform to compare the complexity of different planning problems by restricting/extending the model or the logical language.

### References

1 Wang, Y. and Cao, Q. (2013). On axiomatizations of public announcement logic. *Synthese*, 190(1):103–134.
2 Wang, Y. and Aucher, G. (2013). An alternative axiomatization of DEL and its applications. In Proceedings of *IJCAI* 2013:1147–1154.
3 Wang, Y. and Li, Y. (2012). Not all those who wander are lost: dynamic epistemic reasoning in navigation. In *Proceedings of Advances in Modal Logic 2012*, volume 9, pp. 559–580, College Publications.

## 3.4 Automata Techniques for Epistemic Protocol Synthesis

*Guillaume Aucher (INRIA Rennes – Bretagne Atlantique, FR)*



In this work we aim at applying automata techniques to problems studied in Dynamic Epistemic Logic, such as epistemic planning. To do so, we first remark that repeatedly executing ad infinitum a propositional event model from an initial epistemic model yields a

relational structure that can be finitely represented with automata. This correspondence, together with recent results on uniform strategies, allows us to give an alternative decidability proof of the epistemic planning problem for propositional events, with as by-products accurate upper-bounds on its time complexity, and the possibility to synthesize a finite word automaton that describes the set of all solution plans. In fact, using automata techniques enables us to solve a much more general problem, that we introduce and call epistemic protocol synthesis.

### References

1 Aucher, G.; Maubert, B.; Pinchinat, S., *Automata Techniques for Epistemic Protocol Synthesis*, Proc. of the 2nd Int'l Workshop on Strategic Reasoning, Grenoble, France, pp. 97–103, 2014.

## 3.5 What does it mean to know a number

*Jan van Eijck (CWI – Amsterdam, NL)*



The talk is about the application of DEL model checking to the analysis of cryptographic security protocols. As a step towards that, I will analyze the question of the title. I will make a proposal for how to represent knowledge of large numbers in Kripke models, and I will use this representation to model a protocol for secret key distribution over an insecure network in DEL. In the conclusion, I will make a connection with epistemic planning.

## 3.6 Epistemic modal logic with robots and cameras

*François Schwarzentruber (IRISA – Rennes, FR)*



In this talk, we present a concrete version of epistemic modal logic where agents are located in the plane. Formulas of the language are formulas of epistemic modal logic, that is to say we have construction of the form "$K_a\phi$" meaning that "agent a knows the property $\phi$". Atomic propositions are grounded in the sense that they denote physical properties as "agent a sees agent b", "agent a wears a hat", etc. We may add dynamic extensions to the language as public announcements or action models of dynamic epistemic logic. We evaluate formulas in a given world described by the positions of agents in the plane. From a world, we infer a Kripke structure and this Kripke structure is used to give a semantics to formulas. The model checking problem is defined as follows: given a description w of the positions of the agents and a formula phi, does phi hold in w? We show that the model checking problem is decidable [1] but the exact complexity is an open issue. In the case where there is common knowledge of the positions of the agents (but not the angle of view), we show that the model checking problem is PSPACE-complete [2].

### References

1 Balbiani, P.; Gasquet, O.; Schwarzentruber, F., *Agents that look at one another*, Logic Journal of the IGPL, 21:3, 438–467, 2013.
2 Gasquet, O.; Goranko, V.; Scharzentruber, F., *Big Brother Logic: Logical modeling and reasoning about agents equipped with surveillance cameras in the plane*, Proc. of AAMAS'2014, 2014.

## 3.7 Epistemic Protocols for Gossip

*Maduka Attamah (University of Liverpool, GB)*



The focus of much research in dynamic epistemic logic, and more generally in epistemic and temporal modal logics, is Analysis: given a well-specified input epistemic state, and some well-specified dynamic process, compute the output epistemic state. We focus on Synthesis: given a well-specified input epistemic state, and desirable output, find the process transforming the input into the output. The process found is the epistemic protocol. In this talk we discuss protocol synthesis within a specific epistemic planning problem scenario, namely, the gossip problem.

The gossip problem describes a scenario with a group of $N$ agents where each agent knows a unique secret. The agents can *call* each other, that is a pairwise communication in which the calling pair reveal all the secrets they know to each other. The goal is such that by means of a sequence of calls, all the agents learn the secret of every other agent. Studies in literature include variations of this problem in graph theory and matrix algebra , with results about the number of calls to needed to fully distribute the secrets[1]. We focus on variations of the gossip problem in which an agent communicates with some other agent based on some knowledge (or epistemic) property it has, or on some epistemic property of another agent in the group of agents. Hence we study epistemic protocols for gossip. These kinds of protocols are relevant in many real world peer-to-peer networking applications involving autonomous agents, for which a clear understanding of the underlying information and knowledge dynamics is needed.

In our work [2] we introduce a variant of Dynamic Epistemic Logic (DEL) for describing epistemic gossip protocols, we introduce semantic objects and semantics for interpreting such protocols and give various logical properties of these protocols. We investigate properties such as expected execution length; termination (that is, whether the protocol succeeds in fully distributing all the secrets among the agents); and we develop a framework under which such properties can be described and checked.

### References

1 S.M. Hedetniemi, S.H.; Hedetniemi, S.T.; Liestman, A.L., *A survey of gossiping and broadcasting in communication networks*, Networks, 18, 319–349, 1988.
2 Attamah, M.; van Ditmarsch, H.; Grossi, D.; van der Hoek, W., *Knowledge and Gossip*, under submission.

# 4 Working Groups

## 4.1 Working Group on "Applications of epistemic planning in the real world" (APPL)

*Ron Petrick and Christian Becker-Asano*



### 4.1.1 Epistemic planning applications: theory and practice

One of the key difficulties surrounding the current state of epistemic planning and its ability to be a useful tool for developing real-world applications is the inherent cultural differences that appear to exist between researchers in the formal logic and automated planning communities. On the one hand, the planning community should become more familiar with the variety of formal methods coming from the epistemic logic community, and what they offer to the traditional planning concerns of representation and computation, while on the other hand, the epistemic logic community should look to the representational and algorithmic concerns of the planning community in order to identify what features of formal theories could be useful in practice. Without this deeper understanding of concepts and methods a successful collaboration seems less likely.

Looking beyond the core formal logic and planning communities, it may also be helpful to engage or foster connections with other disciplines where epistemic planning can potentially play a key role:

- robotics, especially in the areas of human-robot interaction and social robotics
- computer gaming and simulation, especially the area of interactive storytelling
- multiagent systems
- cognitive science
- (social) psychology

Engagement with industry may also be possible, however, the wider "usefulness" of epistemic planning may rest on our ability to design, develop, and maintain tools that have the potential to be adopted beyond the core epistemic planning research community.

### 4.1.2 Relevant features for applied epistemic planning

The following list of key features was identified as most relevant to the goal of applying epistemic planning to real-world applications:

#### 4.1.2.1 High-order reasoning

Epistemic planning becomes increasingly relevant in cases which necessitate some level of nested reasoning (e.g., $K_a K_b P$). However, reasoning with nested knowledge operators becomes more difficult as the depth of such operators increases, with most complex reasoning in everyday situations often limited to at most depth three. Identifying and reconciling the distinction between such situations, and those of classical logical reasoning problems such as the Muddy Children example, is an important difference between theory and practice.

#### 4.1.2.2 Tractable

From a computational complexity perspective, a real-world problem must be modelled in such a way that lends itself to tractable reasoning and execution. This is particularly important if

the application is meant to be real-time compatible and involves an aspect of human-robot interaction where response time is a key concern. At the same time, the application scenario needs to be non-trivial from both a formal and practical view, to be of interest to both the logic and planning communities.

#### 4.1.2.3 Multi-agent planning

In the best case, the application scenario should support epistemic planning that involve multiple agents, since such scenarios are more likely to lend themselves to complex situations involving non-trivial epistemic properties. At the same time, with arbitrary groups of agents (e.g., beyond 5 agents) the problem become more difficult from an implementation point of view, especially if all agents are autonomous and have different sets of capabilities. It is also important to note that the case of a single agent should not be completely dismissed, and single agent epistemic planning scenarios may still provide a valuable testbed for both the theoretic and algorithmic approaches.

In addition to these central features, the following features should also be considered.

#### 4.1.2.4 Coalition logics

In order for multiple agents to reach a shared goal the integration and application of coalition logics in epistemic planning seems necessary. This also gives rise to questions of mechanism design and other game theoretic concepts.

#### 4.1.2.5 Distributed vs. centralised planning and coordination of multiple agents

From a distributed systems point of view, questions of synchronisation and knowledge sharing over the network might come into play. This might further complicate the development of robust, practical solutions due to the additional complications of assuming or establishing properties like common knowledge over a network.

#### 4.1.2.6 Knowledge of capability

Modelling first and, especially, higher order knowledge of capability opens up new application scenarios of interest to both core communities. Knowing that someone has the capability of doing a certain action enables new and interesting types of goals in planning which could be used in practical applications.

#### 4.1.2.7 Long-term interaction

The robotics community's increasing interest in long-term interaction with robotic systems provides a potentially useful testbed for the epistemic planning community. In particular, the type of repeated, multi-agent interaction required in such applications would clearly benefit from the availability of a solid framework built on epistemic planning.

#### 4.1.2.8 Adaptability

Any framework that is developed as a joint outcome of the two communities should be generic enough to be easily adaptable to new application scenarios, both within and without the core epistemic planning community.

Table 1 displays some of these relevant features and identifies particular properties of the application design process which are potentially affected by these features.

**Table 1** Relevant features and particular properties of the application design process which are potentially affected by these features.

| Desired Features | | | | |
|---|---|---|---|---|
| | Represent. | Reasoning | Comput. | Strategic |
| Multiple agents | X | | | |
| Common/distrib. knowledge | X | X | | |
| High order (e.g. depth 3 nesting) | X | | | |
| Complex goals (e.g. nesting negation, etc.) | X | | | |
| Knowledge (e.g. capability) | X | X | | |
| Coalitions | | X | | |
| Tractability | | | X | |
| Decidability | | | X | |
| Computability | | | X | |
| Long-term interaction | | | | X |
| Adaptability | | | | X |

### 4.1.3 Applications

A set of possible applications were identified as potential testbeds for epistemic planning, which attempt to look beyond the type of standard toy benchmark problem which is common in many communities:

- Strategy development
  - Narrative generation
  - Strategic game playing
  - Business strategy planning
- Navigating social domains
  - Social robots
  - Conversational agents
  - Tutoring systems
- Security and enforcement
  - Network security
  - International and organised crime detection
- Internet-of-things
- Role playing games
- Algorithmic game theoretic applications

### 4.1.4 Example

As a specific example we consider an international crime scenario and the type of reasoning that is common in this case.

- Jones is trying to track down the many people who were accomplices in the latest bombing at a Baghdad market.
- What Jones knows:
  - X or Y detonated the bomb.
  - W1 shipped the detonator to X or W2 shipped the detonator to Y
- What Jones needs to know:
  - X or Y?
  - W1 or W2?

- How Jones can find out:
  - Check database to find out the characteristics of bomb which will let Jones know whether X or Y.
  - Visit a chat room to find out whether W1 or W2 has been shipping detonators.
  - Constraint: Jones must not let anyone know that he has found out (knows) that W1 (W2) has shipped detonator

**The challenge**: Can we develop a language or theory that supports this type of reasoning? Can we apply automated planning to this problem?

## 4.2 Working Group on "Establishing benchmarks and concrete goals for epistemic planning" (BENCH)

*Bernhard Nebel and Hans van Ditmarsch*



The goal of this working group was to identify a set of (scalable) examples and benchmarks that could be used to demonstrate the usefulness of epistemic planning, to inspire work on adjusting the expresseiveness of epistemic planning formalisms, and most importantly to inspire the design of new methods and systems that are able to solve epistemic planning problems. In the long run, we hope that our work leads to the creation of a new branch in the international planning competition.

Of course, it would be useful to spell out the benchmarks using a formal logical/planning language. However, for now, we decided to have simple sketches of such problems using natural language.

As epistemic planning problems, we understand here planning problems which involve reasoning about epistemic states, even when the goals of such planning problems are not epistemic. For example, when playing Cluedo, winning the game does not necessarily involve an epistemic goal (similarly for other games). However, for solving such games, reasoning about epistemic states is either necessary or profitable.

### 4.2.1 Classification / Dimensions

Differently from classical planning, there is a large number of dimensions along which benchmarks could be classified. These dimensions include:

- **Number of agents**: The main consus was that for epistemic planning one would need at least 2 agents, whereby perhaps only one agent is able to act. However, during the plenary discussion also arguments came up for the need of epistemic reasoning in the single-agent case.
- **Cooperative versus adversarial setting**: Scenarios could be that a group of agents tries to achieve a common goal or that each agent might want to achieve its own goal, which might conflict with the goals of the other agents.
- **Centralized versus distributed**: Plans can be global in the sense that they spell out the actions for all the agents (and are generated by a central system) or plans could be generated decentralized by each agent in isolation. As mentioned in the discussion, this difference might vanish when we consider online planners that only determine the next action. Furthermore, a question is how many agents we are able to control; the other agents (e.g. a customer in a Helpdesk setting) choose their actions non-deterministically.

- **Turn-taking versus simultaneous action**: Actions might take place simultaneously or there might be a turn-taking mechanism.
- **Simple epistemic knowledge, higher order knowledge, common knowledge**: The nesting depth of epistemic propositions can be limited to 1, to any fixed n, or it can be unbounded. Furthermore, we may have common knowledge (or even relativized common knowledge).
- **Communication & physical actions**: Actions can be only physical actions, which can be observed by the other agents, or only communicating actions (private or public announcements), or both.
- **Lying vs. only truthful statements**: Most work in DEL assumes that all communications by agents are truthful. However, there are, of course, interesting cases where (potential) lying plays an important role and where the goal can be to uncover lies of one agent.
- **Puzzle-mode versus shallow reasoning**: While all the dimensions mentioned so far can be connected to language features or restrictions in the formulation of the problem, this dimension is not as easy to capture. What is meant here is that an epistemic planning problem might be a difficult to solve puzzle even for humans (such as the muddy children puzzle), while other problems can be solved by any ordinary human quite easily. And we might want to address the latter ones in the beginning.

### 4.2.2 Benchmarks

The ideas for benchmarks could be classified into two broad sets. The first set consists of general frameworks, where we believe epistemic planning could be applied and where suitable specializations of the general idea might lead to concrete benchmark instances. In other words, there is probably a lot of work necessary to produce concrete benchmarks. The second set is more concrete, but still need significant work in order to produce a set of scalable benchmark problems.

#### 4.2.2.1 General ideas for benchmarks

- (Video) games with joint action – perhaps even only small portions of the games
- Help desk (modeling beliefs of customers)
- Soap opera planning (or narrative planning in general)
- Agents meeting in a grid with partial observability and communication

#### 4.2.2.2 Specific benchmark ideas

- **(Knowledge-based) gossip protocols**, with shortest or longest execution sequences (where longest execution sequences correspond to a very basic form of soap opera, people exchanging secrets until everybody knows everything, drawing out the process for as long as possible):
  Gossip protocols are protocols for peer-to-peer communication in networks, popularly known as 'telephone calls'. In network theory gossip protocols have been investigated in depth since the 1970s, where the paradigm is that a global scheduler assigns calls to pairs of agents. On the assumption that each agent or process has a local state of information, its 'secret', the goal is for all agents to know all secrets as quickly as possible. When calling each other, agents exchange all the secrets they know. For n agents each holding a secret, the minimum is $2n - 4$, and the maximum (on the assumption that actual information growth takes place during a call) is $n(n-1)/2$. There are also stochastic / probabilistic approaches. Now assume that we switch from a global scheduler to agent-based scheduling

(assuming some random selection of the next agent to make a call, simulating everybody rushing to reach the phone to make a call but only one agent succeeding). For four agents a, b, c, d with secrets A, B, C, D, the shortest protocol to distribute all secrets is ab;cd;ac;bd. For that, it is essential that caller c of the second call calls d and not a or b, which would also be informative. Now if we assume that, as c was not involved in the first call, c cannot distinguish a from b and d (the first call could, from c's perspective, have been between a and b, or a and d, or b and d), c has no reason to prefer calling d over a and b. If the second call is ca instead of cd, the shortest sequence to distribute all secrets now has length 5, not the minimum 4 (observe that d, who now only knows his own secret, has to call or be called by three other agents). We can consider knowledge-based gossip protocols where agents call other agents based on their knowledge, e.g.: (i) call an agent whose secret you do not know, or, (ii) call an agent such that you consider it possible to learn a new secret from that agent, or, (iii) call an agent such that you know that you will learn a new secret from that agent, or (iv) call an agent that you consider possible (or likely) to know many secrets already. (For another example, given the first two calls ab;ac, b now can call a again and learn the secret C in that call. So, b may have a reason to call another agent whose secret he already knows. That other agent may have learnt other secrets in the mean time.) Similar consideration abound in network theory, but not from an agent-based planning perspective. Not much is known about synthesizing such knowledge-based gossip protocols. What is there expected execution length? What is the minimum length of a knowledge-based gossip protocol? A wealth of variations investigated in network theory can be similarly modelled for knowledge-based gossip protocols. Postconditions may be general knowledge of all secrets (as above), but also, on the assumption of common knowledge of the protocol, common knowledge of all secrets. Yet other variations consider different agents following different protocols. Who learns all secrets first? What is the best choice of protocol, and against what other protocols executed by other agents? A standard reference on gossip protocols is [1].

- **Extended Wumpus** world with multiple agents, decentralized, with constraints on communication and sensing abilities
- **BW4T**: Blocksworld for teams
- **Epistemic blocksworld** (Thomas Bolander), where blocks have inscriptions only part of which are visible for each agent.
- **Communicating prisoners**: Agents (prisoners), who want to communicate privately without revealing anything to other agents (guards), where there are constraints on when communication can be tapped into
- **Active muddy children**: A variation of the muddy children problem, where the children are allowed to communicate
- **Cluedo**: Here we may want to consider the game in its full complexity as well as restricted scenarios and simplifications.

  Cluedo (for Americans, Clue), is a murder-mystery board game wherein six partying guests are confronted with a dead body, and they are all suspected of the murder. The game board depicts the different rooms of the house wherein the murder is committed, and there are also a number of possible murder weapons. Six suspects (such as Professor Plum), nine rooms, and six possible murder weapons. These options constitute a deck of 21 cards, one of each kind is drawn and are considered the real murderer, murder weapon, and murder room. The other cards are shuffled (again), and distributed to the players. The game consists of moves that allow for the elimination of facts about card ownership, until the first player to guess the murder cards correctly has won. When on the game

board a room is reached by a player, that player may then voice a suspicion, such as 'I think Ms. Scarlett did it, with a knife, in the kitchen'. This question is addressed to another player and interpreted as a request to admit or deny ownership of these cards for that player. If the addressed player doesn't have any of the requested cards, she says so, but if a player holds at least one of the requested cards, she is obliged to show exactly one of those to the requesting player, and to that player only. The four other players cannot see which card has been shown, but of course know that it must have been one of the three. Both denying ownership (i) and showing a card (ii) are epistemic actions, where the action of showing a card has 'interesting' epistemic postconditions, that may also increase the modal complexity of the underlying model. The question about card ownership is successively addresses to all other players until one of them answer the question (with a (i) or (ii) action.) During his turn a player may also make a guess for the murder cards (iii) – this can be done once in the game only, and is not done by saying it aloud but by writing the three cards on a piece of paper and then checking with the three cards on the table (the murder cards) if the guess is correct. If the guess is incorrect, that player has lost and the game continues. (If the guess had been said aloud, then even with an incorrect guess the game would now be over.) A final epistemic action is that of passing on your turn to the next player (iv), which is, on the assumption of perfect rationality, an admission of ignorance of the actual murder cards. For a knowledge analysis it is common to abstract from the aspect of the game board and thus allow questions about any room, and also only to allow correct guesses of the murder cards, that is, you win when you know what the murder cards are. There are a number of implementations of Cluedo and a Dagstuhl 'Planning with Epistemic Goals' participant currently working on one is Tiago de Lima, CRIL, Lens.

- **PIT**: Another game. Pit is another example of a card games with logical dynamics involving (strict) subgroups of all players. In the Pit game (for trading pit – it's a market simulation card game) the players try to corner the market in coffee, wheat, oranges, or a number of other commodities, and it is like the 'Family Game' in that each of these commodities are distributed over the players in the form of cards, and the first player to gather a full suit of cards (i.e., nine cards) of any commodity, wins. The game moves consist of two players exchanging cards. (In other words, unlike Cluedo, it is a planning problem with information change as well as ontic/factual change.) This goes as follows. A requirement to exchange cards is that they are of the same suit. Players shout the number of cards they wish to exchange, simultaneously, and two playing shouting the same number may then make a change. For example, John has 2 apples, 3 oranges, and some other cards, Mary has 2 oranges and yet other cards. John could have shouted 1, 2, or 3 (changing some but not all of the cards of the same suit is also allowed), but goes for 2, and Mary goes for 2 as well. Shout, shout, ... And they make an exchange. John now has 5 orange cards! Still not 9, but better than before. The exchange action is somewhat similar to the move of showing a card in Cluedo: two players gain subgroup common knowledge, in this case, of the new ownership of the exchanged cards. The other players only learn that two players each have at least two cards of the same suit. This rules out some card distributions. This exchanging of cards continues until somebody gathers his suit of nine cards. This can, in principle, go on forever: it is an extensive game of imperfect information, with infinite (but highly repetitive, there is much symmetry) game tree branches. For more information, see [2, 3].

**References**

1 Hedetniemi, S.M.; Hedetniemi, S.T.; Liestman, A.L., *A survey of gossiping and broadcasting in communication networks*, Networks, 18:4, 319–349, 1988.
2 van Ditmarsch, H., *The Logic of Pit*, Synthese, 149:2, 343–374, 2006.
3 van Ditmarsch, H., *Some game theory of Pit*, Proc. of 8th PRICAI, LNAI 3157, 946–947, 2004.

## 4.3 Working Group on "Planning with Epistemic Goals: Complexity of the task" (COMP)

*Jan van Eijck and Thomas Bolander*



### 4.3.1 Definition of 'Planning with Epistemic Goals'

Paradigm for generic planning task: find your way in a grid towards a specified goal. Paradigm for single-agent generic epistemic planning task: find your way in a grid towards a specified goal, given that your initial position is unknown. Paradigm cases for multi-agent generic epistemic planning tasks: depend on the nature of the other agents (co-operative, non-cooperative, mixed), but can also be modelled as navigation tasks in grids.

Different kind of goals: find your goal in the worst case in the minimum number of steps (optimal planning vs fig).

Background assumption in any planning task is some kind of closed world assumption: the agent knows (or the agents know) the limited repertoire of things that can happen. In an epistemic context, the closed world assumption and your goal without violating any constraints (satisficing planning will need a more sophisticated formulation than that of Ray Raiter ("Whatever is not currently known is false"). Instead of this: "Whatever is not currently known about ... satisfies the following constraints ...."

This is closely related to the classical frame problem in AI: how to model the fact that things in the environment do not change arbitrarily? In an epistemic context, this becomes: how to model the fact that agents know that things in the environment do not change arbitrarily?

### 4.3.2 Looking for fragments with decent complexity

#### 4.3.2.1 Brief Reminder

The plan existence problem for planning with epistemic goals has, more often than not, an unmanageable complexity. Already single-agent S5 planning with partial observability is 2EXP-complete, and multi-agent planning is undecidable (even with only 2 agents and no ontic actions). It is therefore essential to consider various restrictions that allow the complexity of epistemic planning to be controlled, so that it can become practically feasible. Below we will consider restrictions obtained by considering fragments of epistemic planning: restrictions on representation, reasoning, plan types.

#### 4.3.2.2 Representational Constraints

A first way to constrain the planning fragments is by putting constraints on the representation used to specify actions, goals, preconditions and postconditions.

- Possible restrictions on action types: public/private/partially observable announcements, public/private/partially observable propositional assignments, sensing actions (specified in simple terms), propositional actions, actions that don't produce exponential growth of the input epistemic model (**safe** actions).
- Possible restrictions on goal types: goals of limited depth, only positive goals.
- Another possible restriction is by putting a bound on the epistemic depth of reasoning, e.g., by only considering models up to depth $k$.
- Finally, it is possible to impose frame conditions: S5, K45, KD45, and so on.

Question: What action types do we even want to consider? Relates to (APPL) and (BENCH).

#### 4.3.2.3 Limited reasoning

- Limited models of belief.
- Alternative logical bases (e.g. 3-valued logic).

#### 4.3.2.4 Limited plan types

- E.g. existence of plans of *polynomially bounded length*.
- *Protocol restrictions* (what actions are allowed when).
  Some examples: Don't allow actions increasing depth beyond ... Do not allow exact same action twice. Make announcements whenever depth of ignorance is becoming too large.

### 4.3.3 Pragmatic Approaches to Deal with High Complexity

How to get good practical running times despite high complexity of the general task?

- Compact representations (OBDDs, symmetry reduced models, ...)
- Heuristics (delete-relaxation, abstraction, assumption of full observability, pruning techniques/preferred operators, ...)
- Online planning, replanning
- Learning techniques, sampling (UCT, ...)
- Alternative encodings/compilations (SAT, QBF, translation to classical planning, ...)
- Domain-dependent planning (HTNs, Golog, ...)
- Search approaches: forward search (with or without backtracking), regression (backwards from the goal).

### 4.3.4 Methodological Choices

| | qualitative | quantitative |
|---|---|---|
| logical language based | | |
| without logical language | | |

- Logical Approach vs Non-linguistic Approach: do we have a formal language to specify the goal?
- Quantitative methods versus qualitative methods: do we use probability or not?
- Markov decision processes versus conformant/contingent planning (all the logical approaches).
- Logical approaches: syntactic vs semantic.
  - Does the logical language allow talk about probabilities? (E.g., DEL vs Probabilistic DEL)

  - Syntactic approach: Knowledge bases as sets of formulas. In this category: situation calculus.
  - Semantic approach: satisfiability checking and model checking. In the model checking category: DEL update with result model checking.
- Offline planning vs online planning. Online planning may or may not use probabilistic methods. All non-deterministic planning is most often implemented as online planning, because of the combinatorial explosion of action outcomes (except conformant planning which is non-deterministic planning without any observability, and hence no feedback from actions to guide replanning).
- Navigation tasks with uncertainty about the location of the agent in a grid seem to call for online planning: plan consists of a repertoire of actions to carry out, and the plan evolves on the basis of the result of these actions.

### 4.3.5 A First Research Goal

An obvious first item on the list of research goals is the development of a uniform framework to list, classify and compare epistemic planning problems. A possible approach: fit all problems in the same mould by transforming them all into model checking problems using some expressive language $L$. A candidate language for this could be epistemic PDL.

Next, the problems can be classified according to criteria like the following:

- Number of agents, nature of their observational powers
- Properties of the models: classification according to deterministic or not, but also: knowledge or belief models, etc.
- Language to specify the preconditions, postconditions and goal(s): fragments of $L$.
- Format of the plan: just a sequence of actions, or conditionals allowed, or even loops? How does the language for preconditions, postconditions and goals figure in the plan format?

On the basis of this classification, complexity can be properly compared.

## 4.4 Working Group on "Languages for Epistemic Planning" (LANG)

*Thomas Ågotnes*



### 4.4.1 Epistemic planning in two communities

Language plays at least three roles in (epistemic) planning:

1. Describing a domain, in terms of actions and their epistemic pre- and postconditions.
2. Describing an initial (epistemic) state and goal condition. A planning domain together with an initial state and goal condition defines a particular instance of a planning problem to be solved.
3. Describing the result of planning: plans.

This gives us at least two variants of the epistemic planning problem. In type 1 epistemic planning the actions have epistemic pre- and postconditions, but the goal is not epistemic. In type 2 epistemic planning also the goal is epistemic. For planning algorithms this distinction does not seem to be important.

The research goals and cultures of the planning and logic communities differ significantly. Logic research focuses on expressive power and considers a wide range of sometimes complex phenomena; in planning simplicity and computational tractability are important. It highly depends on the requirements of the planning competition (IPC) which language features are widely supported by planning systems.

Planning problems are usually modelled as single-agent problems, rather than multi-agent. There are arguments that single-agent epistemic planning is not interesting, but it appears that there are counter-examples to this (for example, a robot could achieve the goal of knowing whether the vase is broken by dropping it, which is counterintuitive). For the multi-agent case, again, there are several variants. In the simplest variant only one agent can act but these actions can possibly affect the beliefs of other agents. His goals might involve epistemic properties of several agents. In more complex scenarios more than one agent can act.

### 4.4.2 Application Example

Devising epistemic planning languages heavily depends on the kind of applications that one wants to support. In order to decide on which features such a language would have to support, one would have to identify case studies of application scenarios where epistemic planning would be useful. Suppose a domestic robot in an office environment with several people (i.e. other agents) in it. The goal of the agent is to get to know which and how many people there are in the building, introduce itself to them, and achieve that everybody knows (common knowledge?) that lunch will be at 12:00 in the dining hall.

- is this sufficient as an interesting scenario?
- what are the kinds of epistemic goals that we should be able to represent?
- about what parts of the representation there could be possibly uncertainty (planning agent's knowledge, other agent's knowledge, preconditions and effects of actions,...)?
- how to represent/encode sensing, communication etc.?

Multi-agent planning is particularly interesting in adversarial scenarios, which the above is not. Alternatively, a surveillance/security robot could be considered that has to deal with intruders. Moreover, games (General Game Playing?) could be an interesting application area.

### 4.4.3 How can epistemic planning benefit from epistemic logic research?

A systematic approach: logic provides a systematic and generic framework which allows to compare and relate more easily the impact of the different parameters chosen (time dependencies between events, epistemic goals, adding probability,...) by planning researchers on the properties of a planning scenario (memoryless, non-deterministic effects, conditional, looping) as well as on the computational complexity of solving a planning task.

In epistemic logic a number of epistemic concepts that are more sophisticated than individual knowledge of propositions have been developed and studied:

- Higher-level knowledge
- Common knowledge, distributed knowledge and other types of group knowledge

Logic can also provide a more refined representation of uncertainty by means of:

- Probabilities (also interesting in the single-agent case)
- Qualitative representation: plausibilities, possibilities, rankings,... (also interesting in the single-agent case)

Logic studies phenomena which are not yet addressed in the planning community such as awareness.

All these above features can be used to increase the applicability of planning formalisms.

Moreover, the omniscience problem of epistemic logic is likely to arise in reformulations in logic of planning problems. It is likely that one of the large number of solutions dealing with the omniscience problem in epistemic logic will be applicable to planning problems.

These concepts are relevant both for expressing pre- and post-conditions of actions and epistemic goals. The epistemic logic community has developed a theoretically well-founded foundation for these concepts. Their computational properties are also well understood. Even though these concepts have not been considered in existing case studies, the planning community doesn't necessarily know what can be useful!

How does one express goals about things the agent is not aware of? And this is interesting even in the single-agent setting! E.g., the Mars rover must learn about all new green creatures it finds out there! How is finding out what others are trying to achieve expressed as a goal?

Logic has also developed a number of model theoretical tools that can possibly be applied to planning. An example is techniques for comparing models, which can be applied to characterise when planning problem descriptions are the same. It is well known that planning is sensitive to framing, i.e., representation of the problem.

One challenge is that classical planning formalisms are based on states as variable assignments. From a theoretical perspective, this is not expressive enough to model all epistemic multi-agent settings.

### 4.4.4 Some research challenges

- Find case studies with epistemic goals involving complex epistemic concepts (higher-level knowledge, common knowledge, etc.).
- Can existing planning description languages (e.g. PDDL, GDL-II) be extended with useful epistemic features in a natural way? (For instance, single-agent S5 DEL planning can implement NPDDL. What multi-agent features can be added to NPDDL? Can these be implemented in multi-agent S5 DEL?)
- Can these extended languages be implemented in epistemic planning frameworks (DEL-based, SitCalc-based, EventCalc-based)? What is the corresponding expressivity and/or complexity?
- How can we define an appropriate language for action models that would allow an easy specification of actions for planning problems?
- How can the expressive power of planning languages be characterised using the epistemic logic framework?
- How can the concepts and methods dealing with succinctness, compactness and elaboration tolerance from the planning community be imported within the epistemic logic framework?
- Can logical techniques for comparing models (simulations, bisimulations, model comparison games) be used to answer the question of when two planning problem descriptions are the same?
- Can planning formalisms be augmented by external "black-box" epistemic reasoning?
- A 'strong' plan guarantees that it will always achieve the goal, while a 'weak'-plan just says that it is possible. Is a logic of degrees of belief a suitable formalism for grading guarantees, as in 'this plan will achieve the goal most of the time', or 'I believe (though I do not know) that this plan will achieve the goal'? Is the complexity of synthesising a plan that is 'believed to work' better or worse than synthesising a strong or weak plan?
- What is the exact and formal connection between the knowledge-based programs of the epistemic logic community and the plans of the planning community?

## Participants

- Thomas Ågotnes
University of Bergen, NO
- Maduka Attamah
University of Liverpool, GB
- Guillaume Aucher
INRIA Rennes – Bretagne Atlantique, FR
- Christian Becker-Asano
Universität Freiburg, DE
- Mikkel Birkegaard Andersen
Technical University of Denmark – Lyngby, DK
- Thomas Bolander
Technical University of Denmark – Lyngby, DK
- Jens Claßen
RWTH Aachen, DE
- Tiago de Lima
Artois University – Lens, FR
- Carmel Domshlak
Technion, IL
- Hector Geffner
UPF – Barcelona, ES
- Valentin Goranko
Technical University of Denmark – Lyngby, DK
- Malte Helmert
Universität Basel, CH
- Andreas Herzig
Paul Sabatier University – Toulouse, FR
- Jörg Hoffmann
Universität des Saarlandes, DE
- Martin Holm Jensen
Technical University of Denmark – Lyngby, DK
- Gerhard Lakemeyer
RWTH Aachen, DE
- Jerome Lang
University Paris-Dauphine, FR
- Yongmei Liu
Sun Yat-sen University – Guangzhou, CN
- Benedikt Löwe
University of Amsterdam, NL
- Robert Mattmüller
Universität Freiburg, DE
- Sheila McIlraith
University of Toronto, CA
- Leora Morgenstern
New York University, US
- Bernhard Nebel
Universität Freiburg, DE
- Ron Petrick
University of Edinburgh, GB
- Gabriele Röger
Universität Basel, CH
- Francois Schwarzentruber
IRISA – Rennes, FR
- Sunil Easaw Simon
Indian Institute of Technology – Kanpur, IN
- Hans Van Ditmarsch
LORIA – Nancy, FR
- Jan van Eijck
CWI – Amsterdam, NL
- Yanjing Wang
Peking University – Beijing, CN

Report from Dagstuhl Seminar 14041

# Quantitative Models: Expressiveness, Analysis, and New Applications

**Edited by**

**Manfred Droste[1], Paul Gastin[2], Kim Guldstrand Larsen[3], and Axel Legay[4]**

1 **Universität Leipzig, DE,** droste@informatik.uni-leipzig.de
2 **ENS – Cachan, FR,** Paul.Gastin@lsv.ens-cachan.fr
3 **Aalborg University, DK,** kgl@cs.aau.dk
4 **INRIA – Rennes, FR,** alegay@irisa.fr

**Abstract**

From Jan. 19 to Jan. 24, 2014, "Quantitative Models: Expressiveness, Analysis, and New Applications " was held in Schloss Dagstuhl – Leibniz Center for Informatics. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

**Seminar** January 19–24, 2014 – http://www.dagstuhl.de/14041
**1998 ACM Subject Classification** D.2.4 Software/Program Verification, F.1.1 Models of Computation
**Keywords and phrases** quantitative models, quantitative analysis, timed and hybrid systems, probabilistic systems, weighted automata, systems biology, smart grid
**Digital Object Identifier** 10.4230/DagRep.4.1.104
**Edited in cooperation with** Thomas Weidner

## 1 Executive Summary

*Manfred Droste*
*Paul Gastin*
*Kim Guldstrand Larsen*
*Axel Legay*



Quantitative models and quantitative analysis in Computer Science is receiving increased attention in order to meet the challenges from application areas such as Cyber Physical Systems. What is aimed at is a revision of the foundation of Computer Science where Boolean models and analyses are replaced by quantitative models and analyses in order that more detailed and practically useful answers can be provided. Recently, a large number of new models, toolsets, and new application domains have emerged. The theory of weighted automata has also developed, introducing extensions of the models which are motivated by the quantitative analysis of systems.

The first objective of the seminar was to bring the quantitative model checking and weighted automata communities together with the goal of discussing the latest developments


Quantitative Models: Expressiveness, Analysis, and New Applications, *Dagstuhl Reports*, Vol. 1, Issue 1, pp. 104–124
Editors: Manfred Droste, Paul Gastin, Kim Guldstrand Larsen, and Axel Legay

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

in those areas. The second objective of this workshop was to go one major step further. In fact, it has been recently observed an increasing usage (and demand) of quantitative models in a wide range of new application domains. This includes, e.g., systems biology and energy grid. However, these different communities are often not aware of each other. This seminar had the major objective to put those various communities in contact with the hope of creating fruitful long term collaborations.

Quantitative model checking covers extended automata-based models that permit to reason on quantities. The model of timed automata introduced by Alur and Dill in 1989 has by now established itself as a universal formalism for describing real-time systems. The notion of zone has led to a number of tools – e.g. BIP, Kronos, UPPAAL – which support efficient analysis (reachability and model checking) of timed automata. Later the more expressive formalism of hybrid automata was introduced and popularized by Henzinger et al and the introduction of the tool HyTech provided a semi-decision algorithm for analyzing so-called linear hybrid systems. Whereas in timed automata the continuous part of a model is restricted to be clocks (which always evolve with rate 1), linear hybrid automata allow more general continuous variables with evolution rates in arbitrary intervals. The notion of priced (or weighted) timed automata was introduced independently by Alur et al and Larsen et al in 2001, with the surprising result that cost optimal reachability is decidable. Since these initial results, efficient tools were developed and a number of more challenging questions have been considered including multi-priced timed automata, optimal infinite scheduling (both with respect to mean pay-off and discounting), priced timed games and model checking for priced timed automata.

Driven by new needs in areas such as cyber physical systems, a series of recent work have tried to combine real-time with stochastic aspects, leading to new models such as timed stochastic automata. One of the main objectives of the seminar was to study those new models and put them in perspective with similar results in weighted automata. The new notion of energy automata (Larsen, Markey, Bouyer, ...) that extends price timed automata and permits to reason on energy problems was also discussed and put in perspective with similar work done at the weighted automata level.

Weighted automata on finite words were already investigated in seminal work of Schützenberger (1961) and Chomsky and Schützenberger (1963). They consist of classical finite automata in which the transitions carry weights which may model, e.g., the cost, the consumption of resources, or the reliability or probability of the successful execution of the transitions. This concept soon developed a flourishing theory. Recently, motivated by practical examples of energy consumption, new quantitative automata models have been introduced and investigated in which the weights of finite or infinite paths are computed e.g. by the average weights or by the accumulation points of the average weights of their transitions. Colcombet (2009) studied regular cost functions which permit a quantitative extension of classical equivalence results relating automata, expressions, algebraic recognizability, and variants of monadic second-order logic. Gastin et al (2010) introduced weighted pebble automata in order to capture the expressive power of weighted extensions of Xpath for XML documents, or temporal logics for linear behaviors. All these concepts provide totally new models for which weighted automata-theoretic methods can often be applied successfully. It was very profitable therefore to bring these different communities together.

Another main theme of the seminar was to create interaction with researchers working in areas where the theoretical models and techniques may have potential applications. In systems biology, the challenge is not only to find mathematical models, but also to define new efficient quantitative analysis techniques capable of coping with very large size complex

systems. Two promising applications are 1) using SMC-based techniques to monitor complex properties that cannot be expressed in classical temporal logic (e.g., oscillation properties), and 2) using interface theories as a formal characterization of phenomena in the area of synthetic biology. As another application area, the challenge of smart electricity grids is to balance the behavior of all participants (suppliers and consumers) to improve efficiency and stability. Again, quantitative models such as energy automata and analysis are emerging as potential key techniques.

In the seminar, 40 researchers from 13 countries discussed their recent research results and developments for quantitative models and their analysis. Five survey lectures, including two lectures covering the application domains, and 32 talks were organized in eight sessions with centralized themes. From the beginning, all lectures and talks raised questions of members from the other fields, and lively discussions followed. In particular, the surveys presented the fields of weighted automata, formal model checking and simulation methods adopted by industry, programmable single-cell biocomputers, models for smart grid balancing, and asymptotic analysis of weighted automata. The lectures and talks dealt with, e.g., quantitative logics and their semantics, expressiveness of models including quantitative measures for infinite behavior (like discounting, mean payoff, long-run averages), and statistical model checking of stochastic hybrid systems, to name only a few topics.

There are a number of open problems concerning the interplay between these fields. For instance, there are many interesting open questions about the connection between energy automata, energy functions and weighted automata, on weighted specification languages used in more algebraic settings, on energy games, and on the combination of real-time and probabilism. The interplay between priced timed automata and weighted automata also demands further investigation. Due to these open challenges, several researchers decided to meet again later in the year, e.g. during the international workshop in Leipzig on "Weighted Automata: Theory and Applications (WATA 2014)".

During the seminar, there was very much interaction between the participants. In particular, the seminar was successful in attracting academic researchers with contacts to industry; this was felt very positive and should definitely be continued. Generally, it was expressed that a future research collaboration between the different present groups should be highly fruitful and would therefore be very desirable. A Dagstuhl seminar would provide an ideal and unique opportunity for this. The successful collaboration in the present seminar was felt to be due in particular to the superb facilities and excellent organization provided by the Dagstuhl center and its team.

## 2 Table of Contents

# 3 Overview of Talks

## 3.1 Computing Behavioral Distances, Compositionally

*Giorgio Bacci (Aalborg University, DK)*


**Joint work of** Bacci, Giorgio; Bacci, Giovanni; Larsen, Kim G., Mardare, Radu;
**Main reference** G. Bacci, G. Bacci, K. G. Larsen, R. Mardare, "Computing Behavioral Distances, Compositionally" in Proc. of the 38th Int'l Symp. on Mathematical Foundations of Computer Science ( MFCS'13), LCNCs, Vol. 8087, pp. 74–85, Springer, 2013.
**URL** http://dx.doi.org/10.1007/978-3-642-40313-2_9

In the last years, behavioral metrics have received an increased attention as a theoretical tool for approximate reasoning on quantitative models. Realistic models are usually specified compositionally by means of operators that describe the interactions between the subcomponents. These specifications may suffer from an exponential growth of the state space making their analysis difficult to perform in practice. We show recent work on the practical benefits of the compositional reasoning for computing the bisimilarity distance of Ferns et al. between Markov Decision Processes with rewards (MDPs). We identified a well behaved class of operators, called safe, that are guaranteed to be non-extensive w.r.t. the bisimilarity distance on MDPs and we will show that, for MDPs built using safe/non-extensive operators, it is possible to exploit the structure of the system improving the performance on state of the art methods for (exactly) computing such distance.

## 3.2 On-the-Fly Exact Computation of Bisimilarity Distances

*Giovanni Bacci (Aalborg University, DK)*


**Joint work of** Bacci, Giorgio; Bacci, Giovanni; Larsen, Kim G.; Mardare, Radu;
**Main reference** G. Bacci, G. Bacci, K. G. Larsen, R. Madare, "On-the-Fly Exact Computation of Bisimilarity Distances," in Proc. of the 19th Int'l Conf. on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'13), LNCS, Vol. 7795, pp. 1–15, Springer, 2013.
**URL** http://dx.doi.org/10.1007/978-3-642-36742-7_1

We describe recent work on an efficient on-the-fly algorithm for exact computation of bisimilarity distances between discrete-time Markov chains. Unlike other existing solutions, our method is able to exactly compute the distances between given states avoiding an exhaustive exploration of the state space. Given a set of target states, our technique successively refines over- approximations of the distance using a greedy strategy which ensures that the state space is further explored only when this is actually needed for improving the current approximation on the given target. Tests performed on a consistent set of (pseudo)randomly generated Markov chains shows that our algorithm improves, on average, the efficiency of the corresponding iterative algorithms with orders of magnitude.

## 3.3 Fundamental Problems of Fuzzy Automata Theory

*Miroslav Ćirić (University of Niš – Serbia)*


**Joint work of** Ćirić, Miroslav; Ignjatovic, Jelena;

In this talk we will present the main ideas and methodology used in dealing with certain fundamental problems of the theory of fuzzy automata, such as equivalence, simulation and bisimulation, state reduction, determinism and determinization, etc. We have found that fuzzy automata can be successfully studied using a fuzzy relational calculus. In particular, it turned out that the basic problems in the study of simulation, bisimulation and state reduction can be reduced to the problems of solving some particular systems of fuzzy relation equations and inequalities. A key role in solving these systems play residuals of fuzzy relations, which naturally generalize residuals of ordinary Boolean relations introduced by Birkhoff in the 1940s. In order to ensure the existence of residuals it is needed that the structure of truth values is a complete residuated lattice, or a quantale (a complete residuated lattice which might lack commutativity). We will also show that some of the obtained results concerning simulation, bisimulation and state reduction can be extended to weighted automata over some types of semirings which allow residuation (max-plus algebras or min-plus algebras – tropical semirings) or relative residuation (additively idempotent semirings), whereas the results concerning determinization can be extended to weighted automata over arbitrary semirings, and even to weighted automata over strong-bimonoids (semirings which might lack distributivity). Finally, we will show how these results of fuzzy automata theory influenced the development of general methods for solving systems of fuzzy relation equations and inequalities, or even more generally, the development of methods for solving systems of equations and inequalities defined by residuated functions. We will also demonstrate how our methodology can be applied in other fields, such as social network analysis and formal concept analysis.

## 3.4 Approximations of Difficult Problems for Tropical Automata

*Thomas Colcombet (CNRS / Université Paris-Diderot)*



In this presentation I will consider the problem of comparing tropical automata (min-plus or max-plus, with non-negative weights). This problem is undecidable [1]. In this talk I will present three decidable approximations of this question:

- regular cost function, in which only the relative boundedness is considered, yielding a robust theory extending regular languages.
- the epsilon-approximation of the comparison of min-plus automata, in which the exact comparison of min-plus automata is decided up to an epsilon-multiplicative margin of error epsilon (collaboration with Laure Daviaud).
- the asymptotic analysis of max-plus automata, in which the asymptotic worst-case behaviour of a max-plus automaton is analysed, with an application to the analysis of the termination time of algorithms under the Size-Change-Abstraction (collaboration with Laure Daviaud and Florian Zuleger).

These research received funding from the European Union's Seventh Framework Programme (FP7/2007-2013) under grant agreement No. 259454 (project GALE).

### References

1 Daniel Krob. The equality problem for rational series with multiplicities in the tropical semiring is undecidable. Internat. J. Algebra Comput., 4(3):405–425, 1994

## 3.5 Equilibria in Multiplayer Cost Games

*Julie De Pril (University of Mons, BE)*


**Main reference** J. De Pril, "Equilibria in Multiplayer Cost Games," PhD thesis, Université de Mons, 2013.
**URL** http://math.umons.ac.be/perso/DePril.Julie/thesis_Julie_DePril.pdf

In order to model complex interactive computer systems with more than two components, and with quantitative objectives that are not necessarily antagonist, we resort to multiplayer non zero-sum quantitative games played on graphs (also called multiplayer cost games). Many parameters appear when studying cost games: the graph can be enriched with prices on edges or not; there can be two or more players; the objectives of the players can be very various and also complicated; different kinds of rational behaviour can be considered for the players, leading to different concepts of equilibria; ... In this talk, we define different kinds of equilibria in these games. For each kind, we present some existence results and state some open questions.

## 3.6 Quantitative Languages: Weighted Automata and Beyond

*Laurent Doyen (ENS – Cachan, FR)*



Quantitative generalizations of classical languages, which assign to each word a real number instead of a boolean value, have applications in modeling resource-constrained computation. We use nondeterministic weighted automata (finite automata with transition weights) to define classes of quantitative languages over infinite words. We study the natural generalization of decision problems and closure properties from automata theory: threshold-emptiness, threshold-universality, language inclusion, language equivalence, and closure under the pointwise operations $\max(L, L')$, $\min(L, L')$, $1 - L$ (which generalize the boolean operations), and sum $L + L'$ where $L$, $L'$ are quantitative languages. In this survey, we give an overview of the results and open questions about decidability, expressiveness (including determinization), and closure properties of the various classes of quantitative languages defined by weighted automata. As none of these classes enjoys both full decidability and positive closure properties (even if we consider an extension to alternating automata), we present mean-payoff automaton expressions, a new syntax to specify quantitative languages with limit-average value. We show that this class of quantitative languages is robust and decidable: it is closed under the four pointwise operations, and we show that all decision problems are decidable.

## 3.7 Weighted Automata and Quantitative Logics

*Manfred Droste (Universität Leipzig)*


**Joint work of** Paul Gastin resp. Ingmar Meinecke;

We investigate weighted automata and their relationship to weighted logics. For this, we present syntax and semantics of a quantitative logic; the semantics counts "how often" a formula is true in a given word. Our main result, extending the classical result of Büchi, shows that if the weights are taken from an arbitrary semiring, then weighted automata and a syntactically defined fragment of our weighted logic are expressively equivalent. A corresponding result holds for infinite words. Moreover, this extends to quantitative automata investigated by Henzinger et al. with (non-semiring) average-type behaviors, or with discounting or limit average objectives for infinite words.

## 3.8 Some Algebraic Structures for the Behavior of Quantitative Systems

*Zoltán Ésik (University of Szeged)*



I will discuss algebraic structures for the behavior of quantitative systems including Conway and iteration semirings, hemirings, semimodules, etc. Some open problems will be presented.

## 3.9 Kleene Algebras and Semimodules for Energy Problems

*Ulrich Fahrenberg (INRIA Bretagne Atlantique – Rennes)*


**Joint work of** Ésik, Zoltán; Fahrenberg, Ulrich; Legay, Axel; Quaas, Karin;

With the purpose of unifying a number of approaches to energy problems found in the literature, we introduce generalized energy automata. These are finite automata whose edges are labeled with energy functions that define how energy levels evolve during transitions. Uncovering a close connection between energy problems and reachability and Büchi acceptance for semiring-weighted automata, we show that these generalized energy problems are decidable. We also provide complexity results for important special cases.

## 3.10 Functional Weighted Automata

*Emmanuel Filiot (Université Libre de Bruxelles, BE)*


**Joint work of** Filiot, Emmanuel; Gentilini, Raffaella; Raskin, Jean-François;
**Main reference** E. Filiot, R. Gentilini, J.-F. Raskin, "Quantitative Languages Defined by Functional Automata," in Proc. of the 23rd Int'l Conf. on Concurrency Theory, LNCS, Vol. 7454, pp. 132–146, Springer, 2012; pre-print available as arXiv:1111.0862v2 [cs.FL].
**URL** http://dx.doi.org/10.1007/978-3-642-32940-1_11
**URL** http://arxiv.org/abs/1111.0862v2

We study several decision problems for functional weighted automata. To associate values with runs, we consider four different measure functions: the sum, the mean, the discounted sum of weights along edges and the ratio between rewards and costs. On the positive side, we show that the existential and universal threshold problems, the language inclusion problem and the equivalence problem are all decidable for the class of functional weighted automata and the four measure functions that we consider. On the negative side, we also study the quantitative extension of the realizability problem and show that it is undecidable for sum, mean and ratio. We show how to decide if the quantitative language defined by a functional weighted discounted sum automaton can be defined with a deterministic automata (it was already known for sum and mean). Finally, we discuss some extension to k-valued weighted automata.

## 3.11 Probabilistic Regular Expressions

*Paul Gastin (ENS – Cachan)*



We provide a Kleene Theorem for (Rabin) probabilistic automata over finite words. Probabilistic automata generalize deterministic finite automata and assign to a word an acceptance probability. We provide probabilistic expressions with probabilistic choice, guarded choice, concatenation, and a star operator. We prove that probabilistic expressions and probabilistic automata are expressively equivalent. Our result actually extends to two-way probabilistic automata with pebbles and corresponding expressions.

## 3.12 Compositional Metric Reasoning with Probabilistic Process Calculi

*Daniel Gebler (Free University of Amsterdam, NL)*


**Joint work of** Gebler, Daniel; Larsen, Kim G.; Tini, Simone;
**Main reference** D. Gebler, S. Tini, "Compositionality of Approximate Bisimulation for Probabilistic Systems," in Proc. of the Combined 20th Int'l Workshop on Expressiveness in Concurrency and 10th Workshop on Structural Operational Semantics (EXPRESS/SOS'13), EPTCS, Vol. 120, pp. 32–46, 2013.
**URL** http://dx.doi.org/10.4204/EPTCS.120.4

Probabilistic process calculi are algebraic theories to specify and verify probabilistic concurrent systems. Bisimulation metric is a fundamental semantic notion that defines the behavioral distance between probabilistic processes. We study which operators of probabilistic process calculi allow for compositional reasoning with respect to bisimulation metric semantics. We

propose continuity as property of process combinators that capture the essential nature of compositional reasoning for both non-recursive and recursive probabilistic processes. Moreover, we characterize the distance between probabilistic processes composed by standard process algebra operators.

## 3.13 Important Splitting for Statistical Model Checking

*Cyrille Jegourel (INRIA Bretagne Atlantique – Rennes)*



Statistical model checking avoids the intractable growth of states associated with probabilistic model checking by estimating the probability of a property from simulations. Rare properties are often important, but pose a challenge for simulation-based approaches: the relative error of the estimate is unbounded. A key objective for statistical model checking rare events is thus to reduce the variance of the estimator. Importance splitting achieves this by estimating a sequence of conditional probabilities, whose product is the required result. To apply this idea to model checking it is necessary to define a score function based on logical properties, and a set of levels that delimit the conditional probabilities.

In this talk we motivate the use of importance splitting for statistical model checking and describe the necessary and desirable properties of score functions and levels. We illustrate how a score function may be derived from a property and present an (adaptive) importance splitting algorithm that discovers optimal levels adaptively.

## 3.14 Inferring Partially Observed Markov Chains in Biology

*Heinz Koeppl (ETH Zürich)*



I will discuss the problem specifics of working with biochemical stochastic models to describe experimental single-cell data. In particular, this involves the incorporation of cellular heterogeneity into the mathematical formalism and I will show how to overcome the computational burden related to the inference of resulting models. Moreover, I will explain our work and ideas related to estimation of molecular states and its relation to conditional Markov processes and statistical model checking.

## 3.15 Assume-Guarantee Reasoning in Continuous-Time

*Jan Křetínský (TU München, DE)*


**Joint work of** Brázdil, Tomáš; Hermanns, Holger; Krčál, Jan; Křetínský, Jan; Řehák, Vojtěch;
**Main reference** H. Hermanns, J. Krčál, J. Křetínský, "Compositional Verification and Optimization of Interactive Markov Chains," in Proc. of the 24th Int'l Conf. on Concurrency Theory, LNCS, Vol. 8052, pp. 364–379, Springer, 2013.
**URL** http://dx.doi.org/10.1007/978-3-642-40184-8_26

We discuss assume-guarantee reasoning for and compositional verification of interactive Markov chains (a model with non-determinism and stochastic continuous-time). We discuss

several ways to interpret this task and some ideas for solutions specific for continuous-time. We also provide a specification formalism for these systems, namely a continuous and modal extension of timed automata.

### References

1 Tomáš Brázdil, Holger Hermanns, Jan Krčál, Jan Křetínský, Vojtěch Řehák: Verification of Open Interactive Markov Chains. FSTTCS 2012.
2 Holger Hermanns, Jan Krčál, Jan Křetínský: Compositional Verification and Optimization of Interactive Markov Chains. CONCUR 2013.

## 3.16 Patroling Games

*Antonín Kučera (Masaryk University – Brno)*


**Joint work of** Brázdil, Tomáš; Hliněný, Petr; Řehák, Vojtěch;

Patrolling is one of the central problems in operational security. Formally, a patrolling problem is specified by a set $U$ of vulnerable targets and a function $d$ which to every target $u$ assigns the (integer) time $d(u)$ needed to complete an intrusion at $u$. The goal is to design an optimal strategy for a defender who is moving from target to target and aims at detecting possible intrusions. The defender can detect an intrusion at $u$ only by visiting u before the intrusion is completed. The goal of the attacker is to maximize the probability of a successful attack. We assume that the attacker is adversarial, i.e., he knows the strategy of the defender and can observe her moves. We prove that the defender has an optimal strategy for every patrolling problem and every environment, and we show how to construct an optimal strategy efficiently.

## 3.17 Formal Analysis of Resource Contention in Multicore Architectures

*Kai Lampka (Uppsala University, SE)*


**Joint work of** Lampka, Kai; Georgia Giannopoulou; Nikolay Stoimenov; LotharThiele;
**Main reference** G. Giannopoulou, K. Lampka, N. Stoimenov, L. Thiele, "Timed Model Checking with Abstractions: Towards Worst-Case Response Time Analysis in Resource-Sharing Manycore Systems," in Proc. of the 10th ACM Int'l Conf. on Embedded Software (EMSOFT'12), pp. 63–72, ACM, 2012.
**URL** http://dx.doi.org/10.1145/2380356.2380372

The talk is concerned with the analysis of real-time constrained software executing on multicore processors with shared resources like caches, memory and inter-core connections. The presented results have been developed as part of the EU FP-7 STREP CERTAINTY [1]. The challenge inherent to the design of mixed-critical real-time system deployed on multi-core architectures is to organize the system in such a way, that the system behaviour can be analyzed precisely. This precision will help to avoid over- provisioning of the architecture and in turn reduces the waste of resources. This talk focus on the structuring of software systems to achieve time predictability, i.e, to reduce the non-determinism of occurrence of

events. Based on the proposed model of computation, the talk introduces an analysis scheme based on Timed Automata and model checking which is described in [2].

### References

1 EU FP 7 STREP: CERTAINTY (Certification of Real Time Applications designed for mixed criticality) http://www.certainty-project.eu
2 Georgia Giannopoulou, Kai Lampka, Nikolay Stoimenov and LotharThiele: Timed Model Checking with Abstractions: Towards Worst-Case Response Time Analysis in Resource-Sharing Manycore Systems. Proc. International Conference on Embedded Software (EMSOFT) 2012. pp. 63–72, ACM 2012.

## 3.18 Statistical Model Checking of Stochastic Hybrid Systems

*Kim Guldstrand Larsen (Aalborg University, DK)*



Timed automata, priced timed automata and energy automata have emerged as useful formalisms for modeling a real-time and energy-aware systems as found in several embedded and cyber-physical systems. Whereas the real-time model checker UPPAAL allows for efficient verification of hard timing constraints of timed automata, model checking of priced timed automata and energy automata are in general undecidable – notable exception being cost-optimal reachability for priced timed automata as supported by the branch UPPAAL Cora. These obstacles are overcome by UPPAAL-SMC, the new highly scalable engine of UPPAAL, which supports (distributed) statistical model checking of stochastic hybrid systems with respect to weighted metric temporal logic. The talk will review UPPAAL-SMC and some of its applications, e.g. to the domains of energy-harvesting wireless sensor networks, schedulability analysis for mixed criticality systems, as well as smart grids. In the talk I will also contemplate on how other branches of UPPAAL may benefit from the new scalable simulation engine of UPPAAL-SMC in order to improve their performance as well as scope in terms of the models that they are supporting. This includes application of UPPAAL-SMC to counter example generation, refinement checking, controller synthesis, optimization, testing and meta-analysis.

## 3.19 Timing Analysis of Parallel Software Using Abstract Execution

*Björn Lisper (Mälardalen University – Västerås, SE)*


**Joint work of** Gustavsson, Andreas; Gustafsson, Jan; Lisper, Björn;
**Main reference** A. Gustavsson, J. Gustafsson, B. Lisper , "Timing Analysis of Parallel Software Using Abstract Execution," in Proc. of the 15th Int'l Conf. on Verification, Model Checking, and Abstract Interpretation (VMCAI'14), LNCS, Vol. 8318, pp. 59–77, Springer, 2014.
**URL** http://dx.doi.org/10.1007/978-3-642-54013-4_4

A major trend in computer architecture is multi-core processors. To fully exploit this type of parallel processor chip, programs running on it will have to be parallel as well. This means that even hard real-time embedded systems will be parallel. Therefore, it is of utmost importance that methods to analyze the timing properties of parallel real-time systems are

developed. We present an algorithm that is founded on abstract interpretation and derives safe approximations of the execution times of parallel programs. The algorithm is formulated and proven correct for a simple parallel language with parallel threads, shared memory and synchronization via locks.

## 3.20 Definition of Star and Epsilon-Removal in Weighted Automata

*Sylvain Lombardy (Université Bordeaux)*



The removal of epsilon-transitions in weighted automata requires, except in some particular cases, to deal with a definition of a star operator in the semiring of weights. This operator must allow both a sound definition of validity of weighted automata and the computation of the epsilon-removal. Different approachs exist; the axiomatique method states axioms on semirings that guarantee the validity of weighted automata and the sound definition of the star operator. We present here another method based on the topology of the semiring, that allows to define the star as a sum; the semirings that can be handled this way includes all the common semirings, but in this framework, weighted automata may be not valid and decision algorithms must be designed.

## 3.21 Robustness of Timed Models

*Nicolas Markey (ENS – Cachan, FR)*


**Joint work of** Markey, Nicolas; Bouyer, Patricia; Sankur, Ocan; Reynier, Pierre-Alain; Fang, Erwin;

Timed automata are governed by an "idealized" semantics, assuming zero-delay transitions and infinite precision in the measure of time. This is not compatible with real-life systems, and small clock drifts can have important impact on the correctness of real-time systems, even when their abstract model has been proven correct. In this talk, I present several recent attempts that have been proposed to overcome this problem, taking into account perturbations in the semantics. I conclude with general perspectives for robust model checking of real-time systems.

## 3.22 Weighted Hybrid Logics

*Benjamin Monmege (Université Libre de Bruxelles, BE)*


**Joint work of** Bollig, Benedikt; Gastin, Paul; Monmege, Benjamin; Zeitoun, Marc;

The link between weighted automata and weighted logic has been investigated since several years. At first, a weighted extension of monadic second order logic (wMSO) has been introduced by Droste and Gastin, which appears to be much too powerful with respect to

weighted automata. However, a fragment of this logic appears to match the expressive power of weighted automata. In this talk, we will consider a weighted hybrid logic which combines the features of weighted regular expressions and weighted first-order logic (a fragment of wMSO). We show an efficient translation from this weighted hybrid logic to weighted pebble automata, an extension of weighted automata with two-way navigation and pebbles. This translation works over general classes of graphs, like words, trees, nested words, pictures, ... Moreover, weighted pebble automata can be evaluated efficiently. In the overall, this permits an efficient processing of quantitative specifications over general classes of graphs, with possible applications to language processing, speech recognition, or XML querying, e.g. Finally, notice that the expressive power of our weighted hybrid logic meets the one of pebble weighted automata, proving the robustness of our models.

## 3.23 Multi-weighted Automata and MSO Logic

*Vitaly Perevoshchikov (Universität Leipzig, DE)*


**Joint work of** Droste, Manfred; Perevoshchikov, Vitaly;
**Main reference** M. Droste, V. Perevoshchikov, "Multi-weighted automata and MSO logic," in Proc. of the 8th Int'l Computer Science Symposium in Russia (CSR'13), LNCS, Vol. 7913, pp. 418–430, Springer, 2013.
**URL** http://dx.doi.org/10.1007/978-3-642-38536-0_36

Weighted automata are non-deterministic automata where the transitions are equipped with weights. They can model quantitative aspects of systems like costs or energy consumption. The value of a run can be computed, for example, as the maximum, limit average, or discounted sum of transition weights. In multi-weighted automata, transitions carry several weights and can model, for example, the ratio between rewards and costs, or the efficiency of use of a primary resource under some upper bound constraint on a secondary resource. Here, we introduce a general model for multi-weighted automata as well as a multi-weighted MSO logic. In our main results, we show that this multi-weighted MSO logic and multi-weighted automata are expressively equivalent both for finite and infinite words. The translation process is effective, leading to decidability results for our multi-weighted MSO logic.

## 3.24 Formal Reductions of Stochastic Rule-based Models of Biochemical Systems

*Tatjana Petrov (IST Austria – Klosterneuburg, AT)*


**Joint work of** Petrov, Tatjana; Koeppl, Heinz; Feret, Jerome; Henzinger, Tom; Ganguly, Arnab;
**Main reference** A. Ganguly, T. Petrov, H. Koeppl , "Markov chain aggregation and its applications to combinatorial reaction networks," Journal of Mathematical Biology, Nov. 2013 (online).
**URL** http://dx.doi.org/10.1007/s00285-013-0738-7

Intuitively, bisimulation is a measure of behavioural similarity between two transition systems. The classical probabilistic bisimulation on transition systems running in continous-time, on discrete state space, coincides with the concept of lumpability in Markov chain theory. We show that such probabilistic bisimulation can be effectively constructed for models of

biochemical networks written in a rule-based language, and, hence, it can provide a significant state space reduction. Then, we discuss possible further directions in this application domain.

### References

1 Petrov Tatjana. *Formal reductions of stochastic rule- based models of biochemical systems*. PhD thesis, ETH Zürich, 2013
2 Petrov Tatjana, Ganguly Arnab, Koeppl Heinz. *Markov chain aggregation and its applications to combinatorial reaction networks*. Journal of Mathematical Biology, Springer, 2013

## 3.25 Verification of Timed One-Counter Automata

*Karin Quaas (Universität Leipzig, DE)*



A timed one-counter automaton is a timed automaton extended with a counter ranging over the natural numbers. During the execution of a transition, the counter can be incremented or decremented. In this way, a timed one-counter automaton can also be regarded as an extension of a one-dimensional vector addition system with states extended with clocks. By an easy extension of the classical region graph construction for timed automata, one can prove that the reachability problem for timed one-counter nets is decidable. I present some results and open questions about typical verification problems like Metric Temporal Logic-model checking and language inclusion for timed one-counter automata and related models.

## 3.26 Topological RNA Structures

*Christian M. Reidys (University of Southern Denmark – Odense, DK)*


**Joint work of** Reidys, Christian M.; Fenix Huang;
**Main reference** F. Huang, C. Reidys, "Shapes of topological RNA structures," submitted.

In this talk I discuss topological RNA structures. The particular topologization method is quite generic and applicable to other discrete structures, like graphs when additional information is given. Topological RNA structures are fatgraphs, a natural enrichment of the concept of undirected graphs and equivalent to cell-complexes of orientable surfaces. In this talk we describe the model and discuss main results and new perspectives.

## 3.27 Using LARES in Order to Tackle Hierarchically Structured Dependable Systems

*Martin Riedl (Universität der Bundeswehr – München, DE)*



In order to bridge the gap between high-level dependability modelling formalisms and formal modelling languages the LAnguage for REconfigurable dependable Systems (LARES) has been defined. It can serve as both an intermediate language as well as a stand-alone modelling approach. LARES provides means for hierarchical modelling, i.e. it separates between definition of structure and behaviour. Furthermore, it introduces scopes in order to restrict visibility of definitions and named statements, which leads to structured model descriptions. Specific language elements are provided which allow asserting questions on states of subinstances in order to imply a specific reaction. The semantics is defined by means of stochastic process algebra and labelled transition systems. LARES has meanwhile been extended by rewards and non-deterministic decisions. It is still an open question how hierarchical structures can be exploited for improving the analysis and which additional extensions would be useful in order to increase the expressiveness of LARES, thus its field of application.

## 3.28 Decidable Properties for Subfragments of Quantitative Monadic Second-Order Logic

*Cristian Riveros (Pontificia Universidad Catolica de Chile, CL)*


**Joint work of** Riveros, Cristian; Kreutzer, Stephan;
**Main reference** S. Kreutzer, C. Riveros, "Quantitative Monadic Second-Order Logic ," in Proc. of the 28th IEEE/ACM Symp. on Logic in Computer Science (LICS'13), pp. 113–122, IEEE, 2013.
**URL** http://dx.doi.org/10.1109/LICS.2013.16

What is a good logic to define quantitative properties? Quantitative monadic second order logic (i.e. Weighted monadic second order logic, Droste and Gastin 2005) is a good alternative for defining quantitative properties that captures the expressiveness of weighted automata. Unfortunately, this equivalence implies that many interesting problems (e.g. containment or emptyness of formulas) of this logic become undecidable. To avoid these undesirable features of QMSO, one is forced to look at its subfragments and sacrifices expressiveness in favor of decidability. Towards this goal, we show that, by refining the analysis of QMSO, one can obtain subfragments that characterize exactly subclasses of weighted automata defined by the level of ambiguity allowed in the automata. This gives us a starting point in order to find a quantitative logic with good decidability properties while being reasonably expressive. In this talk, I will give a summary of the connection of between subfragments of QMSO and the ambiguity of weighted automata. Towards the end, I will show how this is connected with their decidability properties for different semirings.

## 3.29 Basics of Weighted Automata Theory: An Algebraic Perspective

*Jacques Sakarovitch (Telecom ParisTech, FR)*



In this lecture, I shall try to give a comprehensive introduction to the study of weighted automata along the line I have developed in the chapter I wrote in the Handbook of Weighted Automata edited by Droste, Kuich and Vogler and in more recent works achieved with Sylvain Lombardy. The notions of rationality and recognisability will be distinguished, in particular for the purpose of dealing with weighted transducers. The questions of reduction and morphisms of weighted automata will be tackled with the notion of conjugacy of automata. Reduction is the problem of finding an equivalent automaton of smaller size, hopefully of minimal size. Reduction is effective when the weight semiring is a (subsemiring of a) skew field, and yields the decidability of equivalence in these cases, with a cubic complexity. As established by Harju and Karhumäki, this decision result extends to unambiguous transducers via Malcev-Neumann theorem. The notion of morphisms for weighted automata proves to be equivalent to the notion of bisimulation developed in other contexts.

### References

1 S. Lombardy and J. Sakarovitch, The validity of weighted automata, *Int. J. Algebra and Computation* **23** (2013) 893–913. DOI: 10.1142/S0218196713400146
2 S. Lombardy and J. Sakarovitch, Radix cross-sections for length morphisms, in *Proc. LATIN 2010* (A. Lopez-Ortiz, ed.), Lect. Notes in Comp. Sci. **6034** Springer (2010) 184–195.
3 J. Sakarovitch, Rational and recognisable power series, in *Handbook of Weighted Automata* (M. Droste, W. Kuich and H. Vogler, eds.), Springer (2009) 105–174.
4 J. Sakarovitch, *Elements of Automata Theory*, Cambridge University Press (2009).

## 3.30 Energy-Autonomous Smart Micro-Grids

*Gerard J. M. Smit (University of Twente)*


URL http://www.dagstuhl.de/mat/Files/14/14041/14041.SmitGerard.ExtAbstract.pdf

When enough (renewable) generation like PV panels, biomass installations and wind-turbines in combination with storage assets are installed, it may be possible to create a self-supplying (autonomous) neighbourhood in a so- called energy autonomous smart micro-grid. The main objective of our work is: to develop methods and techniques to support the development of energy-autonomous smart micro-grids. This broad main objective can be decomposed in a number of detailed research questions:

- In an energy-autonomous smart micro-grid demand/supply matching (DSM) has to be done on a local level. How to find local balance of demand/supply/storage. A related research question is: How (and for how long) can a micro-grid continue autonomously without a connection to the main electricity grid?
- What distributed energy management systems can be used for a local micro-grid and a cluster of micro-grids (systems of systems) attached to the smart grid.

- Find and use the flexibility of appliances in a micro-grid e.g. storage, charging time of EV, starting time of dishwashers.
- What kind of (wireless) communication networks will support reliable, real-time and efficient communication in a micro-grid?

## 3.31 Probabilistic Rectangular Hybrid Automata

*Jeremy Sproston (University of Turin, IT)*


**Main reference** J. Sproston, "Discrete-Time Verification and Control for Probabilistic Rectangular Hybrid Automata," in Proc. of the 8th Int'l Conf. on Quantitative Evaluation of Systems (QEST'11), pp. 79–88, IEEE, 2011.
**URL** http://dx.doi.org/10.1109/QEST.2011.18

Hybrid automata provide a modeling formalism for systems characterized by a combination of discrete and continuous components. Probabilistic rectangular hybrid automata generalize the class of rectangular hybrid automata with the possibility of representing random behavior of the discrete components of the system. We consider verification and control problems for probabilistic rectangular hybrid automata. When restricting to the case of a semantics in which discrete control transitions can occur only at integer points in time, both the verification and control problems are decidable. We also consider positive and negative results and open problems with regard to the standard continuous-time semantics of probabilistic rectangular hybrid automata.

## 3.32 Model Checking meets Simulation-Based Analysis – Getting Formal Methods Adopted by Industry

*Bart Theelen (Embedded Systems Institute – Eindhoven, NL)*


**Joint work of** Theelen, Bart; Geilen, Marc; Voeten, Jeroen;

In our mission to advance innovation by industrial adoption of academic results, TNO-ESI performs projects with high-tech industries such as ASML, Océ Technologies, Thales, NXP Semiconductors, Philips Healthcare and TP Vision. Favoring formal methods, we observe a gap between industrial needs in analyzing quantitative properties and the capabilities of formal methods for this goal. After highlighting a few aspects of this gap, we briefly identify some relevant deficiencies of state-of-the-art quantitative analysis techniques (focusing on model checking and simulation). As an ingredient to bridging the gap, we indicate the formal link between our model checking and simulation-based analysis approaches. Although concentrating on infinite horizon properties (i.e., complex forms of long-run averages), our techniques also serve best/worst case and (probabilistic/expected) reachability properties. We illustrate our approaches with an example from the domain of dynamic Digital Signal Processing (DSP) applications in high-tech (embedded / cyber-physical) systems. Based on ingredients of our work, we conclude with some technical thoughts to initiate a discussion on advancing model checking based quantitative analysis to improve its adoption by industry.

## 3.33 Probabilistic Logic and Regular Expressions on Infinite Words

*Thomas Weidner (Universität Leipzig, DE)*



We introduce a probabilistic extension to MSO logic on infinite words. This extended logic adds an expected-value operator to classical MSO. We also provide probabilistic omega-regular expressions, which are based on probabilistic regular expressions introduced by Bollig, Gastin, Monmege, Zeitoun. Both formalisms prove to be expressively equivalent to probabilistic Muller-automata. To obtain better decidability results we restrict probabilistic automata and probabilistic expressions such that probabilistic choices occur almost surely only finitely often. The image of such restricted automata, resp. expressions, can be approximated by a finite, computable set.

## 3.34 Models and Control for Smart Grid Balancing

*Rafael Wisniewski (Aalborg University, DK)*


**Joint work of** Luminita C. Totu, Rafael Wisniewski, John Leth;

This presentation addresses the problem of maintaining the balance between consumption and production in the electricity grid when volatile resources, such as wind and sun, account for a large percentage of the power generation. I will consider a participant in the energy market, who manages a portfolio composed of different units. Specifically, the portfolio includes consumption units with flexible consumption. The units in the portfolio are distributed across geographical areas, and the number of units in the portfolio can make the portfolio very large. Consequently, both collecting and distributing data across the portfolio, as well as optimizing individual power schedules in a centralized manner, may become cumbersome. During the talk, I will address two approaches:

1. deterministic, where distributed optimisation techniques are applied,
2. stochastic, where Forward Kolmogorov equations is used to capture the power consumption behavior at the population level.

## Participants

- Giorgio Bacci
Aalborg University, DK
- Giovanni Bacci
Aalborg University, DK
- Benedikt Bollig
ENS – Cachan, FR
- Miroslav Ćirić
University of Niš – Serbia, RS
- Thomas Colcombet
CNRS / Université Paris-Diderot, FR
- Julie De Pril
University of Mons, BE
- Laurent Doyen
ENS – Cachan, FR
- Manfred Droste
Universität Leipzig, DE
- Zoltán Ésik
University of Szeged, HU
- Javier Esparza
TU München, DE
- Ulrich Fahrenberg
INRIA Rennes – Bretagne Atlantique, FR
- Emmanuel Filiot
Université Libre de Bruxelles, BE
- Paul Gastin
ENS – Cachan, FR
- Daniel Gebler
Free Univ. of Amsterdam, NL
- Cyrille Jegourel
INRIA Rennes – Bretagne Atlantique, FR
- Heinz Koeppl
ETH Zürich, CH
- Jan Křetínský
TU München, DE
- Antonín Kučera
Masaryk University – Brno, CZ
- Dietrich Kuske
TU Ilmenau, DE
- Kai Lampka
Uppsala University, SE
- Kim Guldstrand Larsen
Aalborg University, DK
- Axel Legay
INRIA Rennes – Bretagne Atlantique, FR
- Björn Lisper
Mälardalen University – Västerås, SE
- Sylvain Lombardy
Université Bordeaux, FR
- Jan Madsen
Technical Univ. of Denmark – Lyngby, DK
- Nicolas Markey
ENS – Cachan, FR
- Benjamin Monmege
Université Libre de Bruxelles, BE
- Vitaly Perevoshchikov
Universität Leipzig, DE
- Tatjana Petrov
IST Austria – Klosterneuburg, AT
- Karin Quaas
Universität Leipzig, DE
- Christian M. Reidys
University of Southern Denmark – Odense, DK
- Martin Riedl
Universität der Bundeswehr – München, DE
- Cristian Riveros
Pontificia Universidad Catolica de Chile, CL
- Jacques Sakarovitch
Telecom ParisTech, FR
- Sean Sedwards
INRIA Rennes – Bretagne Atlantique, FR
- Gerard J. M. Smit
University of Twente, NL
- Jeremy Sproston
University of Turin, IT
- Bart Theelen
Embedded Systems Institute – Eindhoven, NL
- Thomas Weidner
Universität Leipzig, DE
- Rafael Wisniewski
Aalborg University, DK

Report from Dagstuhl Seminar 14042

# Do-it-yourself Networking: an Interdisciplinary Approach

**Edited by**

**Panayotis Antoniadis[1], Jörg Ott[2], and Andrea Passarella[3]**

1 **ETH Zürich, CH,** antoniadis@tik.ee.ethz.ch
2 **Aalto University, FI,** jo@netlab.tkk.fi
3 **CNR – Pisa, IT,** a.passarella@iit.cnr.it

## Abstract

This report provides a summary of the organization, program, and outcome of the Dagstuhl Seminar titled "Do-it-yourself networking: an interdisciplinary perspective". We first motivate our interest in wireless networks operating outside the public Internet and the selection of the most relevant areas of expertise. Then we describe the process of bringing together a balanced group of representatives from these areas, and the evolution of the seminar over time. An overview of the interactions during the work in groups on specific application areas and explorations of the concept of failure, edited collectively by the members of the different groups, summarizes the main outcomes of the seminar. Finally, we identify some important lessons learned for facilitating interdisciplinary collaborations and conclude with our plans toward building a DIY networking community of researchers and activists.

**Seminar** January 19–22, 2014 – http://www.dagstuhl.de/14042
**1998 ACM Subject Classification** Mobile Computing, Networks Society, Human-computer Interaction
**Keywords and phrases** Community Wireless Networks, Mobile Networking, Delay-Tolerant Networking, Ad-hoc Networking, Urban informatics, Community informatics, Urban Planning, Urban Art, Interaction design, Interdisciplinarity
**Digital Object Identifier** 10.4230/DagRep.4.1.125

## 1 Executive Summary

*Panayotis Antoniadis*
*Jörg Ott*
*Andrea Passarella*



The key objective of the seminar was to bring together a diverse group of researchers and practitioners to reflect on technological and social issues related to the use of local wireless networks that operate outside the public Internet. We managed to bring together a quite balanced group of 32 people with expertise in the design and implementation of wireless ad hoc networks of various types, human-computer interaction, community informatics, urban interaction design, ethnography, media studies, arts and design.

Interdisciplinary interactions took place successfully around specific application areas for which the use of do-it-yourself networks is meaningful. More specifically, we explored the use of such networks for supporting the creation of transient communities of different size and duration, political activism, and similarity matching. In addition, an in depth exploration


Do-it-yourself Networking: an Interdisciplinary Approach, *Dagstuhl Reports*, Vol. 4, Issue 1, pp. 125–151
Editors: Panayotis Antoniadis, Jörg Ott, and Andrea Passarella
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

of the concept of failure provided a useful framework for addressing various challenges in bridging the gap between theory and practice, scientific and social objectives.

Our main finding was that there are certain assumptions that need to be carefully understood and important requirements that need to be fulfilled in order for DIY networking to become a feasible, and desirable, option for shaping the hybrid space of contemporary cities. That calls for a closer collaboration between experts from different fields and disciplines. For this, the most important achievement of our seminar was the balanced and productive interactions between engineers and social scientists around a concrete topic, and the general feeling that a new interdisciplinary community around the topic of DIY networking is meaningful and a goal worth pursuing. Indeed, concrete plans for facilitating the formation and expansion of such a community through online communication and face-to-face meetings, research visits, and common projects between participants that met in Dagstuhl for the first time are already under way.

*When things get so big, I don't trust them at all*
*You want some control – you've got to keep it small*
*D.I.Y. D.I.Y. D.I.Y. D.I.Y.* – Peter Gabriel

## 2 Table of Contents

## 3 Background and motivation

Wireless technology enables at present the creation of local networks outside the public Internet. Even in cases where the public Internet is easily accessible, such local wireless networks form an interesting alternative, autonomous, option for communication, which

1. ensures that all connected devices are in de facto physical proximity,
2. offers opportunities and novel capabilities for interesting combinations of virtual and physical contact, and appropriation of the hybrid space,
3. enables the serendipitous gathering of diverse people without the need to have any specific application installed or provide any credentials,
4. allows for purely anonymous and privacy-preserving virtual interactions, and
5. can create feelings of ownership and independence.

However, timidity, security issues, and the potential lack of common interests could limit the desire of people to participate in local interactions mediated through ICT. Such psychological barriers and various technical challenges hinder today the creation of plug and play DIY networking solutions with applications specialized for local environments, which can compete with the quality of experience offered by popular Internet applications. Then this fact discourages application developers to invest a lot of effort in building applications, undermining the engineering efforts to solve the corresponding technical challenges, and thus leading to a "chicken and egg" problem.

The vision of developing DIY networking tools could be one toward encouraging more face-to-face communication, information sharing between strangers, and exposure to diversity in contemporary cities. Then more ambitious objectives such as e-participation and e-democracy could be also part of the scope of such an endeavor. This means that the design and deployment of DIY networks and related applications, could touch on areas of expertise and interest of a highly diverse community of researchers, engineers, hackers, practitioners, activists, and artists. More specifically, among others, 1) the research on adhoc, DTN, and packet switched networks, 2) the grassroots initiatives building operational wireless mesh networks in various cities, 3) human-computer interaction (HCI), computer supported collaborative work (CSCW), interaction design, computer mediated communication, 4) sociology, media studies, and other social sciences, 5) the emerging interdisciplinary fields of urban informatics, ubiquitous computing, and community informatics, and related disciplines such as urban planning and urban design.

Although, there are already efforts to create links between some of these areas, there are still many isolated groups of researchers and practitioners. For example, people working on applications and uses of ICT are not always aware of technology's capabilities for building local communication networks. On the other hand, scientists in the field of networking are often indifferent with respect to the actual use and social implications of the technical solutions they devise, as long as they fulfill the minimum academic requirements, and are often abandoned after a few years (e.g., when the PhD student leaves).

Interestingly, the idea of the seminar was born after one of the organizers, Jörg Ott, presented in his keynote talk at the MobiOpp 2012 conference[1] an interesting application, called SCAMPImusic, for sharing music locally anchored on specific locations. This application reminded Panayotis Antoniadis, who was attending the conference, a similar application called Undersound discussed in the book "Divining a digital future" by Dourish and Bell 2011.

[1] http://www.cl.cam.ac.uk/events/mobiopp2012/program.html#keynote2

In the discussion that followed, Jörg and Panayotis, realized that it is a pity that there are not closer interactions between the networking and HCI communities around this type of applications. It was then a matter of a few e-mails to decide together, and also with Andrea Passarella, to apply for the organization of a highly interdisciplinary Dagstuhl seminar titled "DIY networking: an interdisciplinary perspective", with the following objectives:

- The sharing of objectives, values, methodologies, and challenges the different fields of research and practice face today;
- The definition of a research framework that will allow disconnected disciplines to exchange knowledge and interact toward the design of successful do-it-yourself networking applications; and
- The definition of next steps toward a shared experimentation platform and the setting up of a venue for sharing artistic, experimental, and research results.

## 4 Organization

The key first challenge identified during the preparation of the seminar proposal was to manage to build a really balanced mixture of researchers and practitioners and avoid as much as possible power games between disciplines, as for example the treatment of ethnographers by engineers as "tape recorders", as convincingly described by Dourish and Bell 2011, p. 61-88.

More specifically, the invitations aimed to bring together people from two interdisciplinary groups of almost equal size:

1. Adhoc/DTN networking, p2p systems, security, and engineering
2. Community and urban informatics, human-computer interaction (HCI), media and communication studies, ethnography, urban planning, arts and design

However, since none of the organizers had presence in the fields of the second group that seemed like a really difficult task. Our strategy was to try to invite "clusters" of people working closely between them in selected areas, such as community informatics and HCI, in order to avoid as much as possible isolated individuals. Our assumption was that this clustering would make it easier for people to accept the invitation in the first place and, most importantly, make them feel more comfortable and confident during the seminar.

We were very happy to see that our strategy proved to be effective and, together with the help of chance and the reputation of the Dagstuhl seminar series, we managed to gather an impressively diverse mix of researchers and practitioners with backgrounds in engineering, activism, art, sociology, anthropology, urban studies, community informatics, and HCI, coming from many parts of the globe such as Australia, Denmark, Finland, Germany, Greece, Italy, Romania, Sweden, Switzerland, UK, and the US. We could say that such a diverse participation was unique for the Dagstuhl seminar series. Even in terms of gender, our low diversity score, 6 women out of 32 participants, was unusually high and as Kat Jungnickel mentioned at her blog entry about our Dagstuhl seminar [2], the women of the group "were made very welcome". Interestingly, more than half of the participants participated for the first time in a Dagstuhl seminar.

[2] http://www.katjungnickel.com/2014/02/28/dagstuhls-diy-networking-seminar-making-a-failure-machine/

# 5 The seminar

The key challenge was to create a common vocabulary and expose people to different ways of thinking in a productive way. For this, we decided to follow three sequential tasks for which only a first step would be made during the duration of the seminar (and hopefully form a basis for future re-iterations):

1. Getting to know each other
2. Working together
3. Summary and future steps

In the following we give a brief overview of the evolution of the seminar around a draft agenda and in the next section a summary of the key ideas and results produced by the different working groups on the second day.

## 5.1 Getting to know each other

Given that all 32 participants were meeting for the first time a large proportion of our diverse group, the task of getting to know each other was identified as the most important. For this, we tried to optimize the use of the limited available time, by starting with a round table supported by one slide per participant collected and shared beforehand, followed by two seed talks by "representatives" of the two main sub-groups of participants: engineering and urban informatics (roughly speaking).

First, **Gunnar Karlsson** decomposed in a very nice way the do-it-yourself meme and clarified how many differents options for Do (e.g., design, deploy, educate, inspire), It (e.g., spectrum, networks, applications), and Yourself (e.g., individuals, communities, organizations), actually exist. This disambiguation effort provided us with a handy reference during the seminar when there were misunderstandings about key assumptions and the meaning of ambiguous terms, such as "network". Most importantly, Gunnar added one more Y in front of the DIY acronym (Y.D.I.Y) standing for "Why DIY?", which proved to be the most popular and challenging question during the seminar.

It is important to note that a traditional challenge in the field of adhoc and delay tolerant networking has been to provide convincing arguments about the importance of this mode of communication in light of specific applications, especially in situations where access to the Internet is affordable. For this, the second keynote by **Marcus Foth**, was ideal in showcasing a wide range of such applications that go beyond the traditional top-down visions of the smart cities and assume an increased level of engagement by citizens, but which have not until today considered DIY networks as their main communication infrastructure. Ranging from hybrid participation platforms (such as Discussions in Space[3]), to the appropriation of the urban space (like in the SMS guerilla project by the Troika group), and the fabrication of gadgets through 3D printing (like the Maker Bot), these applications offered much inspiration for the following interdisciplinary exchanges centered around the key question, why DIY?, posed by Gunnar.

The introductory part of the seminar concluded with a panel on experiences from the field which gave us a glimpse of real life DIY networking projects seen from the perspective

[3] http://www.urbaninformatics.net/media/dis/


**Figure 1** Part of the outcome of the brainstorming meta-session in which we tried to identify important concepts, examples, and ideas, belonging to different categories, such as infrastructure, platforms, processes, and case studies.

of ethnographers (i.e., the account of **Kat Jungnickel** on DIY WiFi initiatives in Australia), activists (i.e., the description of the on-going RedHook WiFi initiative by **Jonathan Baldwin**), researcher communities (i.e., the ExtremeCom conference series by **Anders Lindgren** and Pan Hui, presented by the former), and entrepreneurs (the university spin-off on liberouter by **Teemu Kärkkäinen**). The key message of the panel was that DIY networking is feasible and there are many disconnected efforts today that would benefit from the creation of a community around this concrete design space.

This left us with only a short time to rehearse on working in groups, preparing the field for the next "working together" day. Three groups were formed after a quick decision process with the help of a google doc which was filled with ideas for possible topics of collaboration. Two somehow focused groups, concentrated on the topics of failure and affordable networks, and the use of crypto-currencies, like BitCoin, for local change.

The third and biggest group focused on a more general meta-discussion on case studies, the role of DIY networking, and interdisciplinarity. In this discussion many of the differences in background, assumptions, and ways of thinking between participants manifested and it was often that we had to go back to Gunnar's Y.D.I.Y to be sure that we were all on the same page. Despite the efforts to converge to a classification of concepts that would help us organize our thoughts and proposed solutions (see Figure 1), it was made clear that our task of working together toward concrete outcomes the following day wouldn't be so easy.

## 5.2 Working together

Before starting working in groups on specific topics, three seed talks were scheduled in the morning of the second day to give us some additional inspiration and put the collaborative work into perspective.

First, **Paul Dourish** talked about the art of interdisciplinary research, the different types of interactions between disciplines (inter, multi, trans) and possible spaces of encounter (solutions, phenomena, epistemes). Paul highlighted the significant amount of time and effort required, the role of language, and the need to "give up something" for interdisciplinary interactions to be successful. He also stressed the inevitability of doing politics through design.

Then, **Jon Crowcroft** introduced the concept of peer-to-peer systems and virtual currencies, a technology central to the implementation of distributed DIY networks, and a rather political one. In this context, he presented on-going research on replacing the energy wasteful BitCoin model with a peer-to-peer storage system, the personal cloud, whose needs for encryption and verification could be used for minting a new virtual currency (DO$H – Decentralized Object Storage Help), which can be used as a basis for a private data economy that allows people to sell their data to service providers and advertisers and buy ad-free services.

Finally, **Doug Schuler**, offered us a number of seeds, starting from asking "Why ask why?" and the key concepts of civic intelligence and the role of the citizens, which are part of the answer, to high-level views of the big picture including how local solutions can grow to a "hyper project", and the role of research readjustment, technology development, and the context (e.g. the increasing power of Facebook and Google and the NSA affair). Note that the seed metaphor is a particularly interesting one in the context of DIY networking, and actually our Dagstuhl seminar was conceived more as a seed for future collaborations than a workshop seeking for concrete results and well-baked ideas.

Following these very interesting and diverse seed talks, a short "walkshop" was organized ad-hoc to give us the chance to relax and prepare for the long-awaited "working in groups" session, whose results are summarized in the following. Judging from the success of these intense collaborations in smaller groups, there was a general feeling that we should have perhaps reserved more time for group work than plenary talks. As Kat Jungnickel mentioned in her blog, "Although discussions were expansive and interesting for the first day and a half, and the walkshop around the local village and forest was great, the event became really productive for me when groups shrunk in size and conversations shifted to more specific topics."

However, given the very high degree of diversity less of "getting to know each other" in plenary might lead to more misunderstandings during the group work and the formation of less diverse groups. It is difficult to know how things would have evolved otherwise, but there is clearly a trade-off, and in future similar events we could try to experiment with even less structure and scheduled talks.

## 5.3 Summary and future steps

The last day of the workshop started with the presentations of the outcomes of the group work, some of which were very animated and created a joyful atmosphere in the room. Then a quick roundtable gave everyone the opportunity to share his take away message and ideas for the future. The general feeling was that we managed to build links between the different communities, as well as that DIY networking is a good "triangulator" for enabling fruitful interdisciplinary collaborations, and allow the combination of research and action for addressing real problems.

The rest of our available time was spent to discuss on ideas that will help us to "keep one or more balls rolling" and take advantage of the momentum created during the seminar. The

organization of a follow-up Dagstuhl seminar, the creation of an interdisciplinary workshop on DIY networking, possibly attached to conferences of different disciplines and other events, but also the organization of various less formal meetups, research visits, invited talks, an e-mail list, were some of them, discussed in more detail in the last section of this report.

# 6 Outcomes of group work

Schematically, there were three groups focusing on the design of applications for three different scenarios: transient environments, political activism, and similarity matching (with focus on social medicine), and one group elaborating on the overarching concept of failure.

## 6.1 Failure machine

*Christian Becker, Jon Crowcroft, Paul Dourish, Kevin Fall, Alison Powell, and Irina Shklovski*



Failure is a rich topic for discussion as it holds multiple shifting meanings, is culturally shaped and manifest in diverse assemblies of practice. We shared very different experiences from our fieldwork and practice, talking about how failure in some contexts was the key to success and the start of innovative journey; failure in the form of disaster sometimes operates as a catalyst for invention (but things need to be available before); success as an exception. We asked:

- What is failure?
- How is it avoided?
- Who is allowed to fail? Who isn't?
- How is failure understood, subverted and explored?
- How is it represented? How has this changed over time/ in different places?
- In what ways/contexts/articulations is failure reviled? Cleaned up? Ignored? Celebrated?

Drawing on previous experiences of building Enquiry Machines[4], Kat Jungnickel suggested to build these ideas into a 'Failure Machine'. Enquiry Machines are a series of performed artefacts made in collaboration with others that explore ideas or methods. The point is less about materializing answers or prototyping ideas and more about formulating new critical approaches and literally seeing and touching methods in new ways. EMs are not meant to be finished or polished objects that speak for themselves. In fact, most fail in some way. They remind us that mistakes and tangents are just as important to our insights as the things that 'work'.

It seemed a good idea in this context as it would help to ground the discussion and unite our wide-ranging discussion into something physical. Also, the delightful thing about working on failure is that anything we made or failed to make would be productive. Plus Enquiry Machines are fun to make.

We started by simply talking more, writing down ideas, quotes and drawing things that popped up in conversation. Then we coded these bits of paper according to themes, creating more consolidated taxonomies. This working session moved into the evening and

[4] http://www.katjungnickel.com/portfolio/enquiry-machines/

**Figure 2** Examples of Châtelaines.

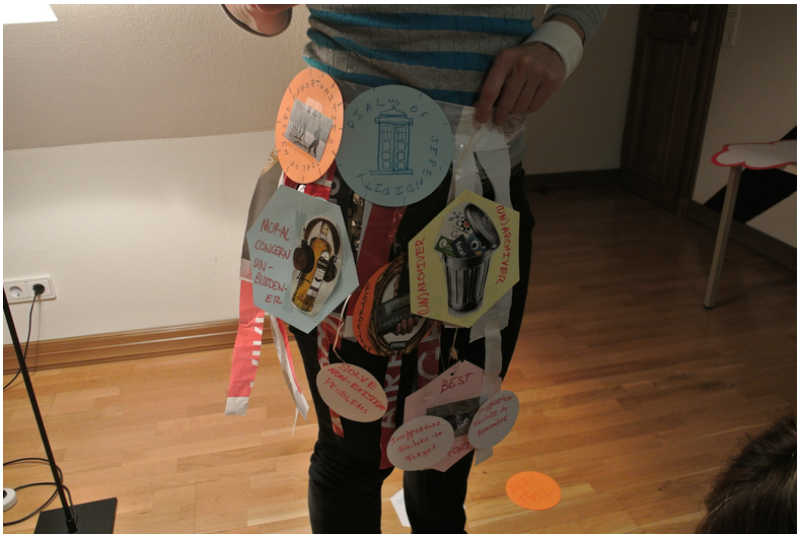

**Figure 3** The failure machine chatelaine, made out of paper, magazines, and sticky tape.

was accompanied by some nice local wine, in fine company and to the background of acoustic guitar played by John and Kevin. It'd be nice to work like this more often. Bits of paper, pens and the whiteboard were the tools of choice. Magazines, coloured paper, tape, string and scissors were soon recruited.

During the session Kat talked a bit about her recent obsession with châtelaine, a fascinating technology introduced to her by Genevieve Bell. Châtelaines were practical and decorative devices worn on the belt and hung with a series of short chains at the end of which were objects related to the task at hand (see Figure 2). They were worn by women from the 16th to 19th Centuries, from lower socio-economic workers to aristocracy. Nurses wore châtelaines with clocks, thermometers, bandages and scissors. Seamstresses had bobbins of thread, thimble cases and needles on the end of their châtelaine chains. Society ladies' châtelaines featured highly decorative perfume bottles, purses, fans and even dance cards.

We decided to make our 'Machine of Enquiry' into a digital châtelaine. We called it 'The Battery Operated Wind-Up Merchant', playing on the ideas about technological lineage, pointing to larger dependent ecologies of use and using humour as a deliberate device to bring to life multiple ideas about failure and also the slightly ridiculous method.

There was a lot of DiY hands-on material adaptation going on. We scoured the castle for string and in its absence made use of tape, scissors, some raffia and a plastic bag. The châtelaine featured a series of filters or apps hanging from each chain that reflected some of the critical themes and ideas generated in our discussions (see Figure 3).


**Figure 4** The plenary presentation of the failure machine, which was much more animated and cheerful than this figure suggests.

We talked about the apps having both independent and potential interrupting characteristics, so they might overlap, tangle and otherwise interfere with one another causing even more noise in the system/process. The apps included 'Dial of serendipity', 'Dial of missed opportunities', 'Lens of temporality', 'Latency creator', '(Un)Archiver', 'Moral concern unburdener' and many more.

The process and presentation of the machine to the larger group was productive and enjoyable (Figure 4). Although making ideas material constituted a different method for some in our group, everyone was buoyed by the experience of collectively approaching the multiplicity and messiness of failure via gendered, historic, cultural and social actors as well as the technical ones. There was even talk of potentially furthering this as an interdisciplinary project and making the Failure Machine again in different, more developed materials.

For more details and photos see http://www.katjungnickel.com/2014/02/28/dagstuhls-diy-networking-seminar-making-a-failure-machine/.

## 6.2 Transient communities

*Panayotis Antoniadis, Jonathan Baldwin, Marcus Foth, Mark Gaved, Paul Houghton, Teemu Kärkkäinen, Jussi Kangasharju, Gunnar Karlsson, Anders Lindgren, Jörg Ott, and Michael Smyth*



This working group focused on relevant applications for DIY networking in transitory environments: people coming together in a particular place for a limited amount of time. It became very quickly apparent that there are different important dimensions that affect the type of applications that make sense and their basic characteristics.

So, our first task was to elaborate on the most important context variables that would affect the choice of application and its characteristics: 1) The number of people involved, 2)

the expected duration of the network-mediated interactions, 3) the size of the target area to be covered, 4) the expected interactions (passive or active, synchronous or asynchronous, broadcasting vs. collecting), 5) the technological dimension regarding the importance of DIY networking: non-existent infrastructure; unwanted infrastructure; insufficient infrastructure, 6) the assumed client devices (smartphones, shared displays, other), 7) the circumstances that would activate the use of the DIY network (disaster, planned event, overlay to existing activities, continuing practice in a space), 8) the legal framework, and 9) the corresponding objectives (pass the time, feedback loops, content sharing and other purposeful activities like alerts, organization of meet-ups, political activism, etc.).

Again, the key question that one would need to answer in many scenarios would be why not just use the 3G network. As one of the participants wrote on our collaborative google doc: "What is the value of very contextualised mobile communication technologies from the standpoint of non-networking guys in this room? I mean: to communicate among us now, is that perfectly fine to use google docs, or any form of DIY network would bring us some additional value? More generally, if one has to design a form of participatory smart citizenship thing, would Haggle-like opportunistic networks (dynamic networks built out of users devices) be useful, or 'global' types of communication services would be enough? And furthermore – how would these things potentially look and feel (to pick up on the mess/materiality point Kat is making)?"

Some of the main reasons for investing on DIY networking instead of relying on the global Internet that we identified during our discussion are the following: Natural localization in space and time, reduced costs (especially relevant for touristic areas and developing countries), privacy issues, inclusive participation (through the use of a captive portal), and feelings of ownership. Interestingly, in our Dagstuhl seminar we had with us three developers of such systems (two of which in our working group, Jonathan Baldwin and Teemu Kärkkäinen):

- http://tidepools.co,runningontopofhttp://commotionwireless.net
- http://www.ict-scampi.eu/results/scampi-liberouter/
- http://occupyhere.org,byDanPhiffer

To test the expressiveness of the set of selected context variables we went through a list of possible examples classified according to the most important variables summarized above.

**Meetings and spontaneous gatherings.** These could concern up to 50–100 people for a duration of several hours to 1–2 weeks. In this scenario the narrative is critical but also a certain level of required attention. Interestingly, although two such DIY networks were available during our Dagstuhl seminar (occupy.here and liberouter), there was limited use, perhaps because of the intense interactions and limited engagement of people in online activities, which was considered actually one of the achievements of the seminar.

**Traveling together (airplane, train).** This is a similar setup as the above in terms of participation and duration, but in addition to the option of ambient, unannounced nodes, as in the case of the L-train notwork, one could consider also the possibility of a more official setup supported by the transport company. In this scenario however, potential participants are assumed to be mostly bored and having access to their mobile devices which makes it relatively easier to attract their attention. So, activities could range from content sharing to recommendations for the destination, chatting, and various short-term games.

**Long events (music festivals, cruiseships, camping sites).** Here the participation might increase significantly, from a hundred to several thousands of people or even 500000 as in the cross country skiing championship. The duration would be also also very variable, typically

from 3–4 days up to a month. Due to the repeated interactions in this set-up, in addition to content sharing more significant social interactions like meet-ups can be initiated. Note also, that in many cases there are international visitors, so relying on 3G would be problematic (both because of individual costs and load). This means that even the important task of sharing official data (such as broadcasting results, announcing on-going events, etc.) would benefit from a DIY networking setup.

**Short events (a music concert, a football match).** In this scenario the whole duration wouldn't be more than a few hours and one would expect different modes of operation during the actual event and during breaks. Contributing to building shared footage (e.g., during a music stadium event, you may be in a seat with a poor view, and would like to see somebody else's view.

**Public transport nodes (bus stop, train station, airport).** In this scenario one would expect limited participation (from 10-20 people in a bus stop to a few hundreds in an airport terminal) and limited duration (from a few minutes to a few hours).

**Communities of practice (construction site).** This is a long-term scenario, at the range of years, which involves a large and possibly changing population of people working together. In this case a DIY network could be used for safety reporting, synchronous communications, and sharing documents at the specific place they are needed.

After exploring the space of possible application areas, we chose to focus on two somehow "extreme" examples: a construction site with 3000 people over three years in one square kilometre; and a bus stop, with up to 20 people for up to ten minutes, in a few square metres.

We identified that a DIY network scenario for the construction site would be the need for workers to anonymously report on unsafe working conditions and bypass the official company network. This is not just a hypothetical scenario since, Jonathan Baldwin informed us that he has been approached by a group of migrant workers that want a way of reporting dangerous work conditions to the national health and safety people but bypassing the employer: if you go through the employer you may lose your job. So, there are many situations in which real safety of the workers may be orthogonal to the company's actual purposes of achieving health and safety inspections, and there is a need for workers to be able to independently report as actual conditions which aren't being seen by the health and safety inspectors, bypassing official channels and company procedures.

In our detailed discussion of this case study, we assumed that there is an official infrastructure (e.g, 3G) but not all workers may have phones, such as migrant workers from another country. The DIY network could have in general both official and unofficial purposes, which may be in conflict as mentioned above:

- *Official*: sharing documents, safety assessments, tracking work, scheduling for the use of specialized tools
- *Unofficial*: social communication (sharing jokes and pictures, chatting), informal learning (getting expert advice, information), anonymous communications alerting on dangerous conditions

While the use of a DIY network was obvious for the unofficial scenario, the possibility for official uses of a DIY network, lead us to reconsider the key question: Why DIY? Why a company should build its own network rather than rely on the use of the existing infrastructure? We identified the following reasons why even in an official scenario DIY networking might prove useful:

**Figure 5** The plenary feedback of the "transient communities" group, whose members literally presented the outcomes of the discussion as written on paper by Mark Gaved, who also animated this unconventional male chorus :-).

- No network coverage in shipyards due to large amount of metal, also true sometimes on construction site when concrete is wet, large amounts of water present.
- Push to talk too expensive without wifi
- Multiple channels (not sure what we discussed here?)
- Supporting small teams/ gangs or workers

Then the discussion shifted to the bus stop scenario. We wanted to think how an occupy.here style ambient and unannounced network node (or phone-to-phone network) could trigger initial, light steps towards community interaction, helping to start interactions between the different people in your neighbourhood, "achieving the initial smile". Michael Smyth referred to the concept of "smirting" – smoking and flirting; caused by the smoking ban – , to highlight the potential of people meeting and engaging with others that they might otherwise not meet. Some ideas on possible applications in this setup included the following:

- Sharing music: letting other people know what music you are listening to, not sharing the actual music but the titles and artists [31], or to listen to music that a group of people brings to a venue [7].
- Bus stop as fabric for displaying some of this information (smart city approach)
- Aiming for the initial shared smile, starting community interaction. Getting different people from the local community interacting: young, old, those who share this space but wouldn't normally interact.
- Situated sharing economy: a local web portal where locals can say if they've food or other things to share (like in freecycle), which would be easy for someone waiting to take the bus back home to carry.
- "Snapchat for buses"
- Community/art approach to triggering community
- Arriving to leave, but maybe on the same bus "the Bus 25 community"

An interesting question that arose during this very creative brainstorming session, was whether some of these applications can make the little interactions at the bus stop so

interesting that people decide to come a little bit early to participate in them. Would a bus company like this? How about the municipality?

The final task of this session was to make a roundtable for all to suggest one thing which would make such a DIY network a success, which resulted to the following list (not all in agreement with them all):

- Unofficial/subverting
- Effortless
- Useful and usable, without requiring a PhD in Computer Science ;-)
- If people who used it told their friend about the network
- Makes you smile, want to do it/use it again
- Help people to open up to strangers and get exposure to diversity
- Making the familiar strange (getting people to think of the place in a different way)

## 6.3 Political Activism

*Ileana Apostol, Fiorella De Cindio, Per Gunningberg, Christian Nold, Dan Phiffer, and Volker Wulf*



This group's discussions could be summarized in two streams pertaining to 1) security, trust, and ownership with respect to political uprisings, and 2) design issues for the case of participatory online platforms. The issues of surveillance, security and trust came forth in response to the question: why the DIY networking model may be a better solution for platforms used in political activism?

The ICT advances enable multiple opportunities for hybrid spatial uses that open up new dimensions of political activism and can strengthen social movements. Recent political uprisings in the Middle East or the Occupy Wall Street movement in New York have shown that, and the same holds with numerous other examples of participatory processes. What is important in the context of DIY networking is that both the online and offline spaces that political activists use for gathering and organizing their actions are subject to different ownership structures, and thus one needs to be aware of their limitations and potential threats.

First, in the online world, private social networks like Facebook may give access to their recorded information to entities of their choice, including governments or secret services, like the NSA surveillance programs uncovered by Edward Snowden. Governments could close down or even take possession of central servers like in the case of the seizure of Indymedia servers by the FBI. They can also limit people's access to certain servers either permanently as in the China case, or temporarily as in the case of the Arab spring, where governments managed in certain circumstances to cut the access to the whole Internet. There are, of course, many places in the world where popular platforms like Facebook, and Twitter may not be easily shut down or censored, as demonstrated by the recent (unsuccessful) effort by Tayyip Erdogan to ban Twitter in Turkey. However, these platforms in addition to their questionable privacy policies and vulnerability to surveillance, they use generic social software which does not allow to customize its design for specific uses, e.g., not allowing to retrieve old information easily. This is an important issue because the large variety of actions implied

in political activism require more options for customization, from reaching out to new people to working inside the movement.

Similarly to the case of the private online spaces used for online interactions between political activists, social movements often use private gathering places as in the case of the Zuccotti Park in New York for the Occupy Wall Street movement. As the spaces for public life are more and more privatized in western cities, one should be aware of the dangers that arise when private spaces are used in terms of surveillance, possibilities for eviction, etc.

DIY networking solutions can provide more flexible solutions both in the virtual and the physical space. For example, local networks hosted in a simple wireless access point, such as occupy.here, can support truly private communications through highly customized user interfaces, and allow for flexible choices of physical spaces, even where Internet connectivity is not available or censored. Additional flexibility, independence, and resilience could be guaranteed if, on top of a DIY network, web-based applications with a decentralized architecture, such as Diaspora, are used to manage the information shared between the participants.

The biggest challenge is that such solutions need to be prepared and configured in advance. That might be a problem due to the spontaneous nature of political activism, which leaves little room for preparation in terms of risk evaluation for communications before or during organized action. The same holds for example, in the case of PGP keys for securing private communication, despite the efforts of the media and institutions like the Electronic Frontier Foundation (EFF) which offer best security practices that may become a prerequisite for better-protected political action.

One suggestion highlighted in the group discussion regards a student or university based model/culture to make technologies for political activists, a DIY sort of network technology that could be build in academic context (despite the problems of wide distribution at the political moment).

The second stream of discussion concerned more mainstream political action, which could give inspiration on the suitable design of the applications built on top of a DIY network like the organization of a 500.000 people Demonstration in Italia by Populo Viola or the Movimiento 5 Stars by B. Grillo. Another example from the US is a CSCW meetup platform serving in the organization of local groups (it enables people to meet in the physical space) was used for US government elections by H. Dean to organize physical meetings. On a different line, applications such as 'fix my street' may be seen as democratic intermediary between citizens and local governments while more sophisticated e-participation tools such as OpenDCN.org developed by the group of Fiorella De Cindio facilitate brainstorming and petitioning and can enable deliberations on complex issues.

## 6.4 Similarity matching and social medicine

*N. Asokan, Ahmed Helmy, Marcin Nagy, and Amalia Sabiescu*



The initial idea was to explore technical solutions for assisting members of communities of place to find each other, interact and share resources based on common interests. Social medicine (what was to become the short name for our group) was proposed as an area of application for these support groups. Social medicine had also been approached in prior discussions among some of the group members.

We started with a general discussion of our core idea and why it was interesting to work on it from the perspective of social science. We used social medicine initially as application area, and then came up with additional applications in education. The most important points that emerged were that local communities have hidden resources that are seldom known to people in close proximity. These resources are nearby people who are experts in a given field, or have a great deal of passion and interest in certain domains. The local community could benefit in many ways by uncovering these hidden resources. We explored several ideas.

Firstly, people could learn from community experts. For instance, a mother wanting her son to learn Portuguese may discover a nearby Portuguese teacher, instead of enrolling her son to a distant language school. This approach may also bring money savings, as such "teachers" may not always be professionals, but also be a language hothead doing his job as a hobby, or favor.

Second, people could form local communities to provide assistance for its members. For instance, a group of working mothers with small children may form a group in which every day one of them stays home to care for all their children, while others may go to work. Such a group may be the answer to existing problems for women that want to join their motherhood and career. It provides money savings in comparison to day care costs, and may also be a more trustworthy solution for mothers that are afraid to leave their children in a daycare facility.

Third, and actually most promising, we explored support groups, people liaising with others who share common passions and interests to get motivation and drive for continuing to nurture their passion, pursue their interests, or solve their problems.

We then went on to bridge the social science side with the technical side, concentrating on finding compelling incentives for deployment, and asking: why and in what community contexts could these types of similarity-based encounters (and associated technical solutions) be needed? Arguably, there are many contexts where people would love to get together, share resources and learn from each other on topics they are interested about. The issue here was to understand why they would need a local network? Why not an Internet-based group using existing infrastructures?

So why a DIY solution? To explore this, we went on to list conditions and constraints by which a community may go for a DIY solution to allow its members to get together based on similar interests. The list included:

- Lack of infrastructure (poor communities)
- Privacy concerns (public sharing of data, avoidance of monitoring, importance of keeping data locally)

In addition, we explored potential technological solutions that could address the main goal: support identification and matching around common interests. Ahmed proposed that one possible road to explore, promising also for research advancement, is the usage of behavioral sensing for creating similarity metrics. There are many challenges to make behavioral sensing usable, so we spent a significant amount of time trying to address usability issues. Firstly, many potential users may be opposed to the idea, as they may be afraid of being tracked and having their privacy violated by personal information reveal. Therefore technology must support user anonymization and if some data are stored on the server, such servers must be well-protected. Ideally, whole computation and data analysis can be run on personal devices to avoid these issues. A second important issue is the problem that people are usually dishonest with themselves and at the same time there is no, and will never be created, an ideal behavioral sensing algorithm. Thus, a successful system must find a right trade-off between proper behavioral sensing and openness for user feedback. The feedback must be

designed in such a way that it is nice, usable and doesn't cause irritation to the users. A user's profile would be fed jointly by feedbacks and through behavioral sensing. The role of behavioral sensing would be, in particular, to allow updates of the user's profile (e.g. potential new interests discovery) and user notification about possible life changes (e.g. lack of social activities as a possible sign of potential bad mood). We explored several scenarios to shape and refine the supporting system, in particular a community yoga class, and support groups for mild cases of depression. Obviously, for any cases relating to illness treatment, the role of professional doctors is irreplaceable, and behavioral sensing works only as a helpful tool.

In a nutshell, the core concept was using similarity metrics for uncovering hidden community resources. Assuming that communities have hidden resources under the form of people with special expertise, interests and passions, we proposed a network for allowing these people to group around shared goals and interests and engage in local activities.

Similarity matching is done based on a combination of user choice and behavioral sensing. A user can complete a profile with their data, goals, interests, etc., and decide which part of the profile they want to make public. When made public, the profile is matched with other people in the community who have similar interests. Matching profiles can be done using cryptographic techniques that do not reveal any information in case there is no match. Behavioral sensing would be used in the private sphere, and its scope of action is regulated by the user. The role of behavioral sensing is to help the user discover aspects of himself that aren't noticeable to him, and which s/he may not be ready to sincerely acknowledge. Also, behavioral sensing is used as feedback provider to allow the user to track progress towards set goals, or nascent tendencies he is trying to fight against.

The main value of such a network is located, however, out of the user's private sphere and also out of the virtual sphere, in the space of real-life sharing and exchange afforded by using the system. By networking with people with similar interests, the user can pursue her/his passion, increase motivation and commitment, join family life with professional career, and benefit from other people that are animated by similar drives, or experts that can provide expert advice or counseling.

Regarding interdisciplinarity, although the group was formed by three computer scientists and a social scientist, this did not entail a heavy orientation towards technical rather than social considerations. The initial idea came from a computer scientist and it was formulated in full consideration of its social value. During discussions, we constantly shifted from social to technical considerations. We started our group work thinking about communities and why our idea could be valuable to people. Then we shifted to a discussion of supporting technologies, and while discussing them, we realized we needed to go back to our discussion of communities to understand in what contexts a DIY solution could be more acceptable than an Internet-based network. The advancement of the concept can be tracked in its continuous transition between social and technical perspectives, until it came out shaped by both sets of considerations. This was a first, visible benefit of interdisciplinarity: conceiving technology that fits in life, and allowing those life areas that need a new technology to speak out their needs.

A second benefit was due to exchanges by which obscure terminology was clarified (especially for the social scientist) and novel perspectives considered (both sides). The exchange also revealed the advantages of employing a fluid process for conceiving a technical solution fit for a real-life context, by drawing jointly on computer and social sciences expertise.

Our proposal generated a lively debate. The concept was disputed from both a social science and a technical perspective. It was questioned whether this type of support groups

were not bordering too much on activities that would be better handled in professional environments (e.g. for depression). Issues of privacy were raised with respect to the use of behavioral sensing, and there were arguments that some users would refrain from using it, and that it could generate feelings of lack of control and agency. We discussed to what extent the benefits of behavioral sensing would make up for this type of concerns, and also how its pitfalls could be avoided by strengthening the sphere of user control and a net distinction between what is kept private and what is shared. There were also positive comments suggesting that local networks can become more important in the future, as public services may become scarce.

## 7 Interdisciplinarity

It seems that there are two main strategies to approach the task of bringing together people from different disciplines to collaborate on a specific topic. The first is to expose their differences in vocabularies, methodologies and objectives. The second is to focus on their commonalities, as for example their interest in specific case studies and applications.

During our Dagstuhl seminar we tried to do both, but it was mainly the latter which proved to be the most productive strategy. As pointed out by one participant "My biggest insight is that even when people come from different disciplines, different vocabularies, and so on, when there is a problem at hand, and there is an issue, and an idea, and people start working together, it just kind of works and if there is a misunderstanding you just solve it because people communicate." Many participants stressed the importance of case studies and application areas as a common ground for interdisciplinary exchanges. For example, "I believe that one of the big things that I got from the workshop has been examples, case studies, and going to the future I would like that we find a way to share these experiences, to share these case studies, ... just to know that one of us see that this case study is interesting because there are so many around, the selection that each one of us can do, look at this because it is interesting, would be a great way to continue the work, especially for real life examples." But it is not only that a common problem fosters efficient collaboration, since "Working with other disciplines helps you uncover problems that are probably worth addressing", as another participant stressed.

### 7.1 Learning from each other

Indeed, our interdisciplinary exchanges provided valuable information to social scientists on the capabilities of technology. For example, a social scientist stated that "This idea that from a network point of view DIY means creating your own channels of communication, I think it is useful of thinking later. I learned something new. That's very cool." And another, that "I really appreciated the opportunity to ask very technical questions to those that actually have the expertise and have better insights and go through a specific case study of a DIY network."" But for engineers it is also very important to interact with social scientists that have a different understanding of the everyday life problems that need to be addressed. As put by an engineer in our group, "I think this is maybe the thing that we lacked during the last years when we were trying to find the applications, trying to find the right use cases without asking other people what they think about it." For another, "the key take away ... was that it is probably time ... to be completely problem driven instead of being technology driven. There were small snippets of very interesting and useful things that I

picked up from people, especially from other fields. Like, if it looks like a microwave people will use it like a microwave."

Similar were the feelings of some of the activists in the group: "My take away now is that I feel I want to make some kind of flow chart for troubleshooting user issues, introducing ethnography and prototyping towards apps and networks, so create these low barriers of entry and immediate relevance for the communities." Or "I think I have slowly come to realize that I need to like step back a little bit, document and talk to more people about how this could be used, and work with people more directly to see how they can not be reliant on technologists for obvious stuff."

Summarizing this process, one of the few representatives from the industry in our group concluded: "It was interesting to watch two very different academic disciplines meet, get to know each other and over a few days start to warm to and learn from each others' approach. The process of working together on simple tasks broke down the initial posturing about the proper way things should be done and ended up working toward a new understanding that I found enriching. It is not often we get such a paradigm shift in how we view what we have been doing from the outside, it takes several days of working closely with highly skilled people looking at the world differently to make that possible."

## 7.2 Breaking the ice

It is not always obvious, however, how to break the ice and make a first step in interdisciplinary collaborations, as well as to integrate diverse preferences and encourage individuals to get out of their comfort zone. For this, there are three important lessons learned during our seminar.

First, it is important to minimize the constraints and allow for self-organization. As described by an engineer, "I would like to congratulate the organizers on the laissez-faire anarchist approach, basically we were self-organized, so that was good." And by a social scientist, "We move from philosophy to real practical cases and this is really good and it has given me some real practical ideas to take away. I was a bit worried about how open it was but I think that openness actually sort of caused some really interesting things happening."

Second, it is very important to create a pleasant atmosphere and allow oneself to have fun while working on complex scientific problems. A regular Dagstuhl participant said: "I have been in six or seven or something Dagstuhls. This is by far the funniest one ever. I realized that working with social scientists, which I did in a bilaterally basis a little bit, behave differently in groups. Like build things they wear, stuff like that. I think that this is a new methodology of implementation, that probably can be used in subsequent meetings. It was quite interesting." And another, "From a personal perspective who didn't join us yesterday cannot imagine how creative this was. This was one of the best experiences of my professional life ... [audience: it was just the wine] ... it was not the wine, it was not the company, it was not the things we talked about, it was the combination of all of it. And this was so awesome. I have the impression that you had similar spirit in the other groups. So, please organize this again. Bring us together again and see what we will follow up."

Third, it is equally important to avoid setting rigid objectives and be overly ambitious, since interdisciplinarity takes time and the phase of "getting to know each other" needs to give enough space and time to all parts to expose their point of view. "Diversity without objectives, just giving out information can make us happy. This is how I feel about DIY networks, that they can create these spaces of sharing. This is actually where ideas come

from. From just putting the information in your head and not trying to do anything specific. I mean . . . the brain does it by itself. So, I think that if we just keep sharing and putting things on the table in a diverse way, we will be happy and everything will be formed by itself." Or put slightly different by another participant, "I have to admit initially when I saw Dagstuhl and the kind of reputation that has, that I felt there was a lot of pressure, a very tight structure, having to deliver something really, something substantial. I think I appreciated and I actually relaxed much more when I realized, no, this is an opportunity to for us to actually to just get together, have very open and creative interdisciplinary discussions, and also I appreciated this kind of agility for us and lots of people going between the theoretical to the applied."

## 7.3 Open challenges

The debate that followed the proposal for a social medicine application revealed the important challenges faced by interdisciplinary research, especially when it tries to bridge the "two cultures", on the one hand, the world of arts, humanities and interpretive social sciences, and on the other hand the world of science and technology, between which, according to C. P. Snow (1959), there seems to exist an unbridgeable gap (see Frodeman et al. 2010, p. 213).

Perhaps a social scientist's request toward the engineers of the group, "don't be too creepy", best summarizes one of the most challenging tensions between the two cultures today, due to the important threats posed by technology on privacy and self-determination. But there are also more fundamental differences, related to vocabularies and methodologies. As mentioned by another social scientist, "I really like interdisciplinary working and collaboration but it is tough and it can be really frustrating at times and it can feel like so much extra work but it also pushes me to think what it is that I do, what it is that I can learn from other people in terms of thinking differently about the words that I use or the words that I don't use. So all of that has been incredibly rewarding." Or as put by an engineer, "I always teach my students to try to think out of the box. I came here and I find that whatever box you are out of, you always find another box . . . there is always room to learn new things."

Another example of tension was related to the recordings of 1-min final statements coming from all participants, from which the quotes in this section are taken from. More specifically, one of the organizers decided to take these recordings without asking the consent of the participants, inspired by a short debate on privacy issues generated by the presentation of the social medicine application and the concept of a failure machine, which could cause accidental or voluntary leakages of information by friends or companies. When participants were informed about the existence of the recording, and asked whether they would wish it to be deleted, a debate started regarding the appropriateness of this action. This debate made for an illustrative case for the various tensions existing between disciplines and individuals, and pointed to differences between theoretical concerns regarding privacy threats posed by technology (or the lack thereof), and the personal engagement in a real situation.[5]

The question of consent and deception in scientific methodology is a rather challenging area where important differences between disciplines and individuals manifest. For example,

[5] One could say that this was an example of an "artistic intervention", which aims to create impact, a "real" emotional reaction either positive or negative, in face of a certain situation. This is a fundamentally different means to evaluate a product of design than the typical rational approach, which establishes concrete performance criteria based on the maximization of a quantitative metric, such as the level of participation, measured as the number of clicks, ratings, etc.

one of the main debates between the fields of behavioural economics and experimental economics is that the former allows "deception" in laboratory experiments, which can ensure "truthful" behaviour, but which according to experimental economists is only a "short-term" achievement. In the long-term, if subjects become aware over time that deception is part of the game, they will not trust the experimenters and the whole scientific methodology will be rendered invalid. Similarly, the requirement of consent for storing and using private information is only meaningful for the short-term. Those giving their consent for their private information to be used for various purposes, assuming they really read the corresponding text or pay attention to the conditions, cannot be aware neither of possible future, different, uses of this information nor of their own respective position in the future (that might change but it will be late to take back the information made available). As an urban planner in the group stressed in relation to the concept of failure, "in planning, the perspective is a long term one, usually we may see the failures in 50 years or so."

So, despite the very positive feelings that our seminar generated regarding the possibility to bridge the gap between the "two cultures" around the design of DIY networks, we are aware that there is still a long way before resolving fundamental differences between disciplines and individuals, in the way of thinking, ethics, and attitude toward critical trade-offs and dualities that new technology brings into our life.

# 8 Interesting ideas to keep in mind

In the following we list some additional ideas exchanged through discussions at the seminar and contributions at our collaborative spaces (wiki and google doc), which we think is worth to keep in mind.

## 8.1 Lessons from the past

**Sometimes less than more ICT is needed.** As Mark Gaved stated: "the two technologies that turned out to be important for social networking were 'tea' and 'cake' ".

**Sustainability is a key challenge.** Experience says that you need to campaign for a sustained period to get stuff adopted, which is why relatively short lived academic projects (3 year PhD) don't typically get adopted much unless they get lucky. This is why we should think in terms of "initiative" (open ended) rather than project (closed time period).

**You need to build before the disaster arrives.** For example, as Jonathan Baldwin informed us, Redhook WiFi was built for community but was mostly used upon Sandy hurricane. As framed by Jon Crowcroft, the carrot is that you get a network you don't have to pay for but it isn't very good most the time, but if the internet is broken (or stick: your data coerced into a government vault) you can fall back on it, and be assured there isn't some operator who have embedded spies (like all the telcos and big cloud providers do:)

## 8.2 Looking toward the future

**Collaborative experience creation.** How/what tools and list of basic service capabilities to put in a community which facilitate unplanned applications to emerge based on local needs, art skills, creative commons enhancement. What elements of this experience are best served

by local networking links to provide sufficient advantage over traditional net tech. It often isn't enough to be "mine", must also be "better" in some way.

**Creativity.** We'd need development tools suitable for this environment, including running on mobiles (rather than always using monster design infrastructure, libraries, etc.). Applications themselves could also serve as building blocks for more complex functions (on individual devices or in the local net). It'd be nice if one could fix or adjust things that don't quite work for her needs.

**Toolkits and hybrid design.** Jon Crowcroft would like to design a toolkit that has both h/w and s/w components – could also have pieces that need to be 3D printed – then we'd have a (liberouter stored decentralized) appstore where people upload stuff they've designed so others can download, as in the Internet of Things project (see hubofallthings). The intention would be to let stuff emerge from what people do with this – so one needs to do a design which deconstrains what people can do (and therefore own) but also constrains interfaces between components, just enough (not too much) so most combinations do something, whether useful or not isn't determined (who are we to say?).

## 9 Conclusion: Toward a DIY networking community

If there was one clear take away message from this seminar, it was the desire to continue our effort and try to build an interdisciplinary community of research and action around the concept of DIY networking. There were various ideas discussed regarding specific next steps and possible meet-ups that would help us advance slowly in a self-organized fashion. The smiles, hugs, and promises for keeping in contact during the farewell phase gave us confidence that there are big chances that the seed placed by this short seminar will eventually produce exciting results.

Our first post-Dagstuhl meeting took place in London, in the context of the IETF meeting, where we tried to refine some of the ideas discussed at the end of the seminar about future steps. One of the key challenges identified was how to give incentives to people to participate in events of different disciplines. One option could be to motivate the collocation of major conferences (as happened recently with the Infocom/CHI collocation in Toronto in 2014) and provide "single registration" options. Since this wouldn't be very easy to implement in practice, the idea to fund specific people that could play the role of "representatives" in conferences of various disciplines was discussed. Another ambitious option could be to set-up a nomadic workshop on DIY networking, which could be collocated every year with a conference of a different related discipline.

Another set of ideas discussed was related to the organization of more action-oriented events targeted to specific locations in cities where workshops, hackathons, etc., could aim to produce specific solutions satisfying local needs. For example, an "urban" ExtremeCom conference taking place in challenged neighbourhoods of big cities, where DIY networking can be more than an alternative option to the Internet. Toward this direction, people from our group participate in a summer school, titled "From Smart Cities to Engaged Citizens", which will explore the design of specific solutions, including DIY networking, targeted for the city of Volos, Greece, in collaboration with local urban researchers and authorities: http://www.internet-science.eu/summer-school-2014.

Finally, Jon Crowcroft proposed a nice metaphor for interdisciplinary exchanges, the Potlach gift-giving feasts, which gave us a playful and ambitious vision to imagine: The organ-

ization of a big dedicated potlach event for interdisciplinary exchanges between researchers and activists a la Burning Man :-).

As an easier, and obvious first step, we decided to build an e-mail list which would allow us to expand our network and help us to exchange related announcements, case studies under progress, etc. Jörg Ott, has already reserved the **diynet.net** domain, which will be inaugurated soon.

## References


1. Alexander, C. (1979). The timeless way of building. Oxford University Press.
2. Antoniadis, P., Le Grande, B., Satsiou, A. Tassiulas, L., Aguiar, R., Barraca, J.P., and Sargento, S. (2008). Community building over Neighborhood Wireless Mesh Networks. IEEE Society and Technology, 27 (1): 48-56.
3. Antoniadis, P. and Apostol, I. (2013). The Neighbourhood Game: from Behavioural Economics to Urban Planning. 1st International Conference on Internet Science.
4. Apostol, I., Antoniadis, P., and Banerjee, T. (2013). Flânerie between Net and Place: Possibilities for Participation in Planning, Journal of Planning Education and Research (SAGE), 33(1): 20-33.
5. Baldwin, J. (2011). TidePools: Social WiFi. Master thesis. Available at http://www.scribd.com/doc/94601219/TidePools-Social-WiFi-Thesis
6. Basagni, S., Conti, M., Giordano, S., and Stojmenovic, I. (2013). Mobile Ad Hoc Networking: The Cutting Edge Directions. Wiley-IEEE Press.
7. Chen, C., Yavuz, E., and Karlsson, G., What a juke! A collaborative music sharing system, IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM), 2012.
8. De Cindio, F., and Schuler, D. (2012). Beyond Community Networks: From Local to Global, from Participation to Deliberation. The Journal of Community Informatics, 8(3).
9. De Cindio, F., Gentile, O., Grew, P., and Redolfi, D. (2003). Community networks: Rules of behavior and social structure. The Information Society, 19(5), 395-406.
10. Dourish, P. (2010). HCI and Environmental Sustainability: The Politics of Design and the Design of Politics. In Proceedings of the ACM Symposium on Designing Interactive Systems: DIS'2010. NY: ACM Press.
11. Dourish, P., and Bell, G. (2011). Divining a Digital Future: Mess and Mythology in Ubiquitous Computing. MIT Press.
12. Fall, K. (2003). A delay-tolerant network architecture for challenged internets, ACM SIGCOMM.
13. Farman, J. (2012). Mobile Interface Theory: Embodied Space and Locative Media. Routledge.
14. Foth, M. (Ed.). (2009). Handbook of Research on Urban Informatics: The Practice and Promise of the Real-Time City. Hershey, PA: IGI. http://eprints.qut.edu.au/13308/
15. Foth, M., Choi, J. H., and Satchell, C. (2011). Urban Informatics. In J. Bardram and N. Ducheneaut (Eds.), Proceedings of CSCW 2011 (pp. 1-8). Hangzhou, China. http://eprints.qut.edu.au/39159/
16. Foth, M., Forlano, L., Satchell, C., and Gibbs, M. (eds.). (2012). From Social Butterfly to Engaged Citizen: Urban Informatics, Social Media, Ubiquitous Computing, and Mobile Technology to Support Citizen Engagement. Cambridge, MA: MIT Press. http://eprints.qut.edu.au/39160/
17. Foth, M., Rittenbruch, M., Robinson, R., and Viller, S. (Eds.) (2014). Street Computing: Urban Informatics and City Interfaces. Abingdon, UK: Routledge. ISBN 978-0-415-84336-2. http://eprints.qut.edu.au/59160/
18. Frodeman, R., Klein, J.T. and Mitcham, C., eds (2010) The Oxford Handbook of Interdisciplinarity. Oxford, UK: Oxford University Press.
19. Gaved, M., and Mulholland, P. (2008). Pioneers, subcultures, and cooperatives: the grassroots augmentation of urban places. In Aurigi, A. and De Cindio, F. (eds.), Augmented urban spaces: articulating the physical and electronic city, England, Ashgate: pp. 171-184.
20. Jacobs, J. (1961). The Death and Life of Great American Cities. Random House, NY.
21. Jungnickel, K. (2014). DIY WIFI: Re-imagining Connectivity, Palgrave Pivot.

22 Kärkkäinen, T., Pitkanen, M. and Ott, J. (2013) Applications in Delay-Tolerant and Opportunistic Networks, in Mobile Ad Hoc Networking: Cutting Edge Directions, Second Edition (eds S. Basagni, M. Conti, S. Giordano and I. Stojmenovic), John Wiley & Sons, Inc., Hoboken, NJ, USA.

23 Lieberman, L.; Paternó, P.; Wulf,V. (eds.): End User Development, Springer, Dordrecht 2006

24 Lindgren, A. and Hui, P. (2011). ExtremeCom: To Boldly Go Where No One Has Gone Before. ACM SIGCOMM Computer Communications Review, 41 (1). pp. 54-59

25 Negroponte, N. (2002). Being wireless. WIRED, 10.10. Available at http://archive.wired.com/wired/archive/10.10/wireless.html

26 Snow, C.P. (1959). The two cultures and the scientific revolution. NY: Cambridge University Press.

27 Pitkänen, N., Kärkkäinen, T., Ott, J., Conti, M., Passarella, A., Giordano, S., Puccinelli, D., Legendre, F., Trifunovic, S., Hummel, K.A, May, M., Hegde, N., Spyropoulos, T. (2012). SCAMPI: service platform for social aware mobile and pervasive computing, Computer Communication Review 42(4): 503-508.

28 Powell, A. (2011). Metaphors, Models and Communicative Spaces: Designing local wireless infrastructure. Canadian Journal of Communication.

29 Schuler, D. (1996). New community networks: Wired for change. New York: Addison-Wesley.

30 Schuler, D. and Day, P. (2004). Shaping the Network Society: The New Role of Civil Society in Cyberspace, MIT Press.

31 Seeburger, J., Foth, M., and Tjondronegoro, D.W. (2012) The sound of music: sharing song selections between collocated strangers in public urban places. In Proceedings of the 11th International Conference on Mobile and Ubiquitous Multimedia, (MUM) 2012.

32 Scott, J., Hui, P., Crowcroft, J., and Diot, C. (2006). Haggle: A networking architecture designed around mobile users, IFIP WONS.

33 Shklovski, I., and de Souza e Silva, A. (2013). An Urban Encounter: Realizing online connectedness through local urban play. Information, Communication & Society, 16(2): 340-361.

34 Schubert, K., Weibert, A., and Wulf, V. Locating Computer Clubs in Multicultural Neighborhoods: How Collaborative Project Work Fosters Integration Processes. International Journal of Human-Computer Studies (2011).

35 Smyth, M., Helgason, I., Brynskov, M., Mitrovic, I.,Zaffiro, G. (2013). UrbanIxD: designing human interactions in the networked city. In: In CHI '13 Extended Abstracts on Human Factors in Computing Systems (CHI EA '13).

36 Whyte, W.H. (1980). The social life of small urban spaces. Washington, D.C.: Conservation Foundation.

37 Wilken, R. (2010). A Community of Strangers? Mobile Media, Art, Tactility and Urban Encounters with the Other. Mobilities, 5(4), 449-478

38 Wulf, V., Misaki, K., Atam, M., Randall, D., Rohde, M. (2013). 'On the Ground' in Sidi Bouzid: Investigating Social Media Use during the Tunisian Revolution. In Proceedings of ACM Conference on Computer Supported Cooperative Work (CSCW 2013).

39 Wulf, V., Aal, K., Abu Kteish, I., Atam, M., Schubert, K., Yerousis, D., Randall, D., Rohde, M. (2013). Fighting against the Wall: Social Media use by Political Activists in a Palestinian Village in: Proceedings of ACM Conference on Computer Human Interaction (CHI 2013).

## Participants

- Panayotis Antoniadis
ETH Zurich, CH
- Ileana Apostol
ETH Zurich, CH
- N. Asokan
University of Helsinki, FI
- Jonathan Baldwin
Interface Foundry, US
- Christian Becker
Universität Mannheim, DE
- Jon Crowcroft
University of Cambridge, GB
- Fiorella De Cindio
University of Milan, IT
- Paul Dourish
Univ. of California – Irvine, US
- Kevin R. Fall
Carnegie Mellon University, US
- Marcus Foth
Queensland University of Technology, AU
- Mark Gaved
The Open University – Milton Keynes, GB
- Per Gunningberg
Uppsala University, SE
- Ahmed Helmy
University of Florida, US
- Paul Houghton
Futurice GmbH – Berlin, DE
- Katrina Jungnickel
University of London/Goldsmiths, GB
- Teemu Kärkkäinen
Aalto University, FI
- Jussi Kangasharju
University of Helsinki, FI
- Gunnar Karlsson
KTH Royal Institute of Technology, SE
- Anders Lindgren
Swedish Institute of Computer Science – Kista, SE
- Marcin Nagy
Aalto University, FI
- Christian Nold
University College London, GB
- Jörg Ott
Aalto University, FI
- Andrea Passarella
CNR – Pisa, IT
- Dan Phiffer
The New Yorker, US
- Alison Powell
London School of Economics, GB
- Amalia Sabiescu
Coventry University, UK
- Douglas Schuler
Evergreen State College – Olympia, US
- Irina Shklovski
IT University of Copenhagen, D
- Michael Smyth
Edinburgh Napier University, GB
- Ersin Uzun
Xerox PARC – Palo Alto, US
- Volker Wulf
Universität Siegen, DE

Report from Dagstuhl Seminar 14051

# Algorithms for Wireless Communication

**Edited by**

**Guy Even[1], Magnús M. Halldórsson[2], Yvonne Anne Pignolet[3], and Christian Scheideler[4]**

1 **Tel Aviv University, IL,** guy@eng.tau.ac.il
2 **Reykjavik University, IS,** mmh@ru.is
3 **ABB Corporate Research, Baden, CH,** yvonne-anne.pignolet@ch.abb.com
4 **Universität Paderborn, DE,** scheideler@uni-paderborn.de

**Abstract**

This report documents the talks and discussions of Dagstuhl Seminar 14051 "Algorithms for Wireless Communication". The presented talks represent a wide spectrum of work on wireless networks. The topic of wireless communication continues to grow in many domains, new applications and deployments of wireless networks in a variety of contexts are being reported. A key focus of the talks and discussions presented here is to discuss models for wireless networks as well as algorithmic results and real world deployments.

**Seminar** January 26–31, 2014 – http://www.dagstuhl.de/14051
**1998 ACM Subject Classification** C.2.1 Network Architecture and Design: Wireless Communication, C.2.2 Network Protocols, C.4 Performance of Systems: Modelling Techniques, E.1 Data Structures: Graphs and Networks, F.2 Analysis of Algorithms and Problem Complexity, G.2.2 Graph Theory: Graph Algorithms
**Keywords and phrases** wireless, algorithms, model, complexity
**Digital Object Identifier** 10.4230/DagRep.4.1.152

## 1 Executive Summary

*Guy Even*
*Magnús M. Halldórsson*
*Yvonne Anne Pignolet*
*Christian Scheideler*



The last decades have seen an ever growing interest in wireless communication networks and their applications. Wireless networks pose many algorithmic challenges for various reasons: Realistic wireless signal propagation and interference models are very complex and therefore hard to use in rigorous algorithmic research, and this is further complicated by emerging technologies such as MIMO (multiple-input and multiple-output). Also, reasonable models for the dynamics and mobility in these networks can be quite complex and are not yet well-understood. Furthermore, standard complexity measures such as time and space are not sufficient any more as energy consumption is also a critical aspect that cannot be neglected. Many protocols for wireless networks have already been proposed by the research community, but most of them have only been studied in simulations or analyzed using rather simple models. So there is doubt whether any of these protocols would actually work in practice.




The purpose of this Dagstuhl seminar was to bring together computer scientists of different backgrounds to review and discuss models and algorithmic approaches in order to obtain a better understanding of the capabilities and limitations of modern wireless networks and to come up with more realistic models and approaches for future research on wireless networks that may then be investigated in joint research projects. The mix of the participating people resulted in fruitful discussions and interesting information exchange. The structure of the seminar took advantage of these different backgrounds by focusing on themed talks and open discussions.

The program included an eclectic mix of algorithmic and systems perspectives, modeling issues and emerging networking techniques, and explorations of the limits and possibilities of fundamental problems.

Discussions of models ranged from simple graph-based communication and interference models, to stochastic models, adversarial interruptions and jamming, dynamic networks and uncertainty formulations, and variations and extensions of signal-strength models.

Presentations from the systems perspective included managing environmental factors affecting measurements, robust predictions of channel capacities, efficiency of backpressure routing, issues in emerging heterogeneous radio environmental contexts, and robots controlled via wireless communication.

New dimensions at different networking layers included MIMO, network coding, interference cancellation, directional antennas and cognitive radio networks.

Finally, new results were presented on various related classic problems including broadcast, local broadcast, game theory, coding, routing, positioning, and connectivity.

## 2 Table of Contents

# 3 Overview of Talks

## 3.1 Jamming-Resistant Learning in Wireless Networks

*Johannes Dams (RWTH Aachen, DE)*


**Joint work of** Johannes Dams, Martin Hoefer and Thomas Kesselheim

We consider capacity maximization in wireless networks under adversarial interference conditions. There are $n$ links, each consisting of a sender and a receiver, which repeatedly try to perform a successful transmission. In each time step, the success of attempted transmissions depends on interference conditions, which are captured by an interference model (e.g. the SINR model). Additionally, an adversarial jammer can render a $(1-\delta)$-fraction of time steps unsuccessful. Our main result is an algorithm based on no-regret learning converging to an $O(1/\delta)$-approximation. It provides even a constant-factor approximation when the jammer exactly blocks a $(1-\delta)$-fraction of time steps. In addition, we consider a stochastic jammer, for which we obtain a constant-factor approximation after a polynomial number of time steps. Using this learning approach and the general proof technique, we can even extend the results to more general settings, in which links arrive and depart dynamically, and where each sender tries to reach multiple receivers. Though these results cannot directly be applied to a setting with multiple channels and stochastic channel availabilities, we can achieve a constant-factor approximation using other learning approaches.

## 3.2 Braess Paradox in Wireless Networks: The Danger of Improved Technology

*Michael Dinitz (Johns Hopkins University – Baltimore, US)*


**Joint work of** Michael Dinitz and Merav Parter

When comparing new wireless technologies, it is common to consider the effect that they have on the capacity of the network (defined as the maximum number of simultaneously satisfiable links). For example, it has been shown that giving receivers the ability to do interference cancellation, or allowing transmitters to use power control, never decreases the capacity and can in certain cases increase it. But there is no reason to expect the optimal capacity to be realized in practice, particularly since maximizing the capacity is known to be NP-hard. In reality, we would expect links to behave as self- interested agents, and thus when introducing a new technology it makes more sense to compare the values reached at game-theoretic equilibria than the optimum values.

In this paper we initiate this line of work by comparing various notions of equilibria (particularly Nash equilibria and no-regret behavior) when using a supposedly "better" technology. We show a version of Braess Paradox for all of them: in certain networks, upgrading technology can actually make the equilibria worse, despite an increase in the capacity. We construct instances where this decrease is a constant factor for power control, interference cancellation, and improvements in the SINR threshold $\beta$, and is $O(\log n)$ when power control is combined with interference cancellation. However, we show that these examples are basically tight: the decrease is at most $O(1)$ for power control, interference cancellation, and improved $\beta$, and is at most $O(\log n)$ when power control is combined with interference cancellation.

## 3.3 Frequency Hopping against a Powerful Adversary

*Yuval Emek (Technion, IL)*


**Joint work of** Yuval Emek and Roger Wattenhofer

Frequency hopping is a central method in wireless communication, offering improved resistance to adversarial interference and interception attempts and easy non-coordinated control in dynamic environments. In this talk, we introduce a new model that supports a rigorous study of frequency hopping in adversarial settings. We then propose new frequency hopping protocols that allow a sender- receiver pair to exploit essentially the full communication capacity despite a powerful adversary that can scan and jam a significant amount of the ongoing transmissions.

## 3.4 Deterministic Rateless Codes for BSC

*Guy Even (Tel Aviv University, IL)*


**Joint work of** Guy Even, Benny Applebaum and Liron David

A rateless code encodes a finite length information word into an infinitely long codeword. The length of the prefix of the noisy codeword required by the decoder depends on signal-to-noise ratio of the communication channel. A rateless code achieves capacity for a family of channels if, for every channel in the family, reliable communication is obtained with a rate that is arbitrarily close to the channel's capacity. The encoder is universal because the same encoding is used for all channels in the family.

We construct the first *deterministic* rateless code for the binary symmetric channel. Our code can be encoded and decoded in $O(\log \log k)$ time per bit and in poly-logarithmic parallel time. Furthermore, the error probability of our code is almost exponentially small $\exp(\Omega(k/poly \log \log(k)))$. Previous rateless codes are probabilistic (i.e., based on code ensembles), require polynomial time per bit for decoding, and have inferior error probabilities.

## 3.5 Robot swarms as mobile sensor networks

*Sándor Fekete (TU Braunschweig, DE)*


**Joint work of** Becker, Aaron; Demaine, Erik; Fekete, Sándor; Habibi, Golnaz; Kroeller, Alexander; Lee, Soung Kyou; McLurkin, James; Schmidt, Christiane

We present two studies of algorithmic methods for swarms of mobile devices as sensor networks. In the first, we show a video describing exploration and mapping of an unknown environment by a swarm of robots with limited capabilities. In the second, we discuss positive and negative results for controlling a massive swarm of particles by uniform external forces in the presence of obstacles.

## 3.6 Arbitrary Transmission Power in the SINR Model

*Fabian Fuchs (KIT – Karlsruhe Institute of Technology, DE)*



Both in the light of energy conservation and the expansion of existing networks, wireless networks face the challenge of nodes with heterogeneous transmission power. However, for more realistic models of wireless communication only few algorithmic results are known. In this talk we propose to consider the transmission power as input. Thus we consider nodes with arbitrary power assignment in the so-called physical or SINR model. Our first result is a bound on the probabilistic interference from all simultaneously transmitting nodes on receivers. This result implies that current local broadcasting algorithms can be generalized to the case of non-uniform transmission power with minor changes. Also, we introduce a new network parameter $\ell$ that measures the length of the longest unidirectional path in the network. After showing that a dependence on $\ell$ is inevitable for distributed node coloring, we present a coloring algorithm its time complexity.

## 3.7 Modeling and Analysis of Wireless Networks using Stochastic Geometry

*Martin Haenggi (University of Notre Dame, US)*



Analyses of specific configurations of wireless networks are tedious and yield results without any generality. In contrast, wireless network ensembles described by spatial stochastic models are fairly tractable and give insight for a general class of networks. Stochastic geometry, in particular point process theory, provides these both these models and the mathematical tools for their analysis. The point process represents the locations of the nodes in a wireless networks. A prominent example is the Poisson point process, for which many interesting closed-form results have been derived. Originally used for sensor and ad hoc networks, it is now increasingly popular also as a model for cellular networks. While simple graphs do not constitute useful models for wireless networks, random geometric graphs based on point processes and the SINR link model are powerful and versatile performance indicators. Many metrics of interest can be extracted from such models, but it needs to be understood that many protocol decisions at the lower layers must be made before a meaningful graph can be defined. In other words, graphs should always be interpreted as the result of certain protocol choices, rather than as the starting point of an analysis or algorithm development.

## 3.8 Network Coding for Multi-Hop Wireless Broadcast

*Bernhard Haeupler (Microsoft Research – Mountain View, US)*



We introduce a simple distributed implementation of random linear network coding (RLNC) [8] and gives several scenarios arising in wireless network broadcast settings in which RLNC achieves large gains over traditional routing strategies.

After explaining an novel, simpler technique for analyzing RLNC [5] we survey several recent results obtaining the first throughput optimal algorithms in different wireless models using RLNC. In particular simple proof (sketches) for a throughput optimal broadcast in the radio network model [3, 4] and the dynamic network model [6, 2, 7] are presented.

Lastly we give some recent results on the network coding gap for broadcast in the radio network model [1].

### References

1 N. Alon, M. Ghaffari, B. Haeupler, and M. Khabbazian. Broadcast throughput in radio networks: Routing vs. network coding. *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1831–1843, 2014.
2 C. Dutta, G. Pandurangan, Z. Sun, and E. Viola. On the Complexity of Information Spreading in Dynamic Networks. *ACM/SIAM Symposium on Discrete Algorithms (SODA)*, 2013.
3 M. Ghaffari and B. Haeupler. Fast structuring of radio networks large for multi-message communications. *International Symposium on Distributed Computing (DISC)*, 8205:492–506, 2013.
4 M. Ghaffari, B. Haeupler, and M. Khabbazian. Randomized broadcast in radio networks with collision detection. *ACM Symposium on Principles of Distributed Computing (PODC)*, pages 325–334, 2013.
5 B. Haeupler. Analyzing Network Coding Gossip Made Easy. *ACM Symposium on Theory of Computing (STOC)*, pages 293–302, 2011.
6 B. Haeupler and D. Karger. Faster Information Dissemination in Dynamic Networks via Network Coding. *ACM Symposium on Principles of Distributed Computing (PODC)*, pages 381–390, 2011.
7 B. Haeupler and F. Kuhn. Lower bounds on information dissemination in dynamic networks. *International Symposium on Distributed Computing (DISC)*, pages 166–180, 2012.
8 T. Ho, M. Médard, R. Koetter, D. R. Karger, M. Effros, J. Shi, and B. Leong. A random linear network coding approach to multicast. *IEEE Transactions on Information Theory (TransInf)*, 52(10):4413–4430, 2006.

## 3.9 Extending SINR to Realistic Environments

*Magnús M. Halldórsson (Reykjavik University, IS)*


Joint work of Marijke Bodlaender and Magnús M. Halldórsson

Signal-strength models of wireless communications capture the gradual fading of signals and the additivity of interference. As such, they are closer to reality than other models. However, nearly all theoretic work in the SINR model depends on the assumption of smooth geometric decay, one that is true in free space but is far off in actual environments. The challenge is to

model realistic environments, including walls, obstacles, reflections and anisotropic antennas, without making the models algorithmically impractical or analytically intractable.

We present a simple solution that allows the modeling of arbitrary static situations by moving from geometry to arbitrary decay spaces. The complexity of a setting is captured by a "metricity" parameter $\zeta$ that indicates how far the decay space is from satisfying the triangular inequality. All results that hold in the SINR model in general metrics carry over to decay spaces, with the resulting time complexity and approximation depending on $\zeta$ in the same way that the original results depends on the path loss term $\alpha$. For distributed algorithms, that to date have appeared to necessarily depend on the planarity, we indicate how they can be adapted to arbitrary decay spaces at a cost in time complexity that depends on a fading parameter of the decay space. In particular, for decay spaces that are doubling, the parameter is constant-bounded.

## 3.10 Online Independet Set with Stochastic Adversaries

*Martin Hoefer (Universität des Saarlandes, DE)*


**Joint work of** Goebel, Oliver; Hoefer, Martin; Kesselheim, Thomas; Schleiden, Thomas; Vöcking, Berthold

We investigate online algorithms for maximum (weight) independent set on graph classes with bounded inductive independence number $\rho$ like interval and disk graphs with applications to, e.g., task scheduling, spectrum allocation and admission control. In the online setting, nodes of an unknown graph arrive one by one over time. An online algorithm has to decide whether an arriving node should be included into the independent set.

Traditional (worst-case) competitive analysis yields only devastating results. Hence, we conduct a stochastic analysis of the problem and introduce a generic sampling approach that allows to devise online algorithms for a variety of input models. It bridges between models of quite different nature – it covers the secretary model, in which an adversarial graph is presented in random order, and the prophet-inequality model, in which a randomly generated graph is presented in adversarial order.

Our first result is an online algorithm for maximum independent set with a competitive ratio of $O(\rho^2)$ in all considered models. It can be extended to maximum-weight independent set by losing only a factor of $O(\log n)$, with n denoting the (expected) number of nodes. This upper bound is complemented by a lower bound of $\Omega(\log n / \log^2 \log n)$ showing that our sampling approach achieves nearly the optimal competitive ratio in all considered models. In addition, we present various extensions, e.g., towards admission control in wireless networks under SINR constraints.

## 3.11 Multichannel Information Dissemination

*Stephan Holzer (MIT – Cambridge, US)*


**Joint work of** Stephan Holzer, Thomas Locher, Yvonne Anne Pignolet, Roger Wattenhofer

We study the information exchange problem on a set of multiple access channels: $k$ arbitrary nodes have information they want to distribute to the entire network via a shared medium partitioned into channels. We present algorithms and lower bounds on the time and channel complexity for disseminating these $k$ information items in a single-hop network of $n$ nodes. More precisely, we devise a deterministic algorithm running in asymptotically optimal time $O(k)$ using $O(n^{\log(k)/k})$ channels if $k \leq \frac{1}{6}\log n$ and $O(\log^{1+\rho}(n/k))$ channels otherwise, where $\rho > 0$ is an arbitrarily small constant. In addition, we show that $\Omega(n^{\Omega(1/k)} + \log_k n)$ channels are necessary to achieve this time complexity.

## 3.12 Nearly Optimal Asynchronous Blind Rendezvous Algorithm for Cognitive Radio Networks

*Qiang-Sheng Hua (Tsinghua University – Beijing, CN)*



Rendezvous is a fundamental process in Cognitive Radio Networks, through which a user establishes a link to communicate with a neighbor on a common channel. Most previous solutions use either a central controller or a Common Control Channel (CCC) to simplify the problem, which are inflexible and vulnerable to faults and attacks. Some blind rendezvous algorithms have been proposed that rely on no centralization. Channel Hopping (CH) is a representative technique used in blind rendezvous, with which each user hops among the available channels according to a pre-defined sequence. However, no existing algorithms can work efficiently for both symmetric (both parties have the same set of channels) and asymmetric users. In this paper, we introduce a new notion called Disjoint Relaxed Difference Set (DRDS) and present a linear time constant approximation algorithm for its construction. Then based on the DRDS, we propose a distributed asynchronous algorithm that can achieve and guarantee fast rendezvous for both symmetric and asymmetric users. We also derive a lower bound for any algorithm using the CH technique. This lower bound shows that our proposed DRDS based distributed rendezvous algorithm is nearly optimal. Extensive simulation results corroborate our theoretical analysis.

## 3.13 Distributed Protocols for SINR

*Tomasz Jurdziński (University of Wroclaw, PL)*



In the advent of large-scale multi-hop wireless technologies, it is of utmost importance to devise efficient distributed protocols to provide basic communication. Unlike in the graph-based radio network model, algorithmic issues for multi-hop communication in the SINR model are not well understood. Especially, very few deterministic solutions are known.

The presentation addresses this issue for ad hoc SINR networks in the Euclidean space. Especially, the goal is to compare efficiency of determinstic and randomized solutions as well as to address the issue of the impact of knowledge of positions on complexity of basic communication problems.

## 3.14 Three wireless systems problems

*Holger Karl (Universität Paderborn, DE)*



The talk briefly highlights three problems in wireless communication. First, the streaming of video to mobile devices exploiting predictions of cellular capacity and user behavior. Second, the optimization of mobile backhaul networks to support wireless techniques like cooperative multipoint. Third, an optimization problem that combines cooperative and frequency diversity techniques, using OFDMA in a relay-extended cellular context.

## 3.15 Algorithmic Problems that Arise from Shifting to Directional Antennas

*Matthew Katz (Ben Gurion University – Beer Sheva, IL)*



Directional antennas have several noticeable advantages over omni-directional antennas. For example, when using a directional antenna, less power is required to transmit to a given receiver and less interference is caused by this transmission. Despite these advantages, the vast majority of the papers dealing with algorithmic problems motivated by wireless networks, consider omni-directional antennas, whose coverage area is often modeled by a disk. In this talk we discuss several basic algorithmic problems that arise when switching to directional antennas.

Let P be a set of points in the plane representing transceivers, and assume that each transceiver is equipped with a directional antenna. The coverage area of a directional antenna at point p of angle alpha, is a circular sector of angle alpha centered at p, where the orientation and range of the antenna can be adjusted. For a given assignment of orientations and ranges, the induced symmetric communication graph (SCG) of P is the undirected graph, in which there is an edge between two vertices (i.e., points) u and v if and only if v lies in u's sector and vice versa. The induced asymmetric communication graph (DCG) of P is the directed graph, in which there is a directed edge from u to v if and only if v lies in u's sector.

We consider several problems arising in wireless networks with directional antennas, under both the symmetric and asymmetric models, including orientation and range assignment, power assignment, and interference reduction. In the orientation and power assignment problem, one needs to assign an orientation and range to each of the antennas, so that the resulting communication graph is connected (in the symmetric model) or strongly connected (in the asymmetric model). The goal is to do so while keeping the ranges short.

The interference of a network is defined as the maximum in-degree of a node in DCG, i.e., the maximum number of transceivers covering a receiver. We address the following question

under both models: What is the minimum interference I, such that for any set P of points in the plane, representing transceivers equipped with a directional antenna of angle alpha, one can assign orientations and ranges to the points in P, so that the induced communication graph G is either connected or strongly connected and the interference of G does not exceed I. We show that in the symmetric model, the advantage of directional antennas with respect to omni- directional antennas is less obvious.

## 3.16 Techniques for SINR Approximation Algorithms

*Thomas Kesselheim (Cornell University, US)*



We consider the capacity-maximization problem in the SINR model: Given pairs of senders and receivers, select a maximum subset of these such that all transmissions can be carried our simultaneously. For the problem variant with uniform powers, there is a constant-factor approximation due to Halldorsson and Mitra [SODA 2011]. We discuss a simplified analysis in spirit of [Kesselheim, ESA 2012]. The algorithm is similar in structure to the one for the problem variant where powers have to be chosen [Kesselheim, SODA 2011]. This allows us to derive an abstraction based on edge-weighted conflict graphs, which is useful for advanced problems.

## 3.17 Backpressure Routing in Wireless Networks: From Theory to Practice

*Bhaskar Krishnamachari (University of Southern California, US)*


**Joint work of** Bhaskar Krishnamachari, Scott Moeller, Majed Alresaini, Mike Neely, Andrea Gasparri, Shangxing Wang, Longbo Huang, Avinash Ridharan, and Omprakash Gnawali

Since the work by Tassiulas and Ephremides in 1992, Backpressure scheduling has been a focus of intense research by network control theorists. I describe our work on translating this theory into practice in the form of Backpressure Collection Protocol (BCP), the first implementation of dynamic bacpressure routing. I also discuss our work on backpressure with adaptive redundancy (BWAR), which utilizes multi-copy routing to improve the delay of backpressure routing for intermittently connected mobile networks. Finally, I present ongoing work on using backpressure scheduling to control the motion of message ferrying robots.

The work described was funded in part by the U.S. National Science Foundation via CNS-1049541.

## 3.18 Probabilistic Analysis of Power Assignments

*Bodo Manthey (University of Twente, NL)*


**Joint work of** Bodo Manthey and Maurits de Graaf

We consider the problem of assigning transmission powers to the devices of a wireless network such that the resulting communication graph is connected and the total transmit power is minimized. Our goal is a probabilistic analysis of this power assignment (PA) problem.

We prove complete convergence of PA for arbitrary combinations of the dimension d and the distance-power gradient $p$. In particular, we prove complete convergence for the case $p \geq d$. As far as we are aware, complete convergence for $p > d$ has not been proved yet for any Euclidean functional.

Furthermore, we prove that the expected approximation ratio of a simple spanning tree heuristic is strictly less than its worst-case ratio of 2.

## 3.19 Continuous Local Strategies for Robotic Formation Problems

*Friedhelm Meyer auf der Heide (Universität Paderborn, DE)*


**Joint work of** Bastian Degener, Barbara Kempkes and Peter Kling

I study a scenario in which n mobile robots with a limited viewing range are distributed in the plane and have to solve a formation problem. The formation problems considered are the GATHERING problem and the CHAIN- FORMATION problem. In the GATHERING problem the robots have to gather in one (not predefined) point, while in the CHAIN-FORMATION problem they have to form a connected chain of minimal length between two stationary base stations. Each robot has to base its decisions where to move only on the current relative positions of neighboring robots within its viewing range. Variants of these problems (especially for the GATHERING problem) have been studied extensively in different discrete time models. In this talk, I focus on a continuous time model where robots continuously sense the positions of other robots within their viewing range and continuously adapt their speed and direction according to some simple, local rules. Hereby, I assume that the robots have to obey a speed limit of one. I present strategies for both problems with $O(n)$ runtime bounds. Moreover, I show that their runtimes are at most by a factor $O(\log(OPT))$ and $O(\log(n))$, resp., away from the optimal time OPT necessary by a global algorithm, for GATHERING and CHAIN-FORMATION, resp.

## 3.20 Thoughts on Models for Wireless Networks

*Calvin Newport (Georgetown University – Washington, US)*



In my talk, I presented three thoughts about wireless algorithm research. The first thought was that it is useful to make a distinction between basic science research (exploring the capabilities and limits of wireless communication) and applied research (developing algorithms for use in practice). My second thought was that for applied research, it is often useful to

introduce uncertainty to the relevant model. My third thought was that for applied research, we should also consider studying problems at the network layer, focusing less on contention (which is hard to formally model) and more on other aspects of wireless that show up at higher layers (uncertainty in delay and receivers, lack of network knowledge, etc.).

## 3.21 The topology of wireless communication and applications

*Merav Parter (Weizmann Institute – Rehovot, IL)*


**Joint work of** Chen Avin, Asaf Cohen, Yoram Haddad, Erez Kantor, Zvi Lotker, Merav Parter and David Peleg

We study the topological properties of wireless communication maps and their usability in algorithmic design. We consider the SINR model, which compares the received power of a signal at a receiver against the sum of strengths of other interfering signals plus background noise. To describe the behavior of a multi-station network, we use the convenient representation of a *reception map*. In the SINR model, the resulting *SINR diagram* partitions the plane into reception zones, one per station, and the complementary region of the plane where no station can be heard. We consider the general case where transmission energies are arbitrary (or non-uniform). Under that setting, the reception zones are not necessarily convex or even connected. This poses the algorithmic challenge of designing efficient point location techniques as well as the theoretical challenge of understanding the geometry of SINR diagrams. We achieve several results in both directions. One of our key results concerns the behavior of a (d+1)-dimensional map. Specifically, although the d-dimensional map might be highly fractured, drawing the map in one dimension higher "heals" the zones, which become connected. In addition, we study the topology of reception regions when reception points are allowed to decode messages using interference cancellation. Our final note concerns the connection between SINR diagrams and Voronoi diagrams.

## 3.22 Enhancing Future Networks with Radio Environmental Context

*Marina Petrova (RWTH Aachen, DE)*



As the demand for ever high data rates is continuing to increase driven by the global growth in smart phones and tablets, it is of great importance to re-think the spectrum allocation strategies and and at the same time enable cognitive radio technology to optimally use the available spectrum resources. Better understanding the radio environment in time and space and being able to dynamically predict interference could potentially help in providing capacity whenever needed and improving the wireless connectivity and service to the user. A Radio Environment Map (REM) is an advanced knowledge base that stores live multi-domain information on the entities in the network as well as the environment historically. The main functionality of a REM is the construction of dynamic interference map for each frequency at each location of interest by collecting spectrum measurements. By using advanced data processing methods and with a help of geographical terrain models, propagation environment, and regulations it can also estimate the state of locations where there is no measurement

data. In this talk we will talk about the use of such REMs for predicting network coverage holes and discuss how using spatial statistics methods such as Kriging can produce reliable predictions under certain data accuracy and measurement density requirements.

## 3.23 Environmental-Aware Protocols for Networked Embedded Systems

*Kay Römer (TU Graz, AT)*



Sensor networks prodive a substrate to realize applications in several domains of utmost importance for our society, including surveillance of critical infrastructures, smart cities, smart grids, and smart healthcare. However, many of these applications are only possible if sensor networks provides dependable performance. Application-specific guarantees on network performance parameters such as data delivery reliability and latency must be given for all system operation conditions. Failure to meet these requirements at all times may lead to reduced user satisfaction, increased costs, or to critical system failures. Unfortunately, existing sensor network technologies mostly follow a best effort approach and do not offer guaranteed performance.

The major hurdle to providing dependable sensor is that their operation is deeply affected by their surrounding environment. Environmental properties such as electromagnetic (EM) radiation, ambient temperature, and humidity have significant impact on achievable network performance. In particular, environmental temperature deeply affects the operation of sensor networks. We have shown in previous work that temperature variations in a deployment may lead to failing transmissions during hot periods. Not only are these environmental conditions hard to predict for a given deployment site, they also may largely vary from one deployment site to another, thus hindering scalable deployment of applications as every new deployment site requires costly customization.

In order to design sensor networks that can provide certain performance guarantees despite changing environmental conditions, there is a need for testbeds with realistic environmental effects, where protocols and applications can be run on real sensor network hardware under repeatable and realistic environmental conditions. In this talk, we first present TempLab, a testbed where user-defined temperature conditions can be created. For this, sensor nodes are equipped with infrared heating lamps that can be controlled via wireless dimmers to create an accurate temperature profile that varies over time and space. Using this testbed, we systematically study the impact of temperature on the signal-to-noise ratio of a wireless link and provide a model that accurately predicts SNR for given transmitter and receiver temperatures. We also study the impact of temperature on the frequency of processor clocks, finding that a temperature change of 30 degrees slows down the processor clock by more than 10 percent. Finally we investigate the impact of temperature on the RPL routing protocol and find that the topology of the routing tree is significantly affected, even leading to network partitions.

## 3.24 Flow/back-pressure/slide approaches for routing and scheduling: from wired to wireless networks

*Adi Rosén (University Paris-Diderot, FR)*



We shortly review (older) results about flow/back-pressure/slide techniques for routing and scheduling in stable and dynamic networks. In particular we review results that show that such protocols, albeit being online and distributed, maintain stability in any network as long as the traffic injected into the network allows stability (roughly speaking, as long as there exist paths for the injected packets which imply injection rate at most 1). We then discuss the challenges of adapting these protocols to wireless networks and some perhaps promising directions for such adaptation.

## 3.25 SINR with adversarial noise

*Christian Scheideler (Universität Paderborn, DE)*


**Joint work of** Adrian Ogierman, Andrea Richa, Christian Scheideler, Stefan Schmid, and Jin Zhang

In my talk I consider the problem of how to efficiently share a wireless medium which is subject to harsh external interference or even jamming. While this problem has already been studied intensively for simplistic single- hop or unit disk graph models, almost no work is known for the SINR interference model.

The talk consists of two parts. First, I introduce a new adversarial SINR model which captures a wide range of interference phenomena. Concretely, I consider a powerful, adaptive adversary which can jam nodes at arbitrary times and which is only limited by some energy budget. In the second part I present a distributed MAC protocol which provably achieves a constant competitive throughput in this environment: I show that, with high probability, the protocol ensures that a constant fraction of the non-blocked time periods is used for successful transmissions. The result also highlights an inherent difference between the SINR model and unit disk graph models.

## 3.26 The Power of MIMO

*Christian Schindelhauer (Universität Freiburg, DE)*


**Joint work of** Christian Schindelhauer and Thomas Janson

We present a wireless network model for MIMO, i.e. coordinated senders and receivers. First, we show that this model reduces to the standard SINR- model in the case of random senders and uncoordinated receivers. Then, we concentrate on the case of senders and receivers on a line. In our free-space model we coordinate the sender phase and amplitude in such a way that on the left we produce a beam while on the other side the path loss coefficient can be arbitrarily small. For equidistant ad-hoc nodes, which only coordinate themselves by

the receipt times of messages we present a broadcasting algorithm with time $\mathcal{O}(\log n)$ for $n$ nodes. In this setting the energy and SINR ratio only allows communication to the next node. So, cooperated to SINR we improve the time from $n-1$ to $\mathcal{O}(\log n)$.

## 3.27 Deterministic Blind Rendezvous in Cognitive Radio Networks

*Ravi Sundaram (Northeastern University – Boston, US)*


**Joint work of** Chen, Sixia; Russell, Alexander; Samanta, Abhishek; Sundaram, Ravi

Blind rendezvous is a fundamental problem in cognitive radio networks. The problem involves a collection of agents (radios) that wish to discover each other in the blind setting where there is no shared infrastructure and they initially have no knowledge of each other. Time is divided into discrete slots; spectrum is divided into discrete channels, $\{1, 2, ..., n\}$. Each agent may access a single channel in a single time slot and we say that two agents rendezvous when they access the same channel in the same time slot. The model is asymmetric: each agent $A_i$ may only use a particular subset $S_i$ of the channels and different agents may have access to different subsets of channels. The goal is to design deterministic channel hopping schedules for each agent so as to guarantee rendezvous between any pair of agents with overlapping channel sets. Two independent sets of authors, Shin et al. and Lin et al., gave the first constructions guaranteeing asynchronous blind rendezvous in $O(n^2)$ and $O(n^3)$ time, respectively. We present a substantially improved construction. Our results are the first that achieve nontrivial dependence on the size of the set of available channels. This allows us, for example, to save roughly a quadratic factor over the best previous results in the important case when channel subsets have constant size. We also achieve the best possible bound of O(1) time for the symmetric situation; previous works could do no better than O(n). Using the probabilistic method and Ramsey theory we provide evidence in support of our suspicion that our construction is asymptotically optimal for small size channel subsets.

## 3.28 Algorithmic Aspects of Geometric Routing on Mobile Ad Hoc Network

*Takeshi Tokuyama (Tohoku University, JP)*



In this talk, we discuss algorithms for geometric routing on a mobile ad hoc network. Although geometric routing on a static ad hoc network is well studied, the maintenance of the underlying network such as Delaunay graph causes difficulty if nodes move freely.

We present a simple and efficient strategy of geometric routing on moving nodes without much effort to update the underlying network. We also show the gap of theory and practice on mobile ad hoc network, and propose a research direction of geometric routing so that it can be used in practice.

## Participants

- Eyjólfur Ingi Asgeirsson
Reykjavik University, IS
- Marijke Bodlaender
Reykjavik University, IS
- Johannes Dams
RWTH Aachen, DE
- Michael Dinitz
Johns Hopkins University – Baltimore, US
- Yuval Emek
Technion, IL
- Guy Even
Tel Aviv University, IL
- Sándor Fekete
TU Braunschweig, DE
- Fabian Fuchs
KIT – Karlsruhe Institute of Technology, DE
- Jie Gao
SUNY – Stony Brook, US
- Seth Gilbert
National Univ. of Singapore, SG
- Martin Haenggi
University of Notre Dame, US
- Bernhard Haeupler
Microsoft Corp. – Mountain View, US
- Magnus M. Halldorsson
Reykjavik University, IS
- Martin Hoefer
Universität des Saarlandes, DE
- Stephan Holzer
MIT – Cambridge, US
- Qiang-Sheng Hua
Tsinghua Univ. – Beijing, CN
- Thomas Janson
Universität Freiburg, DE
- Tomasz Jurdzinski
University of Wroclaw, PL
- Holger Karl
Universität Paderborn, DE
- Matthew J. Katz
Ben Gurion University – Beer Sheva, IL
- Thomas Kesselheim
Cornell University, US
- Christian Konrad
Reykjavik University, IS
- Bhaskar Krishnamachari
Univ. of Southern California, US
- Fabian Daniel Kuhn
Universität Freiburg, DE
- Bodo Manthey
University of Twente, NL
- Friedhelm Meyer auf der Heide
Universität Paderborn, DE
- Calvin Newport
Georgetown University – Washington, US
- Merav Parter
Weizmann Institute – Rehovot, IL
- Marina Petrova
RWTH Aachen, DE
- Yvonne-Anne Pignolet
ABB – Baden-Dättwil, CH
- Dror Rawitz
Bar-Ilan University, IL
- Kay Römer
TU Graz, AT
- Adi Rosén
University Paris-Diderot, FR
- Alexander Russell
University of Connecticut – Storrs, US
- Christian Scheideler
Universität Paderborn, DE
- Christian Schindelhauer
Universität Freiburg, DE
- Ravi Sundaram
Northeastern Univ. – Boston, US
- Takeshi Tokuyama
Tohoku University, JP
- Roger Wattenhofer
ETH Zürich, CH

Report from Dagstuhl Seminar 14052

# Ethics in Data Sharing

**Edited by**

**Julie Cohen[1], Sven Dietrich[2], Aiko Pras[3], Lenore D. Zuck[4], and Mireille Hildebrand[5]**

1 **Georgetown University – Washington, DC, US,** jec@law.georgetown.edu
2 **Stevens Institute of Technology, Hoboken, NJ, US,** spock@cs.stevens.edu
3 **University of Twente, NL,** a.pras@utwente.nl
4 **University of Illinois at Chicago – Chicago, IL, US,** lenore@cs.uic.edu
5 **Vrije Universiteit Brussel – Brussels, BE**

**Abstract**

This report documents the program and the outcomes of Dagstuhl Seminar 14052 "Ethics in Data Sharing". The seminar brought together computer scientists, an ethicist and legal scholars to discuss the topic of "ethics in data sharing."

**Seminar** January 26–31, 2014 – http://www.dagstuhl.de/14052
**1998 ACM Subject Classification** K.4.1 Public Policy Issues, K.6.5 Security and Protection, K.7.4 Professional Ethics
**Keywords and phrases** Ethics, Data Sharing Practices, Data Dissemination, Ethics across borders, Anonymization, Sanitization
**Digital Object Identifier** 10.4230/DagRep.4.1.170

## 1 Executive Summary

*Julie Cohen*
*Sven Dietrich*
*Aiko Pras*
*Lenore D. Zuck*
*Mireille Hildebrandt*



ACM's ethical guidelines (as well as IEEE's) are almost two decades old. The most relevant points to data sharing it makes are "Avoid harm to others" and "Respect the privacy of others." The consequences of not complying with the code are "Treat violations of this code as inconsistent with membership in the ACM" while "Adherence of professionals to a code of ethics is largely a voluntary matter."

In fact, in the current legal system, ethical behavior "doesn't pay." Such guidelines are insufficient for the numerous professionals working for corporations where privacy policies are dictated more by a company than by its employees. Nowadays, we have little control who receives our Personally Identifiable Information (PII), what PII they receive, where collected PII is transferred to, and what is the source of (mis?)information others have on us. This is especially alarming with the rapid progress of data mining, the constant discovery of flaws in data anonymization/sanitization techniques, and the vast amount of electronic data that exists. It is beyond the ability of a layperson to understand the privacy policy of organizations and their consequences on the individual.


Ethics in Data Sharing, *Dagstuhl Reports*, Vol. 4, Issue 1, pp. 170–183
Editors: Julie Cohen, Sven Dietrich, Aiko Pras, Lenore D. Zuck, and Mireille Hildebrand
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

The situation is even more serious when data is shared and disseminated among different countries that naturally have different ethical codes and policies for dealing with privacy issues concerning data sharing. Data transfer has no borders, hence, neither does data sharing, which renders ethical data sharing all the more challenging.

However, the recent EU proposals to update the legal framework of the Fair Information Principles, precisely with an eye to the emergence of hyperconnectivity and ubiquitous data analytics, has introduced the notion of Data Protection by Design. This may provide strong incentives to introduce purpose binding, informed consent, minimal disclosure and profile transparency into the design of the relevant computing systems.

The seminar brought in researchers from all disciplines that involve data sharing across borders with ethical implications. The main focus was on Computer System Security data, with consideration for Electronic Medical Records. We derived a basic model for data sharing, and came up with some suggestions of code of ethics for computer professionals (including researchers) that will elaborate on existing codes in terms of data sharing.

## 2 Table of Contents

# 3 Overview of Talks

## 3.1 Privacy, Surveillance, and Ethics

*Julie E. Cohen (Georgetown University – Washington, DC, US)*


**Main reference** Julie E. Cohen, "What Privacy Is For," Harvard Law Review, 126.7 (May 2013):1904–1933.
**URL** http://harvardlawreview.org/2013/05/what-privacy-is-for/
**URL** http://cdn.harvardlawreview.org/wp-content/uploads/pdfs/vol126_cohen.pdf

The ways that privacy and surveillance are understood make the task of formulating ethical guidelines for data sharing a complex one.

Privacy: Particularly in the U.S., but also to an extent in Europe, legal and political discourse conceptualizes privacy as a form of protection for the liberal self. So understood, however, privacy's principal function is defensive and reactionary. It preserves negative space around individuals who are already autonomous and fully formed, providing shelter from the pressures of societal and technological change. In fact, the liberal self who is the subject of privacy theory and privacy policymaking does not exist. The selves who are the real subjects of privacy law- and policy-making are socially constructed, emerging gradually within situated cultures and networks of relationships. Properly understood, privacy's function is dynamic. It shelters emergent subjectivity from the efforts of commercial and government actors to render individuals and communities fixed, transparent, and predictable. Privacy exists in the gaps within frameworks of social control, and this makes it an indispensable structural feature of liberal democratic political systems. Because privacy shelters play, experimentation, and critical reflection, it is also foundational to the capacity for innovation.

Surveillance: The understanding of surveillance and its relationship to government is changing. Traditional discourses about surveillance have emphasized discipline and control. Today, crowd-sourcing and gamification have become important strategies for collection and processing of personal information, and legal strategies for open access and open innovation also have emerged as important drivers. Many development projects that rely on personal information are framed as open access projects, and seek to exploit and profit from the intellectual cachet that rhetorics of openness can confer. Meanwhile, participants in legal and policy discourses about privacy and data processing work to position privacy and innovation as opposites, and to align data processing with the exercise of economic and innovative liberty. The resulting template for surveillance is light and politically nimble. Commentators have long noted the existence of a surveillance-industrial complex: a symbiotic relationship between state surveillance and private-sector producers of surveillance technologies. The emerging surveillance-innovation complex represents a new phase of this symbiosis, one that casts surveillance in an unambiguously progressive light and paints it as a modality of economic growth. At the same time, the rhetorics of openness and innovation work to keep the regulatory state at arms length.

Translating Critique into Practice: Privacy is a function of structure (built environment, networked environment, etc.) and practice. Both effective privacy regulation and effective design for privacy must mind the gaps. Design practices favoring seamless interoperability and scalability probably need to be rethought from the ground up.

## 3.2 People can't control access to their data

*Jürgen Schönwälder (Jacobs Universität – Bremen, DE)*



A couple of years ago, we started getting interested in the question how far it is possible to identify users from network flow traces. We believe this is actually feasible for a large set of diverse users (although clear bounds on the error rate still need to be determined). But of course, all this raises concerns whether this research is ethical to do. We argued that other organizations are likely doing this work behind closed walls and we prefer to do it in an open fashion and to write about it in order to raise awareness. Others argued that doing this kind of research (a) violates current laws and (b) research results enables even more privacy infractions. So what is the right thing to do? Should researchers actively develop techniques to break privacy? Is this like in cryptology where you have people specializing into cryptography and cryptanalysis?

Concerning privacy, I believe (and did so years before PRISM) that we have lost privacy and that in particular normal people have not even a slight idea of what is possible with today's networked systems. If we start from the premise that ordinary people can't control their privacy anymore, will we see new businesses offering privacy services to ordinary people (privacy as a service)? Is there any chance that laws will protect privacy in a globalized networked world? What is the privacy incident that is big enough for the public sensitivity to privacy related questions to change?

People can't control access to their data. Even if they try, today's systems make it easy to correlate data and to identify individuals. What are the solutions? (a) Produce so much data that correlation systems fail? (b) Design special systems like e.g., Tor to hide communication? (c) Have 3rd parties involved that manage personal data? (d) Give up, we have lost privacy, it won't come back...

## 3.3 Sound experimentation for computer security

*Darren Shou (Symantec Research Labs – Culver City, CA, US)*



In 2010 we created a platform to enable sound experimentation for computer security. We hoped that we could tip the balance of the security arms race and unlock strategic insights. In 2011 we tried to address ethical considerations to share and of sharing data. Now I am concerned with big data and the impact on the expectation of privacy, second chances, and a new social network economy.

Important Topics:

- What are our options for sharing security research data (or other data) that enable multi-national enterprises and what might be the unintended consequences?
- What are the long term cultural and ethical considerations enabled by big data, that is, what if there is no privacy?
- How would a new economy such as the current multi-level business model evolve to handle privacy issues (i.e. how can a company that provides a free service to users by selling their data to advertisers evolve?)
- Should users control the flow of their information?

## 3.4 Ethics in Networking Research

*Burkhard Stiller (Universität Zürich, CH)*



The talk on "Ethics in Networking Research"is based on 2 definitions, 4 main observations, and follows with 3 open and 6 even wider open issues, which determine at least a few lines of inter-disciplinary research challenges.

Firstly, the definition of ethics as a system of moral principles, guided by a set of rules of conduct recognized in respect to a particular class of human actions or a partial group, culture, education, location leads to the moral principles as of an individual. In turn, ethics guide our life and it remains a very relative metric across borders. Note that there is a clear difference between ethics and legal statutes, since not all ethics is coded into laws and they may vary very much according country-specific perceptions. And, an evolution of ethics, laws, and rules is happening on a daily basis.

Secondly, determining data as raw data, combined with context, meta, and additional data, defines the basis for the process of data sharing, which enables the analysis of real life happenings of networking. However, besides the pure networking data, multiple application-related data (such as health data or banking data) incur an ethics dimension, especially in terms of knowing about their details, their possible use, and their exploitation.

Consequently, neither complete nor in full, but driven by major networking views, 3 observation tend to occur. (1) The handling, analysis, usage, interpretation, and relevance simply changes upon technology changes from hard drives, to clusters, clouds, and tomorrows extraterrestrial storage locations. (2) The main areas of conflict cover the (a) researcher's view, (b) the operator's view, and (c) the personal (user's) view, each of which with valid targets on their own, such as (a) in search of the real reason and knowledge gains, (b) improving reliability, performance, compliance, . . . , and (c) respecting privacy and providing access to functionality demanded for. Independent of each of those categories, all stakeholders are formally bound by legal and regulative constraints, in which ethics are not part of explicitly (at least not in many cases), however, in which many ethical dimensions play a role in, sometimes formalized by ethical boards or discussion fora. (3) Finally, the handling, analysis, and usage of data is the key for networking research, which should be based on a (public) study plan and a respective protocol. This has developed over time to a certain standards level in the medical discipline, however, "standardization" there has been acknowledged only to a certain level. In any case, in networking research the raw data needs to be shared, including the meta data, especially to ensure reproducability and proof of findings related to those data. Such work may be based on best practices (which change often, too). Since such networking data contains personal data, the violation of privacy always happens, either in a foreseen or in an unforeseen manner. As one very simple example in that dimension utilities' usage and its measurements serves as an optimization approach for personal behaviors (privacy aspect) and its reveal to the utility provider, who can access detailed, time-based usage information of a single house-hold.

Due to these very few observation taken – there are hundreds other ones not listed explicitly – the 3 challenges of next steps in research, public policy, and regulations driven are addressing the basics:

1. What are the impact factors on ethics applied to networking research, especially on data and information? Does this list include legality, regulation, usability, reproducability, practicability, responsibility, commercialization, sustainability, auditability, . . .

2. What are the mandatory demands derived from those for an ethical networking research?
3. Is there any general way to provide practical guidelines besides use cases only?

Even wider open issues in a more case-specific perspective include the following 6 ones:

1. Can networking research achieve the same level of best practices in study protocols, as e.g., in the medical field? Do we want to see that happening anyway?
2. How to deal with inter-domain, multi-domain scenarios, aspects, usages, analysis, and violations as well as court cases? While being "locally" enforceable, the remote case may not provide any "access" to any enforcement technology (if it even exists), a legal body may not be attainable, a foreign court may consider itself as being not responsible, and once data leaked, they leaked.
3. Even is the use of specific data is unethical in its core form, how to prevent the misuse of those data by third parties? Is an NDA sufficient, how to enforce authorized accesses only, how to handle the case of whistleblowers, is there a guarantee on data deletion, is there a chance to ensure the handling of data only according to purpose, self-destroying otherwise?
4. Is the use of methods, which are unethical and sometimes even illegal in themselves, but which attackers use, ethical, especially to prevent harm and disasters? In such a case legal aspects have to be discussed on a separate track, which needs to include country-specific views and intelligence agencies, too.
5. How to act as a personal individual, as a researcher? What about the person, the employee, the government? How to deal with or encourage loyalties and whistleblowing?
6. If a technical experiments is illegal, it may be still relevant ethically, is that a path to promote? For example the African Spring uproars made changes happen (to the good or the bad – that's not discussed here), which were based on the violation of law in force.

Thus, concluding, the dimensions at which ethics affect networking research are larger than someone initially may assume, but they do show at the same time that flexible borders exist and that some of those or neither clear nor viable for that many technologies and applications in use as of today.

## 3.5 Unintended Consequences of Data Sharing Laws and Rules

*Sam Weber (Software Engineering Institute – Arlington, VA, US)*



Sharing of cybersecurity and health care data involves ethical and legal issues, as well as concerns about the utility of said data. My talk will focus on the unintended consequences of well-meaning privacy policies and regulations. In particular, I will discuss three real-life examples where seemingly sensible regulations have negative consequences.

## 3.6 Packets don't know about ethics

*Roland van Rijswijk-Deij (Radboud University Nijmegen, NL)*


**Main reference** Blog article on Dagstuhl Seminar "Ethics in Data Sharing"
**URL** https://blog.surfnet.nl/?p=3174

Being a researcher and working at a National Research and Education Network (NREN) is like permanently living in a cookie jar brimming with exquisite chocolate chip cookies. There is unlimited access to a sea of big data that can be used for research on networks. Also, because of their not-for-profit and academic nature, NRENs are inclined to share this big data not only within their organisations but also with academic groups with which they partner to perform research on networks.

While having access to all this data is ideal for research, it comes with a significant ethical problem. Most of the data used for research on networks is highly privacy sensitive; it can be traced back to individual users in many cases and can be used to build extensive profiles of these users. Some of the data has no direct content, but is in stead considered "meta data" (for example what we call flow data, information about who communicates with whom at what time in with which protocol). But even based on this meta data, highly invasive personal profiles can be constructed that predict traits of an individual users with a high probability. An example of this is research performed by Katikulapudi et al. [1], in which the authors analysed the Internet use of a group of students based on network flow data. Prior to monitoring Internet use, study participants completed a Center for Epidemiologic Studies Depression (CES-D) depression scale survey. Katikulapudi et al. then show a correlation between certain characteristics of the study participants' Internet use and a high score for depression on the CES-D survey. This is only one example of how bulk analysis of Internet use patterns can reveal highly private information about individuals.

What is worrying is that the kind of data that researchers who perform research on networks use allow these kinds of analyses on the behaviour of individual users, even if the goal of their research probably is not to do so (but rather, for instance, to learn something about the behaviour of certain network services in the face of a severe denial-of-service attack).

Note: I have published an extensive report about my participation in this seminar on SURFnet's innovation blog. The URL for this report is http://blog.surfnet.nl/?p=3174.

### References

1 Raghavendra Katikalapudi, Sriram Chellappan, Frances Montgomery, Donald Wunsch, and Karl Lutzen. Associating Internet Usage with Depressive Behavior Among College Students. *IEEE Technology and Society Magazine*, 31(4):73–80, 2012.

## 3.7 Advancing NREN to Researcher Data Sharing: following up on Dagstuhl

*Roland van Rijswijk-Deij (Radboud University Nijmegen, NL)*


**Main reference** Blog article on Dagstuhl Seminar "Ethics in Data Sharing"
**URL** https://blog.surfnet.nl/?p=3174

### Introduction

As a follow-up to the Dagstuhl Seminar on Ethics in Data Sharing, SURFnet has started a project to professionalise the way it shares data with researchers both within as well as outside of it's constituency. This abstract briefly describes the steps we are taking to follow up on the seminar.

### Current practice

To understand our starting point, the current practice for sharing data between SURFnet, as an NREN, and researchers is summarised below.

The current practice is that we share data under a non-disclosure agreement that covers:

- What data is shared
- For what purpose the data may be used
- Who has access to the data
- How long the data may be stored and when it must be destroyed
- Conditions of publication (e.g. references to individual IP addresses must be anonymised)

The collective participants of the Dagstuhl seminar agreed that this was a very prudent practice, but there are some downsides to this approach. First of all, current practice is that we usually have a personal relationship with the researchers that forms a foundation of trust that if we share data it is treated ethically and with respect. Good as this may be, it gives us little foundation for sharing data with researchers we do not know personally. And from a scientific perspective, it would be better if some of the data we share would also be made available to other researchers so they can reproduce research.

### Following up on Dagstuhl

To follow up on the fruitful discussions at the Dagstuhl Seminar on Ethics in Data Sharing, SURFnet has initiated a project in which some of the attendees to the Dagstuhl Seminar participate. The goal of this project is to come up with a comprehensive policy for data sharing between SURFnet as an NREN and research within as well as outside of its constituency. The policy must cover both the ethics side (how to review research proposals, who will review proposals on both sides of the data sharing agreement, training of reviewers, ...) as well as the legal side (impact of privacy law, contracts, ...).

A first project meeting took place on March 25th 2014 with participants from SURFnet, the research community, ethics and law. The outcome of this meeting is that work will start on a number of action items:

- An analysis of applicable law (focusing on EU and Dutch law)
- An outline for a booklet on ethics for researchers as well as proposal reviewers (review board members)
- A proposal for training of researchers and review board members
- Case studies of past and possible future data sharing requests

A follow-up meeting where the first results of these actions will be discussed is scheduled for June 19th 2014.

We believe that the outcome of this project will have broader applicability than just within SURFnet's constituency. We plan to reach out to other European NRENs and will also sollicit input from the security and network industry.

## 3.8 To use or not to use:When and for what should researchers use data obtained from social networking sites.

*Aimee van Wynsberghe (University of Twente, NL)*


**Joint work of** van Wynsberghe, Aimee; Been, Henry; van Keulen, Maurice
**Main reference** A. van Wynsberghe, H. Been, M. van Keulen, "To use or not to use: guidelines for researchers using data from online social networking sites".
**URL** http://responsible-innovation.org.uk/torrii/resource-detail/1471

In the current age of abundant information sharing and gathering, social networking sites (SNSs) are now thought of as incredible resources for collecting data on individuals. To date, such data is collected in a variety of ways (e.g. passively or aggressively), by a variety of researchers (e.g. academic, industry, governmental) for a variety of purposes (e.g. detecting fraudulent behaviors, detecting consumer patterns, studying user patterns). Given this range in collection methods and uses of the data, the question of importance for ethicists, researchers and citizens alike has to do with when and for what can such data be used? In other words, what are the ethics of using data obtained from social networking sites for research purposes? Even when researchers make attempts at protecting the privacy of subjects, things can go wrong, e.g. Facebook's Tastes, ties and Time project and the release of data in 2008. Given the lack of best practices in terms of applying ethics to this field of research, this presentation aims to present a variety of issues related to the use of data obtained through SNSs along with guidelines or points for debate, regarding how to construct a best practice for ethical research. This guideline is created using suggestions from the current literature as well as first-hand experience as an ethics adviser for a research institute dealing specifically with the research and design of ICT systems.

## 3.9 Ethical considerations of using information obtained from online file sharing sites the case of the PirateBay.

*Aimee van Wynsberghe (University of Twente, NL)*


**Joint work of** van Wynsberghe, Aimee; van der Ham, Jeroen
**Main reference** A. van Wynsberghe, J. van der Ham, Jeroen, "Ethical considerations of using information obtained from online file sharing sites – the case of the PirateBay," in Proc. of ETHICOMP 2014, to appear.

Since the creation of Napster back in the late 1990s for the sharing and distribution of MP3 files across the Internet, the entertainment industry has struggled to deal with the regulation of information sharing at large. From an ethics perspective, the practice of file sharing over the internet presents an interesting value conflict between the protection of intellectual property on the one hand (Von Lohmann, 2003), and fairness on the other (DeVoss and

Porter, 2006). On the one hand, the entertainment industry wishes to uphold their exclusive copyrights to the content, to maintain their business model and their distribution methods. On the other hand, users are demanding easy access to music, television and movie files, and will resort to file-sharing when it is not available at a fair price, or at all. With this in mind, the aim of this paper is to investigate the ethical issues arising from the collection of data from an online sharing site. Most notably the website known as ThePirateBay (TPB), founded in Sweden in 2003, facilitates peer-to-peer file-sharing of movies, music, television programs and more. In different countries the entertainment-industry lobbying organizations are taking different approaches to combat this issue. For a variety of European countries, access to TPB is blocked but users are finding ways around this. In the Netherlands, the entertainment industry has successfully won a court case against TPB forcing them to block users from the Netherlands. The website, however, has not recognized the ruling and has not taken action to block any users in the Netherlands from accessing the site. As a next step the Internet service providers (ISPs) have been taken to court to force them to block access to TPB website. This situation has provided a unique opportunity to study the effects of a website blockade on the file sharing behavior of consumers. It seemed common knowledge for Dutch Internet users that there were ways around the blockade, but the net effect was never measured objectively. A study to this effect is relevant for the Internet Society Netherlands Transparency Working Group, but also for both sides in the court cases (the entertainment industry as well as the Internet providers). To that end, van der Ham et al. (2012) began to measure whether preventing access to these sources of links (i.e. ThePirateBay) has an impact on file sharing behavior of users in the Netherlands. To accomplish this goal the researchers created a tool to measure file sharing activity resulting from links shared through TPB website. The measurement tool takes advantage of the fact that file sharing users use a peer-to-peer sharing mechanism, where part of this mechanism is 'gossiping' IP addresses of other users. This uniquely identifies file sharers and impacts the privacy of these users. Thus, aside from the value conflict when one considers the use of file sharing sites like TPB there exists an additional value conflict when one considers the development of a tool to collect and use the data obtained from such a site. Only after developing this tool and performing the measurements did the researchers realize an ethical analysis of this approach may have been in order. With this, they sought the advice of an ethics adviser (van Wynsberghe). The aim of this paper is two-fold:

1. to conduct an ethical analysis on the collection and use of data obtained through online file sharing sites, and
2. to explore the role and utility of ad-hoc ethics advice.

Both goals use the example of this TPB research as a case study for analysis. For the former, the data collection described here is closely in line with the collection of data from online social networking sites. Thus, to address the ethical issues we will use the framework developed by van Wynsberghe and Been (2013) for the collection of data from online social networking sites. This framework entails an analysis of decision variables and choices of the researcher rather than a study of the ethical intentions (Chen, Y-C et al. 2008) or decision making choices of file sharers (Shang, R-A, 2008). As such the framework is concerned with: the context of use and the privacy concerns for this context; the type and method of data collection; the intended use of information and the amount of information collected; and, analysis of values (e.g. to make explicit and scrutinize designer values). Questions pertaining to the use of the information collected in this research revealed that the stakeholders have a vested interest in proving the effectiveness or non-effectiveness of a file sharing block. This information concerning (non-)effectiveness can be used for different purposes depending on

the stakeholder. For instance, the entertainment industry can use the measurements to identify file sharers in support of their copyright infringement case. Alternatively, the Internet Society Netherlands is an organization that sees the censoring of Internet traffic as a threat to the core of the Internet technology itself. Research regarding (non- )effectiveness is important for their lobbying activities. Consequently, the researchers are left asking who should have access to this information and what are the limits to such access? These questions are in line with those asked by security researchers in cases where vulnerabilities are discovered and the responsible party is unresponsive, or unwilling to fix the vulnerability. Corresponding to the initiatives of responsible research and innovation (RRI) in ICT, engaging an ethicist earlier in this research would have been ideal but in practice this ideal is not often met. Accordingly, alongside the ethical analysis described above, this paper aims to show the utility of an ad-hoc ethical appraisal as a means for steering future research of a similar nature along with pointing out areas of concern for improvement.

## 3.10 Ethical considerations in using information from monitoring file-sharing

*Jeroen van der Ham (University of Amsterdam, NL)*


**Joint work of** van Wynsberghe, Aimee; van der Ham, Jeroen
**Main reference** A. van Wynsberghe, J. van der Ham, "Ethical considerations of using information obtained from online file sharing sites – the case of the PirateBay," in Proc. of ETHICOMP 2014, to appear.
**URL** http://www.aimeevanwynsberghe.com/uploads/1/4/6/0/14604548/van_wynsberghe_piratebay.pdf

In the Netherlands there have been court-cases which have resulted in Internet Service Providers having to block access to The PirateBay website. It quickly became common knowledge that there were many ways around this blockade. The situation has provided a unique opportunity to measure the effect of block a website.

The measurements have been performed by monitoring the file-sharing process. The participants were identified and categorized per country and ISP. The distribution of peers was measured at different time-points and subsequently analyzed and compared.

The ethical analysis in this case is very complex because of the many different parties involved. In the first place there is the tension between the entertainment industry and consumers who may be obtaining and sharing content illegally. Then, through the court cases and their (in)actions, ThePirateBay and Internet Service Providers became involved in the consideration. Finally the Internet Society Netherlands wanted to protect the general openness and neutrality of the Internet.

Amidst all this the research has been performed while possibly identifying participants in the file-sharing process. I am working with Aimee van Wynsberghe to use this case to improve an ethical analysis framework.

# 4 Working Groups

The seminar brought together computer scientists, an ethicist and legal scholars to discuss the topic of "ethics in data sharing." After a set of presentations by the participants, some of which are documented in the previous abstracts, there were three main themes requiring ethical attention that were identified by this group of researchers:

- Best Practices and Institutional Review Boards (IRBs) for Ethics in Computer Science,
- Models of Ethics in Producer-Consumer Relations in Data Sharing for Research and Operations, and
- Building Ethical Technology.

The discussions on the first two themes eventually converged. The participants developed a first model for best practice for data sharing.

An online blog or diary was kept by one of the participants (see Section 3.7), and a report describing the model, based on the discussions and case studies, was published in May 2014 [1].

## References

1 Sven Dietrich, Jeroen van der Ham, Aiko Pras, Roland van Rijswijk-Deij, Darren Shou, Anna Sperotto, Aimee van Wynsberghe, and Lenore D. Zuck. Ethics in Data Sharing – Developing a Model for Best Practice. In *Proceedings of the 2nd Cyber-security Research Ethics Dialog & Strategy Workshop, IEEE CS Security and Privacy Workshops 2014*, San Jose, CA, May 2014.

## Participants


- Jon Callas
Silent Circle – San Jose, CA, US
- Georg Carle
TU München, DE
- Julie E. Cohen
Georgetown University – Washington, DC, US
- Sven Dietrich
Stevens Institute of Technology, NJ, US
- Ronald Leenes
Tilburg University, NL
- Aiko Pras
University of Twente, NL
- Volker Roth
FU Berlin, DE
- Peter Y. A. Ryan
University of Luxembourg, LU
- Jürgen Schönwälder
Jacobs Universität – Bremen, DE
- Darren Shou
Symantec Research Labs – Culver City, CA, US
- Anna Sperotto
University of Twente, NL
- Radu State
University of Luxembourg, LU
- Burkhard Stiller
Universität Zürich, CH
- Jeroen van der Ham
University of Amsterdam, NL
- Roland van Rijswijk-Deij
Radboud Univ. Nijmegen, NL
- Aimee van Wynsberghe
University of Twente, NL
- Da-Wei Wang
Academica Sinica – Taipei, TW
- Sam Weber
Software Engineering Institute – Arlington, VA, US
- Lenore D. Zuck
University of Chicago, IL, US