



DAGSTUHL REPORTS

Volume 7, Issue 4, April 2017

Probabilistic Methods in the Design and Analysis of Algorithms (Dagstuhl Seminar 17141) <i>Bodo Manthey, Claire Mathieu, Heiko Röglin, and Eli Upfal</i>	1
Formal Methods of Transformations (Dagstuhl Seminar 17142) <i>Emmanuel Filiot, Sebastian Maneth, and Helmut Seidl</i>	23
Ambient Notification Environments (Dagstuhl Seminar 17161) <i>Lewis Chuang, Sven Gehrung, Judy Kay, and Albrecht Schmidt</i>	38
Online Privacy and Web Transparency (Dagstuhl Seminar 17162) <i>Nataliia Bielova, Nikolaos Laoutaris, Arvind Narayanan, and Nick Nikiforakis</i> ...	83
Computational Geometry (Dagstuhl Seminar 17171) <i>Otfried Cheong, Anne Driemel, and Jeff Erickson</i>	107

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/dagpub/2192-5283>

Publication date

January, 2018

Bibliographic information published by the Deutsche Nationalbibliothek

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

License

This work is licensed under a Creative Commons Attribution 3.0 DE license (CC BY 3.0 DE).



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

- Attribution: The work must be attributed to its authors.

The copyright is retained by the corresponding authors.

Aims and Scope

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

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

- an executive summary of the seminar program and the fundamental results,
- an overview of the talks given during the seminar (summarized as talk abstracts), and
- summaries from working groups (if applicable).

This basic framework can be extended by suitable contributions that are related to the program of the seminar, e. g. summaries from panel discussions or open problem sessions.

Editorial Board

- Gilles Barthe
- Bernd Becker
- Stephan Diehl
- Hans Hagen
- Reiner Hähnle
- Hannes Hartenstein
- Oliver Kohlbacher
- Stephan Merz
- Bernhard Mitschang
- Bernhard Nebel
- Bernt Schiele
- Albrecht Schmidt
- Raimund Seidel (*Editor-in-Chief*)
- Arjen P. de Vries
- Klaus Wehrle
- Verena Wolf

Editorial Office

Michael Wagner (*Managing Editor*)
Jutka Gasirowski (*Editorial Assistance*)
Dagmar Glaser (*Editorial Assistance*)
Thomas Schillo (*Technical Assistance*)

Contact

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

<http://www.dagstuhl.de/dagrep>

Digital Object Identifier: 10.4230/DagRep.7.4.i

Probabilistic Methods in the Design and Analysis of Algorithms

Edited by

Bodo Manthey¹, Claire Mathieu², Heiko Röglin³, and Eli Upfal⁴

1 University of Twente, NL, b.manthey@utwente.nl

2 ENS – Paris, FR, cmathieu@di.ens.fr

3 Universität Bonn, DE, roeglin@cs.uni-bonn.de

4 Brown University – Providence, US, eliezer_upfal@brown.edu

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 17141 “Probabilistic Methods in the Design and Analysis of Algorithms”.

Probabilistic methods play a central role in theoretical computer science. They are a powerful and widely applied tool used, for example, for designing efficient randomized algorithms and for establishing various lower bounds in complexity theory. They also form the basis of frameworks like average-case and smoothed analysis, in which algorithms are analyzed beyond the classical worst-case perspective. The seminar was on probabilistic methods with a focus on the design and analysis of algorithms.

The seminar helped to consolidate the research and to foster collaborations among the researchers who use probabilistic methods in different areas of the design and analysis of algorithms.

Seminar April 2–7, 2017 – <http://www.dagstuhl.de/17141>

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems, G.1.6 Optimization, G.2 Discrete Mathematics, G.3 Probability and Statistics

Keywords and phrases analysis of algorithms, average-case analysis, random graphs, randomized algorithms, smoothed analysis, sub-linear algorithms

Digital Object Identifier 10.4230/DagRep.7.4.1

Edited in cooperation with Stefan Klotzwijk

1 Executive Summary

Bodo Manthey

Claire Mathieu

Heiko Röglin

Eli Upfal

License  Creative Commons BY 3.0 Unported license
© Bodo Manthey, Claire Mathieu, Heiko Röglin, and Eli Upfal

Probabilistic methods play a central role in theoretical computer science. They are a powerful and widely applied tool used, for example, for designing efficient randomized algorithms and for establishing various lower bounds in complexity theory. They also form the basis of frameworks like average-case and smoothed analysis, in which algorithms are analyzed beyond the classical worst-case perspective. The seminar was on probabilistic methods with a focus on the design and analysis of algorithms.

Probabilistic methods are often used in algorithm analysis when worst-case analysis does not provide useful or empirically accurate results. For example, worst-case analysis suggests



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

Probabilistic Methods in the Design and Analysis of Algorithms, *Dagstuhl Reports*, Vol. 7, Issue 04, pp. 1–22

Editors: Bodo Manthey, Claire Mathieu, Heiko Röglin, and Eli Upfal



Dagstuhl Reports

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

that the simplex method is an exponential-time algorithm for linear programming, while in fact it runs in near-linear time on almost all inputs of interest. For the simplex method and many other algorithms worst-case inputs are often rather contrived and occur hardly ever in practical applications. The last decade has seen much interest in the development of a more realistic and robust algorithmic theory that is not entirely based on worst-case performance. One very successful line of research studies the performance of algorithms on inputs that are to some extent random. Besides average-case analysis, in which inputs are generated randomly according to some fixed distribution, also more sophisticated semi-random models have gained momentum.

Another area in which probabilistic methods play a central role is stochastic optimization. Here uncertainty in the data is modeled by probability distributions and the actual data is only revealed over time. For example, in a scheduling problem one might know the probability distribution of a job's length but one learns its actual length only by executing it.

Probabilistic methods are also central in algorithm design. For many optimization problems, the most efficient known algorithms rely essentially on randomization. In other areas, like sublinear algorithms and hashing, one can even prove that randomization is necessary to obtain good algorithms.

The seminar covered recent progress in the context of probabilistic methods.

2 Table of Contents

Executive Summary	
<i>Bodo Manthey, Claire Mathieu, Heiko Röglin, and Eli Upfal</i>	1
Introduction	5
Overview of Talks	7
Stochastic Control via Entropy Compression	
<i>Dimitris Achlioptas</i>	7
Friend or Foe? Population Protocols can perform Community Detection	
<i>Luca Becchetti</i>	7
Ignore or Comply? On Breaking Symmetry in Consensus	
<i>Petra Berenbrink</i>	8
Information-theoretic thresholds	
<i>Amin Coja-Oghlan</i>	8
Making Banaszczyk’s Bound Constructive for the Komlos Problem	
<i>Daniel Dadush</i>	9
Dual-Pivot Quicksort: Optimality, Analysis and Zeros of Lattice Paths	
<i>Martin Dietzfelbinger</i>	9
Randomized Rumor Spreading Revisited	
<i>Benjamin Doerr</i>	10
Thompson Sampling For Stochastic Bandits with Graph Feedback	
<i>Devdatt Dubhashi</i>	11
New (and old) results about the Moran Process	
<i>Leslie Ann Goldberg</i>	11
Resilience of the chromatic number	
<i>Ross Kang</i>	12
Matching Prophets and Secretaries	
<i>Thomas Kesselheim</i>	12
Probabilistic Analysis of Facility Location on Random Shortest Path Metrics	
<i>Stefan Klotwijk</i>	12
Towards $(1 + \epsilon)$ -Approximate Flow Sparsifiers	
<i>Robert Krauthgamer</i>	13
Algorithms and mechanisms for two-sided bayesian auctions.	
<i>Stefano Leonardi</i>	13
Distance in the Forest Fire Model – How far are you from Eve?	
<i>Frederik Mallmann-Trenn</i>	14
The noisy skyline problem	
<i>Claire Mathieu</i>	14
Advances in Streaming Matching	
<i>Morteza Monemizadeh</i>	15

Multi-label classification with pairwise relations <i>Seffi Naor</i>	16
The distortion of Locality Sensitive Hashing <i>Alessandro Panconesi</i>	16
MapReduce and Streaming Algorithms for Diversity Maximization in Metric Spaces of Bounded Doubling Dimension <i>Andrea Pietracaprina</i>	17
Efficient Diameter Approximation for Large Graphs in MapReduce <i>Geppino Pucci</i>	17
Rademacher Averages: Theory and Practice <i>Matteo Riondato</i>	18
Smoothed Analysis of Local Search for the Maximum-Cut Problem <i>Heiko Röglin</i>	18
On coalescence time in graphs – When is coalescing as fast as meeting? <i>Thomas Sauerwald</i>	19
Open Problem: k -means++ in data streams <i>Melanie Schmidt</i>	19
Dependent Rounding <i>Aravind Srinivasan</i>	20
Greedy Algorithms for Stochastic Unrelated Machine Scheduling <i>Marc Uetz</i>	20
Reconstructing Hidden Permutations Using the Average-Precision (AP) Correlation Statistic <i>Eli Upfal</i>	20
Randomized Shared Memory Algorithms <i>Philipp Woelfel</i>	21
Participants	22

3 Introduction

Average-case Analysis. When worst-case analysis fails to predict the performance of an algorithm, average-case analysis is often the first alternative. Average-case analysis has been the subject of a huge amount of research in the last decades, building on insights from, for instance, random graph theory. However, often very simple probability distributions, such as Erdős-Rényi random graphs, have been used in order to keep the average-case analysis tractable. Only recently, average-case analysis has adopted more complex input models, inspired by the study of complex networks. Examples are the usage of first-passage percolation to generate random metric spaces (Bringmann et al., *Algorithmica*, 2015), incorporating gender in social network models to model the glass ceiling effect (Avin et al., *ITCS* 2015), and the preferential attachment process to generate random graphs whose node degrees follow a power-law distribution.

Smoothed Analysis. The drawback of average-case analysis is that purely random instances often have very special properties that distinguish them from typical instances. In smoothed analysis one considers adversarial inputs that are slightly perturbed at random. Spielman and Teng, who invented this notion in 2001, considered adversarial linear programs in which every coefficient is perturbed by Gaussian noise with standard deviation σ . They proved that the expected running time of a certain variant of the simplex algorithm on such linear programs is polynomially bounded in the input size and $1/\sigma$. Hence, even if the amount of randomness is small, the expected running time of the simplex method is polynomially bounded.

Smoothed analysis can be viewed as a less pessimistic worst-case analysis, in which the randomness rules out pathological worst-case instances that are rarely observed in practice but dominate the worst-case analysis. After its invention it has attracted a great deal of attention and it has been applied in a variety of different contexts, e.g., in multi-objective optimization, local search, clustering, and online algorithms.

Semi-random Graph Models. Semi-random input models can be considered as analogues of smoothed analysis for graph problems and they even predate smoothed analysis by a couple of years. There is a variety of semi-random graph models that go beyond the classical Erdős-Rényi random graphs. In most of these models graphs are generated by a noisy adversary – an adversary whose decisions (whether or not to insert a particular edge) have some small probability of being reversed. Another well-studied class of semi-random models are *planted models*, in which a solution (e.g., an independent set or a partitioning of the vertices in color classes) is chosen and then edges are added randomly or by an adversary of limited power in such a way that the given solution stays a valid solution for the given problem.

Similar to smoothed analysis, semi-random models have been invented in order to better understand the complexity of NP-hard graph problems because Erdős-Rényi random graphs often do not reflect the instances one encounters in practice – many graph problems are efficiently solvable on such random graphs.

Online Problems with Random Order. The competitive ratio approach to evaluating online algorithms can be separated into two parts; the sequence of items given as input to an online problem can form a worst-case set of items; and the order in which they are revealed can be a worst-case order. An intermediate measure, between worst-case competitive ratio and average-case competitive ratio when each item in the sequence is drawn independently from some distribution, occurs when the set of input items is worst-case but the arrival order is

random, as might happen if the requests come from sources that do not coordinate their timing. Many problems can be solved more efficiently when the arrival order is random.

Algorithm Design based on Probabilistic Analysis. It is an interesting question if insights from probabilistic analyses of algorithms can give rise to improved algorithms. Recently, some steps in this direction have been taken. The smoothed analysis of the shadow vertex simplex method has, for example, led to efficient algorithms to find short paths on certain polyhedra and improved diameter bounds (Dadush, Hähnle, *SoCG*, 2015).

Stochastic Optimization. In stochastic optimization one has to cope with uncertainty in the inputs, which is usually modeled by probability distributions. One illustrative example are scheduling problems in which the processing times of the tasks follow known probability distributions and the actual processing time of a job is not known until it is completed. In recent years, there has been significant progress in stochastic optimization, in particular, in the area of stochastic scheduling.

Sample Complexity and Sublinear Algorithms. Sampling is a powerful technique, which is at the core of many randomization techniques. Using a finite, often small, set of observations, we attempt to estimate properties of an entire sample space. How good are estimates obtained from a sample? Any rigorous application of sampling requires an understanding of the *sample complexity* of the problem – the minimum size sample needed to obtain the required results. The problem becomes particularly difficult when we try to use one sample set (one set of observations) to detect, or estimate the probabilities of a large (possibly infinite) set of events. Sample complexity has been extensively studied in the context of statistical learning theory, however recent applications have shown the importance of these issues in related fields such as data mining and big data analysis. In particular, there has been significant interest in applying Rademacher averages to improve progressive sampling, and VC-dimension to the study of social network properties.

Related to sample complexity is the area of sublinear algorithms. For huge data sets, even linear-time algorithms are impractical. Hence, the area of sublinear algorithms has developed. Sublinear algorithms inherently need randomness to sample from the input.

Program of the Seminar

The program of the seminar consisted of 29 talks, including the following survey talks:

- Secretary problems and prophet inequalities, Thomas Kesselheim;
- New (and old) results about the Moran Process, Leslie Ann Goldberg;
- Randomized shared memory algorithms, Philipp Wölfel;
- Rademacher averages, Matteo Riondata.

The rest of the talks were 25-minute presentations on recent research of the participants. The time between lunch and the afternoon coffee break was mostly left open for individual discussions and collaborations in small groups. One of the main goals of the seminar was to foster collaborations among the researchers using probabilistic methods in different ways in the area of algorithms research. The feedback provided by the participants shows that the goals of the seminar, namely to circulate new ideas and create new collaborations, were met to a large extent. The organizers and participants wish to thank the staff and the management of Schloss Dagstuhl for their assistance and support in the arrangement of a very successful and productive event.

4 Overview of Talks

4.1 Stochastic Control via Entropy Compression

Dimitris Achlioptas (University of California – Santa Cruz, US)

License © Creative Commons BY 3.0 Unported license
© Dimitris Achlioptas

Joint work of Dimitris Achlioptas, Fotis Iliopoulos, Nikos Vlassis

Main reference D. Achlioptas, F. Iliopoulos, N. Vlassis, “Stochastic Control via Entropy Compression”, in Proc. of the 44th International Colloquium on Automata, Languages, and Programming (ICALP 2017), LIPIcs, Vol 80, pp. 83:1–83:13, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2017.

URL <http://dx.doi.org/10.4230/LIPIcs.ICALP.2017.83>

We consider an agent trying to bring a system to an acceptable state by repeated probabilistic action (stochastic control). Specifically, in each step the agent observes the flaws present in the current state, selects one of them, and addresses it by probabilistically moving to a new state, one where the addressed flaw is most likely absent, but where one or more new flaws may be present. Several recent works on algorithmizations of the Lovász Local Lemma (LLL) have established sufficient conditions for such an agent to succeed. Motivated by the paradigm of Partially Observable Markov Decision Processes (POMDPs) we study whether such stochastic control is also possible in a noisy environment, where both the process of state-observation and the process of state-evolution are subject to adversarial perturbation (noise). The introduction of noise causes the tools developed for LLL algorithmization to break down since the key LLL ingredient, the sparsity of the causality (dependence) relationship, no longer holds. To overcome this challenge we develop a new analysis where entropy plays a central role, both to measure the rate at which progress towards an acceptable state is made and the rate at which the noise undoes this progress. The end result is a sufficient condition that allows a smooth trade-off between the intensity of the noise and the amenability of the system, recovering an asymmetric LLL condition in the noiseless case. To our knowledge, this is the first tractability result for a nontrivial class of POMDPs under stochastic memoryless control.

4.2 Friend or Foe? Population Protocols can perform Community Detection

Luca Becchetti (Sapienza University of Rome, IT)

License © Creative Commons BY 3.0 Unported license
© Luca Becchetti

Joint work of Luca Becchetti, Andrea Clementi, Emanuele Natale, Francesco Pasquale, Prasad Raghavendra, Luca Trevisan

Main reference L. Becchetti, A. E. F. Clementi, E. Natale, F. Pasquale, P. Raghavendra, L. Trevisan, “Friend or Foe? Population Protocols can perform Community Detection”, arXiv:1703.05045 [cs.DM], 2017.

URL <https://arxiv.org/abs/1703.05045>

We present a simple distributed algorithm that, given a regular graph consisting of two communities (or clusters), each inducing a good expander and such that the cut between them has sparsity polylogarithmic recovers the two communities.

More precisely, upon running the protocol, every node assigns itself a binary label of logarithmic length m , so that with high probability, for all but a small number of outliers, nodes within the same community are assigned labels with Hamming distance $o(m)$, while nodes belonging to different communities receive labels with Hamming distance at least $m/2 - o(m)$. We refer to such an outcome as a *community sensitive labeling* of the graph.

Our algorithm uses polylogarithmic local memory and computes the community sensitive labeling after each node performs polylogarithmic steps of local work.

Our algorithm and its analysis work in the (*random*) *population protocol* model, in which anonymous nodes do not share any global clock (the model is asynchronous) and communication occurs over one single (random) edge per round. We believe, this is the first provably-effective protocol for community detection that works in this model.

4.3 Ignore or Comply? On Breaking Symmetry in Consensus

Petra Berenbrink (Universität Hamburg, DE)

License © Creative Commons BY 3.0 Unported license

© Petra Berenbrink

Joint work of Petra Berenbrink, Andrea Clementi, Robert Elsässer, Peter Kling, Frederik Mallmann-Trenn, Emanuele Natale

Main reference P. Berenbrink, A. E. F. Clementi, R. Elsässer, P. Kling, F. Mallmann-Trenn, E. Natale, “Ignore or Comply? On Breaking Symmetry in Consensus”, to appear in PODC 2017.

URL <https://arxiv.org/abs/1702.04921>

We study consensus processes on the complete graph of n nodes. Initially, each node supports one from up to n opinions. Nodes randomly and in parallel sample the opinions of constant many nodes. Based on these samples, they use an update rule to change their own opinion. The goal is to reach consensus, a configuration where all nodes support the same opinion. We compare two well-known update rules: 2-Choices and 3-Majority. In the former, each node samples two nodes and adopts their opinion if they agree. In the latter, each node samples three nodes: If an opinion is supported by at least two samples the node adopts it, otherwise it randomly adopts one of the sampled opinions. Known results for these update rules focus on initial configurations with a limited number of colors, or typically assume a bias, where one opinion has a much larger support than any other. For such biased configurations, the time to reach consensus is roughly the same for 2-Choices and 3-Majority. Interestingly, we prove that this is no longer true for configurations with a large number of initial colors. In particular, we show that 3-Majority reaches consensus with high probability in $O(n^{3/4} \log^{7/8} n)$ rounds, while 2-Choices can need $\Omega(n/\log n)$ rounds. We thus get the first unconditional sublinear bound for 3-Majority and the first result separating the consensus time of these processes. Along the way, we develop a framework that allows a fine-grained comparison between consensus processes from a specific class. We believe that this framework might help to classify the performance of more consensus processes.

4.4 Information-theoretic thresholds

Amin Coja-Oghlan (Goethe-Universität – Frankfurt am Main, DE)

License © Creative Commons BY 3.0 Unported license

© Amin Coja-Oghlan

Joint work of Amin Coja-Oghlan, Florent Krzakala, Will Perkins, Lenka Zdeborová

Main reference A. Coja-Oghlan, F. Krzakala, W. Perkins, L. Zdeborová, “Information-theoretic thresholds from the cavity method”, arXiv:1611.00814 [cs.DM], 2016.

URL <https://arxiv.org/abs/1611.00814>

Vindicating a sophisticated but non-rigorous physics approach called the cavity method, we establish a formula for the mutual information in statistical inference problems induced by random graphs and we show that the mutual information holds the key to understanding

certain important phase transitions in random graph models. We work out several concrete applications of these general results. For instance, we pinpoint the exact condensation phase transition in the Potts antiferromagnet on the random graph, thereby improving prior approximate results [1]. Further, we prove the conjecture from [3] about the condensation phase transition in the random graph coloring problem for any number $q \geq 3$ of colors. Moreover, we prove the conjecture on the information-theoretic threshold in the disassortative stochastic block model [2]. Additionally, our general result implies the conjectured formula for the mutual information in Low-Density Generator Matrix codes [4].

References

- 1 P. Contucci, S. Dommers, C. Giardinà, S. Starr. *Antiferromagnetic Potts model on the Erdős-Rényi random graph*. Communications in Mathematical Physics 323 (2013) 517–554.
- 2 A. Decelle, F. Krzakala, C. Moore, L. Zdeborová. *Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications*. Phys. Rev. E 84 (2011) 066106.
- 3 F. Krzakala, A. Montanari, F. Ricci-Tersenghi, G. Semerjian, L. Zdeborová. *Gibbs states and the set of solutions of random constraint satisfaction problems*. Proc. National Academy of Sciences 104 (2007) 10318–10323.
- 4 A. Montanari. *Tight bounds for LDPC and LDGM codes under MAP decoding*. IEEE Transactions on Information Theory 51 (2005) 3221–3246.

4.5 Making Banaszczyk’s Bound Constructive for the Komlos Problem

Daniel Dadush (CWI – Amsterdam, NL)

License © Creative Commons BY 3.0 Unported license
© Daniel Dadush

Joint work of Nikhil Bansal, Daniel Dadush, Shashwat Garg

We consider the problem of finding a low discrepancy coloring for sparse set systems where each element lies in at most t sets. We give an efficient algorithm that finds a coloring with discrepancy $O((t \log n)^{1/2})$, matching the best known non-constructive bound for the problem due to Banaszczyk. The previous algorithms only achieved an $O(t^{1/2} \log n)$ bound. The result also extends to the more general Komlos setting, where each vector has norm at most 1, and gives an algorithmic $O(\log^{1/2} n)$ bound.

Joint work with Nikhil Bansal and Shashwat Garg.

4.6 Dual-Pivot Quicksort: Optimality, Analysis and Zeros of Lattice Paths

Martin Dietzfelbinger (TU Ilmenau, DE)

License © Creative Commons BY 3.0 Unported license
© Martin Dietzfelbinger

Joint work of Martin Aumüller, Martin Dietzfelbinger, Daniel Krenn, Clemens Heuberger, Helmut Prodinger
Main reference M. Aumüller, M. Dietzfelbinger, C. Heuberger, D. Krenn, H. Prodinger, “Dual-Pivot Quicksort: Optimality, Analysis and Zeros of Associated Lattice Paths”, arXiv:1611.00258 [math.CO], 2016.
URL <https://arxiv.org/abs/1611.00258>

Quicksort is a venerable sorting algorithm, it is taught in basic algorithms classes, and it is routinely used in practice. Can there be anything new about Quicksort today? *Dual-pivot*

quicksort refers to variants of classical quicksort where in the partitioning step two pivots are used to split the input into three segments. Algorithms of this type had been studied by Sedgewick (1975) and by Hennequin, but they received new attention starting from 2009, when a dual-pivot algorithm due to Yaroslavskiy, Bentley, and Bloch replaced the well-engineered quicksort algorithm in Oracle’s Java 7 runtime library. An analysis of a simplified variant of this algorithm by Nebel and Wild from 2012 gave $1.9 \ln n$ comparisons on average for n input numbers. (Other works ensued. Standard quicksort has $2n \ln n$ expected comparisons.) In the center of the analysis is the partitioning procedure, which splits the input keys in “small” (smaller than small pivot), “medium” (between the two pivots), “large” (larger than large pivot). We identify a partitioning strategy with the minimum average number of key comparisons. It keeps count of how many large and small elements were seen before and prefers the corresponding pivot. The comparison count is closely related to a “random walk” on the integers which keeps track of the difference of large and small elements seen so far. (An alternative way of understanding what is going on is a Pólya urn with three colors.) For the fine analysis it is essential to understand the expected number of times this random walk hits zero. The expected number of comparisons can be determined exactly and as a formula up to lower terms. It is $1.8n \ln n - 2.38\dots n + 1.675 \ln n + O(1)$.

4.7 Randomized Rumor Spreading Revisited

Benjamin Doerr (Ecole Polytechnique – Palaiseau, FR)

License  Creative Commons BY 3.0 Unported license
© Benjamin Doerr

We develop a simple and generic method to analyze randomized rumor spreading processes in fully connected networks. In contrast to all previous works, which heavily exploit the precise definition of the process under investigation, we only need to understand the probability and the covariance of the events that uninformed nodes become informed. This universality allows us to easily analyze the classic push, pull, and push-pull protocols both in their pure version and in several variations such as messages failing with constant probability or nodes calling a random number of others each round. Some dynamic models can be analyzed as well, e.g., when the network is a $G(n, p)$ random graph sampled independently each round [Clementi et al. (ESA 2013)].

Despite this generality, our method determines the expected rumor spreading time precisely apart from additive constants, which is more precise than almost all previous works. We also prove tail bounds showing that a deviation from the expectation by more than an additive number of r rounds occurs with probability at most $\exp(-\Omega(r))$.

We further use our method to discuss the common assumption that nodes can answer any number of incoming calls. We observe that the restriction that only one call can be answered leads to a significant increase of the runtime of the push-pull protocol. In particular, the double logarithmic end phase of the process now takes logarithmic time. This also increases the message complexity from the asymptotically optimal $\Theta(n \log \log n)$ [Karp, Shenker, Schindelhauer, Vöcking (FOCS 2000)] to $\Theta(n \log n)$. We propose a simple variation of the push-pull protocol that reverts back to the double logarithmic end phase and thus to the $\Theta(n \log \log n)$ message complexity.

4.8 Thompson Sampling For Stochastic Bandits with Graph Feedback

Devdatt Dubhashi (Chalmers UT – Göteborg, SE)

License © Creative Commons BY 3.0 Unported license
© Devdatt Dubhashi

Joint work of Aristide Tossou, Christos Dimitrakakis, Devdatt Dubhashi

Main reference A. C. Y. Tossou, C. Dimitrakakis, D. P. Dubhashi, “Thompson Sampling for Stochastic Bandits with Graph Feedback”, in Proc. of the 31st AAAI Conf. on Artificial Intelligence (AAAI 2017), pp. 2660-2666, 2017.

URL <https://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14768>

We present a novel extension of Thompson Sampling for stochastic sequential decision problems with graph feedback, even when the graph structure itself is unknown and/or changing. We provide theoretical guarantees on the Bayesian regret of the algorithm, linking its performance to the underlying properties of the graph. Thompson Sampling has the advantage of being applicable without the need to construct complicated upper confidence bounds for different problems. We illustrate its performance through extensive experimental results on real and simulated networks with graph feedback. More specifically, we tested our algorithms on power law, planted partitions and Erdős-Rényi graphs, as well as on graphs derived from Facebook and Flixster data. These all show that our algorithms clearly outperform related methods that employ upper confidence bounds, even if the latter use more information about the graph.

4.9 New (and old) results about the Moran Process

Leslie Ann Goldberg (University of Oxford, GB)

License © Creative Commons BY 3.0 Unported license
© Leslie Ann Goldberg

Joint work of Andreas Galanis, Andreas Goebel, Leslie Ann Goldberg, John Lapinskas, Johannes Lengler, Florian Meier, Konstantinos Panagiotou, Pascal Pfister, David Richerby

Main reference L. A. Goldberg, J. Lapinskas, J. Lengler, F. Meier, K. Panagiotou, P. Pfister, “Asymptotically Optimal Amplifiers for the Moran Process”, arXiv:1611.04209 [math.PR], 2016.

URL <https://arxiv.org/abs/1611.04209>

Main reference A. Galanis, A. Goebel, L. A. Goldberg, J. Lapinskas, D. Richerby, “Amplifiers for the Moran Process”, Journal of the ACM (JACM) 64(1):5, 2017.

URL <http://dx.doi.org/10.1145/3019609>

The Moran process, as adapted by Lieberman, Hauert and Nowak, is a model of a population on a graph, evolving in discrete time. Individuals in the population are associated with the vertices of the graph. Certain individuals, called “mutants” have fitness r and other individuals, called “non-mutants” have fitness 1. The state of being a mutant or not can be spread from vertices to neighbours. We focus on the situation where the mutation is advantageous, in the sense that $r > 1$. If the graph is strongly connected then, with probability 1, the process will either reach the state where there are only mutants (known as fixation) or it will reach the state where there are only non-mutants (known as extinction). A set of (directed or undirected) graphs is said to be strongly amplifying if the extinction probability tends to 0 when the Moran process is run on graphs in this set, starting from the state with a single mutation, at a uniformly-chosen vertex. It turns out that strong amplifiers exist, even in the undirected case. This talk will tell you what is known about them, including joint work with Galanis, Goebel, Lapinskas and Richerby and also joint work with Lapinskas, Lengler, Meier, Panagiotou and Pfister.

4.10 Resilience of the chromatic number

Ross Kang (Radboud University Nijmegen, NL)

License © Creative Commons BY 3.0 Unported license
© Ross Kang

Joint work of Ross Kang, Colin McDiarmid

For p bounded away from 1 and $np \rightarrow \infty$, the asymptotic value of the chromatic number of the binomial random graph $G(n, p)$ is resilient to the removal of a subgraph H with $o(n \log(np))$ edges. This is best possible. Our work is counterpoint to work of Sudakov and Vu (2008) and Alon and Sudakov (2010), which was on chromatic resilience for $G(n, p)$ subject to the addition, rather than removal, of edges. Joint work with Colin McDiarmid.

4.11 Matching Prophets and Secretaries

Thomas Kesselheim (TU Dortmund, DE)

License © Creative Commons BY 3.0 Unported license
© Thomas Kesselheim

Main reference T. Kesselheim, K. Radke, A. Tönnis, B. Vöcking, “An Optimal Online Algorithm for Weighted Bipartite Matching and Extensions to Combinatorial Auctions”, *Algorithms – ESA 2013, LNCS*, Vol. 8125, pp. 589–600, Springer, 2013.

URL http://dx.doi.org/10.1007/978-3-642-40450-4_50

Main reference M. Feldman, N. Gravin, B. Lucier, “Combinatorial Auctions via Posted Prices”, in *Proc. of the 26th Annual ACM-SIAM Symp. on Discrete Algorithms (SODA 2015)*, pp. 123–135, 2015.

URL <http://dx.doi.org/10.1137/1.9781611973730.10>

We discuss two different approaches for probabilistic input models in online settings. In the first model, the optimization instance is determined by an adversary but it is revealed in random order. The secretary problem is a common representative of this class. In the second model, the optimization instance is drawn from known probability distributions and presented in adversarial order, like for the prophet inequality. For both settings, we survey different online problems with a focus on edge-weighted bipartite matching, for which optimal guarantees are known.

4.12 Probabilistic Analysis of Facility Location on Random Shortest Path Metrics

Stefan Klootwijk (University of Twente, NL)

License © Creative Commons BY 3.0 Unported license
© Stefan Klootwijk

Joint work of Stefan Klootwijk, Bodo Manthey

Main reference S. Klootwijk, “Probabilistic Analysis of Facility Location on Random Shortest Path Metrics”, Master Thesis, University of Twente, 2016.

URL <http://purl.utwente.nl/essays/71002>

We investigate the performance of the facility location problem using the model of random shortest path metrics. We analyze some probabilistic properties for a simple heuristic which provides a solution to the facility location problem: opening a certain number of arbitrary facilities (with that certain number only depending on the facility opening cost). We show that, for almost any facility opening cost, this heuristic yields a $1 + o(1)$ approximation in expectation. In the remaining few cases we show that this heuristic yields an $O(1)$ approximation in expectation.

4.13 Towards $(1 + \varepsilon)$ -Approximate Flow Sparsifiers

Robert Krauthgamer (Weizmann Institute – Rehovot, IL)

License © Creative Commons BY 3.0 Unported license
© Robert Krauthgamer

Joint work of Alexandr Andoni, Anupam Gupta, Robert Krauthgamer

Main reference A. Andoni, A. Gupta, R. Krauthgamer, “Towards $(1 + \varepsilon)$ -Approximate Flow Sparsifiers”, in Proc. of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2014), pp. 279–293, 2014.

URL <http://dx.doi.org/10.1137/1.9781611973402.20>

A key challenge in designing graph algorithms is to “compress” a graph G so as to preserve some of its basic properties, such as distances and cuts. I will discuss vertex-sparsification, where given a graph G with k terminal vertices, the goal is to construct a small graph G' containing the same terminals, such that all cuts/flows between the terminals in G' approximate those in G within factor $1 + \varepsilon$.

I will present a randomized construction for bipartite graphs, and discuss some open questions.

4.14 Algorithms and mechanisms for two-sided bayesian auctions.

Stefano Leonardi (Sapienza University of Rome, IT)

License © Creative Commons BY 3.0 Unported license
© Stefano Leonardi

Joint work of Riccardo Colini-Baldeschi, Paul Goldberg, Bart de Keijzer, Stefano Leonardi, Tim Roughgarden, Stefano Turchetta

Main reference R. Colini-Baldeschi, B. de Keijzer, S. Leonardi, S. Turchetta, “Approximately Efficient Double Auctions with Strong Budget Balance”, in Proc. of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2016), pp. 1424–1443, 2016.

URL <https://doi.org/10.1137/1.9781611974331.ch98>

Main reference R. Colini-Baldeschi, P. W. Goldberg, B. de Keijzer, S. Leonardi, T. Roughgarden, S. Turchetta, “Approximately Efficient Two-Sided Combinatorial Auctions”, in Proc. of the 2017 ACM Conf. on Economics and Computation (EC'17), pp. 591–608, ACM, 2017.

URL <https://doi.org/10.1145/3033274.3085128>

Mechanism design for one-sided markets is an area of extensive research in economics and, since more than a decade, in computer science as well. Two-sided markets, on the other hand, have not received the same attention despite the numerous applications to Internet advertisement. We present algorithms and mechanisms for two-sided markets in which both buyers and sellers act strategically.

An ideal goal in two-sided markets is to maximize the social welfare of buyers and sellers with individually rational (IR), incentive compatible (IC) and budget-balanced mechanisms (BB), which requires that the mechanism does not subsidize the market or make an arbitrary profit from the exchange. Unfortunately, Myerson and Satterthwaite proved in 1983 that this goal cannot be achieved even in the bayesian setting and for the simple case of only one buyer and one seller. In this talk I'll discuss meaningful trade-offs and algorithmic approximations of the above requirements for the bayesian setting.

I'll present simple two-sided sequential posted price mechanisms which provides an $O(1)$ approximation to the optimal social welfare while obeying the IR, IC and BB constraints. The mechanisms work for any number of buyers and sellers with arbitrary, independent distributions and matroid constraints for the case of unit supply buyers and sellers. These mechanisms can be extended to two-sided combinatorial auctions with additive and XOS valuations on subsets of items.

4.15 Distance in the Forest Fire Model – How far are you from Eve?

Frederik Mallmann-Trenn (ENS – Paris, FR)

License © Creative Commons BY 3.0 Unported license
© Frederik Mallmann-Trenn

Joint work of Claire Mathieu, Frederik Mallmann-Trenn, Varun Kanade, Reut Levi, Zvi Lotker
Main reference V. Kanade, R. Levi, Z. Lotker, F. Mallmann-Trenn, C. Mathieu, “Distance in the Forest Fire Model How far are you from Eve?”, in Proc. of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2016), pp. 1602–1620, 2016.
URL <http://dx.doi.org/10.1137/1.9781611974331.ch109>

Leskovec, Kleinberg and Faloutsos (2005) observed that many social networks exhibit properties such as shrinking (i.e. bounded) diameter, densification, and (power-law) heavy tail degree distributions. To explain these phenomena, they introduced a generative model, called the Forest Fire model, and using simulations showed that this model indeed exhibited these properties; however, proving this rigorously was left as an open problem.

In this paper, we analyse one of these properties, shrinking diameter. We define a restricted version of their model that incorporates the main features that seem to contribute towards this property, and prove that the graphs generated by this model exhibit shrinking distance to the seed graph. We prove that an even simpler model, the random walk model, already exhibits this phenomenon.

4.16 The noisy skyline problem

Claire Mathieu (ENS – Paris, FR)

License © Creative Commons BY 3.0 Unported license
© Claire Mathieu

Joint work of Frederik Mallmann-Trenn, Claire Mathieu

Given a set of points in a d -dimensional space, we seek to compute the skyline, i.e., those points which are not strictly dominated by any other point, using few comparisons between elements. We study the setting where comparisons can fail with constant probability. This version of the problem was introduced and studied by Groz & Milo. We provide two algorithms computing the skyline with query complexity $O(nd \log(dk))$ and $O(ndk \log(k))$, where k is the size of the skyline. These results improve on the bounds of Groz and Milo by a factor of k . Furthermore, we prove that the obtained bounds are (up to constants) tight when d is not too large.

4.17 Advances in Streaming Matching

Morteza Monemizadeh (Charles University – Prague, CZ)

License © Creative Commons BY 3.0 Unported license
© Morteza Monemizadeh

- Joint work of** Rajesh Chitnis, Graham Cormode, Hossein Esfandiari, Elena Grigorescu, Mohammad Taghi Hajiaghayi, Hossein Jowhari, Vahid Liaghat, Andrew McGregor, Morteza Monemizadeh, S. Muthukrishnan, Krzysztof Onak, Sofya Vorotnikova, Samson Zhou
- Main reference** G. Cormode, H. Jowhari, M. Monemizadeh, S. Muthukrishnan, “The Sparse Awakens: Streaming Algorithms for Matching Size Estimation in Sparse Graphs”, in Proc. of the 25th Annual European Symposium on Algorithms (ESA 2017), LIPIcs, Vol. 87, pp. 29:1–29:15, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2017.
- URL** <https://arxiv.org/abs/1608.03118>
- Main reference** R.H. Chitnis, G. Cormode, H. Esfandiari, M. T. Hajiaghayi, A. McGregor, M. Monemizadeh, S. Vorotnikova, “Kernelization via Sampling with Applications to Finding Matchings and Related Problems in Dynamic Graph Streams”, in Proc. of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2016), pp. 1326–1344, 2016.
- URL** <http://dx.doi.org/10.1137/1.9781611974331.ch92>
- Main reference** R.H. Chitnis, G. Cormode, M.T. Hajiaghayi, M. Monemizadeh, “Parameterized Streaming: Maximal Matching and Vertex Cover”, in Proc. of the 26th Annual ACM-SIAM Symp. on Discrete Algorithms (SODA 2015), pp. 1234–1251, 2015.
- URL** <http://dx.doi.org/10.1137/1.9781611973730.82>
- Main reference** H. Esfandiari, M.T. Hajiaghayi, V. Liaghat, M. Monemizadeh, K. Onak, “Streaming Algorithms for Estimating the Matching Size in Planar Graphs and Beyond”, in Proc. of the 26th Annual ACM-SIAM Symp. on Discrete Algorithms (SODA 2015), pp. 1217–1233, 2015.
- URL** <http://dx.doi.org/10.1137/1.9781611973730.81>

As noted by Lovász and Plummer in their classic book, “Matching Theory is a central part of graph theory, not only because of its applications, but also because it is the source of important ideas developed during the rapid growth of combinatorics during the last several decades.”

We consider variants of matching in data streams when the insertion and the deletion of edges are revealed in a streaming fashion. In particular,

1. When the input graph is planar, we present a simple and elegant streaming algorithm that with high probability estimates the size of a maximum matching within a constant factor using $O(n^{2/3})$ space, where n is the number of vertices. The approach generalizes to the family of graphs that have bounded arboricity, which include graphs with an excluded constant-size minor.
2. We further reduce the required memory size to $O(\sqrt{n})$ for three restricted settings: (i) when the input graph is a forest; (ii) when we have 2-passes and the input graph has bounded arboricity; and (iii) when the edges arrive in random order and the input graph has bounded arboricity.
3. We present a simple but powerful subgraph sampling primitive that is applicable in a variety of computational models including dynamic graph streams (where the input graph is defined by a sequence of edge/hyperedge insertions and deletions) and distributed systems such as MapReduce. In the case of dynamic graph streams, we use this primitive to prove that there exists an $O(k^2)$ space algorithm that returns the edges of a maximum matching on the assumption the cardinality is at most k . The best previous algorithm used $O(kn)$ space. We prove our result is optimal up to logarithmic factors. Our algorithm has $O(1)$ update time.
4. We also show that there exists an $O(n^2/\alpha^3)$ space algorithm that returns an α -approximation for matchings of arbitrary size. In independent work, Assadi et al. (SODA 2016) proved this approximation algorithm is optimal and provided an alternative algorithm. We generalize our exact and approximate algorithms to weighted matching. For graphs with low arboricity such as planar graphs, the space required for constant approximation can be further reduced. While there has been a substantial amount of work on approximate matching in insert-only graph streams, these are the first non-trivial results in the dynamic setting.

4.18 Multi-label classification with pairwise relations

Seffi Naor (*Technion – Haifa, IL*)

License  Creative Commons BY 3.0 Unported license
© Seffi Naor

Joint work of Shahar Chen, Dotan Di Castro, Zohar Karnin, Liane Lewin-Eytan, Seffi Naor, Roy Schwartz

Motivated by various applications in machine learning, e.g., classification of web pages, semantic tagging of images, and functional genomics, we introduce the metric multi-labeling problem. The input is a set of objects with pairwise relations, given as a weighted graph over the objects, together with a label set and assignment costs of labels to objects. Departing from classical work on the metric labeling problem, assignment costs may be either positive or negative, reflecting a recommendation given by a local learning process which infers label preferences of objects. However, the learning process ignores pairwise relations of objects. Clearly, the labeling cost for completely agreeing with this recommendation is the minimum possible, and this is our benchmark labeling. The objective is to find an assignment of (multiple) labels to objects, at least one for each object, minimizing the sum of the following two terms: (1) the deviation cost from the benchmark labeling, and (2) the separation cost between pairs of similar objects receiving different label sets. Thus, the goal is to globally optimize the assignment of label sets to objects, while taking into account local preferences of objects. We obtain a (tight) 2-approximation for this problem.

Additionally, in many natural applications a bound is given on the number of labels that can be assigned to an object. This leads to a very interesting problem of simultaneous rounding of many (fractional) points contained in a uniform matroid polytope, while bounding pairwise distances. We achieve a true $O(\log k)$ -approximation for this problem (k is the number of labels).

4.19 The distortion of Locality Sensitive Hashing

Alessandro Panconesi (*Sapienza University of Rome, IT*)

License  Creative Commons BY 3.0 Unported license
© Alessandro Panconesi

Joint work of Flavio Chierichetti, Ravi Kumar, Alessandro Panconesi, Erisa Terolli

Main reference F. Chierichetti, R. Kumar, A. Panconesi, E. Terolli, “The Distortion of Locality Sensitive Hashing”, in Proc. of the 8th Innovations in Theoretical Computer Science Conf. (ITCS 2017), LIPIcs, Vol. 67, pp. 54:1–54:18, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2017.

URL <http://dx.doi.org/10.4230/LIPIcs.ITCS.2017.54>

Given a pairwise similarity notion between objects, locality sensitive hashing (LSH) aims to construct a hash function family over the universe of objects such that the probability two objects hash to the same value is their similarity. LSH is a powerful algorithmic tool for large-scale applications and much work has been done to understand LSHable similarities, i.e., similarities that admit an LSH.

In this paper we focus on similarities that are provably non-LSHable and propose a notion of distortion to capture the approximation of such a similarity by a similarity that is LSHable. We consider several well-known non-LSHable similarities and show tight upper and lower bounds on their distortion.

4.20 MapReduce and Streaming Algorithms for Diversity Maximization in Metric Spaces of Bounded Doubling Dimension

Andrea Pietracaprina (University of Padova, IT)

License © Creative Commons BY 3.0 Unported license
© Andrea Pietracaprina

Joint work of Matteo Ceccarelo, Andrea Pietracaprina, Geppino Pucci, Eli Upfal

Main reference M. Ceccarelo, A. Pietracaprina, G. Pucci, E. Upfal, “MapReduce and Streaming Algorithms for Diversity Maximization in Metric Spaces of Bounded Doubling Dimension”, *PVLDB* 10(5):469–480, 2017.

URL <http://doi.org/10.14778/3055540.3055541>

Given a dataset of points in a metric space and an integer k , a diversity maximization problem requires determining a subset of k points maximizing some diversity objective measure (e.g., the minimum or the average distance between a pair of points in the subset). Diversity maximization problems are computationally hard, hence only approximate solutions can be hoped for. Although its applications are mostly in massive data analysis, most of the past research on diversity maximization has concentrated on the standard sequential setting. In this talk, we present novel space and pass/round-efficient approximation algorithms for diversity maximization in the Streaming and MapReduce models, which can handle very large inputs. We show that for point sets belonging to metric spaces of bounded doubling dimension, our algorithms attain an $(\alpha + \varepsilon)$ -approximation ratio, for any constant $\varepsilon > 0$, where α is the best approximation ratio achieved by a polynomial-time, linear-space sequential algorithm. We also provide extensive experimental evidence of the effectiveness of our algorithms on both real world and synthetic datasets, scaling up to over a billion points. While our algorithms are deterministic, we show how randomization can help attain lower space requirements, and believe this to be a promising direction for future improvements. Also, we will briefly discuss recent results on the generalization of our approach to diversity maximization problems under matroid constraints.

4.21 Efficient Diameter Approximation for Large Graphs in MapReduce

Geppino Pucci (University of Padova, IT)

License © Creative Commons BY 3.0 Unported license
© Geppino Pucci

Joint work of Matteo Ceccarelo, Andrea Pietracaprina, Geppino Pucci, Eli Upfal

Main reference M. Ceccarelo, A. Pietracaprina, G. Pucci, E. Upfal, “Space and Time Efficient Parallel Graph Decomposition, Clustering, and Diameter Approximation”, *ACM SPAA 2015*: 182–191.

Main reference M. Ceccarelo, A. Pietracaprina, G. Pucci, E. Upfal, “A Practical Parallel Algorithm for Diameter Approximation of Massive Weighted Graphs”, *IEEE IPDPS 2016*: 12–21.

URL <http://dx.doi.org/10.1109/IPDPS.2016.61>

We present a space and time efficient practical parallel algorithm for approximating the diameter of massive weighted undirected graphs on distributed-memory platforms supporting a MapReduce-like abstraction. The core of the algorithm is a randomized clustering strategy generating disjoint clusters of bounded weighted radius. Theoretically, our algorithm uses linear space and yields a polylogarithmic approximation guarantee; moreover, for the important class of graphs of bounded doubling dimension, it runs in a number of rounds asymptotically smaller than those required by the state-of-the-art Delta-stepping algorithm, which is the only practical linear-space competitor on distributed-memory platforms. We complement our theoretical findings with an extensive experimental analysis on large benchmark graphs, which demonstrates that our algorithm attains substantial improvements on a number of

key performance indicators with respect to the aforementioned competitor, while featuring a similar approximation ratio (a small constant less than 1.4, as opposed to the polylogarithmic theoretical bound).

We also point out some recent interesting applications of clustering for the efficient sequential approximation of diameter and node centralities in graphs of low doubling dimension.

4.22 Rademacher Averages: Theory and Practice

Matteo Riondato (Two Sigma Investments LP – New York, US)

License © Creative Commons BY 3.0 Unported license

© Matteo Riondato

Joint work of Matteo Riondato, Eli Upfal

Main reference M. Riondato, E. Upfal, “ABRA: Approximating Betweenness Centrality in Static and Dynamic Graphs with Rademacher Averages”, in Proc. of the 22nd ACM SIGKDD Int’l Conf. on Knowledge Discovery and Data Mining (KDD 2016), pp. 1145–1154, ACM, 2016.

URL <http://dx.doi.org/10.1145/2939672.2939770>

Rademacher averages are fundamental concepts from statistical learning theory. They allow to study the maximum deviation of a family of sample averages from their expectations, using only information obtained from the sample itself. They were for long time considered to be only of theoretical interest, but in recent works we showed how to use them to speed up important tasks from data analysis, such as frequent pattern mining and centrality estimation from large graphs.

In this talk I first introduce the fundamental definitions and results involving Rademacher averages, including the celebrated symmetrization lemma and Massart’s finite class lemma. Then I show how to use these results in practice to compute high-quality approximations of the betweenness centralities of all nodes in a graph by using progressive random sampling. This second part is based on joint work with Eli Upfal, appearing in the proceedings of ACM KDD’16.

4.23 Smoothed Analysis of Local Search for the Maximum-Cut Problem

Heiko Röglin (Universität Bonn, DE)

License © Creative Commons BY 3.0 Unported license

© Heiko Röglin

Joint work of Michael Etscheid, Heiko Röglin

Main reference M. Etscheid, H. Röglin, “Smoothed Analysis of Local Search for the Maximum-Cut Problem,” ACM Transactions on Algorithms, 13(2):25, 2017.

URL <http://dx.doi.org/10.1145/3011870>

Even though local search heuristics are the method of choice in practice for many well-studied optimization problems, most of them behave poorly in the worst case. This is in particular the case for the Maximum-Cut Problem, for which local search can take an exponential number of steps to terminate and the problem of computing a local optimum is PLS-complete. To narrow the gap between theory and practice, we study local search for the Maximum-Cut Problem in the framework of smoothed analysis in which inputs are subject to a small amount of random noise. We show that the smoothed number of iterations is quasi-polynomial, i.e., it is bounded from above by a polynomial in $n^{\log n}$ and ϕ where n denotes the number of nodes and ϕ denotes the perturbation parameter. This shows that worst-case instances are fragile and it is a first step in explaining why they are rarely observed in practice.

4.24 On coalescence time in graphs – When is coalescing as fast as meeting?

Thomas Sauerwald (University of Cambridge, GB)

License © Creative Commons BY 3.0 Unported license
© Thomas Sauerwald

Joint work of Frederik Mallmann-Trenn, Varun Kanade, Thomas Sauerwald

Main reference V. Kanade, F. Mallmann-Trenn, T. Sauerwald, “On coalescence time in graphs – When is coalescing as fast as meeting?”, arXiv:1611.02460 [cs.DM], 2016.

URL <https://arxiv.org/abs/1611.02460>

Coalescing random walks is a fundamental stochastic process, where a set of particles perform independent discrete-time random walks on an undirected graph. Whenever two or more particles meet at a given node, they merge and continue as a single random walk. The coalescence time is defined as the expected time until only one particle remains, starting from one particle at every node. Despite recent progress such as by Cooper et al. and Berenbrink et al., the coalescence time for graphs such as binary trees, d -dimensional tori, hypercubes and more generally, vertex-transitive graphs, remains unresolved.

We provide a powerful toolkit that results in tight bounds for various topologies including the aforementioned ones. The meeting time is defined as the worst-case expected time required for two random walks to arrive at the same node at the same time. As a general result, we establish that for graphs whose meeting time is only marginally larger than the mixing time (a factor of $\log^2 n$), the coalescence time of n random walks equals the meeting time up to constant factors. This upper bound is complemented by the construction of a graph family demonstrating that this result is the best possible up to constant factors. For almost-regular graphs, we bound the coalescence time by the hitting time, resolving the discrete-time variant of a conjecture by Aldous for this class of graphs. Finally, we prove that for any graph the coalescence time is bounded by $O(n^3)$. By duality, our results give bounds on the voter model.

4.25 Open Problem: k -means++ in data streams

Melanie Schmidt (Universität Bonn, DE)

License © Creative Commons BY 3.0 Unported license
© Melanie Schmidt

The k -means++ algorithm is a simple yet powerful method to compute good solutions for the k -means problem. It beats the immensely popular heuristic for k -means, Lloyd’s algorithm, in practice while providing a (moderate) approximation guarantee. However, it shares an issue with Lloyd’s algorithm: it requires multiple passes over the data and does not scale too well. Modern clustering offers very strong theoretical results for k -means in the one pass setting and also practical algorithms that (heuristically) compute good solutions on the fly. However, it lacks an algorithm that is as elegant as the k -means++ algorithm.

4.26 Dependent Rounding

Aravind Srinivasan (University of Maryland – College Park, US)

License  Creative Commons BY 3.0 Unported license
© Aravind Srinivasan

Rounding algorithms that introduce carefully-chosen dependencies amongst the variables being rounded, have emerged as a useful technique in randomized (approximation) algorithms over the last twenty years or so. I will give a very brief overview of some aspects of this approach.

4.27 Greedy Algorithms for Stochastic Unrelated Machine Scheduling

Marc Uetz (University of Twente, NL)

License  Creative Commons BY 3.0 Unported license
© Marc Uetz

Joint work of Varun Gupta, Ben Moseley, Marc Uetz, Qiaomin Xie
Main reference V. Gupta, B. Moseley, M. Uetz, Q. Xie, “Stochastic Online Scheduling on Unrelated Machines”, in Proc. of the 19th Int’l Conf. on Integer Programming and Combinatorial Optimization (IPCO 2017), LNCS, Vol. 10328, pp. 228–240, Springer, 2017.
URL http://dx.doi.org/10.1007/978-3-319-59250-3_19

We derive the first performance guarantees for an online algorithm that schedules stochastic, nonpreemptive jobs on unrelated machines to minimize the expectation of the total weighted completion time. The performance guarantees depend linearly on the squared coefficient of variation of the underlying random variables, and therefore –asymptotically– match previously known results. Prior work on unrelated machine scheduling was either restricted to the offline case, and/or required sophisticated linear or convex programming relaxations for the assignment of jobs to machines. In contrast, our algorithm is purely combinatorial, and therefore it works for the online setting, too. As to the techniques applied, we obtain our results by dual fitting. In that sense, we also demonstrate how this technique be put to work for stochastic and nonpreemptive scheduling problems.

4.28 Reconstructing Hidden Permutations Using the Average-Precision (AP) Correlation Statistic

Eli Upfal (Brown University – Providence, US)

License  Creative Commons BY 3.0 Unported license
© Eli Upfal

Joint work of Lorenzo De Stefani, Alessandro Epasto, Eli Upfal, Fabio Vandin
Main reference L. De Stefani, A. Epasto, E. Upfal, F. Vandin, “Reconstructing Hidden Permutations Using the Average-Precision (AP) Correlation Statistic”, in Proc. of the 31st AAAI Conf. on Artificial Intelligence (AAAI 2017), pp. 1526–1532, 2016.
URL <http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12046>

We study the problem of learning probabilistic models for permutations, where the order between highly ranked items in the observed permutations is more reliable (i.e., consistent in different rankings) than the order between lower ranked items, a typical phenomena observed in many applications such as web search results and product ranking. We introduce and study a variant of the Mallows model where the distribution is a function of the widely used Average-Precision (AP) Correlation statistic, instead of the standard Kendall’s tau distance.

We present a generative model for constructing samples from this distribution and prove useful properties of that distribution. Using these properties we develop an efficient algorithm that provably computes an asymptotically unbiased estimate of the center permutation, and a faster algorithm that learns with high probability the hidden central permutation for a wide range of the parameters of the model. We complement our theoretical analysis with extensive experiments showing that unsupervised methods based on our model can precisely identify ground-truth clusters of rankings in real-world data. In particular, when compared to the Kendall's tau based methods, our methods are less affected by noise in low-rank items.

4.29 Randomized Shared Memory Algorithms

Philipp Woelfel (University of Calgary, CA)

License © Creative Commons BY 3.0 Unported license
© Philipp Woelfel

Joint work of George Giakkoupis, Philipp Woelfel

Main reference G. Giakkoupis, P. Woelfel, “Randomized Mutual Exclusion with Constant Amortized RMR Complexity on the DSM”, in Proc. of the 55th Annual Symposium on Foundations of Computer Science (FOCS 2014), pp. 504–513, IEEE, 2014.

URL <http://dx.doi.org/10.1109/FOCS.2014.60>

Main reference G. Giakkoupis, P. Woelfel, “On the time and space complexity of randomized test-and-set”, in Proc. of the 2012 ACM symposium on Principles of distributed computing (PODC 2012), pp. 19–28, ACM, 2012.

URL <http://dx.doi.org/10.1145/2332432.2332436>

Randomization has been recognized for decades as a powerful tool in sequential algorithm design. But for concurrent algorithmic problems, probabilistic choices by processes have been mainly used to solve otherwise unsolvable problems (most famously the consensus problem). Only in recent years, randomization has received attention to improve the efficiency of concurrent algorithms.

In this talk, I will give an overview of the asynchronous shared memory model, and how randomization can be employed to find solutions to some fundamental problems, which are more efficient than the best known deterministic algorithms. I will give two main examples of recently devised randomized algorithms:

- A randomized leader election algorithm that has $O(\log \log n)$ expected step complexity; and
- A randomized mutual exclusion algorithm that has amortized constant expected RMR complexity.

Participants

- Dimitris Achlioptas
University of California –
Santa Cruz, US
- Aris Anagnostopoulos
Sapienza University of Rome, IT
- Luca Becchetti
Sapienza University of Rome, IT
- Petra Berenbrink
Universität Hamburg, DE
- Amin Coja-Oghlan
Goethe-Universität –
Frankfurt am Main, DE
- Artur Czumaj
University of Warwick –
Coventry, GB
- Daniel Dadush
CWI – Amsterdam, NL
- Martin Dietzfelbinger
TU Ilmenau, DE
- Benjamin Doerr
Ecole Polytechnique –
Palaiseau, FR
- Devdatt Dubhashi
Chalmers UT – Göteborg, SE
- Matthias Englert
University of Warwick –
Coventry, GB
- Leslie Ann Goldberg
University of Oxford, GB
- Paul W. Goldberg
University of Oxford, GB
- Ross Kang
Radboud University
Nijmegen, NL
- Thomas Kesselheim
TU Dortmund, DE
- Stefan Klotzwijk
University of Twente, NL
- Robert Krauthgamer
Weizmann Institute –
Rehovot, IL
- Marvin Künnemann
MPI für Informatik –
Saarbrücken, DE
- Stefano Leonardi
Sapienza University of Rome, IT
- Frederik Mallmann-Trenn
ENS – Paris, FR
- Bodo Manthey
University of Twente, NL
- Claire Mathieu
ENS – Paris, FR
- Nicole Megow
Universität Bremen, DE
- Morteza Monemizadeh
Charles University – Prague, CZ
- Seffi Naor
Technion – Haifa, IL
- Alessandro Panconesi
Sapienza University of Rome, IT
- Andrea Pietracaprina
University of Padova, IT
- Kirk Pruhs
University of Pittsburgh, US
- Geppino Pucci
University of Padova, IT
- Matteo Riondato
Two Sigma Investments LP –
New York, US
- Heiko Röglin
Universität Bonn, DE
- Thomas Sauerwald
University of Cambridge, GB
- Melanie Schmidt
Universität Bonn, DE
- Christian Sohler
TU Dortmund, DE
- Aravind Srinivasan
University of Maryland –
College Park, US
- Marc Uetz
University of Twente, NL
- Eli Upfal
Brown University –
Providence, US
- Tjark Vredeveld
Maastricht University, NL
- Philipp Woelfel
University of Calgary, CA



Formal Methods of Transformations

Edited by

Emmanuel Filiot¹, Sebastian Maneth², and Helmut Seidl³

1 Free University of Brussels, BE, efiliot@ulb.ac.be

2 Universität Bremen, DE, maneth@uni-bremen.de

3 TU München, DE, seidl@in.tum.de

Abstract

The goal of this Dagstuhl seminar was to gather researchers working on the theory and practice of transformations (also known as transductions) of word and tree structures, which are realised by transducers (automata with outputs). This seminar was motivated by recent advances and breakthrough results, both in the settings of words and trees.

Seminar April 2–5, 2017 – <http://www.dagstuhl.de/17142>

1998 ACM Subject Classification F.1.1 Models of Computation, F.4.3 Formal Languages

Keywords and phrases string transducers, tree transducers, expressiveness, complexity

Digital Object Identifier 10.4230/DagRep.7.4.23

Edited in cooperation with Ismaël Jecker

1 Executive Summary

Emmanuel Filiot

Sebastian Maneth

Helmut Seidl

License  Creative Commons BY 3.0 Unported license
© Emmanuel Filiot, Sebastian Maneth, and Helmut Seidl

The Dagstuhl seminar 17142 “Formal Methods of Transformations” was a short two and half day seminar that took place from April 3rd to 5th, 2017. The aim of this seminar was to bring together researchers working on theory and applications of formal models of transformations (also known as transductions) of strings and trees. A model of transformation which has been central in this seminar is that of a transducer, i.e., an automaton extended with output. Transducers were introduced in the 1960s as formal models for linguistics and syntax-directed translation in compilers. Today, research on string and tree transducers is an active field with various new applications in databases, document processing, natural language processing, software engineering, and verification. To make the seminar more focused, we had identified six research directions as key topics for the seminar:

- Canonical Normal Forms: it is well-known that regular languages admit a unique minimal (and canonical) deterministic finite automaton. Similarly, is it possible to characterize classes of transformations by canonical normal form for transducers?
- Transducer-Logic Relationships: there are well-known automata-logic correspondences in the theory of languages. What is known for transformations and can we obtain similar connections?



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

Formal Methods of Transformations, *Dagstuhl Reports*, Vol. 7, Issue 04, pp. 23–37

Editors: Emmanuel Filiot, Sebastian Maneth, and Helmut Seidl



DAGSTUHL
REPORTS

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

- Subclass Definability: a fundamental question, which requires a deep understanding of the manipulated objects, is that of subclass definability: given (an effective description of) an object in some class C , does this object belong to a given subclass C' of C ? What are the recent results and open problems with respect to subclass definability for transformations?
- New Decidability Results: what are the recent decidability breakthroughs in the theory of transformations and what are the important open problems?
- New Transducer Models: is there a need for new transducer models tailored to new applications?
- Model-Checking Data-Centric Systems: what are the potential applications of transducer theory to the verification of systems that transform data?

The seminar gave a large overview of recent results and open problems with respect to these research directions. It was a follow-up of Dagstuhl seminar 13192 “Tree Transducers and Formal Methods”, which now included researchers on string transducers. String transducers have indeed received a lot of attention in the recent years and new important results have been obtained. The aim of this seminar was also to gather researchers from the string and tree transducer communities. There were 31 participants from 10 countries (Sweden, France, Poland, Germany, US, Belgium, UK, Japan, Portugal, and India). These participants were invited by the organizers Emmanuel Filiot, Sebastian Maneth, and Helmut Seidl to give survey talks and shorter talks on their current research.

The seminar started with a survey talk by Emmanuel Filiot on recent results on string transformations, intended to motivate the need to address the theory of string transformations in this seminar, and to introduce some of its main recent breakthroughs. Mikolaj Bojanczyk then gave a survey talk on the notion of *origins* in transformations. Origins are inherent to all known transducer models and making them explicit in the semantics of transducers gave rise to new decidability and definability results. Mikolaj’s talk was followed by a session on string transducers. Luc Dartois presented a new and expressive logic to define string transformations. The logic offers good decidability properties with respect to satisfiability and equivalence (under the origin semantics). String transformations, under the origin semantics, can be seen as sets of graphs, called origin graphs. Bruno Guillon presented a characterization of the class of origin graphs generated by known string transducer models. Finally, Jacques Sakarovitch closed this session by presenting a contribution, and an open problem, on rational base numeration systems and their analysis by means of automata and transducers.

The afternoon session started with a survey talk by Christof Löding on the automatic synthesis of deterministic transducers of strings and trees, from specifications given by non-deterministic transducers. This problem, called the uniformization problem, can be seen as a variant of the classical Church synthesis problem. Two short talks given by Nathan Lhote and Michaël Cadilhac presented their recent contributions on definability problems for string transducers. Nathan Lhote presented decidability results for checking whether a given rational function is first-order definable, and, more generally, whether it is definable by some string transducer whose (input) transition monoid (which disregards the outputs) belongs to some given class, thus lifting to transformations well-known results from the theory of regular languages. Michaël Cadilhac then introduced the notion of C -continuity for string transformations and showed effectiveness of this notion for particular classes C . This notion gives an alternative and machine-independent way, which also takes outputs into account, of defining meaningful subclasses of transformations.

Day 2 (morning) was a session devoted to applications. It started with a survey talk on symbolic transducers, by Margus Veanes. Symbolic transducers have been introduced to

address practical issues when dealing with very large alphabets (even infinite). Transition labels are replaced by predicates defining sets of allowed labels. Fundamental decidability results and practical applications were presented. Based on symbolic string transducers, Loris d'Antoni then presented a fully automatic method to invert a practical class of functional (and injective) transformations, implemented in the tool GENIC. Adrien Boiret then addressed the problem of checking equivalence for symbolic top-down tree transducers. Keisuke Nakano presented some fundamental results on B-terms, which model function composition. Finally, the morning session ended with a talk by Olivier Carton about the notion of infinite word compression by transducers and its relation to normality.

Andreas Maletti started the afternoon session with a survey talk on tree transducers and their applications in linguistics. He presented transducer models well-suited to some linguistic applications such as automatic natural language translation, and showed a comparison with deep learning approaches. Frank Drewes then discussed the notion of graph transformations, and especially of DAGs, and proposed in his talk a definition of DAG transducers. Motivated by natural-language interface applications, Johanna Björklund presented a study of the expressivity of checking stack transducers, which extend with outputs the well-known model of checking stack automata. Day 2 ended with a talk by Helmut Seidl on a result obtained together with Sebastian Maneth and Gregor Kemper on the decidability of equivalence for deterministic top-down tree-to-string transducers, which had been a long standing open problem.

The last day was dedicated to results on the class of string transformations defined by two-way transducers, mostly definability problems. There were two talks, one by Olivier Gauwin and one by Jean-Marc Talbot, presenting two different techniques, for deciding whether a transformation in this class can be defined by a one-way transducers, i.e., is rational. It was open whether this problem was solvable in elementary time, and the two presented approaches answer this question positively. Another definability problem about the minimization of the number of registers in subclasses of SSTs (a model introduced by Alur and Cerny, with the same expressivity as two-way transducers), was discussed by Pierre-Alain Reynier. Ismaël Jecker presented a class of transducers with strong constraints on their structures, but still as expressive as two-way transducers, and at the same time enjoying many good algorithmic properties, in particular improving by one exponential a famous result by Hopcroft and Ullman. Finally, Krishna S. concluded the seminar by sketching an alternative proof of a result by Alur et. al. on regular expressions for string transformations.

We had long after-lunch breaks (till 4pm) to give the opportunity to the participants to discuss. It was greatly appreciated by the participants, and some of them initiated new collaborations. The discussions inspired new ideas and we hope that joint papers will be published by the participants.

We warmly thank Schloss Dagstuhl for making this fruitful event possible, and for their help in the organization. It is highly appreciated as organizers, and allowed us to focus only on the scientific aspects of the seminar.

2 Table of Contents

Executive Summary

<i>Emmanuel Filiot, Sebastian Maneth, and Helmut Seidl</i>	23
--	----

Overview of Talks

Deterministic Stack Transducers – On the Regularity of Stack Languages <i>Johanna Björklund</i>	28
Transducers with origin information <i>Mikolaj Bojanczyk</i>	28
Continuity and Rational Functions <i>Michaël Cadilhac</i>	29
Compression with transducers <i>Olivier Carton</i>	29
Automatic Program Inversion using Symbolic Transducers <i>Loris d’Antoni</i>	29
A decidable logic for string transductions <i>Luc Dartois</i>	30
Recent Results on Word Transductions <i>Emmanuel Filiot</i>	30
Untwisting two-way transducers in elementary time <i>Olivier Gauwin and Anca Muscholl</i>	31
Which classes of origin graphs are generated by transducers? <i>Bruno Guillon</i>	31
On Reversible Transducers <i>Ismaël Jecker</i>	32
Towards algebraic and logical characterizations of rational functions <i>Nathan Lhote, Emmanuel Filiot, and Olivier Gauwin</i>	32
Tree Automata in Parsing and Machine Translation <i>Andreas Maletti</i>	33
On repetitive right application of B-terms <i>Keisuke Nakano</i>	33
Two twinning properties with applications to register minimisation in streaming string transducers <i>Pierre-Alain Reynier</i>	33
Transducers as odometers in rational base numeration systems <i>Jacques Sakarovitch</i>	34
Equivalence of Deterministic Tree-to-String Transducers is Decidable <i>Helmut Seidl, Gregor Kemper, and Sebastian Maneth</i>	35
Combinator Expressions for Regular Transformations <i>Krishna Shankaranarayanan</i>	35
From Two-way to One-way Finite State Transducers: a Shepherdson’s Approach <i>Jean-Marc Talbot</i>	35

Participants 37

3 Overview of Talks

3.1 Deterministic Stack Transducers – On the Regularity of Stack Languages

Johanna Björklund (University of Umeå, SE)

License © Creative Commons BY 3.0 Unported license
© Johanna Björklund

Joint work of Suna Bensch, Johanna Björklund, Martin Kutrib

Main reference S. Bensch, J. Björklund, M. Kutrib, “Deterministic Stack Transducers”, in Proc. of the Int’l Conf. on Implementation and Application of Automata (CIAA 2016), pp. 27–38, Springer, 2016.

We introduce and investigate stack transducers, which are one-way stack automata with an output tape. A one-way stack automaton is a classical pushdown automaton with the additional ability to move the stack head inside the stack without altering the contents. For stack transducers, we distinguish between a digging and a non-digging mode. In digging mode, the stack transducer can write on the output tape when its stack head is inside the stack, whereas in non-digging mode, the stack transducer is only allowed to emit symbols when its stack head is at the top of the stack. These stack transducers have a motivation from natural-language interface applications, as they capture long-distance dependencies in syntactic, semantic, and discourse structures. We study the computational capacity for deterministic digging and non-digging stack transducers, as well as for their non-erasing and checking versions. We finally show that even for the strongest variant of stack transducers the stack languages are regular.

3.2 Transducers with origin information

Mikolaj Bojanczyk (University of Warsaw, PL)

License © Creative Commons BY 3.0 Unported license
© Mikolaj Bojanczyk

Joint work of Laure Daviaud, Bruno Guillon, Vincent Penelle

Main reference M. Bojanczyk, “Transducers with Origin Information”, in Proc. of the 41st Int’l Colloquium on Automata, Languages, and Programming (ICALP 2014), LNCS, Vol. 8573, pp. 26–37, Springer, 2014.

URL http://dx.doi.org/10.1007/978-3-662-43951-7_3

This talk is about origin semantics for transducers. The idea is to change the semantics of a transducer, by adding information about which positions in the output originate from which positions in the input. Under such semantics, called origin semantics, fewer transducers are equivalent to each other, e.g. the reverse and identity transformations over a one letter alphabet are not equivalent under origin semantics, but they are equivalent under standard semantics. I will discuss how, with origin semantics, many technical and combinatorial problems go away. There is also an implicit and debatable philosophical point, which is that origin semantics better reflects the “true essence” of a transducer, because when thinking of a transducer one also typically has in mind some origin information.

3.3 Continuity and Rational Functions

Michaël Cadilhac (Universität Tübingen, DE)

License © Creative Commons BY 3.0 Unported license
© Michaël Cadilhac

Joint work of Michaël Cadilhac, Olivier Carton, Charles Paperman

Main reference M. Cadilhac, O. Carton, C. Paperman, “Continuity and Rational Functions”, in Proc. of the 44th Int’l Colloquium on Automata, Languages, and Programming (ICALP 2017), LIPIcs, Vol. 80, pp. 115:1–115:14, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2017.

A word-to-word function is continuous for a class of languages V if its inverse maps V -languages to V . This notion provides a basis for an algebraic study of transducers, and was integral to the characterization of the sequential transducers computable in some circuit complexity classes.

Here, we report on the decidability of continuity for functional transducers and some standard classes of regular languages. Previous algebraic studies of transducers have focused on the structure of the underlying input automaton, disregarding the output. We propose a comparison of the two algebraic approaches through two questions: When are the automaton structure and the continuity properties related, and when does continuity propagate to superclasses?

3.4 Compression with transducers

Olivier Carton (University Paris-Diderot, FR)

License © Creative Commons BY 3.0 Unported license
© Olivier Carton

Joint work of Verónica Becher, Olivier Carton, Pablo Ariel Heiber

We strengthen the theorem that establishes that deterministic finite transducers can not compress normal infinite words. We prove that, indeed, non-deterministic finite transducers, even augmented with a fixed number of counters, can not compress normal infinite words. However, there are push-down non-deterministic transducers that can compress normal infinite words. We also obtain new results on the preservation of normality with automata selectors. Complementing Agafonov’s theorem for prefix selectors, we show that suffix selectors also preserve normality. However, there are simple two-sided selectors that do not preserve normality.

3.5 Automatic Program Inversion using Symbolic Transducers

Loris d’Antoni (University of Wisconsin – Madison, US)

License © Creative Commons BY 3.0 Unported license
© Loris d’Antoni

URL <http://pages.cs.wisc.edu/~loris/papers/pldi17inversion.pdf>

We propose a fully-automated technique for inverting functional programs that operate over lists such as string encoders and We propose a fully-automated technique for inverting functional programs that operate over lists such as string encoders and decoders. We consider programs that can be modeled using symbolic extended finite transducers (s-EFTs), an expressive model that can describe complex list-manipulating programs while retaining

several decidable properties. Concretely, given a program P expressed as an s-EFT, we propose techniques for: 1) checking whether P is injective and, if that is the case, 2) building an s-EFT P^{-1} describing its inverse. We first show that it is undecidable to check whether an s-EFT is injective and propose an algorithm for checking injectivity for a restricted, but a practical class of s-EFTs. We then propose an algorithm for inverting s-EFTs based on the following idea: if an s-EFT is injective, inverting it amounts to inverting all its individual transitions. We leverage recent advances in program synthesis and show that the transition inversion problem can be expressed as an instance of the syntax-guided synthesis framework. Finally, we implement the proposed techniques in a tool called GENIC and show that GENIC can invert 13 out of 14 real complex string encoders and decoders, producing inverse programs that are almost identical to manually written ones.

3.6 A decidable logic for string transductions

Luc Dartois (Free University of Brussels, BE)

License © Creative Commons BY 3.0 Unported license
© Luc Dartois

Joint work of Luc Dartois, Emmanuel Filiot, Nathan Lhote

Main reference L. Dartois, E. Filiot, N. Lhote, “Decidable Logics for Transductions and Data Words”, arXiv:1701.03670 [cs.FL], 2017.

URL <https://arxiv.org/abs/1701.03670>

We introduce a logic, called LT, to express properties of transductions, i.e. binary relations from input to output (finite) words. In LT, the input/output dependencies are modeled via an origin function which associates with any position of the output word, the input position from which it originates. The logic LT can express all MSO-definable functions, and is incomparable with MSO-transducers for relations. Despite its high expressive power, we show, among other interesting properties, that LT has decidable satisfiability and equivalence problems. The transduction logic LT is shown to be expressively equivalent to a logic for data words, LD, up to some bijection from transductions with origin to data words (the origin of an output position becomes the data of that position). The logic LD, which is interesting in itself and extends in expressive power known logics for data words, is shown to have decidable satisfiability.

3.7 Recent Results on Word Transductions

Emmanuel Filiot (Free University of Brussels, BE)

License © Creative Commons BY 3.0 Unported license
© Emmanuel Filiot

Main reference E. Filiot, P.-A. Reynier, “Transducers, logic and algebra for functions of finite words”, SIGLOG News, Vol. 3(3), pp. 4–19, ACM, 2016.

URL <https://doi.org/10.1145/2984450.2984453>

This talk surveys some important recent results from the theory of word transductions. It motivates the need to have a Dagstuhl seminar on transductions that includes word transductions, and gives a very brief overview of the talks of this seminar related to word transductions. It is organised around three main pillars for word transductions: computational models (transducers), logical models, and algebraic characterizations. Both the cases of functional and relational transductions are presented.

Some parts of the talks are based on the following surveys:

References

- 1 Emmanuel Filiot and Pierre-Alain Reynier. *Transducers, logic and algebra for functions of finite words*, SIGLOG News, Vol. 3(3), pages 4-19, 2016.
- 2 Emmanuel Filiot. *Logic-Automata Connections for Transformations*. Proceedings of the Indian Conference of Logic and Applications (ICLA). Pages 30-57. 2015.

3.8 Untwisting two-way transducers in elementary time

Olivier Gauwin (University of Bordeaux, FR) and Anca Muscholl (University of Bordeaux, FR)

License © Creative Commons BY 3.0 Unported license
© Olivier Gauwin and Anca Muscholl

Functional transductions realized by two-way transducers (equivalently, by streaming transducers and by MSO transductions) are the natural and standard notion of “regular” mappings from words to words. It was shown recently (LICS’13) that it is decidable if such a transduction can be implemented by some one-way transducer, but the given algorithm has non-elementary complexity. We provide an algorithm of different flavor solving the above question, that has double exponential space complexity. We further apply our technique to decide whether the transduction realized by a two-way transducer can be implemented by a sweeping transducer, with either known or unknown number of passes.

3.9 Which classes of origin graphs are generated by transducers?

Bruno Guillon (University of Warsaw, PL)

License © Creative Commons BY 3.0 Unported license
© Bruno Guillon

Joint work of Mikołaj Bojańczyk, Laure Daviaud, Bruno Guillon, Vincent Penelle
Main reference M. Bojańczyk, L. Daviaud, B. Guillon, V. Penelle, “Which classes of origin graphs are generated by transducers?”, in Proc. of the 44th Int’l Colloquium on Automata, Languages, and Programming (ICALP 2017), LIPIcs, Vol. 80, pp. 114:1–114:13, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2017.
URL <http://dx.doi.org/10.4230/LIPIcs.ICALP.2017.114>

We consider the origin semantics of transducers (2-way automata with outputs, streaming string transducers, string-to-string MSO-transduction) which extends relations by providing a mapping from positions of the output into positions of the input, saying how the output originates from the input. In the other hand, every relation on words can be extended by such a mapping, leading to the notion of origin transduction, viewed as set of particular graph, namely the origin graphs. We exhibit structural properties of origin transductions which are necessary and sufficient for characterising the origin semantics of streaming string transducers.

3.10 On Reversible Transducers

Ismaël Jecker (Free University of Brussels, BE)

License © Creative Commons BY 3.0 Unported license
© Ismaël Jecker

Joint work of Dartois, Luc; Fournier, Paulin; Lhote, Nathan

Main reference L. Dartois, P. Fournier, I. Jecker, N. Lhote, “On Reversible Transducers”, in Proc. of the 44th Int’l Colloquium on Automata, Languages, and Programming (ICALP 2017), LIPIcs, Vol. 80, pp. 113:1–113:12, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2017.

URL <http://dx.doi.org/10.4230/LIPIcs.ICALP.2017.113>

Deterministic two-way transducers define the robust class of regular functions which, among other good properties, is closed under composition. However, the best known algorithms for composing two-way transducers cause a double exponential blow-up in the size of the inputs. We expose a class of transducers for which the composition has polynomial complexity. It is the class of reversible transducers, for which the computation steps can be reversed deterministically. While in the one-way setting this class is not very expressive, we prove that any two-way transducer can be made reversible through a single exponential blow-up. As a consequence, we obtain that the composition of two-way transducers can be done with a single exponential blow-up in the number of states.

3.11 Towards algebraic and logical characterizations of rational functions

Nathan Lhote (University of Bordeaux, FR), Emmanuel Filiot (Free University of Brussels, BE), and Olivier Gauwin (University of Bordeaux, FR)

License © Creative Commons BY 3.0 Unported license
© Nathan Lhote, Emmanuel Filiot, and Olivier Gauwin

Regular languages are characterized by several different formalisms including monadic second-order logic (MSO) and also congruences of finite index. Many links have been established between fragments of MSO and varieties of congruences (sets of congruences with good closure properties). These characterization often yield an effective way to decide if a language can be expressed in a given logic, the most prominent example being the Schutzenberger/McNaughton-Papert theorem which characterizes the first-order languages as the languages recognized by aperiodic congruences. We present how to lift some of these results to word-to-word transductions, with some assumptions on the logical fragment amounting to having access to the linear order over words.

3.12 Tree Automata in Parsing and Machine Translation

Andreas Maletti (*Universität Leipzig, DE*)

License © Creative Commons BY 3.0 Unported license
© Andreas Maletti

Joint work of Fabienne Braune, Joost Engelfriet, Zoltan Fülöp, Daniel Quernheim, Nina Seemann

Main reference J. Engelfriet, Z. Fülöp, A. Maletti, “Composition Closure of Linear Extended Top-down Tree Transducers”, *Theory Comput. Syst.* 60(2):129–171, 2017.

URL <https://doi.org/10.1007/s00224-015-9660-2>

We will discuss how tree automata were rediscovered in the area of statistical parsing of natural language sentences and demonstrate that some techniques developed in that area might also be beneficial in automata theory. On the example of syntax-based machine translation, we will demonstrate the other direction showing how automata theory can provide solutions to problems in natural language processing. With the identification of the exact expressive power in terms of standard models and known closure properties for them, we developed a syntax-based translation system that can translate not only a single parse, but rather the full parse forest delivered by the parser.

3.13 On repetitive right application of B-terms

Keisuke Nakano (*The University of Electro-Communications – Tokyo, JP*)

License © Creative Commons BY 3.0 Unported license
© Keisuke Nakano

Joint work of Mirai Ikebuchi, Keisuke Nakano

Main reference M. Ikebuchi, K. Nakano, “On repetitive right application of B-terms”, arXiv:1703.10938 [cs.LO], 2017.

URL <http://arxiv.org/abs/1703.10938>

B-terms are built from the B combinator alone defined by $B f g x = f (g x)$, which is well-known as a function composition operator. This talk introduces an interesting property of B-terms, that is, whether repetitive right applications of a B-term circulates or not. Conditions for B-terms to and not to have the proposition are discussed. Specifically, we show that some of infinitely many B-terms which do not have the property can be given as a regular tree language.

3.14 Two twinning properties with applications to register minimisation in streaming string transducers

Pierre-Alain Reynier (*Aix-Marseille University, FR*)

License © Creative Commons BY 3.0 Unported license
© Pierre-Alain Reynier

Joint work of Laure Daviaud, Ismaël Jecker, Jean-Marc Talbot and Didier Villevalois

Main reference L. Daviaud, P.-A. Reynier, J.-M. Talbot, “A Generalised Twinning Property for Minimisation of Cost Register Automata”, in *Proc. of the 31st Annual ACM/IEEE Symp. on Logic in Computer Science (LICS 2016)*, pp. 857–866, ACM, 2016.

URL <http://dx.doi.org/10.1145/2933575.2934549>

In this talk, I present recent results on transducers of finite words. In this setting, input-deterministic machines are less expressive than non-deterministic ones. Non-deterministic transducers admitting an equivalent deterministic one are characterized by the twinning property introduced by Choffrut in 1977.

I will present two modifications of this property allowing to solve two different minimization problems :

- first, a simple restriction of the recently introduced model of streaming string transducers, called appending SST, is expressively equivalent to non-deterministic functional transducers. The twinning property of order k , introduced in [1], precisely characterizes the transducers that can be realized by an appending SST with k registers.
- second, though not realizable by an input-deterministic transducer, some functions, called multi-sequential, may be realized as finite unions of such machines. I present the branching twinning property of order k , introduced in [2], allowing to characterize functions that can be realized as a union of k input-deterministic transducers.

References

- 1 Laure Daviaud, Pierre-Alain Reynier, Jean-Marc Talbot. A Generalised Twinning Property for Minimisation of Cost Register Automata. LICS 2016: 857-866
- 2 Laure Daviaud, Ismaël Jecker, Pierre-Alain Reynier, Didier Villevalois: Degree of Sequentiality of Weighted Automata. FoSSaCS 2017: 215-230

3.15 Transducers as odometers in rational base numeration systems

Jacques Sakarovitch (Telecom ParisTech, FR)

License  Creative Commons BY 3.0 Unported license

© Jacques Sakarovitch

Joint work of Shigeki Akiyama, Victor Marsault, Jacques Sakarovitch

The definition of numeration systems with rational base, in a joint work with S. Akiyama and Ch. Frougny, has allowed to make some progress in a number theoretic problem, by means of automata theory and combinatorics of words. At the same time, it raised the problem of understanding the structure of the sets of the representations of the integers in these systems from the point of view of formal language theory.

At first sight, these sets look rather chaotic and do not fit well in the classical Chomsky hierarchy of languages. They all enjoy a property that makes them defeat, so to speak, any kind of iteration lemma. On the other hand, these sets also exhibit remarkable regularity properties.

During the recent years, these regularities have been studied in a series of joint papers with my student V. Marsault. In particular, we have shown that periodic signatures are characteristic of the representation languages in rational base numeration systems. For given rational base p/q , the representation language naturally define an infinite tree $T_{p/q}$ the nodes of which are the integers. Every integer n is then associated with the infinite minimal branch whose label is an infinite word w_n which is characteristic of n . The next remarkable property is that the function that maps w_n onto w_{n+1} is realized by a letter-to-letter transducer whose underlying graph is precisely $T_{p/q}$.

References

- 1 S. Akiyama, Ch. Frougny, and J. Sakarovitch. Powers of rationals modulo 1 and rational base number systems. *Israel J. Math.*, 168:53–91, 2008.
- 2 S. Akiyama, V. Marsault, and J. Sakarovitch. Auto-similarity in rational base number systems. *Proc. of WORDS 2013*, LNCS 8079:34–45, 2013.
- 3 V. Marsault and J. Sakarovitch. Trees and languages with periodic signature *Indagationes Mathematicae* 28:221–246, 2017.

3.16 Equivalence of Deterministic Tree-to-String Transducers is Decidable

Helmut Seidl (TU München, DE), Gregor Kemper, and Sebastian Maneth (University of Edinburgh, GB)

License © Creative Commons BY 3.0 Unported license
 © Helmut Seidl, Gregor Kemper, and Sebastian Maneth
Main reference H. Seidl, S. Maneth, G. Kemper, “Equivalence of Deterministic Top-down Tree-to-String Transducers is Decidable”, in Proc. of the 56th Ann. Symp. on Foundations of Computer Science (FOCS 2015), pp. 943–962, IEEE, 2015.
URL <https://doi.org/10.1109/FOCS.2015.62>

We show that equivalence of deterministic top-down tree-to-string transducers is decidable, thus solving a long standing open problem in formal language theory. For our main result, we prove that equivalence can be certified by means of inductive invariants using polynomial ideals. This allows us to construct two semi-algorithms, one searching for a proof of equivalence, one for a witness of non-equivalence. Furthermore, we extend our result to deterministic top-down tree-to-string transducers which produce output not in a free monoid but in a free group.

3.17 Combinator Expressions for Regular Transformations

Krishna Shankaranarayanan (Indian Institute of Technology – Mumbai, IN)

License © Creative Commons BY 3.0 Unported license
 © Krishna Shankaranarayanan
Joint work of Vrunda Dave, Paul Gastin, Krishna Shankaranarayanan

MSO transductions, functional transductions realized by two-way transducers, as well as streaming string transducers all capture regular transformations of languages. In this work, we show that every regular function captured by a two-way transducer can be constructed from constant functions using the combinators of choice, split-sum and iterated sum, that are analogs of unambiguous union, concatenation and Kleene-star respectively. This is achieved by unambiguously factorizing the input language into forests of bounded height.

Our proof extends to finite as well as infinite words.

3.18 From Two-way to One-way Finite State Transducers: a Shepherdson’s Approach

Jean-Marc Talbot (Aix-Marseille University, FR)

License © Creative Commons BY 3.0 Unported license
 © Jean-Marc Talbot
Joint work of Benjamin Monmege, Pierre-Alain Reynier, Jean-Marc Talbot

We study functions over finite words defined by one-way and two-way finite state transducers (1FST and 2FST, respectively). 2FST are strictly more powerful than 1FST and a natural question is then to decide, given a 2FST describing a function, whether there exists an equivalent 1FST. Filiot, Gauwin, Reynier and Servais showed in 2013 that this problem is decidable, using a procedure based on the construction of Rabin and Scott for transforming two-way into one-way automata, yielding a non-elementary complexity. In this article, we

prove that the problem is in 2-EXPSpace. Our proof technique relies on a new representation of the transducer based on the crossing sequences introduced in Shepherdson's construction (an alternative to the one of Rabin and Scott), and on a normalisation of this transducer. In case of one-way definability, we produce an equivalent 1FST (deterministic with look-ahead) of size at most triply-exponential in the size of the original 2FST.

Participants

- Johanna Björklund
University of Umeå, SE
- Adrien Boiret
INRIA Lille, FR
- Mikolaj Bojanczyk
University of Warsaw, PL
- Michaël Cadilhac
Universität Tübingen, DE
- Olivier Carton
University Paris-Diderot, FR
- Loris d'Antoni
University of Wisconsin –
Madison, US
- Luc Dartois
Free University of Brussels, BE
- Frank Drewes
University of Umeå, SE
- Emmanuel Filiot
Free University of Brussels, BE
- Olivier Gauwin
University of Bordeaux, FR
- Bruno Guillon
University of Warsaw, PL
- Florent Jacquemard
IRCAM & INRIA – Paris, FR
- Ismaël Jecker
Free University of Brussels, BE
- Aurélien Lemay
INRIA Lille, FR
- Nathan Lhote
University of Bordeaux, FR
- Christof Löding
RWTH Aachen, DE
- Andreas Maletti
Universität Leipzig, DE
- Sebastian Maneth
University of Edinburgh, GB
- Anca Muscholl
University of Bordeaux, FR
- Keisuke Nakano
The University of
Electro-Communications –
Tokyo, JP
- Joachim Niehren
INRIA Lille, FR
- Chih-Hao Luke Ong
University of Oxford, GB
- Raphaela Palenta
TU München, DE
- Rogério Reis
University of Porto, PT
- Pierre-Alain Reynier
Aix-Marseille University, FR
- Jacques Sakarovich
Telecom ParisTech, FR
- Helmut Seidl
TU München, DE
- Frédéric Servais
Free University of Brussels, BE
- Krishna Shankaranarayanan
Indian Institute of Technology –
Mumbai, IN
- Jean-Marc Talbot
Aix-Marseille University, FR
- Margus Veanas
Microsoft Corporation –
Redmond, US



Ambient Notification Environments

Edited by

Lewis Chuang¹, Sven Gehring², Judy Kay³, and
Albrecht Schmidt⁴

- 1 Max Planck Institute for Biological Cybernetics, Germany
lewis.chuang@tuebingen.mpg.de
- 2 German Research Center for Artificial Intelligence, Germany
sven.gehring@dfki.de
- 3 The University of Sydney, Australia judy.kay@sydney.edu.au
- 4 University of Stuttgart, Germany albrecht.schmidt@acm.org

Abstract

Direct notifications are on the exponential rise. In our time, numerous personal computing devices and applications vie for limited attention, racing to deliver large amounts of information to us. This often results in users being overwhelmed by notifications and interruptions to their regular schedule, to whom a complete avoidance of technology seems to be the only viable option. In other words, the current approach for notification delivery is unsustainable and will not scale. In the Dagstuhl Seminar 17161 “Ambient Notification Environments” we brought together experts from different fields related to smart homes, ambient intelligence, human-computer interaction, activity recognition, and psychology to discuss a potential alternative approach: ambient notifications. We explored how ambient notifications can support people in their daily activities, by providing relevant information that are contextually embedded in the environment. The objective is to facilitate unobtrusive access to information at the right time and in the right place, hence reducing the disruptions and annoyances that are commonly associated with direct notifications. In this report, we present the numerous ideas and concepts of how the research community could strive toward towards realising ambient notifications. This is based on the presentations and activities conducted during the seminar. Overall, the community is in agreement that current approaches to notifications will not scale and that ambient notifications are a potential solution.

Seminar April 17–20, 2017 – <http://www.dagstuhl.de/17161>

1998 ACM Subject Classification H.5.m. Information interfaces and presentation (e.g., HCI):
Miscellaneous

Keywords and phrases Ambient Notifications, Dagstuhl Seminar, Skill Transfer

Digital Object Identifier 10.4230/DagRep.7.4.38

Edited in cooperation with Pascal Knierim and Alexandra Void

1 Executive Summary

Albrecht Schmidt (University of Stuttgart, DE, albrecht.schmidt@acm.org)

License  Creative Commons BY 3.0 Unported license
© Albrecht Schmidt

Reports indicate that many users interact with their smartphone and wearable devices more than 100 times per day. Oftentimes, these interactions result from direct notifications, presented on the screen or via sounds. New communication applications (e.g. from email to WhatsApp) has increased the frequency of notifications, while social media applications are



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

Lewis Chuang, Sven Gehring, Judy Kay and Albrecht Schmidt, *Dagstuhl Reports*, Vol. 7, Issue 4, pp. 38–82

Editors: Lewis Chuang, Sven Gehring, Judy Kay and Albrecht Schmidt



DAGSTUHL
REPORTS

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

inherently motivated to entice repeat visits and interactions. The end result is that more and more systems as well as applications compete for the user's attention using notifications. Projecting this into the future, it is apparent that current implementation schemes that rely on direct notifications will not scale. A simple extrapolation of the rate in notifications suggests that a near future whereby users will only attend to notifications with no time leftover for productive work. Therefore, a radical restructure of notification delivery is necessary – specifically, one that keeps the user in-the-loop without consuming all of the user's attention. This is a key challenge. If no alternatives to direct notifications can be realized, current visions of ubiquitous computing and smart environments are likely to be unrealistic, given their anticipated undesirability to the end-user. In the Dagstuhl Seminar 17161 "Ambient Notification Environments", we looked at how novel approaches to notification delivery can address the above issue. We brought together researchers and experts that understand the technical, psychological, and social aspects of notification systems. This facilitated a broad discussion that was fuelled by the discussants' joint expertise in mobile and smart home technologies, in ambient sensing and presentation, and psychological models of human attention, to name a few. This discussion brought to the foreground, many of the underlying challenges that deserves further research. On the hand, a technological push in novel communication and smart devices drives an immediacy in user interaction. On the other hand, this will result in a higher demand on human attention. In the current report, we document the diverse approaches that were proposed as an alternative means towards notifications that are more personalized and contextualized, with the express purpose of reducing user effort. The central aim of the seminar was to understand the challenges and questions that we face when designing future interactive systems and to overcome the mounting notification problem. In addition to the ideas raised in the individual presentations and the group sessions, we jointly identified the following research questions and challenges:

1. How can notifications be designed to be simultaneously non-intrusive and yet noticeable?
2. How can artificial intelligence be designed to provide effective context-aware notifications?
3. What is a suitable notification architecture for integrating user devices, smart (shared) environments, and personal data without compromising personal privacy?
4. What is a suitable conceptualization for notification systems and how will a taxonomy for classifying notification delivery that is centered on user preferences and privacy look like?

Overall, the seminar contributed to a common vision of how interactions between human and systems can and ought to progress – one where technological progress does not necessitate an ever growing burden on the user's attention. In this report, we balance an analysis of the underlying drivers and problems, ideas for novel conceptual and technical approach, and most importantly a set of questions and research challenges in this domain that merit the attention of like-minded researchers.

2 Table of Contents

Executive Summary	
<i>Albrecht Schmidt</i>	38
Overview of Pecha Kucha Presentations	
Ambient Notification Techniques to Display Changes of the Ecosystem	
<i>Aaron Quigley</i>	42
Notifications will not scale with Internet of Things	
<i>Albrecht Schmidt</i>	43
Notification Strategies for Mobile and Smart Home Notifications	
<i>Alexandra Voit</i>	44
Peripheral Interaction	
<i>Andreas Butz</i>	45
Notifications and their Impact on the Awareness for Activities	
<i>Aurélien Tabard</i>	46
Embedded device use	
<i>Barry Brown</i>	47
Smart Notification Management for Personal Assistants	
<i>Benjamin Poppinga</i>	47
Intelligent Messages	
<i>Christoph Gebhardt</i>	48
Unobtrusive Interaction for Wearable Computing	
<i>David Dobbstein</i>	49
Augmenting Social Interactions: Realtime Behavioural Feedback using Social Signal Processing Techniques	
<i>Elisabeth André</i>	50
Ambient Episodic Memory Augmentation with Context	
<i>Evangelos Niforatos</i>	51
Ambient Notifications Environments as a Support for Daily Activities	
<i>Frederik Wiehr</i>	53
Blended Interaction and Ambient Notification Environments	
<i>Harald Reiterer</i>	54
Ubiquitous Sensing	
<i>Joseph Paradiso</i>	55
Lifelong Learning	
<i>Judy Kay</i>	56
Exploring the role of notifications relating to Group QoS and proactive behaviour in IoT environments.	
<i>Keith Cheverst</i>	57
Updating without interrupting	
<i>Lewis Chuang</i>	58

Ambient Notifications for Dementia Care <i>Markus Löchtefeld</i>	59
Notifications through Indirect Illumination and Pneumatic Compression <i>Michael Rohs</i>	60
Towards Holistic Notification Management <i>Niels Henze</i>	61
Notifications in Mixed Reality <i>Pascal Knierim</i>	61
Towards Visual Notifications in Multi-User Large-Display-Environments <i>Raimund Dachsel</i>	62
Multimodal Notifications <i>Steve Brewster</i>	64
Guiding Attention in Cyber-Physical Environments <i>Susanne Boll</i>	65
Notification Management is Important to Avoid Information Overload <i>Sven Gehring</i>	66
Towards Attention-Aware Adaptive Notification in Advanced Ubiquitous Computing <i>Tadashi Okoshi</i>	68
Interactive Assistance in Working and Teaching Facilities <i>Thomas Kosch</i>	69
Awareness in Computer Supported Cooperative Work <i>Tom Gross</i>	70
The Role of Task Engagement in Mobile Attention Management <i>Veljko Pejovic</i>	71
Sessions	
Notifications vs. Augmentation <i>Joseph A. Paradiso</i>	72
Multimodal Notifications for Ambient Environments <i>Steve Brewster</i>	73
Augmenting Social Interactions <i>Elisabeth André</i>	74
Debates	75
Speed Dating Session	79
Open Maker Session	80
Conceptual Maker Session	81
Participants	82

3 Overview of Pecha Kucha Presentations

3.1 Ambient Notification Techniques to Display Changes of the Ecosystem

Aaron Quigley (*University of St. Andrews, GB, aquigley@st-andrews.ac.uk*)

License  Creative Commons BY 3.0 Unported license
© Aaron Quigley

Our work has explored a variety of screen, widget and pixel level ambient notification techniques for visualizing display changes in ecosystem of display environments [3]. In such environments, users can easily become overloaded and become unaware of display changes as they alternate their attention towards different displays. We have explored and patented subtle gaze-dependent techniques for visualizing change on unattended displays with techniques including FreezeFrame, PixMap, WindowMap and Aura [1]. We have extended this to explore a range of four notification levels unnoticeable, subtle, intrusive and disruptive. These range from an almost imperceptible visual change to a clear and visually salient change. We explored such ambient notifications in AwToolkit, a novel widget set for developers that supports users in maintaining awareness in multi-display systems [2]. The first set of techniques were studied over consecutive days and the results of the study show that the participant found the system and the techniques useful, subtle, calm and non-intrusive. With AwToolkit the evaluation results reveal a marked increase in user awareness in comparison to the same application implemented without AwToolkit.

Our work indicates that subtle visualization techniques can indeed positively change the display attendance behavior of a working professional who used these techniques as part of their regular work activities. However, many opportunities still exist for taking the technology further, both in terms of refining the techniques, and in terms of identifying cost-effective means of evaluating the efficacy of such techniques. The subtle display of visual change information from unattended displays can alter our interactions and expectations of interfaces. Out of sight no longer means out of mind, when inattention is no longer an inherent cost but a new research and development opportunity.

Finally, our work suggests a visual continuum from effectively invisible ambient notifications to highly disruptive notifications which demand attention. Context-aware systems need to understand when notifications have been perceived and acted upon as required before promoting them through a chain from unnoticeable, to subtle, then intrusive and finally disruptive. Today, too many notifications systems begin with disruptive action, rather than resorting to that as a means of notification as last resort.

References

- 1 J. Dostal, P. O. Kristensson, and A. Quigley. Subtle gaze-dependent techniques for visualising display changes in multi-display environments. In Proceedings of the 2013 international conference on Intelligent user interfaces - IUI '13. Association for Computing Machinery (ACM), 2013.
- 2 J. E. Garrido, V. M. R. Penichet, M. D. Lozano, A. Quigley, and P. O. Kristensson. AwToolkit. In Proceedings of the 2014 International Working Conference on Advanced Visual Interfaces - AVI '14. Association for Computing Machinery (ACM), 2014.
- 3 A taxonomy for and analysis of multi-person-display ecosystems, L Terrenghi, A Quigley, A Dix, Journal of Personal and Ubiquitous Computing 13 (8), 583, 2009

3.2 Notifications will not scale with Internet of Things

Albrecht Schmidt (Universität Stuttgart, DE, albrecht.schmidt@acm.org)

License © Creative Commons BY 3.0 Unported license
© Albrecht Schmidt

With the Internet of Things more and more devices and systems will want to get the attention of the user [1]. Potentially hundreds of devices, sensors, and communication endpoints will create notifications that should be brought to the user's attention. Looking at the current volume of notification users receive and at the expected number of internet enabled devices it becomes clear that classical notifications will not work anymore in the near future.

To empower users to live and interact in a world of ubiquitous computing and making good use of many potentially useful interconnected things, we need to re-think how we do notifications. The key criteria for future ambient notification systems are:

1. **Transient notifications:** the large majority of notifications are transient and do not require user attention. If users do not see them or chose not to interact with them, they will disappear not requiring any interaction or attention by the user. An example for this are ephemeral interfaces [2].
2. **Natural attention management:** notifications are designed that humans can use their natural ability to filter and focus. In a forest we are not overwhelmed by the amount of things we see, even though there may be hundreds of leaves, mushrooms, plants, and insects. Humans are good at directing their attentions to things of important or aspects they are primed for.
3. **Situating information and notifications:** by contextualizing notifications to the current situation, by putting information into context, and by making data and notification available in points where humans need to make decisions [1]. Embedding information and interaction into objects reduces the overall complexity in dealing with the information as it is anchored in the real world and it hopefully reduces the demand on the user's attention.
4. **Ambient, peripheral, and aesthetic presentation:** the notifications should be designed that they can blend in with the environment and that the rather aesthetically enhance an environment than making it cluttered [3]. By these means environments can be made more pleasant and at the same time have them communicate information that can be peripherally perceived.

The long term vision is that we create technologies that make us peripherally aware of things we need to know through technologies that augment and amplify [4] our perception. This can help to overcome temporal and spatial limitations of our perception [5]. Similar to human senses, we are not required to actively seek and process information, we will just know it.

References

- 1 Kranz M, Holleis P, Schmidt A. Embedded interaction: Interacting with the internet of things. *IEEE internet computing*. 2010 Mar;14(2):46-53.
- 2 Döring T, Sylvester A, Schmidt A. A design space for ephemeral user interfaces. In *Proceedings of the 7th International Conference on Tangible, Embedded and Embodied Interaction* 2013 Feb 10 (pp. 75-82). ACM.
- 3 Gellersen HW, Schmidt A, Beigl M. Ambient media for peripheral information display. *Personal Technologies*. 1999 Dec 1;3(4):199-208.
- 4 Schmidt A. Augmenting Human Intellect and Amplifying Perception and Cognition. *IEEE Pervasive Computing*. 2017 Jan;16(1):6-10.
- 5 Schmidt A, Langheinrich M, Kersting K. Perception beyond the Here and Now. *Computer*. 2011 Feb;44(2):86-8.

3.3 Notification Strategies for Mobile and Smart Home Notifications

Alexandra Voit (University of Stuttgart, DE, alexandra.voit@vis.uni-stuttgart.de)

License  Creative Commons BY 3.0 Unported license
© Alexandra Voit

My research aims to investigate notification strategies for various kind of information as mobile or smart home notifications. Today, many users own different kinds of smart devices such as laptops, smartphones, smartwatches or tablets. Displaying notifications is an important feature that all these devices support. The smart devices inform users proactively about incoming messages, upcoming appointments or device specific information, e.g. updates. On the one hand, users appreciate notifications because they do not want to miss such information and they do not have to check the different applications manually. On the other hand, many applications use notifications to keep the user in the loop or to catch the user's attention. In addition, receiving notifications can lead to interruptions and therefore lower the task performance of the user's current primary task.

In most cases, different devices are not aware of each other. If a user uses a laptop, a smartphone, a smartwatch and a tablet, and he receives an email, all devices with an email client installed will try to notify the user. I argue that, in order to respect the user's attention, selecting a specific device to notify the user is a better approach. To gain information about on which device the users want to receive notifications in their current context, we conducted a study using the Experience Sampling Method [1], the user's proximity to the device, whether the user currently uses a device, as well as the user's current location can indicate if the user wants to receive notifications on a specific device. In addition, we investigated notification strategies to display mobile notifications, such as incoming messages or upcoming appointments, within the user's environment. With this approach, the users do not have to check their smartphones anymore. For example, we developed guidelines for displaying notifications on smart TVs [2]. Also, we developed a system to display the user's personal notifications on displays, using ambient lightning, using sounds or speech output [3]. Furthermore, smart wall calendars can be used as ambient display in the users' home to increase the attention regarding upcoming appointments or tasks and to support aging in place [4].

With the raise of the Internet of Things (IoT), smart devices in the user's environment will be able to notify the user as well. For example, current automatic vacuum cleaners notify users on their smartphones about their current state (e.g., when the cleaning process is finished). My work investigates new notification strategies to display smart home notifications [5, 6]. In an online survey, I investigated four notification strategies for four scenarios with different urgencies. The results indicate using ambient notifications at the notifying object itself as well as an ambient notification display can display for non-urgent smart home notifications as e.g., watering plants, feed animals or remind the medication intake. My future work focuses ambient notifications and their representation in the user's home.

References

- 1 Dominik Weber, Alexandra Voit, Philipp Kratzer, and Niels Henze. 2016. In-situ investigation of notifications in multi-device environments. In Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '16). ACM, New York, NY, USA, 1259-1264. DOI: <http://dx.doi.org/10.1145/2971648.2971732>
- 2 Dominik Weber, Sven Mayer, Alexandra Voit, Rodrigo Ventura Fierro, and Niels Henze. 2016. Design Guidelines for Notifications on Smart TVs. In Proceedings of the ACM Inter-

- national Conference on Interactive Experiences for TV and Online Video (TVX '16). ACM, New York, NY, USA, 13-24. DOI: <http://dx.doi.org/10.1145/2932206.2932212>
- 3 Thomas Kubitz, Alexandra Voit, Dominik Weber, and Albrecht Schmidt. 2016. An IoT infrastructure for ubiquitous notifications in intelligent living environments. In Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct (UbiComp '16). ACM, New York, NY, USA, 1536-1541. DOI: <https://doi.org/10.1145/2968219.2968545>
 - 4 Alexandra Voit, Elizabeth Stowell, Dominik Weber, Christoph Witte, Daniel KÄrcher, and Niels Henze. 2016. Envisioning an ambient smart calendar to support aging in place. In Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct (UbiComp '16). ACM, New York, NY, USA, 1596-1601. DOI: <https://doi.org/10.1145/2968219.2968555>
 - 5 Alexandra Voit, Tonja Machulla, Dominik Weber, Valentin Schwind, Stefan Schneegass, and Niels Henze. 2016. Exploring notifications in smart home environments. In Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct (MobileHCI '16). ACM, New York, NY, USA, 942-947. DOI: <https://doi.org/10.1145/2957265.2962661>
 - 6 Alexandra Voit, Dominik Weber, and Stefan Schneegass. 2016. Towards Notifications in the Era of the Internet of Things. In Proceedings of the 6th International Conference on the Internet of Things (IoT'16). ACM, New York, NY, USA, 173-174. DOI: <https://doi.org/10.1145/2991561.2998472>

3.4 Peripheral Interaction

Andreas Butz (LMU Munich, DE, andreas.butz@ifi.lmu.de)

License © Creative Commons BY 3.0 Unported license
© Andreas Butz

The focus of the Dagstuhl Seminar 17161 – Ambient Notification Environments obviously is on ambient notifications. This means, that some computing system (or some human via a computing system) is sending us a certain bit of information in an ambient way. There is work from the last century on these concepts, but already several years ago, this situation struck me a somewhat unbalanced: It seemed nice that systems would notify me in unobtrusive and peripheral ways, but how about the other direction? How about interaction in a similarly lightweight and peripheral way? Tis idea of peripheral interaction was investigated by my former PhD student Doris Hausen at LMU and simultaneously by her colleague Saskia Bakker at TU Eindhoven. The two of them edited a book on peripheral interaction [1], and in its introductory chapter they write: “In everyday life, we perform several activities in our periphery of attention. For example, we are aware of what the weather is like and we can routinely wash our hands without actively thinking about it. However, we can also easily focus on these activities when desired. Contrarily, interactions with computing devices, such as smartphones and tablet computers, usually require focused attention, or even demand it through flashing displays, beeping sounds, and vibrations used to alert people. Hence, these interactions move more unpredictably between periphery and center of attention compared to non- computer-mediated activities. With the number of computers embedded in our everyday environment increasing, inevitably interaction with these computers will move to the periphery of attention. Inspired by the way we fluently divide our attentional resources over various activities in everyday life, we call this type of interaction “peripheral interaction.”

We believe that considering and enabling peripheral interaction with computing technology contributes to more seamlessly embedding of such technology in everyday routines. This chapter briefly explores the history of peripheral interaction as a field of research and lays out how peripheral interaction, in our view, fits into the larger domain of interactive systems and HCI.” My aim when coming to the Dagstuhl Seminar 17161 – Ambient Notification Environments, was to pick up the topic where Doris Hausen had stopped. I was curious to see how others felt about the input direction. However, during the course of the seminar I realized that there are still far more open issues on the output side than I had anticipated. Some of them have only come into life recently because of the scalability issues of notifications in our daily lives. These scalability issues will certainly also apply (probably even stronger) to the input direction, which means that there will be ample research opportunities in that field for years to come.

References

- 1 S. Bakker et al. (eds.), *Peripheral Interaction*, Springer International Publishing Switzerland 2016, Human-Computer Interaction Series, DOI 10.1007/978-3-319-29523-7_1

3.5 Notifications and their Impact on the Awareness for Activities

Aurélien Tabard (University Claude Bernard, FR, aurelien@tabard.fr)

License  Creative Commons BY 3.0 Unported license
© Aurélien Tabard

My research focuses on supporting reflection on individual and group activity. It is structured around two broad questions:

- How can visualization be made more relevant to people’s context in order to better support reflection in-action [4], and
- How Activity Based Computing can accommodate ill defined activities, scaffolded activities and activities spanning across various people, places, devices and time, so that users have better means of reflecting on their past activity [1].

Self- and group-reflection are tightly connected to awareness of personal and group activities. Notifications are an interesting way to increase awareness, however they can also be highly distracting. Recently, I have been exploring how to design slow or calm notifications. These notifications can take the form of physical objects representing data, such as a run [4]. Making notifications physical and passive means that they can be always on, but won’t distract and can blend in the environment. Another strategy we explored was to integrate notification and awareness tools directly into the objects of interest. For instance in a card sorting activity the participation of each user would be integrated into the cards. Rather than prompting one person to participate more or less with an explicit notification, the UI changes smoothly to notify the group of each others participation.

Finally at a larger scale we are currently exploring how to notify people of activities taking place in libraries, a place where calmness and focus are particularly valued, but in which people are increasingly coming to participate to public events. We are exploring by broadcasting notifications through pervasive displays, mobile devices, and dynamic signage.

References

- 1 Aurélien Tabard, Wendy E. Mackay, and Evelyn Eastmond. 2008. From individual to collaborative: the evolution of prism, a hybrid laboratory notebook. In *Proceedings of the*

- 2008 ACM conference on Computer supported cooperative work (CSCW '08). ACM, New York, NY, USA, 569-578.
- 2 Aurélien Tabard, Juan-David Hincapié-Ramos, Morten Esbensen, and Jakob E. Bardram. 2011. The eLabBench: an interactive tabletop system for the biology laboratory. In Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces (ITS '11). ACM, New York, NY, USA, 202-211.
 - 3 Simon Stusak, Aurélien Tabard, Franziska Sauka, Rohit A. Khot and Andreas Butz. 2014. Activity Sculptures: Exploring the Impact of Physical Visualizations on Running Activity. In IEEE Transactions on Visualization and Computer Graphics, vol. 20, no. 12, pp. 2201-2210.
 - 4 Valentin Lachand, Audrey Serna, Aurélien Tabard, and Jean-Charles Marty. 2016. The impact of indexical and symbolic indicators on the regulation of collaborative activities.. In Actes de la 28ième conférence francophone sur l'Interaction Homme-Machine (IHM '16). ACM, New York, NY, USA, 144-154.
 - 5 Adam Rule, Aurélien Tabard, and Jim Hollan. 2017. Using Visual Histories to Reconstruct the Mental Context of Suspended Activities. Human-Computer Interaction, Taylor and Francis. 1-48.

3.6 Embedded device use

Barry Brown (*University of Stockholm, SE barry@dsv.su.se*)

License © Creative Commons BY 3.0 Unported license
© Barry Brown

In this talk I discuss my interest in the ‘embeddedness’ of device use in everyday life. By this I mean the way in which devices are in state of ‘incipient use’ where a individual can start using a device almost at any time without the need to account for this to others. The question of when and where a device is used - before and after, as well as the question of the impact of usage on non-usage time are interesting directions for research. In my own work I have been drawing on video recordings of the use of smartwatches and mobile phones to try and understand this ‘embedded’ relationship - in particular how notifications work to incite usage from a situation of non use, and in contrast how usage can be ‘occasioned’ by the situation - such as waiting for a bus can occasion the use of a device.

References

- 1 Stuart Reeves and Barry Brown (2016) Embeddedness and sequentiality in social media, Proceedings of CSCW, ACM Press

3.7 Smart Notification Management for Personal Assistants

Benjamin Poppinga (*AUDI AG, DE, bp@benjaminpoppinga.de*)

License © Creative Commons BY 3.0 Unported license
© Benjamin Poppinga

In my opinion, a smart user experience will be the key differentiator in a connected world. Nowadays, more and more smart devices enter the market. They are all connected, mutually notifying each other and the users about, e.g., air quality, strong winds, number of people in the house or whether the washing machine is finished. The increasing number of notifications

is leading to information overload. Users who need to deal with that many information simply cannot stay in the loop on everything, and start to miss information or perform worse on their actual tasks. As a result, users feel annoyed or disconnected, which results into a bad reputation for the related products and brands. I started to work on smart notification management for mobiles during my PhD. Now, I am working with Audi and have the vision that the users will have a smart and highly adaptive user experience within our vehicles. To achieve this, I am working on two large research and development areas: driver monitoring and intelligent personal assistants. Driver monitoring is important to understand how busy the driver is with the primary task, e.g., to safely drive the vehicle. Here, deep learning is a breakthrough technology, which allows to have a very detailed impression on what the driver does or which additional information might be helpful in the current situation. Intelligent personal assistants, like Apple's Siri, will be a key user touch point in many digital environments. They have the role to assist us and, thereby, take some burden off the users. In my opinion, an intelligent personal assistant will be the central entity for all highly connected and smart experiences. My impression is that research and development teams worldwide have identified smart and intelligent user experiences as one of the key working streams for the next decades. Consequently, I am not only looking forward to smart devices and smart eco systems. I look forward to a seamless, adaptive, and truly immersive user experience, which feels convenient in all aspects and supports all humans to achieve whatever are their goals or dreams.

3.8 Intelligent Messages

Christoph Gebhardt (ETH Zurich, CH, christoph.gebhardt@inf.ethz.ch)

License © Creative Commons BY 3.0 Unported license
© Christoph Gebhardt

Optimizing the User Experience of Mobile Instant Messengers Mobile instant messages are a huge contributor to the problem of constant interruptions through notifications as perceived by a large number of smartphone users today. To tackle this issue, we work on a system which predicts opportune moments to notify users on received messages in order to maximize receptivity. Receptivity anticipates users' subjective overall reaction to an interruption, encompassing users' interruptability (depending on context) and their experience of interruption (depending on the content and sender of a message) [1]. For that reason, we want to find a model which predicts whether a user wants to be notified about an incoming message with respect to a user's context as well as the sender and the content of the message. Within that task, we are focusing on three particular problems of intelligent notification systems: Due to habit [2] or social pressure [3] users react on messages even if send at interrupting moments, making it impossible to build models solely by observing user reaction times. Related work solves this problem by randomly polling users on the perceived receptivity of a message. This has the disadvantage that the same or similar instances could be requested to be labeled several times, causing an unnecessary burden on users. In contrast, we intend to employ Active Learning to minimize user polling. In this framework, a certain property of a model is exploited to only poll users for messages where the model's prediction is uncertain subject to that property. Research has proven that in the context of notification management personalized models perform better than general models [4]. One consequence of using these individual models is that each new user will experience a

“cold start”, meaning that because of lacking data no proper predictions can be given. The standard approach to overcome this problem is to personalize a general model: a new user starts with a model learned on all existing users which gets individualized online with data collected from that user. To accelerate personalized model learning, we intend to explore user similarities. One conceivable approach could be to find a vector space which describes user preferences, enabling the retrieval of similar users. Related work has shown that content has a significant effect on the receptivity of messages [1, 5]. Nevertheless, nowadays systems do not consider the content of a message when predicting a user’s receptivity. For this reason, we intend to apply semantic clustering on mobile instant messages and use these topic clusters to improve the prediction of receptivity models. Currently, we are collecting data in order to better understand the problem domain and be able to experiment with different models. Therefore, we implemented a smartphone app called MSGSTAZ which is available on Google Play. The application provides statistics and other features centered on mobile messaging. In the background, the app collects sensor data as well as sender and message content for each incoming mobile instant message. In addition, we sample users’ receptivity for random messages.

References

- 1 Fischer et al. (2010). Effects of content and time of delivery on receptivity to mobile interruptions. In Proc. of MobileHCI '10.
- 2 Kushlev, K. et al. (2016). “Silence Your Phones”: Smartphone Notifications Increase Inattention and Hyperactivity Symptoms. In Proc. of CHI '16.
- 3 Pielot, M. et al. (2014). Didn’t you see my message? Predicting Attentiveness to Mobile Instant Messages. In Proc. of CHI '14.
- 4 Pejovic, V., & Musolesi, M. (2014). InterruptMe: Designing Intelligent Prompting Mechanisms for Pervasive Applications. In Proc. of UBICOMP'14.
- 5 Mehrotra, A. et al. (2015). Designing Content-driven Intelligent Notification Mechanisms for Mobile Applications. In Proc. of UbiComp '15.

3.9 Unobtrusive Interaction for Wearable Computing

David Dobbstein (Ulm University, DE, david.dobbstein@uni-ulm.de)

License © Creative Commons BY 3.0 Unported license
© David Dobbstein

Head-worn displays (such as Google Glass) allow for a quick access time to information and by that can serve to augment the user’s memory.

Interaction with such a display, however, is yet a problem. While the input space is limited, the interaction itself can have strong social implications in public spaces. Missing social conventions and technology apprehension can affect the users’ willingness to perform interaction in public. One possible way to address this, is to reach for preferably unobtrusive interactions that do not cast a lot of attention upon the user. Social acceptance may then rise due to continued exposure with the technology.

In our work, we focus on wearable touch input utilizing the user’s clothes. In user studies on where participants would potentially accept wearable touch input, the spatial area at the upper thighs is often mentioned, because it allows for a quick access time from the resting position of a user’s hand with little movement involved. We introduced a touch-enabled belt, allowing the user to rest their hand within the pocket while reaching for the belt with their

thumb for touch interaction. In a user study in public we found that users feel comfortable to interact as long as the interaction is short and thus isn't perceived as interaction but rather random hand movement. For longer interaction, they preferred the area above the trouser pockets for the close proximity to the resting hand. In following work, we integrated touch functionality into the pocket fabric itself. This allows to reach the interface from inside and outside the pocket, which is utilized for dual-sided touch interaction. Most wearable touch interaction systems provide only very limited basic touch gestures such as dimensional swiping, due to a lack of available input space and difficulties providing hand stabilization in mobile scenarios. With our dual-sided touch interaction at the pocket, the hand can be stabilized into position enabling a cursor-based control with a high input expressiveness. In a target selection study, we found that dual-sided interaction utilizing both sided by sliding the thumb into the pocket and using the index finger to tap for selection from outside is significantly faster than familiar single-sided touch interaction using only the index finger to control a cursor (as familiar from laptop computers). This effect was largest when walking.

The dual-sided interaction furthermore allows to use the thumb as a spatial point of reference for the remaining hand which can be used for jointly performed gestures.

3.10 Augmenting Social Interactions: Realtime Behavioural Feedback using Social Signal Processing Techniques

Elisabeth André (University of Augsburg, DE, andre@hcm-lab.de)

License © Creative Commons BY 3.0 Unported license
© Elisabeth André

In our research, we focus on a specific form of ambient notifications that augment social interactions by combining ideas from Augmented Reality with recent advances in Social Signal Processing. The resulting concept called Social Augmentation offers great promise for social skill training. It may be in particular beneficial to people experiencing social phobia or stage fright. By exposing people in-situ to selected stimuli, they may learn cope with negative emotions associated with socially-challenging situations, such as job interviews or group discussions. In [2], we illustrate the idea of Social Augmentation using public speeches as an example. While people engage with a staged or real audience, social cues, such as body posture and movement energy, are analyzed to provide immediate feedback on their behavior and thus help them adapt (see Fig. 1). A particular challenge is to provide feedback in a way that does not interfere with the main task, namely the social interaction with the interlocutor(s). To this end, we have analyzed and evaluated various options, such as visual feedback using head-mounted displays, haptic feedback using a vibrating arm band [1] and acoustic feedback. Social augmentation may also be of benefit to people being able to perceive their social context to a limited extent only due to physical handicaps. In [3], we show how blind and visually impaired users may profit from Social Augmentation by receiving information on their social context via sonification techniques (see Fig. 2).

References

- 1 Damian, I., Elisabeth André, E.: Exploring the Potential of Realtime Haptic Feedback during Social Interactions. Proc. of the Tenth International Conference on Tangible, Embedded, and Embodied Interaction, TEI 2016: 410-416, Eindhoven, The Netherlands, February 14-17, 2016. ACM.



■ **Figure 1** A user wearing an HMD while giving a lecture (left). Using various sensors and social signal processing techniques, the user receives real-time feedback on his behavior superimposed on his field of view (right)



■ **Figure 2** A blind user wearing eye tracking glasses (left). Sonification techniques are used to communicate the facial expression of the person the blind user is directing her attention to (right).

- 2 Damian, I., Tan, C.S.S., Tobias Baur, T., Schöning, J., Luyten, K., André, E.: Augmenting Social Interactions: Realtime Behavioural Feedback using Social Signal Processing Techniques. CHI 2015: 565-574.
- 3 Dietz, M., Maha El Garf, M.E., Ionut Damian, I. Elisabeth André, E.: Exploring Eye-Tracking-Driven Sonification for the Visually Impaired. Proc. of the 7th Augmented Human International Conference, AH 2016:5:1-5:8, ACM, 2016.

3.11 Ambient Episodic Memory Augmentation with Context

Evangelos Niforatos (USI Lugano, CH, evangelos.niforatos@usi.ch)

License © Creative Commons BY 3.0 Unported license
© Evangelos Niforatos

Today's technology is radically altering the nature and scale of information that we can draw on for remembering. Digital photo albums, progress tracking applications, and e-mail archives are only a few examples of technological artefacts that improve our "episodic memory" – our ability to remember past experiences. However, despite this technological proliferation, people still struggle to remember past experiences. In fact, a surprisingly high amount of failing memories concerns one's daily encounters. This might be due to the today's hectic lifestyle that divides one's daily routine in numerous small segments, while often blurring the border between work and personal life. Inevitably, attention, which is important for

creating high quality memories, spans over a large number of often simultaneous tasks. The episodic memory holds contextual information regarding who, what, where, and when of past experiences [3] – a summary of records of our life [1]. Once an episodic memory is formed, its information is said to be "differentially accessible", i.e., different sensory stimuli have different potential in triggering us recalling a particular personal past event. This is called "episodic activation" and describes the idea that in any episodic memory there is a pattern of activation that determines how its details are accessed [2]. Information stored in the episodic memory (i.e., episodic memories) is dominated by visual imagery that can have a "field" or an "observer" perspective, and are usually recalled in a temporal order. For successfully recalling an episodic memory, our memory system utilizes memory traces and encoded contextual information that arrives as input through our five senses. This contextual information is the basis of what is known in Psychology as memory cue, that when replayed help one remember about a past experience, a practise known as "cued recall". Memory cues can come in a wide variety of forms, as long as they help trigger the memory of a certain past event. Memory cues are per definition never "complete" – they do not comprise the entire experience but merely a related context, e.g., a visual capture (picture), a sound (a song), or a smell. A memory cue can also be textual or an abstract visual, hinting at one's activity, location, or social interaction at a certain time (e.g., device and application usage, GPS logs, Facebook posts) – a "personal context" of an episodic memory. Modern technology and trends such as lifelogging and the quantified-self movement, allow us to capture such personal contexts unobtrusively. Personal mobile devices such as laptops, smartphones, digital and wearable cameras (e.g., Narrative Clip), wristbands (e.g., Fitbit) or smart watches can all record – both implicitly and explicitly – a large range of contextual information about one's daily routine. The ambient presentation of memory cues throughout one's daily life hold the potential to significantly improve one's ability to recall past experiences, while sharing many common aspects with ambient notifications. For example, memory cues – same as notifications – can be delivered via a plethora of daily encountered devices such as smartphones, displays, wearables that utilize visual, audible, tactile or even olfactory modalities, while raising challenges about privacy and one's attention and cognitive reserves in general.

References

- 1 Martin A. Conway. 2009. Episodic memories. *Neuropsychologia* 47, 11: 2305–2313. <https://doi.org/10.1016/j.neuropsychologia.2009.02.003>
- 2 Mihaly Racsmany and Martin A. Conway. 2006. Episodic Inhibition. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 32, 1: 44–57. <https://doi.org/10.1037/0278-7393.32.1.44>
- 3 Endel Tulving. 1984. Precipitous elements of episodic memory. *Behavioral and Brain Sciences* 7, 2: 223–68.

3.12 Ambient Notifications Environments as a Support for Daily Activities

Frederik Wiehr (DFKI Saarbrücken, DE, frewie@gmail.com)

License  Creative Commons BY 3.0 Unported license
© Frederik Wiehr

Currently, I am a PhD student in computer science working at the German Research Center for Artificial Intelligence, Saarbrücken in the field of human computer interaction, human factors, activity recognition, and wearables. In my research I explore how to build technology that assists in overcoming physiological and cognitive restrictions. I do so by trying to extract higher-level information from sensors in the environment and on body worn devices that helps to understand the users' current activities and ongoing mental processes. I am further investigating systems that take such information into account and enable individual meaningful assistance enhancing learning and improving social interaction.

More specifically, a current focus of my research are ambient notifications environments that can support people in their daily activities, providing courses of action and relevant information in context. A key aspect of ambient notifications is the unobtrusive provision of information, embedded into the users' environment. The number of services which are providing notifications on numerous devices, both mobile and embedded into the environment, is increasing. Therefore, is important to find unobtrusive ways for notification, avoiding redundancy, and the appropriate detail of information and alert level.

Another aspect of ambient notification environments is its spatial aspect. Kevin A. Lynch describes the 'need to reorganize and pattern our surroundings' driven by the fear of disorientation, or, positively speaking, by the wish for emotional security, which can arise from a harmonious relationship between the self and the outside world [1]. The intrinsic need for emotional security or respectively the fear to get lost stays in contrast with the design of current assistive systems. These momentarily ease our life but holds us back from learning and may also have an effect on general planning and problem-solving capabilities over longer terms. Spatial knowledge acquisition generally enables us to plan and take actions. The spatial relationships of the apprehended environment, as for example patterns, shape, connection, and separation, remain unchanged regardless of the scale at which it is viewed. Montello introduced the classification of psychological spaces in figural, vista, environmental and geographical space [2]. Effectively, the projective size of the spaces in relation the size of the human body serves as a basis for this classification. In terms of ambient notification environments, I investigate modalities and interaction designs that foster spatial skills and competencies while still maintaining an appropriate level of support.

References

- 1 Kevin Lynch. The image of the city, volume 11. MIT press, 1960.
- 2 Montello, D. R. (1993). Scale and multiple psychologies of space. In Spatial information theory a theoretical basis for gis (pp. 312–321). Springer.

3.13 Blended Interaction and Ambient Notification Environments

Harald Reiterer (University of Konstanz, DE, Harald.Reiterer@uni-konstanz.de)

License  Creative Commons BY 3.0 Unported license
© Harald Reiterer

The goal of my research vision Blended Interaction[3] is to enable a user-centered design of natural interactions in interactive spaces. The framework tries to advance Mark Weiser’s vision [4] to create an “invisible” ubiquitous computing “that provides us with the great powers of digital computation in an unobtrusive manner, so that we are freed to use them without thinking and ‘mental gymnastics’ and to focus beyond computers on new goals”. For that purpose, it draws on the principles of Reality-based Interaction [2] and of Conceptual Blending [1]. Reality-based Interaction attempts to make human-computer interaction similar to the interaction with the real world. By drawing on humans’ pre-existing knowledge and skills, the mental effort required to operate a system is reduced and users are free to focus on the actual task without their cognitive flow being interrupted by cumbersome interactions. In contrast to Reality-based Interaction, Blended Interaction not only applies users’ natural skills and pre-existing knowledge of the real world but also considers digital well-established concepts in the design of new user interfaces. As humans spend more and more time in the digital world, we cannot consider human thinking free from digital influences anymore and need to take them into account when designing new interaction concepts. Conceptual Blending theoretically explains how human thinking subconsciously creates a new concept through projection from two existing input concepts. Therefore, human mind connects the two input concepts on the base of a generic space. The generic space contains basic level concepts which are common to both inputs (e.g. both inputs are containers). On base of these commonalities, human mind blends both input concepts in an output concept that has a new and emergent structure which is not available from the inputs alone. Blended Interaction uses this process of indirect projection to theoretically explain that user interfaces only need to share selected aspects of reality for users to be able to understand and operate a new interaction design. This enables us to use computational power to go beyond what is possible in the real world by keeping a natural and intuitive interaction. During the Dagstuhl Seminar I had the opportunity to discuss how our Blended Interaction Framework could be used to design Ambient Notification Environments. The following questions lead my discussions: What are right blends for this kind of environments? What would we call a natural notification? What are typical design tradeoffs that have to be solved? The very fruitful discussions in different break out groups have shown that Blending Interaction has the potential to support designers of ambient notification environments to find novel design solutions for ambient notifications.

References

- 1 Fauconnier, G., & Turner, M. (2002). *The Way We Think: Conceptual Blending and the Mind’s Hidden Complexities*. New York: Basic Book.
- 2 Jacob, R., & Girouard, A. (2008). Reality-based interaction: a framework for post-WIMP interfaces. In *Proc. of CHI ’08* (pp. 201–210). New York: ACM Press
- 3 Jetter, H., Reiterer, H., & Geyer, F. (2014). Blended Interaction: Understanding Natural Human- Computer Interaction in Post-WIMP Interactive Spaces. *Theme Issue on Designing Collaborative Interactive Spaces, Personal and Ubiquitous Computing*, 18(5), 1139–1158
- 4 Weiser, M. (1991). The computer for the 21st century. *Scientific American*, 265(3), 94–104.

3.14 Ubiquitous Sensing

Joseph Paradiso(MIT – Cambridge, US, joep@media.mit.edu)

License  Creative Commons BY 3.0 Unported license
© Joseph Paradiso

My research has in recent years focussed on ways in which ubiquitous sensing and actuation can connect to people in different ways. At Dagstuhl this year, I presented ways in which this melds onto what can be considered ‘notification’ - in some sense, ambient notification is what we see happening now, but this will eventually be considered to be augmented perception once users gracefully graft onto ubiquitous sensing now deploying under IoT. My team and I have outlined this via high-level review articles [1, 2] outlining the future of perception in the world of dense sensing. More detail has been presented on research projects that focus onto pieces of this - for example, our ‘Tidmarsh’ project [3], where we are building what we call an ‘attention-driven sensory prosthetic’, which funnels information into user perception that is derived from elements of the users’ environment that he or she is focussing on. We have also designed ambient displays [4] that track user location, state, and gesture, providing relevant information accordingly in a fashion relating to ambient notification. My team has also recently explored personal lighting systems as a way to signal others [5], as well as rooms with ambient projection and lighting that can change with user state, context, and potential ‘notification’. There are many other projects running in my group that overlap with the notification theme in interesting and productive ways, hence I was delighted to take part this year.

References

- 1 Dublon, G., Paradiso, J. A., “How a Sensor-Filled World Will Change Human Consciousness,” *Scientific American*, July 2014, pp. 36-41.
- 2 Paradiso, J., “Our Extended Sensoria - How Humans Will Connect with The Internet of Things,” in *The Next Step: Exponential Life*, Open Mind Collection, BBVA Press, 2016, pp. 47-75.
- 3 Mayton, B., et al, ‘The Networked Sensory Landscape: Capturing and Experiencing Ecological Change Across Scales,’ to appear in *Presence*, MIT Press Journal, Special issue on Arts, Aesthetics, and Performance in VR and Telepresence, 2018.
- 4 Gillian, N., Pfenninger, S., Russell, S., and Paradiso, J.A., “Gestures Everywhere: A Multimodal Sensor Fusion and Analysis Framework for Pervasive Displays,” In *Proceedings of The International Symposium on Pervasive Displays (PerDis '14)*, Sven Gehring (Ed.). ACM, New York, NY, June 2014, pp. 98-103.
- 5 Zhao, N. and Paradiso, J.A., “HALO: wearable lighting,” in *Adjunct Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2015 ACM International Symposium on Wearable Computers*, pp. 601-606, Sept. 7-11, 2015.
- 6 Zhao, N., Azaria, A., Paradiso, J.A., “Mediated Atmospheres: a Multimodal Mediated Work Environment,” To appear in *Proc. UbiComp 2017*.

3.15 Lifelong Learning

Judy Kay (The University of Sydney, AU, judy.kay@sydney.edu.au)

License  Creative Commons BY 3.0 Unported license
© Judy Kay

My research aims to support lifelong and lifewide learning, by harnessing people’s digital footprints, the data collected by diverse sensors on devices [1]. There are many ways that this links to ambient notification environments both in formal learning settings and more broadly.

Broadly, my work aims to transform raw sensor data into a user controlled, scrutable user model. The data ranges from activity tracks, to personal loggers, to interfaces for group work. The user model provides an interpretation of the user available data to represent aspects that are important for a person’s goals for learning. Then interfaces, often called Open Learner Models, OLMs, [2] make this information available in useful forms, to support long term metacognitive processes, of self-reflection, self-monitoring and planning. There are two important roles for notifications associated with OLMs: prompts to remind the user to attend to the interface at appropriate times; and scaffolding to help people pay attention to the relevant aspects. An alternative approach we have been exploring is to create ambient interfaces that are carefully situated at the precise places that the user should consider their OLM.

For example, in our tabletop classroom, we created an environment where students worked in 5 groups, each with 4-6 students, at an enhanced tabletop system that tracked each student’s work on the class activity [3]. Then, harnessing the data collected by the tabletops, we created a teacher’s dashboard that provided real-time, continuously updated information about the progress of each group and each student in the group. In this context, notifications proved valuable for ensuring teachers gave more attention to the groups that most needed intervention [4]. So this is one example of an environment where active notifications seems valuable for enabling a teacher to manage complexity.

At a quite different level, we have explored how to help people with challenging behaviour change, in trying to meet the recommended levels of vegetable intake. So we maintain daily OLM of the number of serves consumed. One approach we have explored takes over the lock screen so that the 30-80 times a day many people unlock their mobile phone, they see a “subtle notification” that they intended to eat vegetables and they can easily see their progress on a gamified, pleasing display and very easily log each serve of vegetables.

Continuing this theme, we have been exploring how a SAL logger (Simple, Situated Ambient Logger) can help people use a simple OLM plus logging interface in diverse ways. For example, we have studied people configure their SAL to support behaviour change and maintenance by tracking intake of each food group [5] and to build new tiny habits, such a target number of pauses from desktop computer use for relaxed breathing, eye exercises or drinking water. We have also experimented with deployed use of SAL loggers for medication taking.

References

- 1 Kay, J., & Kummerfeld, B. (2012). Creating personalized systems that people can scrutinize and control: Drivers, principles and experience. *ACM Trans on Interactive Intelligent Systems*, 2(4), 1–42.
- 2 Bull, S., & Kay, J. (2016). SMILI☺: a framework for interfaces to learning data in open learner models, learning analytics and related fields. *International Journal of Artificial Intelligence in Education*, 26(1), 293-331.

- 3 Martinez-Maldonado, R., Dimitriadis, Y., Martinez-Monés, A., Kay, J., & Yacef, K. (2013). Capturing and analyzing verbal and physical collaborative learning interactions at an enriched interactive tabletop. *International Journal of Computer-Supported Collaborative Learning*, 8(4), 455–485.
- 4 Martinez-Maldonado, R., Clayphan, A., Yacef, K., & Kay, J. (2015). MTFeedback: providing notifications to enhance teacher awareness of small group work in the classroom. *IEEE Transactions on Learning Technologies*, 8(2), 187-200.
- 5 Yekeh, F., Kay, J., Kummerfeld, B., Tang, L. M., & Allman-Farinelli, M. A. (2015, December). Can SAL support self reflection for health and nutrition?. In *Proceedings of the Annual Meeting of the Australian Special Interest Group for Computer Human Interaction* (pp. 134-141). ACM.

3.16 Exploring the role of notifications relating to Group QoS and proactive behaviour in IoT environments.

Keith Cheverst (Lancaster University, GB, k.cheverst@lancaster.ac.uk)

License © Creative Commons BY 3.0 Unported license
© Keith Cheverst

My research has considered notification issues across a range of areas including notifications to support users of distributed mobile groupware and notifications relating to proactive behaviour in an IoT environment.

In relation to distributed mobile groupware I was interested in the investigating the usefulness (and the required underlying middleware support) for enabling the specification of Group QoS and the use of notifications signifying Group QoS violations. For example, one might specify a Group QoS relating to the timely delivery of a message to all members of a given group. Assuming that one or more of the group members failed to receive the message in a timely manner, (e.g. because of a network disconnection) then the message sender would receive a visual notification that the specified Group QoS had been violated enabling her to act accordingly. The requirement for supporting Group QoS came from our work with a regional electricity board carried out under the EPSRC funded MOST project. A key requirement was for field engineers to be provided with appropriate ‘Mobile Awareness’ [1] when making changes to the switched power network in a distributed fashion rather than the more traditional centralized approach. The distributed approach required engineers to communicate updates to the state of the power network to the group of other field engineers who might be making their own changes/switches concurrently. Given the safety critical nature of the domain (e.g. ensuring that all engineers held an up-to-date view on which parts of the power network were currently switched to ‘live’) the use of appropriate notifications was critical in ensuring that engineers were made aware of when other field engineers were not receiving updates to the state of the power network in a timely manner, e.g. because of problems of poor wireless connectivity.

The potential role of notifications associated with IoT was studied as part of our research that explored the user experience issues associated with a system known as the Intelligent Office [2]. This system supported the proactive behavior of devices (e.g. the office fan) based on rules inferred from user behavior (e.g. under what contextual conditions the user would turn on their office fan to cool reduce the temperature rather than opening a window). In the developed system a notification would appear prior to the system actuating the given device in order to enable the user to either reject the suggested actuation, view the ‘if-then’ rule

used to trigger the actuation, or even to inspect the learn behavior and associated context history used to infer the rule (and so support the user in scrutinizing the user modeling aspect of the system).

References

- 1 Cheverst, K., Blair, G.S., Davies, N. and Friday, A. The Support of Mobile-awareness in Collaborative Groupware, in ACM Personal Technologies Journal, Vol. 3, No. 1&2, pp 33-42. Springer-Verlag, 1999.
- 2 Cheverst, K., Byun, H.E., Fitton, D., Sas, C., Kray C. and Villar, N. Exploring Issues of User Model Transparency and Proactive Behaviour in an Office Environment Control System, in Special Issue of UMUAI (User Modelling and User-Adapted Interaction) on User Modeling in Ubiquitous Computing, Guest Editors: Anthony Jameson, Antonio Krüger, Kluwer, Vol. 15, No. 3-4, pp.235-273, August 2005.

3.17 Updating without interrupting

*Lewis Chuang (Max Planck Institute for Biological Cybernetics, DE,
lewis.chuang@tuebingen.mpg.de)*

License  Creative Commons BY 3.0 Unported license
© Lewis Chuang

Ambient notifications are designed to update the receiver without interruption. Can this work? This depends on the perceptual or cognitive resources that a notification would demand of the receiver in order to be effectively noticed. After all, a notification is not a notification if it is not noticed. Nonetheless, a receiver can certainly be updated without being behaviourally disrupted. My research has consistently demonstrated that task-irrelevant environmental sounds can elicit brain responses that underlie working memory updating without compromising steering itself [1]. This response weakens as the steering task becomes more difficult. In other words, it is noticed but less so when the receiver is occupied. It requires resources but it does not demand them. Ambient notifications are designed to operate in a similar way. In this seminar, we heard of light environments that change hue according to the user activity [2], of worn clothing that constricted instead vibrating upon receiving notifications [3], of wall panels that created new environments in response to the user's brain state, and more. The brain responds to these changes in the environment (hopefully in positive and the intended way) without stealing from the current task.

References

- 1 Scheer M, Bühlhoff HH and Chuang LL (2016) Steering demands diminish the early-P3, late-P3 and RON components of the event-related potential of task-irrelevant environmental sounds, *Frontiers in Human Neuroscience*, 10(73), 1-15.
- 2 Zhao N., Aldrich M., Reinhart C., Paradiso J., (2015) A Multidimensional Continuous Contextual Lighting Control System Using Google Glass, BuildSys'15, Seoul, South Korea.
- 3 Pohl H., Brandes P., Quang H. N., Rohs M. (2017), Squeezeback: Pneumatic Compression for Notifications, Proceedings of the SIGCHI Conference on Human Factors in Computing Systems

3.18 Ambient Notifications for Dementia Care

Markus Löchtefeld (Aalborg University, DK, mloc@create.aau.dk)

License  Creative Commons BY 3.0 Unported license
© Markus Löchtefeld

Age-related health issues, such as dementia, will become a major cost factor in Denmark as well as many other western countries. The Danish National Research Centre for Dementia (2015) estimates that the number of people with dementia in Denmark today is 84.000 and will increase to 151.000 in 2040. People with dementia have a tendency to wander which includes going to places they should not, night-time wandering, trailing others and attempts to leave the house or care facility [4]. Especially the latter is a big concern for relatives and caregivers. As people with dementia are often unable to take care of themselves there is a serious risk that their straying can lead to dehydration, hypothermia and in the worst-case fatality [1]. The tendency to wander is often paired with previous life routines, e.g., going to work each day. These types of wandering behaviour occur for different reasons and is not always obvious why. Cognitive impairment is often a causative factor for wandering and a correlation between the severity of the impairment and the behaviour has been shown [1]. Other factors such as stress, boredom and environmental influence also affect when, and how often, the wandering behaviour is triggered (e.g., too busy, too much noise) [1, 2]. We want to increase security and reduce the concerns of residents, caregivers and relatives without limiting residents' sense of autonomy, and their opportunities for daily activities. Therefore, electronic tracking and tagging solutions have been suggested and implemented to ensure the safety of the residents. Tracking devices are often used to preserve the independence of a person with dementia and to provide the caregivers with reassurance. One of the most important parts to ensure the effectiveness of such a system is the user interface. It is important that the caregiver receives the right information at the right time in the right manner so the caregivers can take the necessary actions e.g. to avoid that a resident leaves the facility. Especially, as we want to avoid additional temporal or mental demand for the caregivers the user interface needs to blend into their workflow. We will develop novel notification techniques that are embedded into the environment and its architecture. For example, instead of just using mobile devices that the care giver carries around or public displays in the facility, different elements in the environment such as color of the lighting could notify the caregivers of potential danger in an ambient manner.

References

- 1 Algase, D. L. (1999). "Wandering: A dementia-compromised behavior. *Journal of Gerontological Nursing*", 25(9), Sep, 10-16.
- 2 Cohen-Mansfield, J., & Werner, P. (1995). Environmental influences on agitation: An integrative summary of an observational study. *American Journal of Alzheimer's Disease and Other Dementias*, 10(1), 32-39.
- 3 Danish National Research Center for Dementia (2015) (Nationalt Videnscenter for Demens), Forekomst af demens i Danmark, August 2015, <http://www.videnscenterfordemens.dk/statistik/forekomst-af-demensi-danmark>, url-date: 2016-08-08.
- 4 Hope, T., Keene, J., McShane, R. H., Fairburn, C. G., Gedling, K., & Jacoby, R. (2001). "Wandering in dementia: a longitudinal study". *International Psychogeriatrics*, 13(02), Jun, 137-147.

3.19 Notifications through Indirect Illumination and Pneumatic Compression

Michael Rohs (University of Hannover, DE, michael.rohs@hci.uni-hannover.de)

License  Creative Commons BY 3.0 Unported license
© Michael Rohs

We investigated two different notification approaches for wearable devices that aim at less intrusive notification experiences. The first is based on visible light that is scattered through the skin. The second is based on pneumatic compression feedback. This work was done in the HCI group of the University of Hannover and is mostly by Henning Pohl, who just completed his PhD thesis and is now with the University of Copenhagen.

The first approach to mobile notifications uses the underside of a smartwatch for indirect light feedback. A main function of smartwatches is to notify the user. Typically notifications that arrive on the mobile phone are dispatched to the smartwatch and played there as vibration feedback. However, vibration feedback can quickly become annoying as it is quite attention grabbing. This is the reason for some wearers of smartwatches to disable this functionality. We propose an alternative in which eight LEDs on the backside of the watch illuminate the skin surrounding the watch. Human skin transmits light. Only about 5% of the light is directly reflected from the skin's surface. The rest is scattered and subject to dermal and epidermal remittance as well as absorption. This effect becomes stronger near the infrared range. This is the reason we used red LEDs with a wavelength of 623 nm. For manufacturing reasons the SMD-LEDs are soldered on a PCB, embedded in a thin silicone layer and face down.

We performed an in-the-wild user study in which 13 participants wore our watch prototype with indirect light feedback for 24 hours on a weekday. They only took off the watch at night. The light stimuli were played at random times every couple of minutes. The study took part in the summer of 2015 with 6 hours of sunshine per day on average. Participants reacted to 80.3% of the light stimuli and took 16.6 seconds to react to the stimuli. The reaction time distribution has a long tail, so some of the notifications took much longer to be noticed. The participants generally liked the light feedback, but were concerned about others noticing the light feedback in public. Moreover we designed and tested different light patterns, like rotations and oscillations. Participants were generally able to distinguish most of the light patterns. There were no large differences in terms of reaction rate and reaction time.

The second approach to mobile notifications uses pneumatic compression feedback. A repurposed blood pressure meter is used at very low pressure levels to create uniform compression around a limb without generating shear forces. A unique characteristic of this approach is the wide range of haptic feedback that can be generated, from very subtle to very strong, so strong that it inhibits movement. Low pressure levels can be used for continuous background feedback. We conducted an absolute detection threshold experiment and found threshold levels of 0.7 kPa in the lab and 2.3 kPa in the wild. In comparison, blood pressure cuffs are inflated to 16.0 kPa (120 mmHg) for systolic and 10.7 kPa (80 mmHg) for diastolic blood pressure, i.e., 15x-23x the lab detection threshold and 5x-7x times the in-the-wild detection threshold. We did a 1-hour experiment with 9 participants, in which they wore a cuff inflated to 1.3 kPa for during deskwork. They did not report any feeling of annoyance or inhibition.

References

- 1 Pohl, Medrek, Rohs. ScatterWatch: Subtle Notifications via Indirect Illumination Scattered in the Skin. MobileHCI 2016.
- 2 Pohl, Brandes, Quang, Rohs. Squeezeback: Pneumatic Compression for Notifications. CHI 2017.

3.20 Towards Holistic Notification Management

Niels Henze (*University of Stuttgart, GER, niels.henze@vis.uni-stuttgart.de*)

License  Creative Commons BY 3.0 Unported license
© Niels Henze

We currently witness a steadily increasing number of notifications that are delivered to users through a similarly increasing number of notification channels. Today, users can be notified through their smartphones, tablets, PCs, and smart watches. In the near future, Internet-of-Things devices will not only generate additional notifications but will also be able to notify users themselves. In our work, we investigate how users interact with notifications today [1], study users preferences when they have multiple devices that can notify them [2], and explore notifications in smart home environments [3]. As we assume that current approaches for notification management do not scale, we envision a holistic notification system that manages incoming notifications across device boundaries.

References

- 1 Sahami Shirazi, A., Henze, N., Dingler, T., Pielot, M., Weber, D., & Schmidt, A. (2014, April). Large-scale assessment of mobile notifications. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (pp. 3055-3064). ACM.
- 2 Weber, D., Voit, A., Kratzer, P., & Henze, N. (2016, September). In-situ investigation of notifications in multi-device environments. In Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing (pp. 1259-1264). ACM.
- 3 Voit, A., Machulla, T., Weber, D., Schwind, V., Schneegass, S., & Henze, N. (2016, September). Exploring notifications in smart home environments. In Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct (pp. 942-947). ACM.

3.21 Notifications in Mixed Reality

Pascal Knierim (*University of Stuttgart, DE, pascal.knierim@vis.uni-stuttgart.de*)

License  Creative Commons BY 3.0 Unported license
© Pascal Knierim

We already face a vast amount of notifications delivered to our smartphones and tablets, and nowadays to smart TVs and watches, too. This trend is further fueled by more and more Internet of things (IoT) devices. Further, due to the advances in the field of mixed reality, notification output systems will move closer to our eyes to catch our visual attention.

Head mounted displays such as the Google Glass already offer limited input and output space to display notifications and are always in the field of view of the user. New augmented reality (AR) glasses allow displaying high-resolution three-dimensional content located in the real world. There is a potential need for a notification system which is aware of the

user’s current context and location. Instead of using levitating physical tokens [2] to visualize digital information, AR-glasses in combination with context aware notification systems could reduce the notification overload we currently face [3]. Several virtual reality headsets are now available for end users. Users immersed in a virtual world are not penetrated by notifications since they are currently not forwarded into the virtual world. However, we assume that systems will support notifications in VR shortly. Possible concepts to notify the user include well-known visual and auditory, but also haptic notifications [1] methods are imaginable. In all instances of the mixed reality spectrum, we need to investigate new techniques to manage the increasing amount of notifications to unburden the user.

References

- 1 Pascal Knierim, Thomas Kosch, Valentin Schwind, Markus Funk, Francisco Kiss, Stefan Schneegass, and Niels Henze. 2017. Tactile Drones - Providing Immersive Tactile Feedback in Virtual Reality through Quadcopters. In Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '17). ACM, New York, NY, USA, 433-436. DOI: <https://doi.org/10.1145/3027063.3050426>
- 2 Florian Schaub, and Pascal Knierim. Drone-based Privacy Interfaces: Opportunities and Challenges. Twelfth Symposium on Usable Privacy and Security (SOUPS 2016). USENIX Association, 2016.
- 3 Dominik Weber, Alexandra Voit, Philipp Kratzer, and Niels Henze. 2016. In-situ investigation of notifications in multi-device environments. In Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '16). ACM, New York, NY, USA, 1259-1264. DOI: <http://dx.doi.org/10.1145/2971648.2971732>

3.22 Towards Visual Notifications in Multi-User Large-Display-Environments

Raimund Dachsel (TU Dresden, DE, raimund.dachsel@tu-dresden.de)

License  Creative Commons BY 3.0 Unported license
© Raimund Dachsel

Towards Visual Notifications in Multi-User Large-Display-Environments In our research at the Interactive Media Lab Dresden we are concerned in particular with using several interaction modalities – such as multitouch, pen, gaze, gestural, and tangible input – for the interaction of one or multiple people in environments with interactive surfaces. Typically, input modalities can be synergistically combined, like gaze input with multitouch input [6], or offered in parallel for a seamless transition, like body movements, hand gestures, and touch input in BodyLenses [2]. In addition, multiple interactive surfaces of different sizes can be combined to provide effective multi-display environments (MDE), like handhelds with large wall displays [8, 4]. Naturally, such environments are well suited for multi-user interaction and collaboration.

Our research in multi-user multi-display environments illustrates that notifications are no longer limited to personal settings, but need to be managed for multiple users simultaneously. What is the best notification place and technology? Where can people be notified during collaborative work in MDE? In our Sled approach [8] we for example use personal body-worn devices to provide visual feedback and notifications to users at a large display wall. With BodyLenses [2] we contributed personal territories for individual users at interactive wall-sized displays, which, like a personal lens, can be moved around by just moving the body

or interacted with by means of hand gestures and touch input. Another related question is who is actually interacting in an MDE and how to assign both interactions and notifications to individual users. With YouTouch! [7] we developed an approach using depth-cameras and color histograms to identify users even under difficult conditions. Knowledge about the user's current attention using eye gaze [6] can also help supporting the display of appropriate peripheral notifications depending on the user's point of regard.

Another research avenue of the Interactive Media Lab is interactive information visualization using natural user interfaces. In our visualization research, we raise the question how visualization knowledge can help us in designing effective notifications and which modality in general is most appropriate for ambient notifications. The question where to present feedback and notifications has been addressed by our lab in several ways. Examples include tangible everyday objects [1], smart projections in everyday environments [5], and even normal paper being illuminated using electroluminescence in digital pen and paper applications [3].

For the future, we are not only interested in the "how" of ambient notifications – which often is the focus of current research, but increasingly in the question "how much" and the degree of user control. We call for more mindfulness in ambient notifications and a rigorous revealing of dark patterns.

References

- 1 Wolfgang Büschel, Ulrike Kister, Mathias Frisch, and Raimund Dachsel. 2014. T4 - transparent and translucent tangibles on tablespots. In Proceedings of the 2014 International Working Conference on Advanced Visual Interfaces (AVI '14). ACM, New York, NY, USA, 81-88. DOI: <https://doi.org/10.1145/2598153.2598179>
- 2 Ulrike Kister, Patrick Reipschläger, Fabrice Matulic, and Raimund Dachsel. 2015. BodyLenses: Embodied Magic Lenses and Personal Territories for Wall Displays. In Proceedings of the 2015 International Conference on Interactive Tabletops & Surfaces (ITS '15). ACM, New York, NY, USA, 117-126. DOI: <https://doi.org/10.1145/2817721.2817726>
- 3 Konstantin Klamka and Raimund Dachsel. 2017. IllumiPaper: Illuminated Interactive Paper. In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17). ACM, New York, NY, USA, 5605-5618. DOI: <https://doi.org/10.1145/3025453.3025525>
- 4 Ricardo Langner, Ulrich von Zadow, Tom Horak, Annett Mitschick, Raimund Dachsel. Content Sharing Between Spatially-Aware Mobile Phones and Large Vertical Displays Supporting Collaborative Work. In "Collaboration Meets Interactive Spaces". Springer International Publishing, 75-96, 2016.
- 5 Fabrice Matulic, Wolfgang Büschel, Michael Ying Yang, Stephan Ihrke, Anmol Ramraika, Carsten Rother, and Raimund Dachsel. 2016. Smart Ubiquitous Projection: Discovering Surfaces for the Projection of Adaptive Content. In Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '16). ACM, New York, NY, USA, 2592-2600. DOI: <https://doi.org/10.1145/2851581.2892545>
- 6 Sophie Stellmach and Raimund Dachsel. 2013. Still looking: investigating seamless gaze-supported selection, positioning, and manipulation of distant targets. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13). ACM, New York, NY, USA, 285-294. DOI: <https://doi.org/10.1145/2470654.2470695>
- 7 Ulrich von Zadow, Patrick Reipschläger, Daniel Bösel, Anita Sellent, and Raimund Dachsel. 2016. YouTouch! Low-Cost User Identification at an Interactive Display Wall. In Proceedings of the International Working Conference on Advanced Visual Interfaces (AVI '16), Paolo Buono, Rosa Lanzilotti, and Maristella Matera (Eds.). ACM, New York, NY, USA, 144-151. DOI: <https://doi.org/10.1145/2909132.2909258>

- 8 Ulrich von Zadow, Wolfgang Büschel, Ricardo Langner, and Raimund Dachsel. 2014. SleeD: Using a Sleeve Display to Interact with Touch-sensitive Display Walls. In Proceedings of the Ninth ACM International Conference on Interactive Tabletops and Surfaces (ITS '14). ACM, New York, NY, USA, 129-138. DOI: <http://dx.doi.org/10.1145/2669485.2669507>

3.23 Multimodal Notifications

Steve Brewster (University of Glasgow, GB, Stephen.Brewster@glasgow.ac.uk)

License  Creative Commons BY 3.0 Unported license
© Steve Brewster

My research focuses on multimodal feedback and using the range of capabilities that humans have to enable rich interaction with technology. These different feedback modalities can be used to create ambient notifications for users in different settings.

We have developed notifications in several different modalities. For non-speech audio, Earcons can be used to create messages for ambient feedback [1]. By changing auditory parameters such as timbre, rhythm, pitch, spatial location or tempo, different messages can be built up. These can be composed by concatenation or by hierarchically by changing parameters in turn [2].

Tactons are the touch-based equivalent of Earcons. By manipulating parameters of tactual perception, different messages can be created for ambient notifications, for example, waveform, rhythm and body location [3]. Tactons have been used in many different applications including as messages for mobile phones and for in-car feedback. Other types of notifications we have tested have included smell-based [4], ambient visual and thermal notifications [5]. The general process in all of the different modalities is to use an applied psychophysical approach: we first look at the range of human perception to establish human limits and we then develop and test a range of cues within that to see what is perceivable and usable.

Two specific areas of application are in homecare setting and user interfaces for cars. Caring for people in their own homes is a growing area for research due to ageing populations and decreasing healthcare budgets. Due to the wide range of sensory loss that can occur in old age, different forms of feedback and interaction are needed for homecare systems. A user with poor eyesight may need a very different system to someone who has poor hearing. Multimodal notifications can play a significant role as the same messages can be configured to display in different modalities. This makes them accessible for a wide range of users [6].

The key issue for notifications in cars is that they should not distract the driver. As driving requires visual attention on the road ahead, multimodal notifications can be used to inform the driver of information without needing visual attention. We have tested a range of different modality types in the car, from ambient visual, audio, tactile and thermal, plus many different combinations. The results from the studies suggest that more modalities are more attention grabbing [7]. If you want to get a driver's attention for important information then use multiple modalities. If the message is of lower importance then a single modality may be more appropriate.

References

- 1 Brewster, S.A., Wright, P.C. & Edwards, A.D.N. (1992). A detailed investigation into the effectiveness of earcons. In G. Kramer (Ed.), Auditory display, sonification, audification and auditory interfaces. The Proceedings of the First International Conference on Auditory Display, Santa Fe Institute, Santa Fe, NM: Addison-Wesley, pp. 471-498

- 2 Brewster, S.A. (1998). Using non-speech sounds to provide navigation cues. *ACM Transactions on Computer-Human Interaction*, 5(2), pp 224-259.
- 3 Brewster, S.A. and Brown, L.M. Tactons: Structured Tactile Messages for Non-Visual Information Display. In *Proceedings of Australasian User Interface Conference 2004* (Dunedin, New Zealand), Australasian Computer Society, pp15-23.
- 4 Warnock, D., McGee-Lennon, M. and Brewster, S.: Multiple Notification Modalities and Older Users. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '13*, ACM Press, pp. 1091-1094, 2013.
- 5 Wilson, G., Brewster, S., Halvey, M. and Hughes, S.: Thermal Icons: Evaluating Structured Thermal Feedback for Mobile Interaction. In *Proceedings of the 14th International Conference on Human-Computer Interaction with Mobile Devices and Services - MobileHCI '12*, ACM Press, pp. 309-312, 2012.
- 6 Warnock, David and McGee-Lennon, Marilyn and Brewster, Stephen: Older users, multimodal reminders and assisted living technology. In *Health informatics journal* 18(3), pp. 181-90, 2012.
- 7 Politis, I., Brewster, S. and Pollick, F.: Evaluating Multimodal Driver Displays under Varying Situational Urgency. In *Proceedings of ACM CHI 2014*, ACM Press, pp. 4067-4076, 2014.

3.24 Guiding Attention in Cyber-Physical Environments

Susanne Boll (University of Oldenburg, DE, Susanne.Boll@informatik.uni-oldenburg.de)

License © Creative Commons BY 3.0 Unported license
© Susanne Boll

Cyber-physical systems technology blends human attention with intelligent digital systems. Even with greatly increased automation in the future, the human will need to be in the loop for monitoring. However, interaction with cyber-physical systems like cars, ships, robots, smart homes, or emergency rooms will change dramatically: (1) humans will interact much less frequently with the larger automated cyber-physical systems than with today's simpler automated systems; (2) human interaction will be needed to a much greater extent for tasks in which the human is superior to the machine and where automation finds its limitations; (3) humans will be free to dedicate more and more of their cognitive resources to other tasks, with their attention shifting only when needed to interact with the automated Cyber-Physical environments such as Vehicles [1], Intensive Care Units or Ship Bridges.

Interfaces currently addressing these requirements for decision making in complex cyber-physical systems are rather primitive, mostly limited to single, unspecific alerts and auditory cues for gaining and dragging the attention to information and entities. Frequent, often unspecific alerts are leaving the human with the demanding task of identifying and localizing the problem. Different alarms, color changing displays, and sounds are even competing for the individual's attention, leading to irritation, alarm fatigue, and potentially dangerous human errors. While existing research in psychology shows that attention can be cued and that the quality of arousal, notification and attention shift influences the quality of task performance, a translation into human-computer interaction technology is fundamentally lacking.

In our ongoing work, we develop and evaluate pervasive displays for guiding human attention in cyber-physical environments by arousing an individual or several people and directing their attention and cognitive resources to a task or task sequence as they interact with an automated cyber-physical system. Starting out with work with pervasive display

for notifications for example in the field of task switching [2] we moved into the domain of increasingly automated cyber-physical systems. With head-mounted peripheral display technology we are notifying nurses in intensive care units about alarms while reducing acoustic noise [3]. In the maritime domain, we are visualizing critical situations and off-screen objects by peripheral head-mounted displays to increase situation awareness in maritime navigation [4]. We are designing ambient notifications for take-over of a human after a phase of automated driving [5]. Our overall goal is to develop digital technology for guiding and shifting attention by an egocentric display, which is formed by an ensemble of embodied, peripheral, and ambient displays. Our research will result in a fundamental understanding about effectively supporting of human-machine cooperation of tomorrow in which human and machine cooperate as a team supported by digital technology in increasingly automated cyber-physical systems.

References

- 1 Andrew L. Kun, Susanne Boll, Albrecht Schmidt, “Shifting Gears: User Interfaces in the Age of Autonomous Driving,” *IEEE Pervasive Computing* 15 (1), 32-38
- 2 Heiko Müller, Anastasia Kazakova, Wilko Heuten, and Susanne Boll. 2016. Supporting efficient task switching in a work environment with a pervasive display. In *Proceedings of the 5th ACM International Symposium on Pervasive Displays (PerDis '16)*. ACM, New York, NY, USA, 13-19.
- 3 Vanessa Cobus, Wilko Heuten, and Susanne Boll. 2017. Multimodal head-mounted display for multimodal alarms in intensive care units. In *Proceedings of the 6th ACM International Symposium on Pervasive Displays (PerDis '17)*. ACM, New York, NY, USA, Article 26, 2 pages.
- 4 Tim Claudius Stratmann, Susanne Boll. *Demon Hunt-The Role of Endsley’s Demons of Situation Awareness in Maritime Accidents*. International Working Conference on Human-Centred Software Engineering/ International Working Conference on Human Error, Safety, and System Development, 2016
- 5 Shadan Sadeghian Borojeni, Lewis Chuang, Wilko Heuten, and Susanne Boll. 2016. Assisting Drivers with Ambient Take-Over Requests in Highly Automated Driving. In *Proceedings of the 8th International Conference on Automotive User Interfaces and Interactive Vehicular Applications (Automotive'UI 16)*. ACM, New York, NY, USA, 237-244.

3.25 Notification Management is Important to Avoid Information Overload

Sven Gehring (DFKI Saarbrücken, DE, sven.gehring@dfki.de)

License  Creative Commons BY 3.0 Unported license
© Sven Gehring

Digital devices that are capable of providing notifications to their users have become ubiquitous over the last decade. We encounter and interact with such digital systems in our daily lives in different shapes. In our smart homes we have multiple interactive devices providing notifications on the status of the house’s infrastructure, such as the heating system, doorbell, blinds, etc.. It is common that the notifications are provided by a smartphone application. However, smart home sensors of different vendors usually come with different dedicated apps, increasing the number of sources for notifications. Additionally, many of today’s household appliances also communicate their current status through notifications having different modalities. For example, dishwashers project the information onto the floor

or washing machines beep or send a message to a phone. This already results in a huge amount of – often distracting – notifications in our daily lives. Regarding smartphones, the situation becomes even worse [1, 2]. It is quite common to have more than 100 different applications installed on the smartphone. Most of the are displaying notifications to inform the user about new content, incoming messages, or they simple crave for his attention to open the app. If you have multiple smart devices (e.g., smartphone, tablet, smartwatch, notebook, smart TV, etc.), you usually receive the same notifications on each of the devices, causing an information overload that may lead to a notification blindness, where the user either ignores incoming notifications or just turns them off. A similar effect could be observed with public displays, that flood public spaces with digital advertisements [3]. We need to develop smart notification management systems that handle incoming notifications and provide them to the user in an appropriate situation, modality and frequency. If you acknowledged a notification on one device, it should also be acknowledged on all other devices. Another example could be that notifications are only displayed on the device that is the closest to the user. Ideally, in the future, notifications will detached from particular digital devices such as smartphones, tablets or wearables. Notifications should be displayed at the right time, in the right modality and in the right place. They should be ambiently embedded into our environment, which is a challenging task [4]. For example, a water dispenser in a living room could be illuminated to notify a person of drinking more. A further issue that needs to be addressed when developing smart notification management systems is the demographic change. In the future, the number elderly people compared young people will drastically increase. In order to develop usable notification management systems and in order to achieve a good user acceptance, systems need to be designed with elderly users in mind.

References

- 1 Matthias Böhmer, Christian Lander, Sven Gehring, Duncan P. Brumby, and Antonio Krüger. 2014. *Interrupted by a phone call: exploring designs for lowering the impact of call notifications for smartphone users*. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14). ACM, New York, NY, USA, 3045-3054. DOI=<http://dx.doi.org/10.1145/2556288.2557066>
- 2 Luis Leiva, Matthias Böhmer, Sven Gehring, and Antonio Krüger. 2012. *Back to the app: the costs of mobile application interruptions*. In Proceedings of the 14th international conference on Human-computer interaction with mobile devices and services (MobileHCI '12). ACM, New York, NY, USA, 291-294. DOI=<http://dx.doi.org/10.1145/2371574.2371617>
- 3 Müller J. et al.. 2009. *Display Blindness: The Effect of Expectations on Attention towards Digital Signage*. In: Tokuda H., Beigl M., Friday A., Brush A.J.B., Tobe Y. (eds) Pervasive Computing. Pervasive 2009. Lecture Notes in Computer Science, vol 5538. Springer, Berlin, Heidelberg
- 4 Frederik Wiehr, Alexandra Voit, Dominik Weber, Sven Gehring, Christoph Witte, Daniel Kärcher, Niels Henze, and Antonio Krüger. 2016. *Challenges in designing and implementing adaptive ambient notification environments*. In Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct (UbiComp '16). ACM, New York, NY, USA, 1578-1583. DOI: <https://doi.org/10.1145/2968219.2968552>

3.26 Towards Attention-Aware Adaptive Notification in Advanced Ubiquitous Computing

Tadashi Okoshi (Keio University, JP, slash@sfc.keio.ac.jp)

License  Creative Commons BY 3.0 Unported license
© Tadashi Okoshi

There has been an explosion of information available for people to read and act on. However, the amount of attention that can apply to this growing amount of information, has remained constant. Approaches for dealing with this include multitasking or dividing attention among a number of sources, and relying on push notifications to bring information to the forefront of their attention. However, notifications are responsible for an even greater number of interruptions. This is exacerbated by the fact that users are carrying, wearing and using a growing number of mobile and wearable computing devices including notebooks, tablets, smart phones, smart watches or wearable sensors, all of which can deliver interruptive notifications. Making the problem even worse are the growing number of installed applications on each device, each of which can also interrupt the mobile device owner. In particular, communication-based applications that support phone calls, texts chats, and social networking particularly suffer from this. But, games, news and other applications have similar issues as well, leading to a setting where people's everyday lives are significantly impacted from a feeling that they are being constantly interrupted by their computing systems.

Given the ever-increasing degree of information overload, the limited resource of human attention is the new bottleneck in interactive computing. In our research, we particularly focus on interruption overload, a form of distraction caused by the excessive number and inappropriate delivery of notifications from computing systems. All widely-used notification systems deliver notifications as soon as they are received, and this has been shown to negatively affect users' work productivity. Users experience notifications at "random" timings; that is, as they arrive on her devices. In other words, notifications from a variety of applications and services reach Melissa without any consideration of whether she is actually interruptible, causing divided attention and possible having negative impacts on her work productivity. In ubiquitous computing where computer systems promote calmness, they need to behave adaptively with regard to the user's current attention and interruptibility status. To address this problem, we follow the proposal of past literatures, to defer notifications until the user is experiencing a natural "breakpoint", defined as the boundary between two adjacent units of a user's activity. Deferring the interruptive notifications to this point can lower the impact of the interruption on users' cognitive load. Our research software Attelia successfully detected mobile user's breakpoints both during her physical activities and device interactions, in-real-time, by using mobile sensing on the smartphone and smartwatch and machine learning classification technique. Our current significant research challenges firstly include evaluation of our software in the real-world. We integrated the Attelia logic into Yahoo! Japan Android app. Our extensive in-the-wild user study with more than 680,000 users revealed that notifications in breakpoint timings surely improve users' experiences with reduced response time, increased click and engagement level. The second challenge is integration with smart city research. To review citizen's subjective experiences in a designed smart city, we plotted user's affective status (including interruptibility) onto a geographical map called "Affective Map". Our initial evaluation with the map revealed different types of breakpoints and interruptibility are observed in different locations and situations, such as rides in the train and buses.

3.27 Interactive Assistance in Working and Teaching Facilities

Thomas Kosch (University of Stuttgart, DE, thomas.kosch@vis.uni-stuttgart.de)

License  Creative Commons BY 3.0 Unported license
© Thomas Kosch

Disturbances at work and learning environments by external factors interrupt the workflow in which the person was. Besides of decreasing the concentration on the current task, a refocus on the task is required to enter a flow state again. Distributing the current physiological or cognitive states helps to avoid unnecessary interruptions by others. Within the course of our research, we investigate how physiological parameters can be measured effectively and visualized ambiently to notify the person itself and their social environment. Informing social contacts about the current state avoids unnecessary interruptions when doing intensive tasks.

Generally, our work encompasses the measurement of physiological conditions to provide interactive assistance in working and teaching facilities with embedded cognition-awareness. We focus on the analysis of eye movements, brain activity, facial parameters, galvanic skin response, and muscular activity. This is complemented by the training of models, which are suited to individual physiological responses. Such metrics can be used in a variety of use cases: reflecting physiological conditions over a timespan or adapting visualizations are just a few areas of interest. Within this, we investigate how physiological states can be used collaboratively, for instance, when it comes to notify other people or the person itself about his state.

It is important to be aware how to communicate this without drawing too much attention at the same time. A visual approach is to slowly increase the brightness of a lamp according to the stress and mental level of an individual. Heating up the door handle to a reasonable temperature provides a haptic solution. Finding the optimal representation modality states another challenge during the course of this research.

On a visual level, we have looked into how brain activation can be represented in an understandable way [1, 2]. Demanded brain parts are visualized in a color coded way, making the overall cognitive load visible the individual. Transferring these insights ambiently is our scope for future work. This is complemented by a feedback study [3], where we investigate the suitability of ambient tactile, visual, and auditory feedback. The results show a high preference for visual feedback.

References

- 1 T. Kosch, “Real-time brain mapping for treating substance abuse using neurofeedback.” 2015.
- 2 T. Kosch, M. Hassib, and A. Schmidt, “The brain matters: a 3d real-time visualization to examine brain source activation leveraging neurofeedback,” in Proceedings of the 2016 chi conference extended abstracts on human factors in computing systems, New York, NY, USA, 2016, pp. 1570-1576.
- 3 T. Kosch, R. Kettner, M. Funk, and A. Schmidt, “Comparing tactile, auditory, and visual assembly error-feedback for workers with cognitive impairments,” in Proceedings of the 18th international acm sigaccess conference on computers & accessibility, 2016.

3.28 Awareness in Computer Supported Cooperative Work

Tom Gross (*University of Bamberg, DE, email@tomgross.net*)

License  Creative Commons BY 3.0 Unported license
© Tom Gross

Ambient notification environments are an exciting research topic with many unsolved research questions and challenges. At the same time they have some predecessors and related areas that can provide great input. Particularly context-aware systems and event-notification infrastructures have a lot of similar challenges. Also the whole body of research in Computer-Supported Cooperative Work (CSCW) on awareness, which in general addresses questions around providing users with important information of their environment including other team members, has a lot to offer. Great concepts and systems have been developed and ethnomethodologically-informed ethnographic studies have a lot to offer towards a better understanding of how users manage to perform their own tasks and at the same time capturing what is going on around them.

At the Dagstuhl Seminar 17161 on Ambient Notification Environments I gave an overview along my historic survey paper on 25 years of awareness research in CSCW [1], where I characterised the origins of awareness in CSCW both in terms of ethnographic findings as well as in early concepts and technologies. Great concepts and technology have been supporting coexistence awareness with mutual information on users' presence and cooperation awareness with mutual information on users' activities in shared editors and other environments. Base technologies for sensing data and for presenting awareness information have been reported as well as various approaches on modelling awareness information.

In my own work related to ambient notification environments we have been doing surveys on awareness in recent [1] as well as early [2] works. We have also developed our own event notification environments. For instance, in Gross and Prinz [3], we have presented an early approach that addressed challenges that are still relevant today concerning the capturing of data in their context of origin, modelling work contexts of users and keeping those models up-to-date, and presenting in-situ relevant awareness information on the users' computer screens as well as in their physical environment with ambient displays.

References

- 1 Gross, Tom. Supporting Effortless Coordination: 25 Years of Awareness Research. *Computer Supported Cooperative Work: The Journal of Collaborative Computing* 22, 4-6 (Aug.-Dec. 2013). pp. 425-474.
- 2 Gross, Tom, Stary, Christian and Totter, Alex. User-Centred Awareness in Computer-Supported Cooperative Work-Systems: Structured Embedding of Findings from Social Sciences. *International Journal of Human-Computer Interaction (IJHCI)* 18, 3 (June 2005). pp. 323-360.
- 3 Gross, Tom and Prinz, Wolfgang. Modelling Shared Contexts in Cooperative Environments: Concept, Implementation, and Evaluation. *Computer Supported Cooperative Work: The Journal of Collaborative Computing* 13, 3-4 (Aug. 2004). pp. 283-303.

3.29 The Role of Task Engagement in Mobile Attention Management

Veljko Pejovic (University of Ljubljana, SI, Veljko.Pejovic@fri.uni-lj.si)

License  Creative Commons BY 3.0 Unported license
© Veljko Pejovic

Interruptions following mobile notifications are a complex phenomenon stemming from the message arriving at an inappropriate time, when the user is commuting on a bicycle, sleeping, or in a meeting, to name a few such situations. In addition, messages can be inappropriate due to their content or the nature of the sender-receiver relationship. In our work we have built models and practical systems that take the above factors into account and try to predict a user's interruptibility [1, 2, 3]. However, another important aspect that impacts a user's interruptibility is the current task engagement, yet, this has been overlooked in mobile notification management systems. The knowledge of a user's task engagement or the cognitive load, opens up an array of possibilities for a seamless mobile computing device – human interaction, beyond just mobile notification management. Today's most ubiquitous personal sensing devices, such as smartphones, are equipped with an array of sensors that may be used to infer different aspects of human behavior. Still, inferring task engagement using smartphone sensors has not been explored. In our research we investigated the automated task engagement inference using only smartphone sensors [4]. We designed, developed and deployed a mobile sensing application TaskyApp, and collected sensor readings and task engagement labels from eight users in an office setting. Using machine learning we demonstrated that there is weak link between the smartphone sensor data and the task engagement. The most promising are the movement features, sensed by the built-in accelerometers, that show a significant correlation with the task engagement labels provided by the user. However, for full inference, the smartphone is not sufficient. In our future research we plan to use wearable devices, predominantly smartwatches and wristbands, in order to get physiological data pertaining to the user. Information such as heart rate variability, electrodermal activity and skin temperature have been shown to correlate with the user's cognitive load. Yet, these correlations have been shown only in well controlled laboratory settings with specialised equipment. Our research, on the other hand, aims to recognise task engagement using commodity wearable sensors.

References

- 1 V. Pejovic and M. Musolesi InterruptMe: Designing Intelligent Prompting Mechanisms for Pervasive Applications UbiComp'14, Seattle, WA, USA, September 2014.
- 2 A. Mehrotra, M. Musolesi, R. Hendley and V. Pejovic Designing Content-driven Intelligent Notification Mechanisms for Mobile Applications UbiComp'15, Osaka, Japan, September 2015.
- 3 A. Mehrotra, V. Pejovic, J. Vermeulen, R. Hendley and M. Musolesi My Phone and Me: Understanding User's Receptivity to Mobile Notifications ACM CHI'16, San Jose, CA, USA, May 2016.
- 4 G. Urh and V. Pejovic TaskyApp: Inferring Task Engagement via Smartphone Sensing to appear at Ubintention workshop with UbiComp'16, Heidelberg, Germany, September 2016.

4 Sessions

4.1 Notifications vs. Augmentation

Joseph A. Paradiso (MIT – Cambridge, US, joep@media.mit.edu)

License  Creative Commons BY 3.0 Unported license

© Joseph A. Paradiso

URL <https://www.media.mit.edu/people/joep/overview/>

Currently, sensors are embedded in everyday devices. Joe Paradiso presented their developed general decentralized web based sensor framework (CHAIN API) to describe sensors (real sensors, virtual sensors) as RESTful resources. This is a lightweight, adaptive way to describe sensors and to define relations between different services with a shared vocabulary. Combining different sensors with Amazon Echo or Google Home is anticipated to give rise to the "Digital Butler". These devices are expected to inform the users proactively about their own needs. After a transition phase to ubiquitous computing, smart environments will become an extension of the self.

In addition, Joe Paradiso described a need to redefine the user's presence in physical environments. In the future, sensors will be able to deliver different information directly into the user's perception. In this regard, Joe Paradiso presented built environments that affords extra-sensory perception. For example, sensors that are embedded through the MIT building collects multi-modal sensor data that is produced by the building itself or its inhabitants that are invisible to natural perception. To communicate this data, real-time visualization are engineered that presents aggregated information to whoever might require it, e.g. the temperature in different areas (cf. DoppelLab ¹). While natural human perception is incapable of 'seeing' temperature and is limited of seeing people hidden behind walls, such visualizations will allow us to do so.

Another offered example was the Living laboratory ² near Plymouth, Massachusetts with a size of 600 acres. In this area, many sensors were deployed to document the ecological processes of the restoration to natural wetland. In total, four key locations were set up as base-stations (running on solar cells) as well as numerous and varied sensors, e.g. The Mayton 2.0 Tidmarsh Wireless Sensor. Thus, information about the environment could be collected and transmitted in many different ways. For instance, they have set up a website that visualizes real-time aggregated data from the various sensors to allow people to explore the environment from the comfort of their own homes. Besides this, they have developed a DoppelMarsh VR sensor visualization to display the living observatory at the same time as the aggregated sensor information, e.g. temperature. Another VR instantiation provides a holistic view of the living observatory from a high altitude view together with real time sensor visualization. All of these instances afford humans the ability to access and process information about remote environments in ways that extend beyond our natural capabilities.

To allow users to "inhabit" remote environments, it is not sufficient to provide visual information. A separate project seeks to display virtual sounds in the user's environment that are indistinguishable from the user's real, actual surroundings. They use an eye-tracker, a chest strap sensor, a touch and a head tracker as well as a bone conduction headphone to display virtual sounds from a distance. Another framework, they developed is Quantizer a collaboration of ATLAS and the Media Lab. This framework supports diverse musical

¹ DoppelLab: <http://doppelab.media.mit.edu/>

² tidmarsh.media.mit.edu

compositions running on a real time ATLAS/LHC data. It composes music based on the detected particle collisions in the large hydron collider.

The ubiquitous deployment of sensors will not only allow us to modify our felt environments, it will also allow environments to modify themselves in response to our needs. In another project, wearable sensors were used in combination with sensors in the building to adjust the room temperature in a way that saved heating energy as well as optimize user comfort. Similarly, ambient lighting conditions were controlled using synchronized cameras within the infrastructure as well as a Google glass that adjusted reflected light to be optimal at the location that intersected with the user's line-of-sight. Also, the light conditions were regulated with a smartphone app to correspond with the user's moods (e.g., casual, focused). All of these serve a basis for create workspaces that are notified of their occupants' requirements such as to create a conducive ambience, for example mediated workspaces to relax employees (cf. Mindful Photons ³).

4.2 Multimodal Notifications for Ambient Environments

Steve Brewster

(University of Glasgow – Glasgow, GB, stephen.brewster@glasgow.ac.uk)

License © Creative Commons BY 3.0 Unported license

© Steve Brewster

URL <http://www.dcs.gla.ac.uk/stephen/aboutme.shtml>

Stephen Brewster presented a substantial body of research that explored how multi-modal notifications could be effectively combined to create interfaces that display information, particularly in a way that exploit available sensors.

One key example that was provided, which would benefit from this endeavor, was home support for aging. Here, the goal is to support an elderly person with reminders about medication, appointments, and food intake. In this case, it is important to be able to provide discreet and private notifications that are considerate of the context and do not embarrass the user. A key factor for design is the fact that users tend to desire a feeling of being in control. Therefore, notifications has to be delivered to the user with the right notification modality and at the right moment, e.g, during performing everyday activities or potentially concentrating on a current task or interacting with others.

Vision continues to be the most dominant channel for delivering notifications. This is because it affords a large design space (e.g., shape and color). Another form of a notification is the earcon – non-speech audio such as music, structured sound and other sound effects. However, auditory notifications suffer from the disadvantage that they are not typically private. In this regard, tactile notifications represent a more private channel but one that requires near proximity in order to be detected.

More exotic channels of delivery have been explored. For example, olfactory notifications, such as the smell of dark chocolate. Temperature can also communicate strong and intuitive notifications where by values of cold, warm, or hot are arguably related to strong emotional responses (e.g., warm for loving, cold for distant). Typically, the palm is the most receptive body part, but acceptable sensing can be achieved with the wrist and arm. Temperature

³ Mindful Photons: <https://www.media.mit.edu/projects/mindful-photons-context-aware-lighting/overview/>

notifications can be embedded in our environment, whereby a door handle can be either cold, warm, or hot. User studies have shown that this is consistently interpreted as the occupant being, respectively: away, present, or busy.

Several studies showed that combining traditional notification modalities, e.g. tactile and sound, results that notifications were perceived as more urgent and critical. In addition, the use of physical objects within the environment of the user as a notification channel is a completely new approach to display notifications, which should be investigated in future work. Another research questions is how traditional and new notification modalities can be combined to display notifications.

To conclude, Stephen Brewster issued a list of challenges and key questions that beset the design of future notification systems:

1. How can we assure that people who sent the message and those who receive perceive the importance and urgency levels to be equivalent?
2. How can we ensure that notifications are discreet and private?
3. How do we give people a sense of control over their notifications?
4. What do notifications suggest about a person's character (e.g., discipline, attention, control) ?
5. Is this generally true that notifications affect performance (in a memory game), choice of modality did not affect performance.
6. Why is the urgency increasing with increasing redundancy across modalities.
7. Can we physically change an object to create notifications (eg., more easily seen or felt)?
8. How to combine modalities successfully?
9. Are more notification channels always better?
10. What new possibilities are there for other notification strategies?
11. When should we use which channel or strategy for which users?

Altogether, multi-modal design spaces are more varied than would be conventionally imagined. This is especially when dimensions beyond our natural senses are considered (e.g., temperature). A key challenge lies in understanding how they can be effectively combined to communicate notifications in a way that is intuitive to our instinctive interpretations and in a way that can be pragmatically embedded in our environments.

4.3 Augmenting Social Interactions

Elisabeth André

(Augsburg University – Augsburg, DE, andre@informatik.uni-augsburg.de)

License © Creative Commons BY 3.0 Unported license

© Elisabeth André

URL <https://www.informatik.uni-augsburg.de/lehrstuehle/hcm/staff/andre/>

In this talk, Elisabeth André expounded on the state-of-the-art capability of computing systems to sense and interpret social cues, for the purpose of generating appropriate responses to their users. She presented use cases that relied heavily on augmented reality and social signal processing, the latter for which her research has developed numerous algorithms that can be readily implemented with off-the-shelf sensor devices.

With regards to social signal processing, her work targets the decoding of emotions, facial expressions, gestures, speech, physiological measurements, interest, engagement, trust, personality, rapport, group cohesion as well as functional roles. As mentioned above, her

extensive research has resulted in the development of a social signal interpretation framework⁴ that seeks to be readily deployed across research labs and different environments. This framework allows for the synchronization of a large variety of sensory data recordings as well as the interpretation of emotional and social behaviors of observed users.

In one use example, Elisabeth André reported how this allows social interaction to be supported in order to motivate targeted users to behave in a context-appropriate fashion. For example, she noted how the inappropriate behavior of candidates at interviews could have untoward consequences. Therefore, continuous feedback could be provided, based on exhibited behavior to shape the behavior of the person that SSI observes. Project TARDIS is a milestone in this endeavor to inculcate desirable interview behavior⁵. Here, students are trained within a game that provides closed-loop feedback, to generate appropriate gestures and verbal speech that are typical to a successful job interview.

In another example, OpenSSI can be deployed to pick up and aggregate social signals from a target live audience. This can be used to provide real-time feedback to speakers, allowing them to modify their talks in response. In one study, participants gave short presentations with or without the support of such a system. The main findings were that the inappropriate behavior of the participants decreased when speakers were provided with real-time notifications, based on the social signals sensed and aggregated from the environment. In a separate qualitative study, more naturalistic discussions were set up to investigate if real-time feedback, which were based on sensed social behavior, would improve or disrupt social interactions. Participants were divided into groups that discussed a topic for 10 minutes. Different participants within each group provided notifications via a variety of channels (e.g., visual head-mounted-display, tactile and audio output). The study's results demonstrated that real-time notifications based on social sensing resulted in discussions that were more balanced. Furthermore, the notification system were not considered to be particularly disruptive. Participants reported that they appreciated feedback and modified their behavior accordingly: "It was a good feeling to see the icons in green. It lasts longer than a nod [and easier to detect]." Tactile notifications were rated to be the most disturbing and audio as the least.

An interesting outlook of this work would be to investigate whether participants might pay attention under conditions that provide subliminal feedback. Currently, behavior only appears to be responsive to supraliminal notifications. Nonetheless, this could be a matter of its implementation. The rapid advances noted in wearable devices (e.g., Google Glasses) could provide the means whereby notifications could be delivered that are at the same time accessible to their users' consciousness without being obtrusive.

4.4 Debates

During the seminar, three debates were organized on topics that participants voted on. The first debate was on whether user control is an illusion regarding notification systems and their applications. The second debate dealt with the question of whether artificial intelligence will be able to solve all problems regarding notifications in the future. The last debate addressed whether adding a cost for messages is a good idea for restricting the deluge of unnecessary and undesired notifications.

⁴ OpenSSI: <https://hcm-lab.de/projects/ssi/>

⁵ TARDIS: hcm-lab.de/logue



■ **Figure 3** Seminar participants debating controversial topics.

Every debate featured an advocating as well as an opposing team, with four speakers per team. Each speaker spoke for 3 minutes. The summaries are structured to capture the turn-by-turn dialogue between the advocates and their opponents.

4.4.1 Debate 1: User control is an illusion

Advocates:

Control is something that directs behavior or the turn of events. In contrast, illusion is a deceptive idea or belief. The user is under a false impression. In addition, users might not be allowed to make changes for certain applications. If I want the app, I have to give up control (i.e., terms and conditions) Also, developers want to have the control rather than to allow it to users. Users may believe that they are in control but they are not.

Opposition:

The control freaks also looked at some definitions Users can always switch their phones off. That is an ultimate point of control. Humans are able to focus and cut out unwanted information. New methods to present information.

Advocates:

We are not hermits. We are surrounded by devices that we cannot turn off. We have attention that is less than a goldfish. We cannot understand the rules behind these systems so how can we control them? We cannot turn things off because notifications are sent by our loved ones and we do not wish to turn things off and hurt people's feelings. Because of our risk-averse nature, we fear turning off notifications in case something bad happens. We cannot control our participation.

Opposition:

Personal notifications need not be in everybody's face (e.g., vibrotactile cues). Systems can learn from what you do. You are in control of the information that you supply to notification systems. Adjustable parameters allow for control.

Advocates:

The reality is that we are bombarded by notifications. Attention is an ability that we have evolved to control. This ability evolved from many centuries of physical manipulations. It is not developed to deal with modern technology. You cannot switch off your brain.

Opposition:

We are responsible for information overload. Nature holds more information than systems can send us. Stimuli are changing (becoming more salient) than the natural environment. Any notification can be dealt with in different ways. We have a wonderful capacity for ignorance, and we can learn to ignore new stimuli

Advocates:

Providers have a lot of influence their applications. We cannot switch our technology off for the whole week. Opting-out is not an option because it takes too much effort.

Opposition:

These applications are something that we created ourselves. Humans grow up in new systems and are able to adapt to these new systems. Never before have users had more control than now, over things that they have never imagined was possible.

4.4.2 Debate 2: Artificial Intelligence will solve all of our notification problems**Advocates:**

What are the problems in the moment? Incoming notifications grab our attention, therefore the task completion rates are getting lower. Also, the receiving notifications is poorly timed. Thou, users receive notifications in inappropriate situations. In contrast, artificial intelligence (AI) is designed for scheduling. Humans are bad in dealing with notifications and set up their notification systems. However, in the future notifications will adapt to autonomous systems (not humans). In addition, the most irritant notifications come from things that people don't care about which will be acted on by autonomous agents.

Opposition:

AI has consistently been over-rated and failed in its promises. Precision: we want 100% precision this requires sensitivity to context and content. Even if we have the perfect algorithms, the sensing capability will not be there. Also, people want the illusion of control. They will always prefer even their own bad decisions. We will not be able to build these systems and even if we can build them, they will not be accepted.

Advocates:

Today, AI is playing better games than humans (e.g., chess, go,). Why is AI better than the human? Google is working on things to help predict (better than humans) when we need notifications. AI will filter information automatically and help the users dealing with information overload.

Opposition:

If machines are communicating to one another without considering the human, would that not create a new species? We do not govern our life and universe on the basis of mathematics, which is what governs AI.

Advocates:

We do not need a general AI, we need only an AI that deals with notifications. This is a rather small area and such problems are dissolved by AI. Today, the user has to read all notifications and the user had to read emails in the beginning. Today, the spam problem of emails is solved by AI. AI will be fundamental in dealing with filtering notifications according to user's wishes.

Opposition:

We are talking about general area AI because notifications pervade our every day activities. There are no sensors to correctly interpret things like being in love. Allowing AI to make decisions, such as who to marry will create problems since they cannot sense and infer correctly.

Advocates:

The opposition argues that AI will take over our lives. In contrast, the technological gap will narrow. AI will allow us to enjoy our lives by pushing things that we do not care about into the background.

Opposition:

Economics not technology will solve the problems that we face. AI will ask many people about information sensing from the real world. Notifications are the essential thing from our lives in the future.

4.4.3 Debate 3: There should be a cost to send messages**Advocates:**

Systems are not designed to take our attention as a valuable resource. Today, we have costs for sending messages as time, social capital. Also, there is a social opportunity cost when one has to deal with notifications. Not just on the sender but also on the recipient.

Opposition:

Adding cost to messages is a threat to basic human rights. It is discriminating and will divide the world population. Imagine a world where poor people are not able to communicate with others anymore. In addition, some institutions will be disconnected. In total, the world will be more disconnected and people will feel lonely.

Advocates:

There are already existing costs, for example in the way of advertisements. There are negative costs of not sending messages. We have limited resources such as time and attention that we invest when we send messages.

Opposition:

There are issues of practicability and implications of introducing such costs. We currently have systems that allow for free messages. If a cost is introduced, there will be higher

risks of being locked into systems. The only way to do so would be to have a government regulating communications. This poses a risk to net neutrality. In the end, this will reduce connectedness. Introducing costs will address the symptom but not the problem.

Advocates:

We are already paying for communications anyway. Costs be handled in a different way. You are considering the sender and the receiver even though there are many agents in-between. The net is NOT neutral today. You made a confusion about infrastructure and common pool resources (e.g., attention). We suffer from problems of congestants. How do we deal with common pool resources? The best strategy from an economic point of view is to introduce a cost to reduce throughflow.

Opposition:

The government has a too narrow viewpoint. It is not a capitalist industrialized world. The Arab revolution would not have happened if costs were involved. Therefore, costs for messages will kill democracy. Micropayments in developing country are based entirely on free SMSs. In an age of democratic change, we need to look at remote schooling, which depends on no cost messaging. Another important factor is the crisis management. There is a need to inform people in certain situations.

Advocates:

We have no metric or model for the inherent costs of messaging and notifications. Trump was using notifications. Perhaps a cost to him that was imposed based on the number of recipients could have avoided populism.

Opposition:

You just want to receive fewer messages. Adding a cost is a selfish decision. Humans want to communicate with one another. No one has the right to prevent people from speaking to each other. We want everybody to be able to communicate. People will use services that are available and free. Costs will NOT stop people from communicating to one another, just restrict communication.

4.5 Speed Dating Session

Speed dating sessions were conducted throughout the seminar to motivate spontaneous discussions between the participants on “burning questions”. In this activity, participants comprised two rows that faced one other. Every 3 minutes, one row would move to the left while the other remained static. The following are examples of questions that were posed:

- Is the scalability issue of notifications (50 per second in 2030) real or not?
- What are your strategies for dealing with notifications?
- Think of positive/negative examples of notifications!

Each question was discussed twice with different partners. This meant that participants were offered a new topic to discuss with new partners.

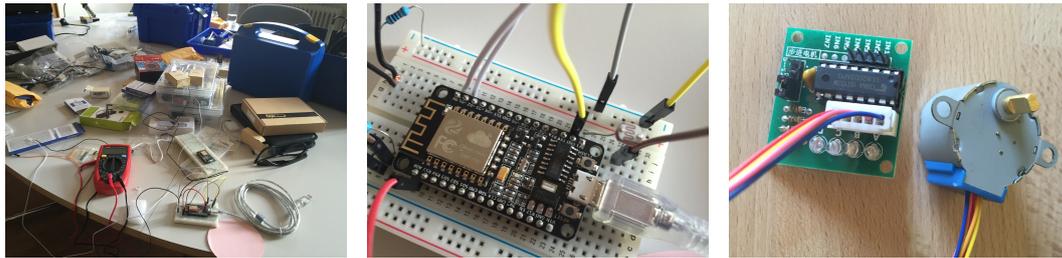


■ **Figure 4** Short but intense discussions during the speed-dating session.

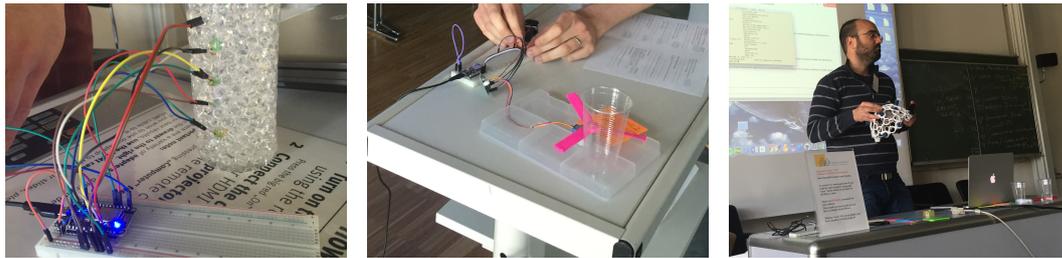
4.6 Open Maker Session

The Open Maker Session started with a short introduction into Arduino development and the ESP8266 chip. A small ultra-low power microcontroller which supports to connect to a Wi-Fi network and make simple TCP/IP connections, read sensor data, control connected device or receive notifications. Afterward, the participants brainstormed in small groups and started prototyping for two hours. Many groups continued developing after the dinner to finalize their prototype.

During the morning session of the final day, participants presented their new concepts for handling the vast amount of notifications we receive daily. The proposed solutions covered a broad range of different approaches. Creative groups worked on a project to investigate how colorful their lives are. Therefore, they used a color sensor to detect colors in their environment and displaying the detected color with an RGB LED. Others build a tangible notification lighthouse which informs the user about urgent or important notifications by different illumination. In figure 6 (center) a destructive notification system is displayed. It notifies the user (e.g. to drink more) by knocking over a drinking glass. Concluding, participants were able to rapid prototype functional notification mechanisms and explored new concepts on how a notification environment could be established in the future.



■ **Figure 5** Plenty of hardware was provided to support creative prototype development.



■ **Figure 6** The prototypes developed during the open maker session were presented at the last day of the seminar.

<p>People</p> <p>Sender/audience 3rd party System / programmer – organisation Collectives</p> <p>Border case: tannoy announcement at airport Design idea: public notifications</p>	<p>Activity</p> <p>System / other initiated Interruption vs ongoing vs parallel Ambient Activity awareness (CSCW concept)</p> <p>Border case: Dropbox notifications Design idea: garment forcing you to go out</p>
<p>Content</p> <p>Style Signal / content Escalation / fading</p> <p>Border case: indexical notifications Design idea: peek into the app</p>	<p>Time</p> <p>Notification tail Zombie notifications Notification speed Very long term notifications</p> <p>Border case: notification history Design idea: clock of the long now</p>

■ **Figure 7** Taxonomy of notifications

4.7 Conceptual Maker Session

The taxonomy for a notification system that handles notifications has to consider the person who sent the message, the current receiver’s activity as well the content of the notification itself as well the time when the notification is received (cf. Figure 7).

Participants

- Elisabeth André
Universität Augsburg, DE
- Susanne Boll
Universität Oldenburg, DE
- Stephen Brewster
University of Glasgow, GB
- Barry Brown
Stockholm University, SE
- Andreas Bulling
MPI für Informatik –
Saarbrücken, DE
- Andreas Butz
LMU München, DE
- Keith Cheverst
Lancaster University, GB
- Lewis Chuang
MPI für biologische Kybernetik –
Tübingen, DE
- Raimund Dachselt
TU Dresden, DE
- Anind K. Dey
Carnegie Mellon University –
Pittsburgh, US
- David Dobbelstein
Universität Ulm, DE
- Christoph Gebhardt
ETH Zürich, CH
- Sven Gehring
DFKI – Saarbrücken, DE
- Hans Gellersen
Lancaster University, GB
- Tom Gross
Universität Bamberg, DE
- Niels Henze
Universität Stuttgart, DE
- Judy Kay
The University of Sydney, AU
- Pascal Knierim
Universität Stuttgart, DE
- Thomas Kosch
Universität Stuttgart, DE
- Bob Kummerfeld
The University of Sydney, AU
- Markus Löchtefeld
Aalborg University, DK
- Evangelos Niforatos
University of Lugano, CH
- Tadashi Okoshi
Keio University – Fujisawa, JP
- Joseph A. Paradiso
MIT – Cambridge, US
- Veljko Pejovic
University of Ljubljana, SI
- Benjamin Poppinga
AUDI AG – Ingolstadt, DE
- Aaron Quigley
University of St. Andrews, GB
- Harald Reiterer
Universität Konstanz, DE
- Michael Rohs
Universität Hannover, DE
- Albrecht Schmidt
Universität Stuttgart, DE
- Aurélien Tabard
University Claude Bernard –
Lyon, FR
- Pauline Vierne
Universität der Künste –
Berlin, DE
- Alexandra Voit
Universität Stuttgart, DE
- Frederik Wiehr
DFKI – Saarbrücken, DE



Online Privacy and Web Transparency

Edited by

Nataliia Bielova¹, Nikolaos Laoutaris², Arvind Narayanan³, and Nick Nikiforakis⁴

1 Université Côte d’Azur, Inria, FR, nataliia.bielova@inria.fr

2 Telefónica Research – Barcelona, ES, laoutaris@gmail.com

3 Princeton University, US, arvindn@cs.princeton.edu

4 Stony Brook University, US, nick@cs.stonybrook.edu

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 17162 “Online Privacy and Web Transparency”. The seminar brought 29 participants in computer science, law and policy together, coming from companies and research institutions across Europe and the US.

The 2.5-days seminar had a well-filled program, with 25 research talks, followed by 7 short panel discussions, and 6 5-minute talks. Online privacy and Web transparency is a broad research field, that includes detection of privacy leaks on the Web and mobiles, measurement of tracking technologies on the Web, transparency tools to detect bias and discrimination, as well as how laws and regulations address these problems from a law research perspective, and how technical solutions can influence standards and laws.

Seminar April 17–20, 2017 – <http://www.dagstuhl.de/17162>

1998 ACM Subject Classification D.4.6 Security and Protection (K.6.5), H.3.4 Systems and Software – User profiles and alert services, H.3.5 Online Information Services – Web-based services, K.4.1 Public Policy Issues – Ethics, Privacy, Regulation

Keywords and phrases Big data, Privacy Enhancing Technologies, Transparency, User Discrimination, Web Privacy

Digital Object Identifier 10.4230/DagRep.7.4.83

1 Executive Summary

Nataliia Bielova

Nikolaos Laoutaris

Arvind Narayanan

Nick Nikiforakis

License © Creative Commons BY 3.0 Unported license
© Nataliia Bielova, Nikolaos Laoutaris, Arvind Narayanan, and Nick Nikiforakis

The Dagstuhl Seminar on *Online Privacy and Web Transparency* was the first seminar at Dagstuhl that gathered together researchers working in web applications, online privacy, transparency on the web, privacy enhancing technologies, privacy measurement, and network economics, as well as several representatives of law and policy discipline.



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

Online Privacy and Web Transparency, *Dagstuhl Reports*, Vol. 7, Issue 04, pp. 83–106

Editors: Nataliia Bielova, Nikolaos Laoutaris, Arvind Narayanan, and Nick Nikiforakis



DAGSTUHL
REPORTS

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

Research context

The web has become an essential part of our society and is currently the main medium of information delivery. Billions of users browse the web on a daily basis, and there are single websites that have reached over one billion user accounts. In this environment, the ability to track users and their online habits can be very lucrative for advertising companies, yet very intrusive for the privacy of users.

Recent research has shown that third-party advertising networks and data brokers use a wide range of techniques in order to track users across the web – these techniques are used to reconstruct browsing sessions and to create profiles of users, inferring, among others, their hobbies, health status, political inclinations, and level of wealth. This information can be used to, not only deliver better targeted advertisements to users, but also to discriminate users, for example by providing customized prices for products based on a user's willingness and ability to pay. To protect users, several solutions have been proposed, ranging from the laws, policies and the W3C Do-Not-Track candidate recommendation, to browser tools developed by companies and volunteers, and other client-side mechanisms proposed by researchers. At the same time, a number of tools have been developed to increase transparency on the web and allow end users to know when they are being tracked and when discrimination happens.

The seminar aimed to address the open questions of how to protect user privacy and how to increase transparency on the web. The key objectives of the seminar are (i) review the state of the art in the field; (ii) identify key technical challenges and brainstorm potential solutions; (iii) understand how computer science research results can influence law and policy; (iv) discuss ethical and legal issues in privacy research.

The seminar brought together scientists from the privacy and transparency communities, as well as policy makers interested in understanding how existing privacy laws and policies can be implemented, and representatives of Internet users organisations. The discussions at this Dagstuhl Seminar were strongly inspired by the following questions and challenges:

Technology

- How can we detect tracking and algorithmic discrimination most effectively? What are the scientific and engineering challenges to overcome? What are the relative merits of automated, semi-automated, and crowdsourced approaches?
- How can we ensure that methodologies, techniques, and tools are shared across different communities working on this topic?
- How can we design the next generation of privacy tools, get users to actively use the tools, and generate data for privacy researchers to scrutinize?
- What are the tracking techniques and data collection practices on mobile devices and how do they compare to those on the web?
- What are the privacy and transparency issues raised by the Internet of Things, and how do we address them?

Law

- Do the current laws and policies cover existing tracking technologies? What is the process for reporting newly discovered tracking techniques to the appropriate Data Protection Authorities?
- Even with the appropriate legislation in place, how can we ensure that companies comply with the law? What can researchers do to help enforce compliance?

Ethical issues

- What is the best way to conduct web privacy research taking ethical issues into account?
- When does a study necessitate ethical review?

Users

- What is the most efficient way to raise user awareness about web tracking and transparency tools?

2 Table of Contents

Executive Summary	
<i>Nataliia Bielova, Nikolaos Laoutaris, Arvind Narayanan, and Nick Nikiforakis . . .</i>	83
Seminar program	88
Overview of Talks	91
How can we reconcile diversity and privacy?	
<i>Benoit Baudry</i>	91
Browser Extension and Login-Leak Experiment	
<i>Nataliia Bielova, Claude Castelluccia, and Gábor Gy. Gulyás</i>	92
A Brief History of Mobile Privacy Leaks	
<i>David Choffnes</i>	94
FDVT: Facebook Data Valuation Tool	
<i>Ángel Cuevas Rumin, Rubén Cuevas Rumin, and Jos'e González Cabañas</i>	94
Analyzing the Impact of Large Scale Online Tracking Measurement	
<i>Steven Englehardt and Arvind Narayanan</i>	95
Sandboxing Trackers with Resource Blocking Lists	
<i>Steven Englehardt</i>	95
The Diffix Framework: Noise Revisited, Again	
<i>Paul Francis</i>	96
Explaining Online Ads	
<i>Krishna P. Gummadi</i>	96
New Faces Of Bias in Online Labor Markets	
<i>Aniko Hannak</i>	96
Web transparency tools demo	
<i>Costas Jordanou</i>	97
Transparency in the era of programmatic, real-time bidding advertising and Cookie Synchronization	
<i>Nicolas Kourtellis</i>	97
How news media use Twitter to attract traffic?	
<i>Arnaud Legout</i>	98
Surveillance as a Regulatory Model	
<i>Timothy Libert</i>	99
Modern Web Privacy: standards, implementations, deployments	
<i>Lukasz Olejnik</i>	100
Ultrasonic Device Tracking for Fun and Profit	
<i>Konrad Rieck</i>	100
Control What You Include ! Server-Side Protection against Third Party Web Tracking	
<i>Dolière Francis Some, Nataliia Bielova, and Tamara Rezk</i>	101

Understanding the privacy risks of browser extensions
Oleksii Starov and Nick Nikiforakis 101

Accountability for Privacy and Discrimination
Michael Carl Tschantz 102

Tracking Classification, a cluster analysis approach to identify missing values
Robbert van Eijk 103

Break-out sessions 103

 Security and privacy trade off 103

 Research feasibility and validity 104

 Bluesky proposals 105

Participants 106

3 Seminar program

The seminar attracted 29 participants, in privacy and transparency, and law and policy together, coming from companies and research institutions across Europe and the US. The group presented a nice mix of participants from computer science community, industry (including researchers from AT&T, Telefónica Research and Cliqz), as well as researchers in law and policy (KU Leuven and Leiden University), and also representatives of data protection authorities (French CNIL and Federal Trade Commission).

The seminar had a well-filled program, with 25 research talks, followed by 7 short panel discussions, 3 break-out sessions and 6 5-minute talks. The organisers tried to keep enough time during breaks and in the evening for informal discussion, however because the seminar was only 2.5 days, we did not have the social event. To compensate for the possible lack of informal discussions, the last morning of the seminar was kept for the conclusion and free discussions.

Research talks

The organizers invited all the participants to take the floor during the seminar, and encouraged them to have rather short 15 minutes talks (differently from standard 25 minutes talks at conferences). This allowed to organise small panel discussions after each session where participants could engage in discussions.

Online privacy and web transparency is a new research field, where researchers from different backgrounds try to propose solutions to the problems. The organizers decided that each session will have a separate topic or open question, where researchers from different domains, as well as companies, can be mixed. As a result, different research results and open problems were presented in 5 working topics:

1. Advertisement and Real-Time Bidding
2. How technical solutions can influence standards/regulations?
3. Transparency and Web tracking
4. Bias and discrimination
5. Privacy by construction

For each of the sessions, the list of talks is presented in this section. For more detailed information about each talk, we refer to the talk abstracts in section 4.

Session 1: Advertisement and Real-Time Bidding

- Krishna Gummadi gave us insights on how to explain online advertisement in Facebook, trying to answer the questions such as what's the best way to explain online ads, what properties that tell us the explanations are good enough and how to automatise explanations?
- Nicolas Kourtellis illustrated how to bring more transparency into online ads by analysing how much advertisers pay to show ads.
- Robbert van Eijk discussed how to classify trackers and ad exchange companies and presented a demo showing how ad networks are connected to each other.
- Panel discussed the problem of utility from explanations (Netflix shows explanations to manipulate users to watch certain movies), what's the cost and benefit for the society from RTB, and how privacy leakage can be measured.

Session 2: How technical solutions can influence standards/regulations?

- Joe Calandrino presented the activities of FTC in investigation, education and research and in particular around topics related to web tracking.
- Lukasz Oleynik gave us insights on new browser APIs that allows to acquire user's behavior from sensors: he gave us history of battery status API leakage that was first revealed in 2011 and finally removed from Firefox in 2017.
- Damian Clifford explained us that the existing Web tracking technologies are covered by the current EU laws (GDPR), and require user consent.
- Timothy Libert presented his webXRy platform for online privacy measurements based on PhantomJS and motivated that chilling effects for OBA industry would be desirable.
- Panel was skeptic about enforceability of laws, due to many technical problems: how to distinguish functional cookies from tracking ones? How to detect other tracking methods that don't have any cookie at all? Participants recalled alternative methods, such as class actions and made analogy with environmental regulations.

Session 3: Transparency and Web tracking

- Günes Acar gave us insights on online crawlers and pointed that one can have different results when using a headless browser. He suspects that companies try to detect crawlers and probably try to evade transparency tools.
- Angel Cuevas presented his FDVT Facebook tool that approximately estimates the revenue users generate for Facebook based on the ads they receive and ads they click on.
- Costas Iordanou demonstrated two tools: \$sheriff that detects price discrimination and eyeWnder that provides information to the user about web advertisements.
- Steven Englehardt presented OpenWPM, a platform for online privacy measurement and gave us insights on statefull and stateless technologies detected on 1 million websites.
- Konrad Rieck discussed the new study on ultrasonic device tracking that is found in three commercial solutions.
- David Choffnes discussed longitudinal measurements of leaks in mobile apps and underlined that we need relations between the third parties for better measurements.
- Oleksii Starov gave us insights that browser extensions have more tracking powers than scripts and showed that some extensions leak browser and search history.
- Josep M. Pujol proposed to distinguish "parties that rely on tracking" from "trackers", because in some cases tracking is unintentional, an may be a by-product of a design choice.

Session 4: Bias and discrimination

- Michael Carl Tschantz discussed the problem of accountability for showing certain ads to certain groups of population. He questioned whether companies are accountable for discrimination of showing ads to a certain group or whether an advertiser should be accountable?
- Vincent Tobiana presented the first analysis on Facebook ethnic affinity and tried to evaluate whether FB users receive different ads based on their ethnic affinity.
- Aniko Hannak gave us insights on bias in job search (LinkedIn, Indeed, etc.) and professional communities (GitHub) and discussed whether online bias is different from offline, and started a discussion whether large-scale crawling may break terms of service of companies that are analysed.

- The panel engaged in discussion on offline vs. online behaviour, whether it should be copied, on discrimination/bias versus targeting. The panel tried to find the fundamental difference between “bias” and “discrimination”.

Session 5: Privacy by construction

- Diana Vlad-Calcic discussed European policy overview and presented existing EU programs to support research on online privacy and web transparency.
- Benoit Baudry gave us insights about the problem of code diversity that leads to efficient fingerprinting and discussed solutions for browser randomisation.
- Francis Some presented a tool that automatically prevents third-party tracking on the server side and protects users from unintentional tracking by third-party content.
- Steven Englehardt explained that sandboxing tracking scripts breaks many sites, and therefore we need other methods to detect and protect from web browser fingerprinting.

Break-out sessions

After the first day of the seminar, the organisers proposed several topics for break-out sessions, that were voted by the participants. As a result of the vote, three break-out sessions were chosen – they enabled participants to discuss selected topics in privacy research in smaller teams (around 10 people per team). The three topics were:

- Security and privacy trade off
- Research feasibility and validity
- Bluesky proposals

The purpose of the break-out sessions was to informally discuss the most important problems in privacy and transparency, state research challenges, and legal problems. As part of the break-out sessions, the teams identified the most relevant problems in the field and main challenges for the specific Web privacy and transparency area. The break-out sessions lasted 90 minutes and were held in parallel, on Wednesday afternoon. Each participant joined a break-out sessions of her choice. The last 30 minutes session on Thursday was used to report back the results of the three break-out sessions to the full group by means of an informal discussion. The reports of the three break-out sessions are summarized in section 5.

5-minute talks

To encourage participants to share their new ideas and results, we had one 5-minute session, where the following speakers presented their work:

- Update on the Data Transparency Lab by Nikolaos Laoutaris
- Browser Extension and Login-Leak Experiment by Nataliia Bielova
- Harvest documentary by Dave Choffnes
- Stealing browsing history using light sensors by Lukasz Olejnik
- Use of browser fingerprinting for security purposes by Gunes Acar
- How news media use Twitter to attract traffic? by Arnaud Legout

Conclusion

The seminar brought together 29 participants in computer science, law and policy, coming from companies and research institutions across Europe and the US. The seminar had a well-filled program, with 25 research talks, followed by 7 short panel discussions, and 6 5-minute talks. Online privacy and Web transparency is a broad research field, and hence a diverse set of recent research results were presented. They covered Web tracking technologies and transparency tools, behavioural advertisement, privacy protection mechanisms and technologies, bias and discrimination. The representatives from FTC and CNIL gave us insights on how to influence standards and regulations, while law and policy researchers explained how the current Web technologies are covered by the EU laws.

The seminar also featured three break-out sessions on Security and Privacy trade offs, Research feasibility and validity (in Web crawling and bias analysis), and Bluesy proposals. The goal of the break-out session was to discuss the most important open problems in the Web privacy and transparency research, and the special brainstorming session on bluesy proposals tried to propose completely new solutions and approaches to improve user's privacy on the Web – the summaries are documented in this report.

Finally, several new collaborations have been created as a result of this seminar, and at least one person has received a job offer due to the discussions that took place in Dagstuhl. A group of participants have organised a *Slack*¹ community to exchange news and ideas in the area.

4 Overview of Talks

4.1 How can we reconcile diversity and privacy?

Benoit Baudry (INRIA – Rennes, FR)

License © Creative Commons BY 3.0 Unported license
© Benoit Baudry

Main reference P. Laperdrix, W. Rudametkin, B. Baudry, “Beauty and the Beast: Diverting Modern Web Browsers to Build Unique Browser Fingerprints”, in Proc. of the IEEE Symposium on Security and Privacy (SP 2016), pp. 878–894, IEEE, 2016.

URL <https://doi.org/10.1109/SP.2016.57>

With this talk, I wish to trigger a discussion about possible ways to reconcile two values, which are fundamentally good, and yet, seem hardly compatible: diversity and privacy.

Diversity is good. Diversity is an essential property of natural systems, it is a characteristic sought in most human organizations and it is a moral value cherished by mankind. Diversity is key for the robustness of complex systems, it is essential to prevent monocultures and the risk of single points of failures. Diversity is also extremely beneficial in software systems for security and safety.

However, diversity threatens online privacy. Meanwhile, individualization is the counterpart of diversity. Consequently, multiple forms of diversity have become threats for the privacy of web users. For example, browser fingerprinting has emerged from the massive diversity of software and hardware components that users can assemble to set their environment. The diversity of online behaviors can be tracked, analyzed and learned to create filter bubbles.

¹ <https://onlineprivacy.slack.com>

Here, I will present our observations about the massive diversity of browser fingerprints. Then, I will discuss some solutions we have investigated to reconcole this with privacy, as well as some of the limitations we encountered.

4.2 Browser Extension and Login-Leak Experiment

Nataliia Bielova (INRIA Sophia Antipolis, FR), Claude Castelluccia, and Gábor Gy. Gulyás

License © Creative Commons BY 3.0 Unported license
 © Nataliia Bielova, Claude Castelluccia, and Gábor Gy. Gulyás
URL <http://extensions.inria.fr>

When a user browses the web, various trackers are spying on her online activities. Even though such trackers are invisible, they collect information about her, such as which pages she visits, which buttons clicks, and what text she types. This information is often used to show her targeted advertisements and may require her to pay a higher price during online shopping depending on the collected information.

Recent studies [6, 12, 5, 7, 3, 4, 1, 13, 10, 2] show that users can be identified based on their device characteristics: this tracking method is called *device fingerprinting*. Such unique collection of device’s properties, or a fingerprint, can often uniquely identify the user who visited the website. Usually, the fingerprint includes technical parameters like what browser and operating system a visitor is using, what timezone she is from, what fonts she has in her system, or what audio card her device supports. Beyond pure technical characteristics, which are not explicitly chosen by the user, a visitor can be also identified by more *behavioral* characteristics, such as the browser extensions she has installed and the websites where she has logged in. Detecting extensions and website logins can clearly make a significant contribution to fingerprinting.

In our new experiment at <http://extensions.inria.fr>, we demonstrate how websites can use *behavioral fingerprinting* and detect two aspects of user’s online behavior: web browser extensions and websites a user has logged in². Using a detection method based on Web Accessible Resources, [11], we are able to detect more than 13,000 Chrome browser extensions, including Adblock, Pinterest, and Ghostery. Our experiment demonstrates an important privacy concern: the more privacy extensions you install, the more identifiable you are!

References

- 1 E. Abgrall, Y. L. Traon, M. Monperrus, S. Gombault, M. Heiderich, and A. Ribault. XSS-FP: browser fingerprinting using HTML parser quirks. *CoRR*, abs/1211.4812, 2012.
- 2 G. Acar, C. Eubank, S. Englehardt, M. Juárez, A. Narayanan, and C. Díaz. The web never forgets: Persistent tracking mechanisms in the wild. In G. Ahn, M. Yung, and N. Li, editors, *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, Scottsdale, AZ, USA, November 3-7, 2014*, pages 674–689. ACM, 2014.
- 3 G. Acar, M. Juárez, N. Nikiforakis, C. Díaz, S. F. Gürses, F. Piessens, and B. Preneel. Fpde-detective: dusting the web for fingerprinters. In A. Sadeghi, V. D. Gligor, and M. Yung, editors, *2013 ACM SIGSAC Conference on Computer and Communications Security, CCS’13, Berlin, Germany, November 4-8, 2013*, pages 1129–1140. ACM, 2013.

² In the experiment, we collect user’s browser fingerprint, together with the browser extensions installed and a list of websites the user has logged in. We only collect anonymous data during the experiment (see our Privacy Policy at <https://extensions.inrialpes.fr/privacy.php>). We securely store the data on an Inria server, use it only for research purpose and not share it with anyone outside of Inria.

- 4 K. Boda, Á. M. Földes, G. G. Gulyás, and S. Imre. User tracking on the web via cross-browser fingerprinting. In P. Laud, editor, *Information Security Technology for Applications – 16th Nordic Conference on Secure IT Systems, NordSec 2011, Tallinn, Estonia, October 26-28, 2011, Revised Selected Papers*, volume 7161 of *Lecture Notes in Computer Science*, pages 31–46. Springer, 2011.
- 5 Y. Cao, S. Li, and E. Wijmans. (cross-)browser fingerprinting via os and hardware level features. In *24th Annual Network and Distributed System Security Symposium, NDSS 2017, San Diego, California, USA, 26 February – 1 March, 2017*, 2017. To Appear.
- 6 P. Eckersley. How unique is your web browser? In M. J. Atallah and N. J. Hopper, editors, *Privacy Enhancing Technologies, 10th International Symposium, PETS 2010, Berlin, Germany, July 21-23, 2010. Proceedings*, volume 6205 of *Lecture Notes in Computer Science*, pages 1–18. Springer, 2010.
- 7 S. Englehardt and A. Narayanan. Online tracking: A 1-million-site measurement and analysis. In E. R. Weippl, S. Katzenbeisser, C. Kruegel, A. C. Myers, and S. Halevi, editors, *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, October 24-28, 2016*, pages 1388–1401. ACM, 2016.
- 8 E. Homakov. Using content-security-policy for evil. <http://homakov.blogspot.fr/2014/01/using-content-security-policy-for-evil.html>, 2014.
- 9 R. Linus. Your social media fingerprint. <https://robinlinus.github.io/socialmedia-leak/>, 2016.
- 10 N. Nikiforakis, A. Kapravelos, W. Joosen, C. Kruegel, F. Piessens, and G. Vigna. Cookieless monster: Exploring the ecosystem of web-based device fingerprinting. In *2013 IEEE Symposium on Security and Privacy, SP 2013, Berkeley, CA, USA, May 19-22, 2013*, pages 541–555. IEEE Computer Society, 2013.
- 11 A. Sjösten, S. Van Acker, and A. Sabelfeld. Discovering browser extensions via web accessible resources. In *Proceedings of the Seventh ACM on Conference on Data and Application Security and Privacy, CODASPY '17*, pages 329–336, New York, NY, USA, 2017. ACM.
- 12 O. Starov and N. Nikiforakis. Extended tracking powers: Measuring the privacy diffusion enabled by browser extensions. In *Proceedings of the 26th International Conference on World Wide Web, WWW 2017, Perth, Australia, April 3 – 7, 2017*, 2017. To Appear.
- 13 N. Takei, T. Saito, K. Takasu, and T. Yamada. Web browser fingerprinting using only cascading style sheets. In L. Barolli, F. Xhafa, M. R. Ogiela, and L. Ogiela, editors, *10th International Conference on Broadband and Wireless Computing, Communication and Applications, BWCCA 2015, Krakow, Poland, November 4-6, 2015*, pages 57–63. IEEE Computer Society, 2015.

4.3 A Brief History of Mobile Privacy Leaks

David Choffnes (Northeastern University – Boston, US)

License © Creative Commons BY 3.0 Unported license
© David Choffnes

Joint work of Jingjing Ren, Ashwin Rao, Martina Lindorfer, Arnaud Legout, Christophe Leung, Christo Wilson
Main reference J. Ren, A. Rao, M. Lindorfer, A. Legout, D. R. Choffnes, “ReCon: Revealing and Controlling PII Leaks in Mobile Network Traffic”, in Proc. of the 14th Annual Int’l Conf. on Mobile Systems, Applications, and Services (MobiSys 2016), pp. 361–374, ACM, 2016.

URL <http://doi.acm.org/10.1145/2906388.2906392>

Mobile devices have unparalleled access to our daily lives, but give us little access to how they gather and share such information over time. In this talk, I summarized some research [1, 2] my group has been doing to address this problem, using a project we call ReCon. I will cover how we identify personally identifiable information exposed by mobile devices to other parties over the Internet, some of the interesting findings from analyzing hundreds of apps and users, how the nature of data collection is changing over time, and what are the implications for users, policymakers, and regulators. As part of ongoing work, we have been analyzing how data collection from apps changes over time, and what are the corresponding privacy and security implications for users.

References

- 1 Jingjing Ren, Ashwin Rao, Martina Lindorfer, Arnaud Legout, David Choffnes. *ReCon: Revealing and Controlling PII Leaks in Mobile Network Traffic*. In Proceedings of MobiSys (MobiSys ’16), Singapore, June 2016.
- 2 Christophe Leung, Jingjing Ren, David Choffnes, and Christo Wilson. *Should You Use the App for That? Comparing the Privacy Implications of Web- and App-based Online Services*. In Proceedings of the 16th ACM Internet Measurement Conference (IMC’16), Santa Monica, CA, November 2016.

4.4 FDVT: Facebook Data Valuation Tool

Ángel Cuevas Rumin (Univ. Carlos III – Madrid, ES), Rubén Cuevas Rumin, and José González Cabañas

License © Creative Commons BY 3.0 Unported license
© Ángel Cuevas Rumin, Rubén Cuevas Rumin, and José González Cabañas

Main reference J. González Cabañas, Á. Cuevas, R. Cuevas, “FDVT: Data Valuation Tool for Facebook Users”, in Proc. of the 2017 CHI Conf. on Human Factors in Computing Systems (CHI’17), pp. 3799–3809, ACM, 2017.

URL <https://doi.org/10.1145/3025453.3025903>

The OECD, the European Union and other public and private initiatives are claiming for the necessity of tools that create awareness among Internet users about the monetary value associated to the commercial exploitation of their online personal information. In this talk we present a recent developed tool addressing this challenge, the Data Valuation Tool for Facebook users (FDVT). The FDVT provides Facebook users with a personalised and real-time estimation of the revenue they generate for Facebook based on the ads they receive and the ads they click on while browsing in this social network. The FDVT has been implemented as a web browser extension available for Google Chrome and Firefox through fdvt.org. Currently, more than 5000 users have installed the FDVT.

4.5 Analyzing the Impact of Large Scale Online Tracking Measurement

Steven Englehardt (Princeton University, US) and Arvind Narayanan (Princeton University, US)

License © Creative Commons BY 3.0 Unported license
 © Steven Englehardt and Arvind Narayanan
Main reference S. Englehardt, A. Narayanan, “Online Tracking: A 1-million-site Measurement and Analysis”, in Proc. of the 2016 ACM SIGSAC Conf. on Computer and Communications Security (CCS 2016), pp. 1388–1401, ACM, 2016.
URL <https://doi.org/10.1145/2976749.2978313>

In this talk I’ll summarize the findings of a 1-million-site measurement of online tracking. I’ll share results of stateful (cookie-based) and stateless (fingerprinting-based) tracking measurements. I’ll examine the impact our measurements have had on the adoption of new techniques by trackers and implementation of defenses by browsers. By extracting lessons from this research and the ensuing impact, I’ll propose several new directions for the tracking measurement field.

4.6 Sandboxing Trackers with Resource Blocking Lists

Steven Englehardt (Princeton University, US)

License © Creative Commons BY 3.0 Unported license
 © Steven Englehardt
Joint work of Steven Englehardt, Tanvi Vyas, Eric Rescorla
URL https://bugzilla.mozilla.org/show_bug.cgi?id=1298207

The most effective consumer privacy tools available today work by blocking advertising and tracking resources from loading in a user’s browser. Resource blocking frequently leads to lost revenue and site breakage. We aim to design a tracking protection feature that minimizes breakage and preserves the revenue stream provided by advertisements while limiting the tracking capabilities of third-party content. We explore several possible client-side solutions and evaluate their effectiveness at preventing tracking and minimizing site breakage, and in their engineering feasibility.

We find that sandboxing tracking resources is either ineffective or infeasible depending on the configuration. Sandbox configurations which don’t provide Javascript sandboxing either continue to block a large percentage of resources or fail to significantly impact the level of tracking. Sandboxing javascript has the potential to both reduce the level of tracking and amount of breakage, but requires a design that’s heavily coupled to current implementation of the tracking scripts. Our work highlights several hurdles to user privacy caused by the “mash-up” nature of the modern web and raises questions to the extent that purely client-side solutions can protect users.

4.7 The Diffix Framework: Noise Revisited, Again

Paul Francis (MPI-SWS – Kaiserslautern, DE)

License  Creative Commons BY 3.0 Unported license
© Paul Francis

A longstanding problem is that of obtaining high-quality statistical data from a dataset about individuals while protecting the privacy of those individuals. We believe that we have broken new ground on this problem. Diffix is an approach to database anonymization that: has minimal distortion (noise with standard deviation of only 2 for counting queries); places no limit on the number of queries; has rich query semantics (most of SQL, many statistical operations); is easy to configure; can comfortably be called anonymous according to European privacy law. In this talk, I will give a short demo of a commercial-quality implementation of Diffix, and give a brief overview of the main concepts.

4.8 Explaining Online Ads

Krishna P. Gummadi (MPI-SWS – Saarbrücken, DE)

License  Creative Commons BY 3.0 Unported license
© Krishna P. Gummadi

Abstract: Online service providers like Netflix and Facebook are beginning to offer automated explanations to their consumers, e.g., Facebook’s “why am I seeing this ad?” feature. While these explanations have the potential to make the inner working of the services transparent to consumers, they raise several questions. Specifically, (i) What are the types of explanations that are being provided by these services today? (ii) What properties or standards should explanations offered to consumers satisfy (in order to be useful and / or meaningful)? (iii) How can we (automatically) construct explanations that meet specific standards? In an ongoing effort, we are attempting to answer these questions in the context of Facebook targeted advertisements. In this presentation, I will discuss our preliminary findings.

4.9 New Faces Of Bias in Online Labor Markets

Aniko Hannak (Central European University – Budapest, HU)

License  Creative Commons BY 3.0 Unported license
© Aniko Hannak

The internet is fundamentally changing the labor economy. Millions of people use sites like LinkedIn, Upwork or Dribbble to find employment. These services are often driven by algorithms that rate, sort, recommend, and match workers and employers. In theory, many of the mechanisms that cause discrimination in traditional labor markets – cognitive bias, network homophily, statistical discrimination – should be absent from online markets. However, recent studies indicate that these mechanisms do transfer to online platforms, where they may be exacerbated by seemingly harmless design choices.

In this talk I will investigate three techniques that online platforms use to match users with content: social network algorithms, search algorithms and public review systems. Specifically, I present case studies of 6 different employment platforms, using large scale user data from

the employers perspective. I show that biases known from traditional labor markets are indeed present in online platforms, although they manifest in new ways. First, I present results that focus on the visibility of users, which directly impacts the chances of being selected for a job or selling a product. I find that women often receive lower visibility either due to their ranking in the sites' search interface, or their positions in the underlying social network. Furthermore, I investigate social feedback and other success measures found on user profiles, another important factor in hiring decisions. Overall, my investigations show that demographic features are often correlated with the attention and the social feedback workers and employees receive. Exploring these new forms of inequalities, understanding where social biases enter systems and which mechanisms reinforce them, can be crucial for developing mitigation strategies.

4.10 Web transparency tools demo

Costas Iordanou (Telefónica Research – Barcelona, ES)

License © Creative Commons BY 3.0 Unported license
© Costas Iordanou

The first tool is related to online price discrimination called \$heriff. The tool allows internet users to check the prices of different products and services available on the web and look for evidence of price discrimination. The tool is available for two web browsers. Google chrome version is available at <https://chrome.google.com/webstore/detail/heriffv2/emobhicogmenmngifjhjbfiohhjijl> and Mozilla Firefox version is available at https://addons.mozilla.org/en-US/firefox/addon/sheriff_v2/.

The second tool is related to web advertisements called eyeWnder. The tool is following the crowdsourced approach and provides information to the user regarding web advertisements in real time. The tool is also able to visualise the web browsing history of the user annotated with the interest categories assigned by advertisers. The tool is available at <http://www.eyewnder.com/>.

4.11 Transparency in the era of programmatic, real-time bidding advertising and Cookie Synchronization

Nicolas Kourtellis (Telefónica Research – Barcelona, ES)

License © Creative Commons BY 3.0 Unported license
© Nicolas Kourtellis

Joint work of Panagiotis Papadopoulos, Pablo Rodriguez, Nikolaos Laoutaris, Evangelos Markatos

Main reference P. Papadopoulos, N. Kourtellis, P. Rodriguez, N. Laoutaris, "If you are not paying for it, you are the product: How much do advertisers pay to reach you?", arXiv:1701.07058 [cs.GT], 2017.

URL <https://arxiv.org/pdf/1701.07058.pdf>

This presentation was divided into two parts. The first part proposed a methodology on evaluating users' data from the advertising ecosystem side using RTB ads. The second part presented a longitudinal study on the cookie synchronization process used by advertisers and trackers to share data of users.

Part 1: Online advertising is progressively moving towards a programmatic model in which ads are matched to actual interests of individuals collected as they browse the web. Letting the huge debate around privacy aside, a very important question in this area, for which little is known, is: How much do advertisers pay for an individual's personal data?

In this study[1], we develop a first of its kind methodology for computing exactly that – the price paid for one’s privacy – and we do that in real time. Our approach is based on tapping on the Real Time Bidding (RTB) protocol to collect cleartext and encrypted prices for winning bids paid by advertisers in order to place targeted ads. Our main technical contribution is a method for tallying winning bids even when they are encrypted. We achieve this by training a model using as ground truth prices obtained by running our own “probe” ad-campaigns. We implement our methodology through a browser addon and a back-end server that provides it with fresh models for encrypted bids. We validate our methodology using a one year long trace of 1600 mobile users and demonstrate that it can estimate a user’s advertising worth with more than 82% accuracy.

Part 2: Cookies are still the dominant user targeting mechanism on the web. Third-party online companies maintain large user data stores for all unique users encountered and identified using anonymous cookies. In order to identify and track users across different publishers and through time, third parties connect with each other and synchronise their cookies. Consequently, Cookie Synchronization (CSync) is one of the de facto tracking mechanisms of modern web. CSync facilitates an information sharing channel between third parties that may or may not have direct access to the website the user visits. With CSync they can not only reconstruct the browsing history of a user by bypassing the same origin policy, but also merge the user data they own, in the background. In this paper [2], we perform a first to our knowledge longitudinal study of CSync in the wild, using a year-long dataset that includes browsing activity from 1600 real users. Through our study, we aim to understand the protocol’s characteristics, growth and the flow graph of user personal information while being leaked to third parties. Our results show that 97% of the regular web users are exposed to Cookie Synchronization: most of them within the first week of their browsing. Our experiments also suggest that the average user is exposed to 63 distinct cookie synchronization events. This implies that despite the fact that all 63 of the trackers thought that they were tracking 63 different users (identities), after cookie synchronization, all 63 identities can point back to the same single user.

References

- 1 Panagiotis Papadopoulos, Nicolas Kourtellis, Pablo Rodriguez Rodriguez, Nikolaos Laoutaris. If you are not paying for it, you are the product: How much do advertisers pay for your personal data? <https://arxiv.org/pdf/1701.07058.pdf>
- 2 Panagiotis Papadopoulos, Nicolas Kourtellis, Evangelos Markatos. Syncing the cookie monster: A longitudinal study of cookie syncing and its impact on user privacy.

4.12 How news media use Twitter to attract traffic?

Arnaud Legout (INRIA Sophia Antipolis, FR)

License  Creative Commons BY 3.0 Unported license

© Arnaud Legout

Joint work of Maksym Gabielkov, Arthi Ramachandran, Arnaud Legout, Augustin Chaintreau
Main reference M. Gabielkov, A. Ramachandran, A. Legout, A. Chaintreau, “Social Clicks: What and Who Gets Read on Twitter?”, in Proc. of the 2016 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Science, pp. 179–192, ACM, 2016.
URL <https://doi.org/10.1145/2964791.2901462>

Online news domains increasingly rely on social media to drive traffic to their website. Yet we know surprisingly little about how social media conversation mentioning an online article actually generates a click to it. Posting behaviors, in contrast, have been fully or partially available and scrutinized over the years. While this has led to multiple assumptions on the

diffusion of information, each were designed or validated while ignoring this important step.

We present a large scale, validated and reproducible study of social clicks – that is also the first data of its kind – gathering a month of web visits to online resources that are located in 5 leading news domains and that are mentioned in the third largest social media by web referral (Twitter). Our dataset amounts to 2.8 million posts, together responsible for 75 billion potential views on this social media, and 9.6 million actual clicks to 59,088 unique resources. We design a reproducible methodology, carefully corrected its biases, enabling data sharing, future collection and validation. As we prove, properties of clicks and social media Click-Through-Rates (CTR) impact multiple aspects of information diffusion, all previously unknown. Secondary resources, that are not promoted through headlines and are responsible for the long tail of content popularity, generate more clicks both in absolute and relative terms. Social media attention is actually long-lived, in contrast with temporal evolution estimated from posts or impressions. The actual influence of an intermediary or a resource is poorly predicted by their posting behavior, but we show how that prediction can be made more precise.

4.13 Surveillance as a Regulatory Model

Timothy Libert (University of Pennsylvania – Philadelphia, US)

License  Creative Commons BY 3.0 Unported license
© Timothy Libert

Bruce Schneier has noted that “surveillance is the business model of the internet”. One must therefore ask, what is the regulatory model of the Internet? Currently, the dominant regulatory stances in regards to online tracking rely on industry self-regulatory guidelines which are largely focused on the paradigm of “notice and choice”. This paradigm is a vastly weakened version of the normatively-grounded Fair Information Practice Principles, is not practiced by industry in any meaningful way, and fails to protect user privacy in sensitive areas such as personal health. To the degree to which industry is regulated, it is often done through a mixture of fines, bad publicity, and user attempts at blocking tracking mechanisms. These scatter-shot approaches have failed to place significant limits on the spread of online tracking.

This talk proposes a new possibility: surveillance as a regulatory model. With many researchers now conducting large-scale censuses of web tracking practices, it is possible to constantly monitor the activities of companies tracking users on the web and provide regulators to both daily and historical reports on the state of tracking. If and when companies engage in deceptive practices it should be possible to spot such practices quickly, apply fines based on the nature of the deception, and multiply fines based on the number of sites affected over time.

While marquee names such as Facebook and Google receive the most media and regulatory attention, such a large-scale approach could be applied to a much larger variety of companies and result in a significantly expanded pool of companies who would be under constant scrutiny. This would facilitate moving from the current model of large fines levied infrequently, to smaller-scale fines levied on a regular basis.

In short, if the business model of the internet is surveillance, the regulatory model for online privacy should follow suit. The purpose of this talk is to be fairly short, and provoke discussion towards the technical requirements of such an approach, the receptiveness of regulators, and overall feasibility.

4.14 Modern Web Privacy: standards, implementations, deployments

Lukasz Olejnik (University College London, GB)

License © Creative Commons BY 3.0 Unported license
© Lukasz Olejnik

Joint work of Lukasz Olejnik, Steven Englehardt, Arvind Narayanan

Main reference Battery Status Not Included: Assessing Privacy in Web Standards, International Privacy Engineering Workshop, 2017

URL <https://blog.lukaszolejnik.com/battery-status-not-included-assessing-privacy-in-w3c-web-standards/>

For majority of users, web browser is the most important computer application. Increasingly complex, exciting and rich, features are standardized by W3C and implemented in web browsers on a normal basis. New browser features introduce interesting privacy challenges for standardization, research and development.

The importance of privacy engineering has become increasingly apparent to standardizers (W3C), implementers (browser vendors) and web developers. Standardized guidelines for privacy assessments exist and are continuously improved in response to research and experience. The ever increasing complexity and richness of the web ecosystem necessitates a continual reevaluation of privacy assessments.

I provide a case study analysis of the evolution of the W3C Battery Status API, and discuss the specification and related implementations through the previous disclosure of several privacy vulnerabilities. We examine how implementations change and adoption shifts in response to these vulnerabilities. To provide context we present new measurement results and usage statistics for the Battery Status API, showing that there is a heavy fingerprinting use on the modern web. I will mention a list of methodologies, design patterns and recommendations to improve the privacy engineering process during the drafting of specifications and preparing of implementations.

I point out the strong features such as new communication channels or access to low-level sensors that the Modern Web is starting to offer and discuss a number of possible consequences to privacy. Are we ready for the web as a part of the Internet of Things and the ensuing challenges? I discuss how to measure privacy in the new web paradigms.

4.15 Ultrasonic Device Tracking for Fun and Profit

Konrad Rieck (TU Braunschweig, DE)

License © Creative Commons BY 3.0 Unported license
© Konrad Rieck

Joint work of Christian Wressnegger, Erwin Quiring, Daniel Arp, Konrad Rieck

Main reference D. Arp, E. Quiring, C. Wressnegger, K. Rieck, “Privacy Threats through Ultrasonic Side Channels on Mobile Devices”, in Proc. of the 2nd IEEE European Symposium on Security and Privacy (EuroS&P), pp. 35–47, IEEE, 2017.

URL <https://doi.org/10.1109/EuroSP.2017.33>

Device tracking is a serious threat to the privacy of users. Recently, several companies have started to use ultrasound for tracking mobile devices. To this end, ultrasonic markers are embedded in an audio signal and unnoticeably tracked using the microphone of mobile devices. This side channel allows an adversary to identify a user’s current location, spy on her TV viewing habits or link together her different mobile devices. In this talk, we explore the capabilities and the current prevalence of this new tracking technique based on three commercial solutions. We discuss detection and mitigation approaches, and present case studies on Web and TV media.

4.16 Control What You Include ! Server-Side Protection against Third Party Web Tracking

Dolière Francis Some (Université Côte d’Azur, Inria, FR), Nataliia Bielova (Université Côte d’Azur, Inria, FR), and Tamara Rezk

License © Creative Commons BY 3.0 Unported license
© Dolière Francis Some, Nataliia Bielova, and Tamara Rezk

Third party tracking is the practice by which third parties recognize users across different websites as they browse the web. Recent studies show that more than 90% of Alexa top 500 websites[1] contain third party content that is tracking its users across the web. Website developers often need to include third party content in order to provide basic functionality. However, when a developer includes a third party content, she cannot know whether the third party contains tracking mechanisms. If a website developer wants to protect her users from being tracked, the only solution is to exclude any third-party content, thus trading functionality for privacy. We describe and implement a privacy-preserving web architecture that gives website developers a control over third party tracking: developers are able to include functionally useful third party content, the same time ensuring that the end users are not tracked by the third parties.

References

- 1 Franziska Roesner and Tadayoshi Kohno and David Wetherall. Detecting and Defending Against Third-Party Tracking on the Web. Proc. of the 9th NSDI. pages 155-168, 2012.

4.17 Understanding the privacy risks of browser extensions

Oleksii Starov (Stony Brook University, US) and Nick Nikiforakis (Stony Brook University, US)

License © Creative Commons BY 3.0 Unported license
© Oleksii Starov and Nick Nikiforakis

Main reference O. Starov, N. Nikiforakis, “Extended Tracking Powers: Measuring the Privacy Diffusion Enabled by Browser Extensions,” in Proceedings of the 26th International Conference on World Wide Web (WWW 2017), pp. 1481–1490, ACM, 2017.

URL <https://doi.org/10.1145/3038912.3052596>

Main reference O. Starov, N. Nikiforakis, “XHOUND: Quantifying the Fingerprintability of Browser Extensions,” in Proc. of the 2nd IEEE European Symposium on Security and Privacy (EuroS&P), pp. 941–956, IEEE, 2017.

URL <https://doi.org/10.1109/SP.2017.18>

The increased popularity of browser extensions can have serious negative consequences to a user’s online privacy. In this talk, we describe our findings regarding extension privacy from two recent studies. First, we will describe the threat of browser extensions leaking a user’s browser history, on purpose or accidentally, to third parties. In our recent work (published at WWW 2017), we built a dynamic analysis system and discovered that above 6% of browser extensions leak sensitive information to one or more third parties. Second, we discuss the issue of extension fingerprinting. In our recent paper at IEEE S&P 2017, we showed that it is possible to fingerprint browser extensions automatically and at scale. Specifically, we present the details of our dynamic analysis system (XHound) which stimulates browser extensions in order to identify their on-page DOM side-effects which can be abused by a web page to infer an extension’s presence or absence. Our tool was able to automatically extract fingerprintable vectors from 9.2% of the top 10K Chrome extensions that run on any page, and 16.6% of

extensions that run on popular domains like Google, YouTube, and Facebook. We explain why the threat of extension fingerprinting is more serious than traditional fingerprinting and discuss possible solutions, both for data-leaking extensions as well as for protecting users against extension fingerprinting.

4.18 Accountability for Privacy and Discrimination

Michael Carl Tschantz (ICSI – Berkeley, US)

License  Creative Commons BY 3.0 Unported license
© Michael Carl Tschantz

My prior work, Datta et al. [1], presented AdFisher, an automated tool that explores how user behaviors, Google’s ads, and its Ad Settings interact. AdFisher found that setting a simulated user’s gender to female resulted it in getting fewer instances of an ad related to high paying jobs than setting the gender to male. It also found that visiting webpages associated with substance abuse changed the ads shown but not the settings page. However, we cannot determine who or what caused these findings due to our limited visibility into the ad ecosystem, which includes interactions between Google, advertisers, websites, and users.

I believe Google owes us an account of why such discrimination occurred and that such accounts form the basis of “accountability”, as opposed to responsibility or punishment. However, providing such accountability has its difficulties.

One difficulty is interpreting vague policies that refer to what information is “about” or the “subject” of data. As big data analytics find increasingly unexpected associations between unexpected features, it is becoming increasingly difficult to delimit data along these lines. Motivated by this difficulty, we are exploring ways of identifying data for protection based upon where the data comes from instead of what it is about.

Another difficulty is the need to account for responsibility when multiple actors interact to produce a result. Datta et al. [2] developed a theory of causal responsibility, Quantitative Input Influence, that assigns responsibility to the participants in a blackbox system based upon cooperative game theory. I will discuss the difficulties, both practical and conceptual, with applying this theory to large systems.

References

- 1 Amit Datta, Michael Carl Tschantz, and Anupam Datta. Automated Experiments on Ad Privacy Settings: A Tale of Opacity, Choice, and Discrimination. Proceedings on Privacy Enhancing Technologies (PoPETs), 2015.
- 2 Anupam Datta, Shayak Sen, Yair Zick. Algorithmic Transparency via Quantitative Input Influence: Theory and Experiments with Learning Systems. Proceedings of 37th IEEE Symposium on Security and Privacy, 2016.

4.19 Tracking Classification, a cluster analysis approach to identify missing values

Robbert van Eijk (Leiden University Dual PhD Centre, NL)

- License** © Creative Commons BY 3.0 Unported license
© Robbert van Eijk
- Main reference** R. van Eijk, “Tracking Classification: A Cluster Analysis Approach to Identify Missing Values” (Presentations Slides).
URL <http://dx.doi.org/10.2139/ssrn.2954946>
- Main reference** R. van Eijk, “A Brief Introduction to Real Time Bidding (RTB)”
URL <https://ssrn.com/abstract=2856941>
- Main reference** R. van Eijk, “Web Tracking Detection System (TDS): An Effective Strategy to Reduce Systematic Monitoring and Profiling of User Habits Across Websites”
URL <https://ssrn.com/abstract=2696690>

Presentation part of Dagstuhl Seminar ‘Online Privacy and Web Transparency’. This talk is motivated by the changing legal framework in the European Union, i.e., the General Data Protection Regulation and the (draft) ePrivacy Regulation. Transparency and fairness are key elements when it comes to informed consent. The problem with real-time bidding however is that many third-parties are unknown at the time of asking consent from a user. The talk will explore a methodology for classification of third-parties by looking at information flows between real-time bidding networks using cluster analysis. Interactive HTML-widgets will be used to guide the exploration. The difference in effectiveness of ad-blockers in the EU versus the US will also be discussed.

The presentation leads to two conclusions against the background of regulatory changes in the EU (GDPR and ePrivacy Regulation). Conclusion on technology: with cluster analysis on a standard referer graph it is possible to differentiate between actors and their role in interconnected RTB systems. The following cluster analysis approaches have proven to be useful: (a) node betweenness (b) cluster-edge betweenness (c) eigenvector centrality. Conclusion on policy: cookie enforcement can be effective as the example in Slovenia shows.

Keywords: web tracking, big data, privacy, data protection

5 Break-out sessions

5.1 Security and privacy trade off

This break-out session focused on a stateless web tracking technique called browser fingerprinting. The main question on the trade off between security and privacy was: fingerprinting is used to track users, but in the same time is it very useful for security – for example, an attacker who stole user’s credentials would not be able to access the system since his fingerprint is different than those of the legitimate user.

Law and policy researchers pointed out that the **usage of fingerprinting can also be bounded by law**: it may be used only under the conditions of necessity and proportionality. Moreover, with upcoming EU laws, such as GDPR and ePrivacy, any non-functional tracking will require user consent, while functional tracking (like the one for security) will not require consent. The group has also concluded that personalised ads are not considered necessary for the provision of the service.

From a technical perspective, the group discussed and compared using **fingerprinting as a security measure** versus using security questions. By using security questions, an attacker who knows enough information about the user could break in. Some researchers concluded

that negative authentication may be a good option: instead of whitelisting “good” (known) fingerprints, the system may blacklist “bad” fingerprints thus blocking bot fingerprinting.

Researchers have agreed that it’s **difficult to detect fingerprinting** since there are no “bad” APIs. The state-of-the-art literature doesn’t propose 100% accurate detection of fingerprinting. Fingerprinting is also used in malware operations when an attacker finds a vulnerable target because of the browser’s fingerprint. Therefore proposing solutions to detect and/or prevent fingerprinting is an important goal not only for privacy reasons, but also for security of web applications.

It was pointed out that the usage of **fingerprinting for authentication** may be profitable, but in this case two fingerprints must be collected (one to use, another to store) to beat replay attacks. It can be considered as “functional tracking” if it is ensured that fingerprinting is always used together with authentication. However, if fingerprinting is used on other landing pages, then it’s not functional.

Finally, researchers have raised a question: **Is fingerprinting a lost cause?** Recent communications with Mozilla underline that the company will not expand the APIs of the Firefox browser without purpose, however they do not believe that they can have an unfingerprintable browser. Some preliminary research shows that fingerprinting may not work very well in the regular population – there are no studies that show how fingerprintable are regular users, who are not computer science experts. Another reason for these results is that regular population uses more mobile devices, which are less fingerprintable and less unique than the desktop devices.

We would like to thank Nick Nikiforakis for taking notes during this break-out session, and presenting the break-out results to the full group.

5.2 Research feasibility and validity

This break-out session discussed the ethical problems of doing research in the field of privacy and transparency. Researchers have agreed that repeatability has become a standard way of performing research in this field, **most of the results in online privacy and transparency are reproducible**, and researchers share voluntarily their findings and data with the community.

Many of the researchers in this session have experience in interacting with companies, while trying to obtain some data from them. However, the companies are precocious and not willing to share data, and the main reason is that researchers have shown that any **data anonymisation algorithm is not providing 100% guarantee**, and therefore companies don’t take the risk of revealing user’s data.

Researcher discussed that our research may **impact users**, and therefore requires IRB or another ethical committee to validate the study that involves users. Research that involves automatic crawling of web services may also **impact companies**, for example fake profiles may generate revenues for advertisers, and therefore companies may lose some amount of money due to the automatic crawls. We, as researchers, should at least estimate how much our experiments may cost to the companies.

The participants also discussed the differences between **legal and ethical issues** related to transparency research. Creation of fake profiles may violate terms of service, and causes a legal issue. We may also ask a company, whose service we are evaluating, for a permission to perform our study, and thus behave ethically. However, this may influence the practices of the company, and all “unethical”, “biased” or “discriminative” behaviour may be removed before

we even start our analysis. For automatic crawling, there is already a convention “robots.txt” that describes how a service may be crawled. Violating this convention is considered an ethical violation, and not a legal one.

An interesting ethical question raised during the discussion was: **Is it ok for us, researchers, to analyse data that is collected unethically/illegally by someone else?** Is there a difference on using data collected unethically from leaked data (when we don't know how it was collected)? Even the law researcher in the group could not answer this question immediately. More concrete question raised was: Should we delete tweets from our tweeter dataset if they are deleted in twitter?

We would like to thank Nataliia Bielova for taking notes during this break-out session, and presenting the break-out results to the full group.

5.3 Bluesky proposals

The break-out session started from a provocative question in order to initiate lively discussions: **How would you design a clean slate web, both technology and business model, to do all its doing now and solve all privacy problems?**

The participants discussed whether **we should try to mimic the offline world**, despite its imperfections, or should we go for utopia, for an online world that is even more private than the offline we have now. To answer this question, the group mentioned “Evolution vs Revolution” by T. Khun, and raised another question: How close are we to a revolution? If the conditions for a revolution are near, do people have the tools to realise it? Revolutions are usually local not global. Can you revolt against the tool that has been used for revolutions?

Researchers brainstormed on **a clean state web**: Shall anti-monopoly by design be a property of a clean slate web? If you we to break a current monopoly through regulator intervention, how would we do it? They discussed micro payments, and whether they should be a basic ingredient of the clean slate design. Music industry did it. We would never have iTunes and Spotify without Napster and BitTorrent. AdBlockers seem to play the role of Napster forcing the AdTech to sit on the table and discuss alternatives. If not micro-payments, what about an all you can eat subscription on the browser. Can we feed all the publishers with 5 euro per month a la Spotify?

Finally, should we have a data levy imposed on data processors? Similar to what is being imposed on storage devices or tobacco to compensate for harm done? An ideal solution would be clean slate design without middlemen and with the user in the loop.

We would like to thank Nikolas Laoutaris for taking notes during this break-out session, and presenting the break-out results to the full group.

Participants

- Günes Acar
KU Leuven, BE
- Benoit Baudry
INRIA – Rennes, FR
- Nataliia Bielova
INRIA Sophia Antipolis, FR
- Joe Calandrino
Federal Trade Commission –
Washington, US
- David Choffnes
Northeastern University –
Boston, US
- Damian Clifford
KU Leuven, BE
- Angel Cuevas Rumin
Univ. Carlos III – Madrid, ES
- Steven Englehardt
Princeton University, US
- Paul Francis
MPI-SWS – Kaiserslautern, DE
- Krishna P. Gummadi
MPI-SWS – Saarbrücken, DE
- Aniko Hannak
Central European University –
Budapest, HU
- Costas Iordanou
Telefónica Research –
Barcelona, ES
- Nicolas Kourtellis
Telefónica Research –
Barcelona, ES
- Balachander Krishnamurthy
AT&T Labs Research –
New York, US
- Nikolaos Laoutaris
Telefónica Research –
Barcelona, ES
- Arnaud Legout
INRIA Sophia Antipolis, FR
- Timothy Libert
University of Pennsylvania –
Philadelphia, US
- Arvind Narayanan
Princeton University, US
- Nick Nikiforakis
Stony Brook University, US
- Lukasz Olejnik
University College London, GB
- Josep M. Pujol
Cliqz – München, DE
- Konrad Rieck
TU Braunschweig, DE
- Dolière Francis Somé
INRIA Sophia Antipolis, FR
- Oleksii Starov
Stony Brook University, US
- Vincent Toubiana
CNIL – Paris, FR
- Michael Carl Tschantz
ICSI – Berkeley, US
- Peggy Valcke
KU Leuven, BE
- Robbert van Eijk
Leiden University Dual PhD
Centre, NL
- Diana Vlad-Calcic
European Commission –
Brussels, BE



Report from Dagstuhl Seminar 17171

Computational Geometry

Edited by

Otfried Cheong¹, Anne Driemel², and Jeff Erickson³

1 KAIST – Daejeon, KR, otfried@kaist.edu

2 TU Eindhoven, NL, a.driemel@tue.nl

3 University of Illinois – Urbana-Champaign, US, jeffe@illinois.edu

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 17171 “Computational Geometry”. The seminar was held from 23rd to 28th April 2017 and 47 participants from various countries attended it. Recent advances in computational geometry were presented and new challenges were identified. The report collects the abstracts of talks and open problems presented in the seminar.

Seminar April 23–28, 2017 – <http://www.dagstuhl.de/17171>

1998 ACM Subject Classification F.2 Analysis of Algorithms and Problem Complexity, G.2 Discrete Mathematics, G.4 Mathematical Software

Keywords and phrases algorithms, applications, combinatorics, complexity, geometric computing, high-dimensional computational geometry, implementation, monitoring and shape data

Digital Object Identifier 10.4230/DagRep.7.4.107

Edited in cooperation with Marcel Roeloffzen

1 Executive Summary

Anne Driemel

Otfried Cheong

Jeff Erickson

License © Creative Commons BY 3.0 Unported license
© Anne Driemel, Otfried Cheong, and Jeff Erickson

Computational Geometry

Computational geometry is concerned with the design, analysis, and implementation of algorithms for geometric and topological problems, which arise naturally in a wide range of areas, including computer graphics, CAD, robotics, computer vision, image processing, spatial databases, GIS, molecular biology, sensor networks, machine learning, data mining, scientific computing, theoretical computer science, and pure mathematics. Computational geometry is a vibrant and mature field of research, with several dedicated international conferences and journals and strong intellectual connections with other computing and mathematics disciplines.

Seminar Topics

The emphasis of the seminar was on presenting recent developments in computational geometry, as well as identifying new challenges, opportunities, and connections to other fields of computing. In addition to the usual broad coverage of new results in the field, the



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

Computational Geometry, *Dagstuhl Reports*, Vol. 7, Issue 04, pp. 107–127

Editors: Otfried Cheong, Anne Driemel, and Jeff Erickson



Dagstuhl Reports

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

seminar included broad survey talks on monitoring and shape data and on high-dimensional geometric computing, two focus areas that have seen exciting recent progress and that present numerous opportunities for further cross-disciplinary impact.

Computational geometry for monitoring and shape data

The combination of movement and geometry has always been an important topic in computational geometry, initially motivated by robotics and resulting in the study of kinetic data structures. With the advent of widely available location tracking technologies such as GPS sensors, trajectory analysis has become a topic in itself, which has connections to other classical topics in computational geometry such as shape analysis. Still, efficient technologies to perform the most basic operations are lacking. We need data structures supporting similarity queries on trajectory data and geometric clustering algorithms that can handle the infinite-dimensional geometry inherent in the data. A related type of data, namely time series data, has not received much attention in the computational geometry community, despite its universality and its close relation to trajectory data. Shedding light on the interconnections of these topics will promote new results in the field which will address these timely questions.

Computing in high-dimensional and infinite-dimensional spaces

The famous “curse of dimensionality” prevents exact geometric computations in high-dimensional spaces. Most of the data in science and engineering is high-dimensional, rendering classical geometric techniques, such as the sweepline approach, insufficient. One way to address this issue is to use sparsity, but it is not always easy to find a sparse representation of the data. The search of the most efficient representation and how to exploit this representation leads to dimension-reduction techniques, metric embeddings, and approximation algorithms. This line of research has strong ties to machine learning and discrete mathematics as well as computational geometry.

Participants

Dagstuhl seminars on computational geometry have been organized in a two year rhythm since a start in 1990. They have been extremely successful both in disseminating the knowledge and identifying new research thrusts. Many major results in computational geometry were first presented in Dagstuhl seminars, and interactions among the participants at these seminars have led to numerous new results in the field. These seminars have also played an important role in bringing researchers together, fostering collaboration, and exposing young talent to the seniors of the field. They have arguably been the most influential meetings in the field of computational geometry. The organizers held a lottery for the third time this year; the lottery allows to create space to invite younger researchers, rejuvenating the seminar, while keeping a large group of senior and well-known scholars involved. The seminar has now a more balanced attendance in terms of seniority and gender than in the past. This year, 47 researchers from various countries and continents attended the seminar, showing the strong interest of the community for this event. The feedback from participants was very positive. No other meeting in our field allows young researchers to meet with, get to know, and work with well-known and senior scholars to the extent possible at the Dagstuhl Seminar. We warmly thank the scientific, administrative and technical staff at Schloss Dagstuhl! Dagstuhl allows people to really meet and socialize, providing them with a wonderful atmosphere of a unique closed and pleasant environment, which is highly beneficial to interactions. Therefore, Schloss Dagstuhl itself is a great strength of the seminar.

2 Table of Contents

Executive Summary

Anne Driemel, Otfried Cheong, and Jeff Erickson 107

Overview of Talks

Graph Embedding While Preserving Pairwise Distances <i>Tetsuo Asano</i>	111
Improved Time-Space Trade-offs for Computing Voronoi Diagrams <i>Bahareh Banyassady, Matias Korman, Wolfgang Mulzer, Marcel Roeloffzen, Paul Seiferth, Yannik Stein, and André van Renssen</i>	111
The Morse theory of Čech and Delaunay Complexes <i>Ulrich Bauer</i>	112
Conditional Lower Bounds for Similarity Measures on Curves and Strings <i>Karl Bringmann</i>	112
Untangling and Unwinding Curves <i>Jeff Erickson</i>	112
A Nearly Quadratic Bound for the Decision Tree Complexity of k-SUM <i>Esther Ezra</i>	113
Computing Optimal Flight Patterns with Minimum Turn Cost <i>Sándor Fekete</i>	113
Faster Algorithms for the Geometric Transportation Problem <i>Kyle Jordan Fox</i>	114
Don't Collect too Much: Geometric Approaches for Protecting Location and Trajectory Privacy <i>Jie Gao</i>	114
Spatio-Temporal Analysis of Team Sports and Computational Geometry <i>Joachim Gudmundsson</i>	115
High-dimensional Theta Numbers <i>Anna Gundert</i>	115
Minkowski Sums of Polyhedra with Holes <i>Dan Halperin</i>	115
Self-Aligning Shapes <i>David G. Kirkpatrick</i>	116
On the Computational Bottleneck in Sampling-Based Robot Motion Planning <i>Michal Kleinbort</i>	116
Computing Wave Impact in Self-Organised Mussel Beds <i>Maarten Löffler</i>	116
The Art Gallery Problem is ETR-complete <i>Tillmann Miltzow</i>	117
TSP With Locational Uncertainty: The Adversarial Model <i>Joseph S. B. Mitchell</i>	117

Dynamic Planar Voronoi Diagrams for General Distance Functions <i>Wolfgang Mulzer, Haim Kaplan, Liam Roditty, Paul Seiferth, and Micha Sharir</i> . . .	118
Maximizing Volume Subject to Combinatorial Constraints <i>Aleksandar Nikolov</i>	118
Colorful simplicial depth, Minkowski sums, and generalized Gale transforms <i>Zuzana Patáková</i>	119
Approximate Range Maximization <i>Jeff M. Phillips</i>	119
A Treehouse with Custom Windows: Minimum Distortion Embeddings into Bounded Treewidth Graphs <i>Benjamin Raichel</i>	120
Spatial data processing with SAP HANA <i>Alejandro Salinger</i>	120
A Framework for Algorithm Stability <i>Kevin Verbeek</i>	121
Towards Spectral Sparsification of Simplicial Complexes based on Generalized Effective Resistance <i>Bei Wang</i>	121
Open problems	
Nash equilibria of spanner games <i>Mohammad Ali Abam</i>	122
Negative cycles on surface embedded graphs <i>Jeff Erickson</i>	122
Maximum-Area Triangle in a Convex Polygon <i>Maarten Löffler</i>	123
Is G area-universal? <i>Tillmann Miltzow</i>	123
Finding small polygons with a fixed number of triangulations; simple polygons with a unique triangulation <i>Joseph S. B. Mitchell</i>	123
Minimum k dimensional cut of an n dimensional polytope <i>Aleksandar Nikolov</i>	124
VC-dimension of inflated polynomials <i>Jeff M. Phillips</i>	124
Stability of Topological Signatures of Graphs <i>Bei Wang</i>	125
Participants	127

3 Overview of Talks

3.1 Graph Embedding While Preserving Pairwise Distances

Tetsuo Asano (JAIST – Ishikawa, JP)

License © Creative Commons BY 3.0 Unported license
© Tetsuo Asano

Main reference T. Cox, M. Cox, “Multidimensional Scaling Second Edition”, Chapman & Hall CRC, 2001.

In this talk we consider a distance preserving graph embedding problem: Given a weighted graph, embed vertices into points in d -space so that for each edge the distance between their corresponding endpoints is as close as possible to the weight of the edge. If it is known that there is an exact embedding (without any error) for a full matrix for all pairs of vertices, then the existing algorithm known as Principal Coordinate Analysis (PCO) can find such an exact embedding in polynomial time. Although it is believed that PCO almost always gives a good solution, it is not true. We show a worst example for PCO and compare it with a heuristic algorithm. We also consider a special case where every pair of points has the same distance. We show the problem can be solved in one dimension. We also show some experimental results in two dimensions which reflect the results in one dimension.

3.2 Improved Time-Space Trade-offs for Computing Voronoi Diagrams

Bahareh Banyassady (FU Berlin, DE), Matias Korman (Tohoku University – Sendai, JP), Wolfgang Mulzer (FU Berlin, DE), Marcel Roeloffzen, Paul Seiferth, Yannik Stein, and André van Renssen (National Institute of Informatics – Tokyo, JP)

License © Creative Commons BY 3.0 Unported license
© Bahareh Banyassady, Matias Korman, Wolfgang Mulzer, Marcel Roeloffzen, Paul Seiferth, Yannik Stein, and André van Renssen

Joint work of Bahareh Banyassady, Matias Korman, Wolfgang Mulzer, André van Renssen, Marcel Roeloffzen, Paul Seiferth, Yannik Stein

Main reference B. Banyassady, M. Korman, W. Mulzer, A. van Renssen, M. Roeloffzen, P. Seiferth, Y. Stein, “Improved Time-Space Trade-offs for Computing Voronoi Diagrams”, in Proc. of the 34th Symposium on Theoretical Aspects of Computer Science (STACS 2017), LIPIcs, Vol. 66, pp. 9:1–9:14, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2017.

URL <http://dx.doi.org/10.4230/LIPIcs.STACS.2017.9>

In this talk, we are interested in computing various Voronoi diagrams for a given set of sites P in the memory-constrained model. In this model, we assume that the input is in a read-only array and the algorithm may use an additional workspace memory of size $O(s)$ words, $s \in \{1, \dots, n\}$, for reading and writing intermediate data, and the output is write-only. Clearly, when s increases, the running time of the algorithm decreases. In this model, we provide a time-space trade-off for computing the nearest site Voronoi diagram and the farthest site Voronoi diagram with running time $O(n^2/s \log s)$ using $O(s)$ words of workspace. Furthermore, we extend this result to compute the family of all higher-order Voronoi diagrams, up to a given order $K \in \{1, \dots, O(\sqrt{s})\}$, in a pipelined fashion in $O\left(\frac{n^2 K^5}{s} (\log s + K \log K)\right)$ time using $O(s)$ words of workspace. The main idea is to use edges of each diagram to compute the edges of the next diagram. However, this needs to be coordinated carefully, in order to keep the bound on the used memory and to prevent edges from being reported multiple times.

3.3 The Morse theory of Čech and Delaunay Complexes

Ulrich Bauer (TU München, DE)

License © Creative Commons BY 3.0 Unported license
© Ulrich Bauer

Joint work of Herbert Edelsbrunner, Ulrich Bauer

Main reference U. Bauer, H. Edelsbrunner, “The Morse theory of Čech and Delaunay complexes”, *Trans. Amer. Math. Soc.* 369(2017), pp. 3741–3762, 2016.

URL <http://dx.doi.org/10.1090/tran/6991>

Given a finite set of points in R^n and a radius parameter, we study the Čech, Delaunay-Čech, Delaunay (or alpha), and Wrap complexes in the light of generalized discrete Morse theory. Establishing the Čech and Delaunay complexes as sublevel sets of generalized discrete Morse functions, we prove that the four complexes are simple-homotopy equivalent by a sequence of simplicial collapses, which are explicitly described by a single discrete gradient field.

3.4 Conditional Lower Bounds for Similarity Measures on Curves and Strings

Karl Bringmann (MPI für Informatik – Saarbrücken, DE)

License © Creative Commons BY 3.0 Unported license
© Karl Bringmann

Joint work of Marvin Künnemann, Wolfgang Mulzer, Amir Abboud, Karl Bringmann

Main reference K. Bringmann, “Why Walking the Dog Takes Time: Frechet Distance Has No Strongly Subquadratic Algorithms Unless SETH Fails”, in *Proc. of the 55th Annual Symp. on Foundations of Computer Science (FOCS 2014)*, pp. 661–670, IEEE, 2014.

URL <https://doi.org/10.1109/FOCS.2014.76>

This talk is an introduction to the area of Fine-grained Complexity, where we show running time lower bounds conditional on certain conjectures on classic problems, such as the Strong Exponential Time Hypothesis for the Satisfiability problem. Specifically, we survey lower bounds known for similarity measures on curves (e.g. Frechet distance) and strings (e.g. longest common subsequence and edit distance), and we discuss many recent extensions.

3.5 Untangling and Unwinding Curves

Jeff Erickson (University of Illinois – Urbana-Champaign, US)

License © Creative Commons BY 3.0 Unported license
© Jeff Erickson

Joint work of Hsien-Chih Chang, Jeff Erickson

Main reference H.-C. Chang, J. Erickson, “Untangling Planar Curves”, in *Proc. of the 32nd Int’l Symp. on Computational Geometry (SoCG 2016)*, *LIPICs*, Vol. 51, pp. 29:1–29:16, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016.

URL <https://doi.org/10.4230/LIPICs.SocG.2016.29>

Any closed curve in the plane can be transformed into a simple closed curve using a finite sequence of local transformations called homotopy moves. We prove that $\Theta(n^{3/2})$ homotopy moves are necessary and sufficient in the worst case, improving the previous best $O(n^2)$ upper bound due to Steinitz in 1916 and the previous best $\Omega(n)$ lower bound, which is trivial. We also prove that $\Omega(n^2)$ moves are necessary in the worst case to simplify a contractible curve in the annulus, and therefore in any surface with non-positive Euler characteristic; a matching $O(n^2)$ upper bound follows from Steinitz’s planar results and more recent work by Hass and Scott.

3.6 A Nearly Quadratic Bound for the Decision Tree Complexity of k -SUM

Esther Ezra (Georgia Institute of Technology – Atlanta, US)

License © Creative Commons BY 3.0 Unported license
© Esther Ezra

Joint work of Esther Ezra, Micha Sharir

We show that the k -SUM problem can be solved by a linear decision tree of depth $O(n^2 \log^2 n)$, improving the recent bound $O(n^3 \log^3 n)$ of Cardinal et al. Our bound depends linearly on k , and allows us to conclude that the number of linear queries required to decide the n -dimensional Knapsack or SubsetSum problems is only $O(n^3 \log n)$, improving the currently best known bounds by a factor of n . Our algorithm extends to the RAM model, showing that the k -SUM problem can be solved in expected polynomial time, for any fixed k , with the above bound on the number of linear queries. Our approach relies on a new point-location mechanism, exploiting “Epsilon-cuttings” that are based on vertical decompositions in hyperplane arrangements in high dimensions. A major side result of our analysis is a sharper bound on the complexity of the vertical decomposition of such an arrangement (in terms of its dependence on the dimension).

3.7 Computing Optimal Flight Patterns with Minimum Turn Cost

Sándor Fekete (TU Braunschweig, DE)

License © Creative Commons BY 3.0 Unported license
© Sándor Fekete

Joint work of Aaron Becker, Mustapha Debboun, Dominik Krupke, An Nguyen, Sándor Fekete

URL <https://youtu.be/SFyOMDgdNao>

We present results arising from the problem of sweeping a mosquito-infested area with an Unmanned Aerial Vehicle (UAV) equipped with an electrified metal grid. Planning good trajectories is related to a number of classic problems of geometric optimization, in particular the Traveling Salesman Problem, the Lawn Mower Problem and, most closely, Milling with Turn Cost. We describe how planning a good trajectory can be reduced to considering penalty and budget variants of covering a grid graph with minimum turn cost. On the theoretical side, we show the solution of a problem from The Open Problems Project that had been open for more than 15 years, and hint at approximation algorithms. On the practical side, we describe an exact method based on Integer Programming that is able to compute provably optimal instances with over 500 pixels. These solutions are actually used for practical trajectories, as demonstrated in a video.

3.8 Faster Algorithms for the Geometric Transportation Problem

Kyle Jordan Fox (Duke University – Durham, US)

License © Creative Commons BY 3.0 Unported license
© Kyle Jordan Fox

Joint work of Pankaj K. Agarwal, Kyle Fox, Debmalya Panigrahi, Kasturi R. Varadarajan, Allen Xiao
Main reference P. K. Agarwal, K. Fox, D. Panigrahi, K. R. Varadarajan, A. Xiao, “Faster algorithms for the geometric transportation problem”, in Proc. of the 33rd International Symposium on Computational Geometry (SoCG 2017), LIPIcs, Vol. 77, pp. 7:1–7:15, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017.

URL <https://doi.org/10.4230/LIPIcs.SoCG.2017.7>

Let $R, B \subset R^d$, for constant d , be two point sets with $|R| + |B| = n$, and let $\lambda : R \cup B \rightarrow N$ such that $\sum_{r \in R} \lambda(r) = \sum_{b \in B} \lambda(b)$ be demand functions over R and B . Let $d(\cdot, \cdot)$ be a suitable distance function such as the L_p distance. The transportation problem asks to find a map $\tau : R \times B \rightarrow N$ such that $\sum_{b \in B} \tau(r, b) = \lambda(r)$, $\sum_{r \in R} \tau(r, b) = \lambda(b)$, and $\sum_{r \in R, b \in B} \tau(r, b) d(r, b)$ is minimized. We present three new results for the transportation problem when $d(\cdot, \cdot)$ is any L_p metric:

- For any constant $\epsilon > 0$, an $O(n^{1+\epsilon})$ expected time randomized algorithm that returns a transportation map with expected cost $O(\log^2(1/\epsilon))$ times the optimal cost.
- For any $\epsilon > 0$, a $(1 + \epsilon)$ -approximation in $O(n^{3/2} \epsilon^d \text{polylog}(U) \text{polylog}(n))$ time, where $U = \max_{p \in R \cup B} \lambda(p)$.
- An exact strongly polynomial $O(n^2 \text{polylog}(n))$ time algorithm, for $d = 2$.

3.9 Don't Collect too Much: Geometric Approaches for Protecting Location and Trajectory Privacy

Jie Gao (Stony Brook University, US)

License © Creative Commons BY 3.0 Unported license
© Jie Gao

Joint work of Jiaxin Ding, Chien Chun Ni, Mengyu Zhou, Xiaotian Yin, Wei Han, Dengpan Zhou, David Gu, Jie Gao
Main reference J. Ding, C.-C. Ni, M. Zhou, J. Gao, “MinHash hierarchy for privacy preserving trajectory sensing and query”, in Proc. of the 16th ACM/IEEE Int'l Conf. on Information Processing in Sensor Networks, pp. 17–28, ACM, 2017.

URL <http://doi.acm.org/10.1145/3055031.3055076>

Main reference X. Yin, C.-C. Ni, J. Ding, W. Han, D. Zhou, J. Gao, X. D. Gu, “Decentralized human trajectories tracking using hodge decomposition in sensor networks”, in Proc. of the 23rd SIGSPATIAL Int'l Conf. on Advances in Geographic Information Systems (SIGSPATIAL '15), pp. 54:1–54:4, ACM, 2015.

URL <http://doi.acm.org/10.1145/2820783.2820844>

Large amounts of geometric data are becoming available due to recent technology advances in sensing, communication and computation. GPS traces, location registrations, sensor data in smart environments can be used to infer human behaviors and patterns. While we celebrate the enormous learning opportunities these have enabled, the learned patterns may reveal sensitive or personal identifying information. In this talk I will describe geometric methods to protect location and trajectory information. Our methods focus on reducing location/trajectory data collected by sensors to meet the target privacy requirements using differential forms, Hodge decomposition and MinHash schemes.

3.10 Spatio-Temporal Analysis of Team Sports and Computational Geometry

Joachim Gudmundsson (The University of Sydney, AU)

License © Creative Commons BY 3.0 Unported license
© Joachim Gudmundsson

Joint work of Michael Horton, Joachim Gudmundsson

Main reference J. Gudmundsson, M. Horton, “Spatio-Temporal Analysis of Team Sports”, ACM Computing Surveys, Vol. 50(2): 22:1-22:34, ACM, 2017.

URL <http://doi.acm.org/10.1145/3054132>

Team-based invasion sports such as football, basketball and hockey are similar in the sense that the players are able to move freely around the playing area; and that player and team performance cannot be fully analysed without considering the movements and interactions of all players as a group. State of the art object tracking systems now produce spatio-temporal traces of player trajectories with high definition and high frequency, and this, in turn, has facilitated a variety of research efforts, across many disciplines, to extract insight from the trajectories. In this talk we focus on some of the geometric problems that arise in the area.

3.11 High-dimensional Theta Numbers

Anna Gundert (Universität Köln, DE)

License © Creative Commons BY 3.0 Unported license
© Anna Gundert

Joint work of Christine Bachoc, Anna Gundert, Alberto Passuello

Main reference C. Bachoc, A. Gundert, A. Passuello, “The Theta Number of Simplicial Complexes”, arXiv:1704.01836 [math.CO], 2017.

URL <https://arxiv.org/abs/1704.01836>

We introduce a generalization of the celebrated Lovász theta number of a graph to simplicial complexes of arbitrary dimension. Our generalization takes advantage of real simplicial cohomology theory, in particular combinatorial Laplacians, and provides a semidefinite programming upper bound of the independence number of a simplicial complex. We consider properties of the graph theta number such as the relationship to Hoffman’s ratio bound and to the chromatic number and study how they extend to higher dimensions. Furthermore, we analyze the value of the theta number for dense random simplicial complexes.

3.12 Minkowski Sums of Polyhedra with Holes

Dan Halperin (Tel Aviv University, IL)

License © Creative Commons BY 3.0 Unported license
© Dan Halperin

Joint work of Alon Baram, Efi Fogel, Dan Halperin, Michael Hemmer, Sebastian Morr

The Minkowski sum of two sets P and Q in Euclidean space is the result of adding every point (position vector) in P to every point in Q . Considering the Minkowski sum of two polyhedra with holes, we show that one can always fill up the holes in one of the summand polyhedra and still get the same Minkowski sum as of the original summands. We present a simple proof of this observation, improving on (our) earlier rather involved proof of a more restricted claim. As we explain, this observation helps in speeding up the computation of Minkowski sums in practice. We also review additional recent results in computing and using Minkowski sums.

3.13 Self-Aligning Shapes

David G. Kirkpatrick (University of British Columbia – Vancouver, CA)

License © Creative Commons BY 3.0 Unported license
© David G. Kirkpatrick

Joint work of Ashwin Gopinath, David G. Kirkpatrick, Paul W.K. Rothmund and Chris Thachuk

Main reference A. Gopinath, D. G. Kirkpatrick, P. W. K. Rothmund, Chris Thachuk, “Progressive Alignment of Shapes”, in Proc. of the 28th Canadian Conference on Computational Geometry (CCCG 2016), pp. 230–236, Simon Fraser University, 2016.

URL <http://www.cccg.ca/proceedings/2016/proceedings2016.pdf>

A planar shape S is said to be self-aligning if any two overlapping, but otherwise arbitrarily placed, copies of S , can be brought into a unique configuration of maximum overlap by a continuous motion that monotonically increases their overlap. The identification of self-aligning shapes is motivated by applications of self-assembly, driven by molecular forces, in nano-fabrication processes. We are interested in the design and certification of shapes that are self-aligning, but also satisfy certain other constraints that arise due to fabrication issues.

3.14 On the Computational Bottleneck in Sampling-Based Robot Motion Planning

Michal Kleinbort (Tel Aviv University, IL)

License © Creative Commons BY 3.0 Unported license
© Michal Kleinbort

Joint work of Michal Kleinbort, Oren Salzman, Dan Halperin

Main reference M. Kleinbort, O. Salzman, D. Halperin, “Collision detection or nearest-neighbor search? On the computational bottleneck in sampling-based motion planning”, arXiv:1607.04800v3 [cs.RO], 2016.

URL <https://arxiv.org/abs/1607.04800v3>

Many sampling-based motion planning algorithms rely heavily on two main components: (i) collision detection and (ii) nearest-neighbor search. The complexity of the latter dominates the asymptotic running time of such algorithms. However, collision detection is often considered to be the computational bottleneck in practice. We describe settings in which the practical computational role of nearest-neighbor search is far from being negligible, i.e., the portion of running time taken up by nearest-neighbor search is comparable to, or sometimes even greater than the portion of time taken up by collision detection. We show that by using efficient, specially-tailored nearest-neighbor data structures, the overall running time of certain motion-planning algorithms in such settings can be significantly reduced.

3.15 Computing Wave Impact in Self-Organised Mussel Beds

Maarten Löffler (Utrecht University, NL)

License © Creative Commons BY 3.0 Unported license
© Maarten Löffler

We model the effects of byssal connections made by mussels within patterned mussel beds on bed stability as a disk graph, and propose a formula for assessing which mussels, if any, would get dislodged from the bed under the impact of a wave. We formulate the computation as a flow problem, giving access to efficient algorithms to evaluate the formula. We then analyse the geometry of the graph, and show that we only need to compute a maximum flow in a restricted part of the graph, giving rise to a near-linear solution in practice.

3.16 The Art Gallery Problem is ETR-complete

Tillmann Miltzow (Free University of Brussels, BE)

License © Creative Commons BY 3.0 Unported license
© Tillmann Miltzow

Joint work of Mikkel Abrahamsen, Anna Adamaszek, Tillmann Miltzow

Main reference M. Abrahamsen, A. Adamaszek, T. Miltzow, “The Art Gallery Problem is $\exists\mathbb{R}$ -complete”, arXiv:1704.06969 [cs.CG], 2017.

URL <https://arxiv.org/abs/1704.06969>

We prove that the *art gallery problem* is equivalent under polynomial time reductions to deciding whether a system of polynomial equations over the real numbers has a solution. The art gallery problem is a classical problem in computational geometry, introduced in 1973 by Viktor Klee. Given a simple polygon P and an integer k , the goal is to decide if there exists a set G of k *guards* within P such that every point $p \in P$ is seen by at least one guard $g \in G$. Each guard corresponds to a point in the polygon P , and we say that a guard g *sees* a point p if the line segment pg is contained in P .

The art gallery problem has stimulated a myriad of research in geometry and in algorithms. However, despite extensive research, the complexity status of the art gallery problem has not been resolved. It has long been known that the problem is NP-hard, but no one has been able to show that it lies in NP. Recently, the computational geometry community became more aware of the complexity class $\exists\mathbb{R}$. The class $\exists\mathbb{R}$ consists of problems that can be reduced in polynomial time to the problem of deciding whether a system of polynomial equations with integer coefficients and any number of real variables has a solution. It can be easily seen that $\text{NP} \subseteq \exists\mathbb{R}$. We prove that the art gallery problem is $\exists\mathbb{R}$ -complete, implying that (1) any system of polynomial equations over the real numbers can be encoded as an instance of the art gallery problem, and (2) the art gallery problem is not in the complexity class NP unless $\text{NP} = \exists\mathbb{R}$. As a corollary of our construction, we prove that for any real algebraic number α there is an instance of the art gallery problem where one of the coordinates of the guards equals α in any guard set of minimum cardinality. That rules out many geometric approaches to the problem.

3.17 TSP With Locational Uncertainty: The Adversarial Model

Joseph S. B. Mitchell (Stony Brook University, US)

License © Creative Commons BY 3.0 Unported license
© Joseph S. B. Mitchell

Joint work of Gui Citovsky, Tyler Mayer, Joseph S. B. Mitchell

In this paper we study a natural special case of the Traveling Salesman Problem (TSP) with point-locational-uncertainty which we will call the *adversarial TSP* problem (ATSP). Given a metric space (X, d) and a set of subsets $R = \{R_1, R_2, \dots, R_n\} : R_i \subseteq X$, the goal is to devise an ordering of the regions, σ_R , that the tour will visit such that when a single point is chosen from each region, the induced tour over those points in the ordering prescribed by σ_R is as short as possible. Unlike the classical locational-uncertainty-TSP problem, which focuses on minimizing the expected length of such a tour when the point within each region is chosen according to some probability distribution, here, we focus on the *adversarial model* in which once the choice of σ_R is announced, an adversary selects a point from each region in order to make the resulting tour as long as possible. In other words, we consider an offline problem in which the goal is to determine an ordering of the regions R that is optimal

with respect to the “worst” point possible within each region being chosen by an adversary, who knows the chosen ordering. We give a 3-approximation when R is a set of arbitrary regions/sets of points in a metric space. We show how geometry leads to improved constant factor approximations when regions are parallel line segments of the same lengths, and a polynomial-time approximation scheme (PTAS) for the important special case in which R is a set of disjoint unit disks in the plane.

3.18 Dynamic Planar Voronoi Diagrams for General Distance Functions

Wolfgang Mulzer (FU Berlin, DE), Haim Kaplan, Liam Roditty, Paul Seiferth, and Micha Sharir

License © Creative Commons BY 3.0 Unported license

© Wolfgang Mulzer, Haim Kaplan, Liam Roditty, Paul Seiferth, and Micha Sharir

Main reference H. Kaplan, W. Mulzer, L. Roditty, P. Seiferth, M. Sharir, “Dynamic Planar Voronoi Diagrams for General Distance Functions and their Algorithmic Applications”, in Proc. of the 28th ACM-SIAM Symp. on Discrete Algorithms (SODA 2017), pp. 2495–2504, ACM, 2017.

URL <http://dx.doi.org/10.1137/1.9781611974782.165>

We describe a new data structure for dynamic nearest neighbor queries in the plane with respect to a general family of distance functions that includes L_p -norms and additively weighted Euclidean distances, and for general (convex, pairwise disjoint) sites that have constant description complexity (line segments, disks, etc.). Our data structure has a polylogarithmic update and query time, improving an earlier data structure of Agarwal, Efrat and Sharir that required $O(n^\epsilon)$ time for an update and $O(\log n)$ time for a query.

3.19 Maximizing Volume Subject to Combinatorial Constraints

Aleksandar Nikolov (University of Toronto, CA)

License © Creative Commons BY 3.0 Unported license

© Aleksandar Nikolov

The maximum volume j -simplex problem asks to compute the j -dimensional simplex of maximum volume inside the convex hull of a given set of n points in d -dimensional space. We discuss a deterministic approximation algorithm for this problem which achieves an approximation ratio of $e^{j/2+o(j)}$ in time polynomial in n and d . The problem is known to be NP-hard to approximate within a factor of c^j for some constant $c > 1$. Our algorithm also gives a factor $e^{j+o(j)}$ approximation for the problem of finding the principal $j \times j$ submatrix of a rank d positive semidefinite matrix with the largest determinant. In the final part of the talk we sketch how these results have been generalized to maximizing volume or determinants subject to more complicated combinatorial constraints. Such problems appear naturally in machine learning and experimental design when dealing with models of diversity.

3.20 Colorful simplicial depth, Minkowski sums, and generalized Gale transforms

Zuzana Patáková (IST Austria – Klosterneuburg, AT)

License © Creative Commons BY 3.0 Unported license
© Zuzana Patáková

Joint work of Karim Adiprasito, Philip Brinkmann, Arnau Padrol, Pavel Paták, Zuzana Patáková, Raman Sanyal
Main reference K. Adiprasito, P. Brinkmann, A. Padrol, P. Paták, Z. Patáková, R. Sanyal, “Colorful simplicial depth, Minkowski sums, and generalized Gale transforms”, arXiv:1607.00347 [math.CO], 2016.

URL <https://arxiv.org/abs/1607.00347>

Given $d + 1$ sets S_1, S_2, \dots, S_{d+1} (called color classes) in \mathbb{R}^d , a simplex is called *colorful*, if all its vertices are in different color classes. The number of colorful simplices containing a point $p \in \mathbb{R}^d$ is known as the *colorful simplicial depth of p* . The first result concerning colorful simplicial depth in discrete geometry was the colorful Carathéodory’s theorem by Imre Bárány in 1982: “Any point $p \in \mathbb{R}^d$ contained in the convex hull of all color classes has a non-zero simplicial depth provided that each color class has at least $d + 1$ points.”

In 2006 Deza, Huang, Stephen, and Terlaky asked for the minimal and maximal values of the colorful simplicial depth of the point p in colorful Carathéodory’s theorem. We use methods from combinatorial topology to prove a tight upper bound of the form $1 + \prod_{i=1}^{d+1} (|S_i| - 1)$.

The second goal of this talk is to highlight a connection between colorful configurations and faces of Minkowski sums. Considering the Gale transform of the Cayley embedding, we define *colorful Gale transforms* associated to a collection of convex polytopes that capture the facial structure of Minkowski sums in the combinatorics of colorful configurations. This dictionary between Minkowski sums and colorful configurations allows us to resolve a conjecture of Ben Burton from the *normal surface* theory.

3.21 Approximate Range Maximization

Jeff M. Phillips (University of Utah – Salt Lake City, US)

License © Creative Commons BY 3.0 Unported license
© Jeff M. Phillips

Joint work of Michael Matheny, Raghendra Singh, Kaiqiang Wang, Liang Zhang, Jeff M. Phillips
Main reference M. Matheny, R. Singh, K. Wang, L. Zhang, J. M. Phillips, “Scalable Spatial Scan Statistics through Sampling”, in Proc. of the 24th ACM SIGSPATIAL Int’l Conf. on Advances in Geographic Information Systems, (GIS 2016), pp. 20:1–20:10, ACM, 2016.

URL <http://doi.acm.org/10.1145/2996913.2996939>

Consider a geometric range space (X, A) where each data point $x \in X$ has two or more values (say r and b). Also consider a function $\phi(A)$ defined on any subset C in (X, A) on the sum of values in that range e.g., $r_C = \sum_{x \in C} r(x)$ and $b_C = \sum_{x \in C} b(x)$. The maximum range is $A^* = \arg \max_{X \text{ in } (X, A)} \phi(A)$. Our goal is to find some \hat{A} such that $|\phi(\hat{A}) - \phi(A^*)| \leq \varepsilon$. We develop algorithms for this problem for range spaces defined by balls, halfspaces, and axis-aligned rectangles; it has applications in many areas including discrepancy evaluation and spatial scan statistics.

3.22 A Treehouse with Custom Windows: Minimum Distortion Embeddings into Bounded Treewidth Graphs

Benjamin Raichel (University of Texas – Dallas, US)

License © Creative Commons BY 3.0 Unported license
© Benjamin Raichel

Joint work of Amir Nayyeri, Benjamin Raichel

Main reference A. Nayyeri, B. Raichel, “A treehouse with custom windows: Minimum distortion embeddings into bounded treewidth graphs”, in Proc. of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2017), pp. 724–736, ACM, 2017.

URL <https://doi.org/10.1137/1.9781611974782.46>

We describe a $(1 + \varepsilon)$ -approximation algorithm for finding the minimum distortion embedding of an n -point metric space X into the shortest path metric space of a weighted graph G with m vertices. The running time of our algorithm is

$$m^{O(1)} \cdot n^{O(\lambda)} \cdot (\delta_{opt} \Delta)^{\lambda \cdot (1/\varepsilon)^{\lambda+2} \cdot \lambda \cdot (O(\delta_{opt}))^{2\lambda}}$$

parametrized by the values of the minimum distortion, δ_{opt} , the spread, Δ , of the points of X , the treewidth, λ , of G , and the doubling dimension, λ , of G .

In particular, our result implies a PTAS provided an X with polynomial spread, and the doubling dimension of G , the treewidth of G , and δ_{opt} , are all constant. For example, if X has a polynomial spread and δ_{opt} is a constant, we obtain PTAS’s for embedding X into the following spaces: the line, a cycle, a tree of bounded doubling dimension, and a k -outer planar graph of bounded doubling dimension (for a constant k).

3.23 Spatial data processing with SAP HANA

Alejandro Salinger (SAP SE – Walldorf, DE)

License © Creative Commons BY 3.0 Unported license
© Alejandro Salinger

URL http://help.sap.com/hana/SAP_HANA_Spatial_Reference_en.pdf

The rapid development of geospatial technologies in recent years has led to an increase in the types and amount of enterprise data containing geospatial information. The ability to store and process spatial data can add a valuable dimension to business analytics.

In this talk we give an overview of the spatial data processing capabilities of SAP HANA, SAP’s in-memory data platform. We describe the spatial data types and operations supported and give examples of SQL queries using spatial data. These include using simple predicates such as determining if a point is within a given geometry as well as more complex queries that use clustering methods.

We then give some insights about the optimizations used to reduce computational time in window queries, including the use of parallelism, space-filling curves, dictionaries, and indices.

Finally, we show a live demo based on an application for proactive pipeline maintenance used by a major natural gas provider. The application allows visualizing pipelines in a map and finding those that are within a given distance to buildings. We briefly describe the underlying algorithms and data structures used to compute spatial joins that can process millions of geometries in a just a few seconds.

3.24 A Framework for Algorithm Stability

Kevin Verbeek (TU Eindhoven, NL)

License  Creative Commons BY 3.0 Unported license
© Kevin Verbeek

We say that an algorithm is stable if small changes in the input result in small changes in the output. Algorithm stability plays an important role when analyzing and visualizing time-varying data. However, so far, there are only few theoretical results on the stability of algorithms, possibly due to a lack of theoretical analysis tools. In this talk we present a framework for analyzing the stability of algorithms. We focus in particular on the tradeoff between the stability of an algorithm and the quality of the solution it computes. Our framework allows for three types of stability analysis with increasing degrees of complexity: event stability, topological stability, and Lipschitz stability. We demonstrate the use of our stability framework by applying it to kinetic Euclidean minimum spanning trees.

3.25 Towards Spectral Sparsification of Simplicial Complexes based on Generalized Effective Resistance

Bei Wang (University of Utah – Salt Lake City, US)

License  Creative Commons BY 3.0 Unported license
© Bei Wang
Joint work of Braxton Osting, Sourabh Palande, Bei Wang

As a generalization of the use of graphs to describe pairwise interactions, simplicial complexes can be used to model higher-order interactions between three or more objects in complex systems. There has been a recent surge in activity for the development of data analysis methods applicable to simplicial complexes, including techniques based on computational topology, higher-order random processes, generalized Cheeger inequalities, isoperimetric inequalities, and spectral methods. In particular, spectral learning methods (e.g. label propagation and clustering) that directly operate on simplicial complexes represent a new direction emerging from the confluence of computational topology and machine learning. Similar to the challenges faced by massive graphs, computational methods that operate on simplicial complexes are severely limited by computational costs associated with massive datasets.

To apply spectral methods in learning to massive datasets modeled as simplicial complexes, we work towards the sparsification of simplicial complexes based on preserving the spectrum of the associated Laplacian operators. We show that the theory of Spielman and Srivastava for the sparsification of graphs extends to the generality of simplicial complexes via the up Laplacian. In particular, we introduce a generalized effective resistance for simplexes; provide an algorithm for sparsifying simplicial complexes at a fixed dimension; and give a specific version of the generalized Cheeger inequalities for weighted simplicial complexes under the sparsified setting. In addition, we demonstrate via experiments the preservation of up Laplacian during sparsification, as well as the utility of sparsification with respect to spectral clustering.

4 Open problems

4.1 Nash equilibria of spanner games

Mohammad Ali Abam (Sharif University of Technology – Tehran, IR)

License  Creative Commons BY 3.0 Unported license
© Mohammad Ali Abam

Starting from a t -sink spanner, that is a directed graph, where from every vertex there is a path to a sink where the length of the path is at most a factor t more than the cost of a direct edge. The vertices (players) take turns and in each turn a player can replace its outgoing edge by a different edge, but it needs to maintain the t -sink spanner property of the graph. Each player pays for its outgoing edge. Does this game have a nash-equilibrium?

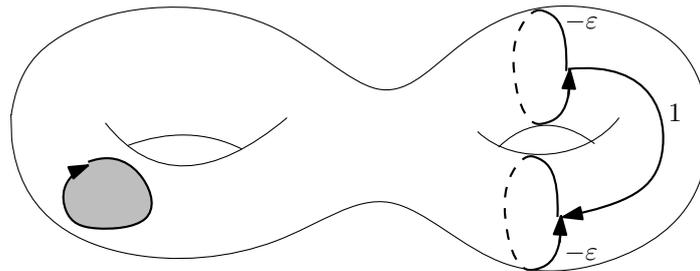
A similar problem can be posed for a general t -spanner, where there should be spanning paths between any pair of vertices. In this case during each turn, a vertex can change its set of outgoing edges under the restriction that the graph must remain a spanner. The question is then if this is a process that terminates and if so, with how many edges?

4.2 Negative cycles on surface embedded graphs

Jeff Erickson (University of Illinois – Urbana-Champaign, US)

License  Creative Commons BY 3.0 Unported license
© Jeff Erickson

Given a directed graph on a (possibly non-compatible) surface, where edges have weights that may be negative. The problem is to determine if there is a negative contractible cycle, that is, a closed walk which may reuse edges and vertices. What is the complexity of this decision problem? Is this problem in \mathbb{P} or even in \mathbb{NP} (the cycle might have exponential length), is it \mathbb{NP} -hard?



4.3 Maximum-Area Triangle in a Convex Polygon

Maarten Löffler (*Utrecht University, NL*)

License © Creative Commons BY 3.0 Unported license
© Maarten Löffler

Main reference D. P. Dobkin, L. Snyder, “On a general method for maximizing and minimizing among certain geometric problems”, in 20th Annual Symposium on Foundations of Computer Science (FOCS 1979), pp. 9–17, IEEE, 1979.

URL <https://doi.org/10.1109/SFCS.1979.28>

Given a convex polygon P , find the largest-area inscribed triangle. We revisit the linear-time algorithm proposed by Dobkin and Snyder and ask if there is a counter-example to their claimed proof.

Update: There has been progress on this problem since the problem was posed: Vahideh Keikha, Maarten Löffler, Jérôme Urhausen, Ivor van der Hoog: “Maximum-Area Triangle in a Convex Polygon, Revisited.” CoRR abs/1705.11035, 2017.

4.4 Is G area-universal?

Tillmann Miltzow (*Free University of Brussels, BE*)

License © Creative Commons BY 3.0 Unported license
© Tillmann Miltzow

Joint work of Linda Kleist, Pawel Rzazewski, Michael Gene Dobbins, Tillmann Miltzow

Main reference L. Kleist, P. Rzazewski, T. Miltzow, “Is Area Universality $\forall\exists\mathbb{R}$ -complete?”, in Proc. of the 33rd European Workshop on Computational Geometry (EuroCG 2017), Malmö University, 2017.

URL <http://csconferences.mah.se/eurocg2017/proceedings.pdf>

During the Dagstuhl Seminar, I asked whether the graph G (drawn on the blackboard) is area universal. See the EuroCG abstract for details (page 181).

4.5 Finding small polygons with a fixed number of triangulations; simple polygons with a unique triangulation

Joseph S. B. Mitchell (*Stony Brook University, US*)

License © Creative Commons BY 3.0 Unported license
© Joseph S. B. Mitchell

Given an integer k , let $N(k)$ denote the minimum number n , so that there is a polygon of n vertices that has exactly k triangulations. This number always exists for any k . The first problem is to find $N(k)$.

A second problem is to detect if a given polygon P has a unique triangulation. This can be done in $O(n)$ time from a starting triangulation. Using Chazelle’s linear triangulation algorithm this yields linear time in total; however, since no implementation for linear-time triangulation exists, an interesting question is if there is a different, simple linear-time algorithm.

4.6 Minimum k dimensional cut of an n dimensional polytope

Aleksandar Nikolov (University of Toronto, CA)

License © Creative Commons BY 3.0 Unported license
© Aleksandar Nikolov

Main reference A. Pajor, N. Tomczak-Jaegermann, “Volume ratio and other s -numbers of operators related to local properties of Banach spaces”, *J. Funct. Anal.*, 87(2): 273–293, Elsevier, 1989.

URL [https://doi.org/10.1016/0022-1236\(89\)90011-6](https://doi.org/10.1016/0022-1236(89)90011-6)

Given a convex polytope P in \mathbb{R}^n that is symmetric around the origin (i.e. $P = -P$), specified in H-representation as $P = \{x : b \leq Ax \leq b\}$ for a matrix A and a vector b with non-negative entries. Note that P can have exponentially many vertices in the size of its representation. More generally, we are given a convex body K symmetric around the origin, specified via a membership or a separation oracle. Can we efficiently find a k -dimensional subspace through the origin that minimizes its cut-volume with P ? That is, can we find or approximate $\min\{\text{vol}(W \cap P)^{1/k} : W \text{ is a subspace of dimension } k\}$?

The problem is easy for centrally symmetric ellipsoids E , specified by $E = \{x : x^\top Bx \leq 1\}$, for a positive semidefinite matrix B . Then the minimizing subspace is spanned by eigenvectors associated with the largest k eigenvalues of M . Using this fact, and Milman’s M-ellipsoid theorem, we can get a factor $C^{n/k}$ -approximation for an arbitrary symmetric convex body K , specified by a membership oracle, and an absolute constant C , independent of n , k , and K . However, this is unsatisfying when k is sublinear in n . On the other hand, for a polytope P in H-representation and $k = 1$, the problem just reduces to computing the closest (in Euclidean distance) facet to the origin. Intermediate values of k seem challenging.

This problem is related to computing volume ratio numbers, which can be used to compute estimates on the Gaussian mean width of a convex body (see Pajor and Tomczak-Jaegermann, *JFA* 1989).

4.7 VC-dimension of inflated polynomials

Jeff M. Phillips (University of Utah – Salt Lake City, US)

License © Creative Commons BY 3.0 Unported license
© Jeff M. Phillips

Main reference J. M. Phillips, Y. Zheng, “Subsampling in Smoothed Range Spaces”, in *Proc. of the 26th International Conference on Algorithmic Learning Theory (ALT 2015)*, LNCS, Vol. 9355, pp. 224–238, Springer, 2015.

URL https://doi.org/10.1007/978-3-319-24486-0_15

Is there a bound on the VC-dimension of the range space defined by shapes formed by the Minkowski sum of a ball and a polynomial curve? In more detail, let \mathcal{S}_p be any family of subsets of \mathbb{R}^d which are defined by a shape whose boundary is a polynomially curve of degree p . Simple examples are when $H \in \mathcal{S}_1$ is a halfspace, or when $D \in \mathcal{S}_2$ is a disk; but the problem seems more challenging when we only restrict the degree of the polynomial. Then define the set of subsets $\mathcal{M}_p = \{S \oplus B \mid S \in \mathcal{S}_p\}$ where B is any ball and \oplus is the Minkowski sum. What is the VC-dimension of $(\mathbb{R}^d, \mathcal{M}_p)$?

4.8 Stability of Topological Signatures of Graphs

Bei Wang (University of Utah – Salt Lake City, US)

License  Creative Commons BY 3.0 Unported license
© Bei Wang

This open problem arises from a joint work with Mustafa Hajij, Carlos Scheidegger and Paul Rosen [4], where we are interested in quantifying and visualizing structural changes of a time-varying graph using persistent homology. However for simplification purposes here, we describe the open problem by focusing on a single snapshot G of a time-varying graph.

Let $G = (V, E, w)$ be a weighted, undirected graph. To study the topological signatures of G , we first embed G into a metric space (M, d) , using the shortest-path distance or other distance metrics based on the graph Laplacian [3] (such as commute-time distance, discrete biharmonic distance [5] and diffusion distance [2]). The topological signatures of such a (metric-space-embedded) graph is then extracted, by computing persistent homology of its corresponding Vietoris-Rips filtration; the topological features of the Vietoris-Rips filtration are captured by its persistence diagrams PD . Such persistence diagrams encoding the topological signatures of G can be compared and structural changes among time-varying snapshots may be detected and visualized.

The question is: *given the above analysis pipeline, can we quantify the stability of topological signatures of graphs?* In other words, if we add random perturbation to a graph, how stable (and therefore potentially trustworthy) are its topological signatures?

The stability of topological signatures of a given graph G can be influenced by at least three factors: the perturbation model, the stability of its metric structure embedding (M, d) , and the stability of its corresponding persistence diagram PD . In terms of random perturbation models, we can consider deletion-only and insertion-only perturbation, as well as rewiring (while maintaining degree or joint degree distributions). We would like to know, under a particular perturbation model, and for a certain metric space embedding, whether we can obtain stable topological signatures. As topological methods can be coordinate and deformation invariant, under certain restrictive settings, can we hope for stable topological signatures even if the metric structures are not stable?

The metric structure behind perturbed graphs in the setting of shortest path metric and Erdős-Rényi type perturbation has been studied recently [6]. Given a “noisy” observation G of a true graph G^* , Parthasarathy et al. [6] propose a de-noising procedure to recover (approximately) the “true” shortest path metric of G^* from G . Since the shortest path metric is known to be sensitive to random perturbations, such a de-noising procedure can be used as a way to recover the stable metric structure of a graph. On the other hand, persistent homology stability result [1] implies that if M and N are two finite metric space, then the distance between persistence diagrams constructed by Vietoris-Rips filtration is bounded by twice of their Gromov-Hausdorff distance [7]. Therefore, the shortest-path de-noising procedure of [6] can be combined with persistent homology stability result (e.g. [1, 7]) to potentially address an instance of this open problem.

References

- 1 Frédéric Chazal, David Cohen-Steiner, Leonidas J Guibas, Facundo Méemoli, and Steve Y. Oudot. Gromov-hausdorff stable signatures for shapes using persistence. *Computer Graphics Forum*, 28(5):1393–1403, 2009.
- 2 Ronald R Coifman, Stephane Lafon, Ann B Lee, Mauro Maggioni, Boaz Nadler, Frederick Warner, and Steven W Zucker. Geometric diffusions as a tool for harmonic analysis and

- structure definition of data: Diffusion maps. *Proceedings National Academy of Sciences of the United States of America*, 102(21):7426–7431, 2005.
- 3 Dragoš M Cvetković, Michael Doob, and Horst Sachs. *Spectra of graphs: theory and application*, volume 87. Academic Press, 1980.
 - 4 Mustafa Hajij, Bei Wang, Carlos Scheidegger, and Paul Rosen. Persistent homology guided exploration of time-varying graphs. [arXiv:1707.06683](https://arxiv.org/abs/1707.06683), 2017.
 - 5 Yaron Lipman, Raif M Rustamov, and Thomas A Funkhouser. Biharmonic distance. *ACM Transactions on Graphics*, 29(3):27, 2010.
 - 6 Srinivasan Parthasarathy, David Sivakoff, Minghao Tian, and Yusu Wang. A quest to unravel the metric structure behind perturbed networks. *Proceedings 33rd International Symposium on Computational Geometry*, 2017.
 - 7 Katharine Turner. Generalizations of the rips filtration for quasi-metric spaces with persistent homology stability results. [arXiv:1608.00365](https://arxiv.org/abs/1608.00365), 2016.

Participants

- Mohammad Ali Abam
Sharif University of Technology –
Tehran, IR
- Tetsuo Asano
JAIST – Ishikawa, JP
- Sang Won Bae
Kyonggi University, KR
- Bahareh Banyassady
FU Berlin, DE
- Ulrich Bauer
TU München, DE
- Karl Bringmann
MPI für Informatik –
Saarbrücken, DE
- Otfried Cheong
KAIST – Daejeon, KR
- Man-Kwun Chiu
National Institute of Informatics –
Tokyo, JP
- Jinhee Chun
Tohoku University – Sendai, JP
- Kenneth L. Clarkson
IBM Almaden Center –
San Jose, US
- Michael Gene Dobbins
Binghamton University, US
- Anne Driemel
TU Eindhoven, NL
- Jeff Erickson
University of Illinois –
Urbana-Champaign, US
- Esther Ezra
Georgia Institute of Technology –
Atlanta, US
- Sándor Fekete
TU Braunschweig, DE
- Kyle Jordan Fox
Duke University – Durham, US
- Jie Gao
Stony Brook University, US
- Joachim Gudmundsson
The University of Sydney, AU
- Anna Gundert
Universität Köln, DE
- Dan Halperin
Tel Aviv University, IL
- David G. Kirkpatrick
University of British Columbia –
Vancouver, CA
- Michal Kleinbort
Tel Aviv University, IL
- Christian Knauer
Universität Bayreuth, DE
- Matias Korman
Tohoku University – Sendai, JP
- Amer Krivosija
TU Dortmund, DE
- Maarten Löffler
Utrecht University, NL
- Tillmann Miltzow
Free University of Brussels, BE
- Joseph S. B. Mitchell
Stony Brook University, US
- Wolfgang Mulzer
FU Berlin, DE
- Aleksandar Nikolov
University of Toronto, CA
- Zuzana Patáková
IST Austria –
Klosterneuburg, AT
- Jeff M. Phillips
University of Utah –
Salt Lake City, US
- Benjamin Raichel
University of Texas – Dallas, US
- Abhishek Rathod
TU München, DE
- Saurabh Ray
New York University –
Abu Dhabi, AE
- Marcel J. M. Roeloffzen
National Institute of Informatics –
Tokyo, JP
- Michael Sagraloff
MPI für Informatik –
Saarbrücken, DE
- Alejandro Salinger
SAP SE – Walldorf, DE
- Lena Schlipf
FernUniversität in Hagen, DE
- Raimund Seidel
Universität des Saarlandes, DE
- Fabian Stehn
Universität Bayreuth, DE
- Monique Teillaud
INRIA Nancy – Grand Est, FR
- André van Renssen
National Institute of Informatics –
Tokyo, JP
- Kevin Verbeek
TU Eindhoven, NL
- Bei Wang
University of Utah –
Salt Lake City, US
- David P. Woodruff
IBM Almaden Center –
San Jose, US

