



DAGSTUHL REPORTS

Volume 2, Issue 8, August 2012

Robust Query Processing (Dagstuhl Seminar 12321) <i>Goetz Graefe, Wey Guy, Harumi A. Kuno, and Glenn Paulley</i>	1
Mobility Data Mining and Privacy (Dagstuhl Seminar 12331) <i>Christopher W. Clifton, Bart Kuijpers, Katharina Morik, and Yucel Saygin</i>	16
Verifying Reliability (Dagstuhl Seminar 12341) <i>Görschwin Fey, Masahiro Fujita, Natasa Miskov-Zivanov, Kaushik Roy, and Matteo Sonza Reorda</i>	54
Engineering Multiagent Systems (Dagstuhl Seminar 122342) <i>Jürgen Dix, Koen V. Hindriks, Brian Logan, and Wayne Wobcke</i>	74
Information Flow and Its Applications (Dagstuhl Seminar 12352) <i>Samson Abramsky, Jean Krivine, and Michael W. Mislove</i>	99

ISSN 2192-5283

Published online and open access by

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

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

Publication date

February, 2013

Bibliographic information published by the Deutsche Nationalbibliothek

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

License

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



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

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

The copyright is retained by the corresponding authors.

Aims and Scope

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

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

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

Editorial Board

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

Editorial Office

Marc Herbstritt (*Managing Editor*)

Jutka Gasirowski (*Editorial Assistance*)

Thomas Schillo (*Technical Assistance*)

Contact

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

Digital Object Identifier: 10.4230/DagRep.2.8.i

www.dagstuhl.de/dagrep

Report from Dagstuhl Seminar 12321

Robust Query Processing

Edited by

Goetz Graefe¹, Wey Guy², Harumi A. Kuno³, and Glenn Paulley⁴

1 HP Labs, USA goetz.graefe@hp.com

2 USA weyyuanguy@hotmail.com

3 HP Labs, USA harumi.kuno@hp.com

4 Conestoga College, Kitchener, Ontario, Canada gpaulley@acm.org

Abstract

The 2012 Dagstuhl 12321 Workshop on Robust Query Processing, held from 5–10 August 2012, brought together researchers from both academia and industry to discuss various aspects of robustness in database management systems and ideas for future research. The Workshop was designed as a sequel to an earlier Workshop, Dagstuhl Workshop 10381, that studied a similar set of topics. In this article we summarize some of the main discussion topics of the 12321 Workshop, the results to date, and some open problems that remain.

Seminar 09.–13. July, 2012 – www.dagstuhl.de/12321

1998 ACM Subject Classification H.2 Database Management, H.2.2 Physical Design, H.2.4 Systems, H.2.7 Database Administration

Keywords and phrases robust query processing, adaptive query optimization, query execution, indexing, workload management, reliability, application availability

Digital Object Identifier 10.4230/DagRep.2.8.1


1 Executive Summary

Goetz Graefe

Wey Guy

Harumi A. Kuno

Glenn Paulley

License  Creative Commons BY-NC-ND 3.0 Unported license
© Goetz Graefe, Wey Guy, Harumi A. Kuno, and Glenn Paulley

Introduction

In early August 2012 researchers from both academia and industry assembled in Dagstuhl at the 2012 Dagstuhl Workshop on Robust Query Processing, Workshop 12321. An earlier Workshop—Dagstuhl Workshop 10381—held in September 2010 [16] had supplied an opportunity to look at issues of Robust Query Processing but had failed to make significant progress in exploring the topic to any significant depth. In 2012, 12321 Workshop participants looked afresh at some of the issues surrounding Robust Query Processing with greater success and with the strong possibility of future publications in the area that would advance the state-of-the-art in query processing technology.



Except where otherwise noted, content of this report is licensed under a Creative Commons BY-NC-ND 3.0 Unported license
Robust Query Processing, *Dagstuhl Reports*, Vol. 2, Issue 8, pp. 1–15
Editors: Goetz Graefe, Wey Guy, Harumi A. Kuno, and Glenn Paulley



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

Background and related research

A considerable amount of query processing research over the past 20 years has focused on improving relational database system optimization and execution techniques for complex queries and complex, ever-changing workloads. Complex queries provide optimization challenges because selectivity and cardinality estimation errors multiply, and so there is a large body of work on improving cardinality estimation techniques and doing so in an automatic fashion: from capturing histogram information at run time [1, 17], to mitigating the effects of correlation on the independence assumption [21], to utilizing constraints to bound estimation error [18, 15, 9, 10], to permitting various query rewritings to simplify the original statement [11, 23, 28, 27, 19, 26]. Studies of the feasibility of query re-optimization [8, 7], or deferring optimization to execution time [24], have until recently largely been based on the premise that the need for such techniques is due either to recovering from estimation errors at optimization time in the former case, or avoiding the problem entirely by performing all optimization on-the-fly, such as with Eddies [6] rather than in a staged, ‘waterfall’ kind of paradigm.

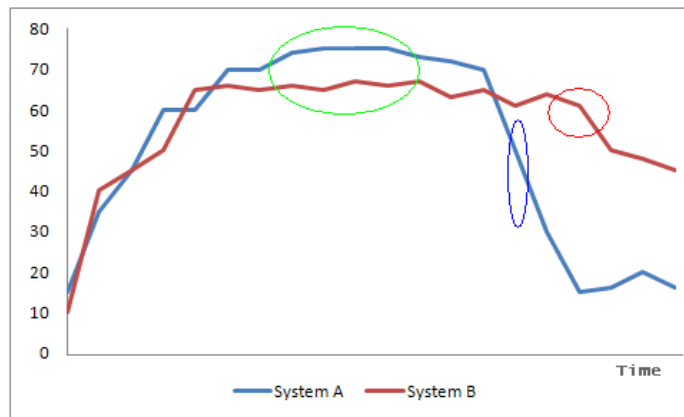
More recent work on adaptive query processing [13, 25, 14, 24] has considered techniques to handle the interaction of query workloads [3, 4, 5], coupled with the realization that changes to environmental conditions can significantly impact a query’s chosen execution plan. These environmental conditions include:

- changes to the amount of memory available (buffer pool, heap memory);
- changes to I/O bandwidth due to concurrent disk activity;
- locking and other waits caused by concurrency control mechanisms;
- detected flaws in the currently executing plan;
- number of available CPU cores;
- changes to the server’s multiprogramming level [2];
- changes to physical access paths, such as the availability of indexes, which could be created on the fly;
- congestion with the telecommunications network;
- contents of the server’s buffer pool;
- inter-query interaction (contention on the server’s transaction log, ‘hot’ rows, and so on).

Background – Dagstuhl seminar 10381

Self-managing database technology, which includes automatic index tuning, automatic database statistics, self-correcting cardinality estimation in query optimization, dynamic resource management, adaptive workload management, and many other approaches, while both interesting and promising, tends to be studied in isolation of other server components. At the 2010 Dagstuhl Workshop on Robust Query Processing (Dagstuhl seminar 10381) held on 19–24 September 2010, seminar attendees tried to unify the study of these technologies in three fundamental ways:

1. determine approaches for evaluating these technologies in the ‘real’ environment where these independently-developed components would interact;
2. establish a metric with which to measure the ‘robustness’ of a database server, making quantitative evaluations feasible so as to compare the worthiness of particular approaches. For example, is dynamic join reordering during query execution worth more than cardinality estimation feedback from query execution to query optimization?



■ **Figure 1** Comparison of Systems A and B in response to increasing workloads over time.

- utilize a metric, or metrics, to permit the construction of regression tests for particular systems. The development of suitable metrics could lead to the development of a new, possibly industry-standard benchmark, that could be used to measure self-managing database systems by industry analysts, customers, vendors, and academic researchers and thus lead to better improvements in robust operation.


At the 2010 Dagstuhl seminar, attendees struggled somewhat with trying to define the notion of robustness, let alone trying to measure or quantify it. Robustness is, arguably, somewhat orthogonal to absolute performance; what we are trying to assess is a system's ability to continue to operate in the face of changing workloads, system parameters and environmental conditions.

An example of the sorts of problems encountered in trying to define robustness is illustrated in Figure 1. Figure 1 shows the throughput rates of two systems, System A (blue line) and System B (red line), over time, for the same workload. The Y-axis represents the throughput rate, and the X-axis is elapsed time. Over time, the workload steadily increases.

Three areas of the graph are highlighted in Figure 1. The first, in green, shows that as the workload is increased, System A outperforms System B by some margin. That peak performance cannot be maintained, however, as the load continues to be increased. The area in blue shows that once System A becomes overloaded, performance drops precipitously. On the other hand, System B shows a much more gradual degradation (circled in red), offering more robust behaviour than System A but with the tradeoff of not being able to match System A's peak performance.

One can argue that Figure 1 mixes the notions of query processing and workload management. In Figure 2 we simplify the problem further, and consider only simple range queries (using two columns) over a single table, where the (logarithmic) X-axis denotes the size of the result set.

In Figure 2, the yellow line illustrates a table scan: it is robust—it delivers identical performance over all result sets—but with relatively poor performance. The dashed red line is a traditional index-to-index lookup plan: that is, search in secondary index, row fetch out of the primary (clustered) index. This plan is considerably faster for very small selectivities, but becomes considerably poorer with only a marginal decrease in selectivity. The solid red line shows, in comparison, substantial-but-imperfect improvements over the index-to-index technique, due to asynchronous prefetch coupled with sorting batches of



■ **Figure 2** Comparison of access plans for a single table range query.

row pointers obtained from the secondary index. This query execution strategy is available in Microsoft SQL Server 2005. While Figure 2 is just one simple query—one table, range predicates on two columns—Figure 2 illustrates both the magnitude of the problem and the opportunity for improving the robustness of such plans.

At the 2010 Dagstuhl seminar, seminar attendees explored a number of different ways in which to define robustness. One idea was to define a metric for robustness as the accumulated variance in the wall clock times of workloads—or particular queries—or, alternatively, some measure of the distribution of that variance, a 2nd level effect. Since this working definition includes wall clock times, it implicitly includes factors such as optimizer quality, since a robustness metric such as this must include statement execution times. However, while the sophistication of query optimization, and the quality of access plans, is a component of a robust database management system, it is not the only component that impacts any notion of robustness.

This working definition of robustness raised as many questions as answers, and many of these were still unresolved by the end of the workshop. Those questions included:

- Sources of variance in query optimization include the statistics model, the cardinality model, and the cost model, with the latter usually being less critical in practice than the former two. One measure of ‘robustness’ is to assess the accuracy between estimates and actuals. What level of importance should the ‘correctness’ of a query optimizer have on a metric of robustness?
- Which offers more opportunity for disaster—and disaster prevention: robust query optimization or robust query execution?
- Is robustness a global property? Does it make sense to measure robustness at the component level? If so, which components, and what weight should be attached to each?
- Several of the attendees at the 2010 Dagstuhl workshop advocated a two-dimensional tradeoff between actual performance and ‘consistency’. But what is ‘consistency’? Is it merely plan quality, or something more?
- Robustness for who? Expectations are different between product engineers and end users; one should not try to define robustness unless one addresses whose expectations you are

trying to satisfy. Both rely on an idealized model of how a system should behave. Can we define that model? At the same time, what expectations can a user have of a really complex query?

- Is adaptivity the only way to achieve robustness?
- What would a benchmark for robustness attempt to measure?

During the workshop we analyzed these questions from various perspectives. Unfortunately we failed to reach consensus on a clear definition of robustness, how to measure it, and what sorts of tradeoffs to include. Our hope, in this, the second Dagstuhl workshop on Robust Query Processing, is to make additional progress towards clarifying the problems, and possibly make some progress towards defining some general—or specific—approaches to improve DBMS robustness.

2 Table of Contents

Executive Summary
Goetz Graefe, Wey Guy, Harumi A. Kuno, and Glenn Paulley 1

Dagstuhl seminar 12321 7
 Approach 7
 Scope 7
 Qualitative and quantitative measures 7

Potential topics for exploration 9

Promising techniques 10

Achievements and further work 12

Participants 15

3 Dagstuhl seminar 12321

3.1 Approach

At the 2012 Dagstuhl seminar, attendees looked at the problem of robust query processing both more simply and more directly. Rather than attempt to define robustness *per se*, instead the attendees were asked to modify a classical model of query processing so that the result would be more ‘robust’, whatever that might mean. The basic idea was to model query processing as a simplistic, generic, cascading ‘waterfall’ system implementation, which included parsing, followed by query rewrite optimization, then join enumeration, and finally query execution. In detail, this set of cascading query processing functions included the following steps:

- logical database design;
- physical database design;
- SQL parsing;
- plan caching: per query, per parameter, and so on, including recompilation control and plan testing;
- query rewrite optimization;
- query optimization proper, including join enumeration;
- query execution;
- database server resource management: memory, threads, cores, and so on;
- database server storage layer, including the server’s buffer pool and transactional support.

Workshop participants were then invited to propose approaches that would modify this ‘strawman’ query processing model by proposing the modification or transposition of query optimization or query execution phases that, it was hoped, would lead to more robust system behaviour.

3.2 Scope

The organizers purposefully chose to keep the focus of discussion relatively narrow and concentrated on more traditional components of relation database query processing. Hence the scope of our discussions included aspects such as Query optimization, query execution, physical database design, resource management, parallelism, data skew, database updates, and database utilities. Outside of the seminar’s scope were arguably more esoteric query processing topics such as text and image retrieval, workload management, cloud-only issues, map-reduce issues, extract, transform, load processing, and stream query processing.

3.3 Qualitative and quantitative measures

To permit workshop participants to focus their ideas on the costs and benefits of their proposals, the workshop organizers developed a list of questions that each work group needed to address with their specific proposal. The questions were:

1. *What is the decision that’s to be moved within the waterfall model?* Within the waterfall framework workshop participants could make several choices; for example, they could decide to introduce a new query processing layer, rather than (simply) move functionality

- from one layer to another. As a concrete example, one might decide to create a new Proactive Physical Database Design layer within the model, thereby treating physical database design as a process and not as a static constraint imposed on server operation.
2. *Where is the decision made in the original waterfall model? What alternative or additional location do you suggest?* To keep the discussions simple, the organizers decided to avoid getting into complex situations such as ‘parallel’ waterfall implementations. For example, the physical database design phase, typically considered an offline task, could be ‘pushed’ into a periodic query processing phase, but we would stop at considering an ‘online’ physical database design task that could greatly perturb the overall model of the system.
 3. *How do you know that this is a real problem?* Industry representatives at the Workshop stated that periodic workload fluctuations are the reality in commercial environments and that manual intervention to deal with issues as they arise is unrealistic. Even if performance analysis tools are available, the resources required to diagnose and solve database server performance problems are too great. Moreover, workload creation and/or modeling remains often too time-consuming for many database administrators, exacerbating diagnosis issues.
 4. *What is the desired effect on robust performance or robust scalability in database query processing?* Not all proposals may lead to better absolute performance. Rather, the target metric of the proposals is to improve the system’s robustness, however that may be defined. For example, a specific proposal may: (a) better anticipate the workload, (b) inform the DBA ahead of the system’s peak workload, (c) permit the system to be better prepared for the system’s expected workload, (d) support better energy efficiency, (e) improve better worst case expected performance of a given workload, or (f) provide additional insights for a DBA to permit better fault diagnosis.
 5. *What are the conditions for moving this decision?* One of the tasks for each group looking at a specific proposal was to determine the parameters for moving a query processing task from one phase to another. In some cases, decision tasks would be designed to move to other phases only under some conditions—for example, due to the periodicity of the workload—whereas in other cases proposals included ideas to permanently move decisions from one query processing phase to another.
 6. *How do the two decision points differ in terms of available useful information?* For example, in many cases it is feasible to consider the migration of a query processing task to another phase only when the metadata surrounding the execution context is known and complete. For example, the system may have a real, captured workload, captured performance indicators, and estimated and actual query execution costs. For other, the amount of metadata required may be significantly less.
 7. *What are the limitations, i.e., when does the shift in decision time not work or not provide any benefit?* For example, is there a type of workload for which a shift in query processing execution will be counter-productive?
 8. *How much effort (software development) would be required?* Is it feasible to consider implementing this proposed technique within a short-term engineering project? Are the risks of implementing the proposal quantifiable? Are the risks real? Can they be mitigated in some way?
 9. *How can the potential benefit be proven for the first time, e.g., in a first publication or in a first white paper on a new feature?* Can the benefits of the new technique be described sufficiently well in a database systems conference paper? What workloads can be used to demonstrate the proposal’s effectiveness, so that the proposal’s benefit can be independently verified by others in the field?

10. *How can the benefit be protected against contradictory software changes, e.g., with a regression test?* This is a complex and difficult problem due in part to subtle changes in underlying assumptions that are made in virtually all complex software systems.
11. *Can regression tests check mechanisms and policies separately?*

To illustrate a potential re-ordering technique, seminar chair Goetz Graefe used an example of moving statistics collection from the physical database design step into the query optimization phase. This would mean that statistics collection would be performed on an on-demand basis during query optimization, rather than as part of a separate, offline ‘performance tuning’ phase. The positive impact of such a change would be to avoid blind cardinality estimation and access plan choices, at the expense of increasing the query’s compile-time cost. Conditions that could impact the effectiveness of the proposal would be missing statistics relevant to the query, or stale statistics that would yield an inaccurate cost estimate. A proof of concept could include examples of poor index and join order choices using an industry standard benchmark, along with mechanisms to control the creation and refresh of these statistics at optimization time.

4 Potential topics for exploration

On the seminar’s first day, Monday, seminar attendees brainstormed a variety of potential ideas to shift query processing functionality from one query processing phase to another. These included:

1. Admission control moved from workload management to query execution: ‘pause and resume’ functionality in queries and utilities.
2. Index creation moved from physical database design to query execution.
3. Materialized view and auxiliary data structures moved from physical database design to query optimization or query execution.
4. Physical database design (vertical partitioning, columnar storage, and column and/or table compression techniques) moved from physical database design to query execution.
5. Join ordering moved from query optimization to query execution.
6. Join and aggregation algorithm: moved from query optimization to query execution, along with set operations, `DISTINCT`, etc.
7. Memory allocation moved from resource management and query optimization to query execution.
8. Buffer pool policies: move policy control from the storage layer to query optimization and/or query execution.
9. Transaction control: move from the server kernel to the application.
10. Serial execution (1 query at a time): consider moving away from concurrent workload management to fixed, serial execution of all transactions.
11. Query optimization a week ago: move physical database design decisions to the query optimizer by pro-actively, iteratively modifying the database’s physical database design through continuous workload analysis.
12. Consider a minimal (in terms of algorithms, implementation effort, and so on) query execution engine that is robust.
13. Query optimization: move from pre-execution to post-execution.

14. Consider the introduction of a robustness indicator during query execution; the idea is to analyze the robustness or efficacy of the query optimization phase—as a simple example, the standard deviation of query execution times.
15. Develop a comparator for pairs of ‘similar’ queries in order to explain what is different between them, both in terms of SQL semantics and in terms of access plans.
16. Holistic robustness.
17. Parallel plan generation, including degree of parallelism: move from query optimization to query execution.
18. Statistics collection and maintenance: move from physical database design to query optimization.
19. Storage formats: move from physical database design to another phase, possibly query optimization or query execution based on plan cache information.
20. Move refresh statistics, predicate normalization and reordering, multi-query awareness, access path selection, join order, join and aggregation algorithms, cardinality estimation, cost calculation, query transformation, common subexpression compilation, parallel planning, and Halloween protection to within the join sequence.
21. Pre-execution during query optimization; for cardinality estimation and for intermediate results, consider the interleaving of query optimization and query execution.
22. Data structure reorganization: move from query execution to query optimization.
23. Develop a measure of robustness and predictable performance as a cost calculation component for a query optimizer.
24. Statistics refresh: move from physical database design or query optimization to query execution
25. Plan cache management: consider augmenting mechanisms and policies set by the query optimization phase with outcomes from the query execution phase.
26. Plan caching decisions: move from the plan cache manager to within the query execution phase.

Over the remaining days of the 12321 seminar, attendees selected topics from the above list and met to consider these ideas in greater detail, and if possible develop a study approach for each that could take place once the Seminar had concluded.

5 Promising techniques

During the seminar’s remaining days, attendees focused on studying the techniques above and developed concrete plans of study for a variety of approaches that each could improve a DBMS system’s robustness. Of those discussed at Dagstuhl, attendees reached consensus that the following ideas were worthwhile exploring in greater depth once the seminar had concluded:

1. *Smooth operations.* The proposal is to implement a new query execution operator, called SmoothScan. At the query execution level, via the SmoothScan operator the system continuously refines the choices made initially by the query optimizer, being able to switch dynamically between index look-up techniques to table scans and vice-versa in a smooth manner.
2. *Opportunistic query processing.* Instead of executing only a single query plan (which may or may not be adaptable), the proposal takes an opportunistic approach that carefully

chooses and executes multiple query plans in parallel, where the fastest plan at any given stage of the query execution determines the overall execution time.

3. *Run-time join order selection.* Investigate techniques to not only re-order joins at execution time, but utilize the construction of small intermediate results to interleave optimization and execution and reduce the degree of errors in cardinality estimation.
4. *Robust plans.* The authors propose a new class of operators and plans that are ‘smooth’ in that their performance degrades gracefully if the expected and the actual characteristics of the underlying data differ. Smooth operators continuously refine the plan choices made by the query optimizer up-front, in order to provide robust performance throughout a wide range of query parameters.
5. *Assessing robustness of query execution plans.* Develop metrics and techniques for comparing two query execution plans with respect to their robustness.
6. *Testing adaptive execution.* Develop an experimental methodology to assess the impact of cardinality estimation errors on system performance and therefore evaluate, by comparison, the impact of adaptive execution methods.
7. *Pro-active physical design.* Using an underlying assumption of periodicity in most workloads, pro-actively, iteratively modify the database’s physical design through continuous workload analysis.
8. *Adaptive partitioning.* In this proposal, the authors seek to continuously adapt physical design considerations to the current workload. Extending the ideas of *database cracking* [20], the authors propose additional ‘cracking’ techniques to both partition physical media (including SSDs) and to handle multi-dimensional queries over multiple attributes using KD-trees and a *Z*-ordering curve to track related data partitions.
9. *Adaptive resource allocation.* In this proposal the authors make the case for dynamic and adaptive resource allocation: dynamic memory allocation at run-time, and dynamic workload throttling, adaptively control concurrent query execution.
10. *Physical database design without workload knowledge.* In this proposal, the authors look at physical database design decisions during bulk loading. The potential set of design choices are page size/format, sorting, indexing, partitioning (horizontal, vertical), compression, distribution/replication within a distributed environment, and statistics.
11. *Weather prediction.* In an analogy to weather prediction, the authors propose a technique to manage user expectations of system performance through analysis of the current workload and prediction parameters analyzed using the current system state.
12. *Lazy parallelization.* Static optimization involving parallelization carries considerable risk, due to several root causes: (1) the exact amount of work to be done is not known at compile time, (2) data skew can have a significant impact on the benefits of parallelism, and (3) the amount of resources available at run time may not be known at compile time. Instead, the proponents of this approach intend to move parallelism choices from the query optimization phase to the query execution phase. In this scenario, the query optimizer would generate ‘generic’ access plans that could be modified on-the-fly during query execution to increase or decrease the degree of parallelism and alter the server’s resource assignment accordingly.
13. *Pause and resume.* This proposal focused on mechanisms that permit stopping and resuming later with the least amount of work repeated or wasted—or even reverted, using some form of undo recovery in order to reach a point from which the server can resume the operation. While similar in intent to mechanisms that permit, for example, dynamic memory allocation, the policies and mechanisms with this proposal are quite different,

and rely on task scheduling and concurrency control mechanisms in order to permit the resumption of a paused SQL request.

14. *Physical database design in query optimization.* The idea behind this proposal is to move some physical database design decisions to the query optimization phase. For example, one idea is to defer index creation to query optimization time, so that indexes are created only when the optimizer can determine that they are beneficial. While this has the benefit of avoiding static physical database design and workload analysis, there are no guarantees that the system can detect cost and benefit bounds for all decisions and all inputs.

6 Achievements and further work

The 2012 Dagstuhl 12321 Workshop made considerable progress towards the goals of developing more robust query processing behaviour. That progress was made, primarily, by maintaining strict focus on the task of shifting a query processing decision from one phase to another, and then answering the questions listed in Section 3.3. The use of that framework enabled more concentrated discussion on the relative merits of the specific techniques, and at the same time reduced debate about the definition of ‘robust’ behaviour, or how to quantify it.

In 2010, Seminar 10381 dwelled on measurement and benchmarks, and subsequent to the seminar two benchmark papers were published. One, which combined elements of both the TPC-C and TPC-H industry-standard benchmarks and entitled the CH-Benchmark, was published in the 2011 DBTEST Workshop, held in Athens the following summer [12]. The other used the metaphor of an agricultural ‘tractor pull contest’ to create a benchmark to examine several measures related to robustness. Like the CH-Benchmark above, the complete tractor pulling benchmark is described in the Proceedings of the 2011 DBTEST Workshop [22].

In 2012, an achievement of the 2012 Dagstuhl seminar was in enumerating various approaches that could either (1) improve the robustness of a database management system under some criteria, or (2) developing metrics that could be used to measure aspects of a system’s behaviour that could be used to optimize a system’s set of decision points. The organizers are confident that this concrete progress will lead to several publications in the very near future.

In particular, two drafts of potential proposals have already been developed: the first for ‘smooth’ scans and ‘smooth’ operators, and the second regarding proactive physical database design for periodic workloads. The seminar’s organizers are confident that several other proposals discussed at the 12321 Workshop will develop into research projects in the coming months.

However, despite the significant progress made at the 12321 Seminar in developing robust query processing techniques, it became clear during both the 10381 and 12321 seminars that there were no known ways of systematically, but efficiently, testing the robustness properties of a database management system in a holistic way. While there is a fairly obvious connection between robustness—however one might try to define it—and self-managing database management system techniques, testing these systems to ensure that they provide robust behaviour is typically limited in practice to individual unit tests of specific self-managing components. Testing the interaction of these various technologies together, with a production-like workload, remains an unsolved and difficult problem. Indeed, testing and

metrics development remain an unknown factor for many, if not all, of the ideas generated at this latest Workshop.

References

- 1 Ashraf Aboulnaga and Surajit Chaudhuri. Self-tuning histograms: Building histograms without looking at data. In *ACM SIGMOD International Conference on Management of Data*, pages 181–192, Philadelphia, Pennsylvania, May 1999.
- 2 Mohammed Abouzour, Kenneth Salem, and Peter Bumbulis. Automatic tuning of the multiprogramming level in Sybase SQL Anywhere. In *ICDE Workshops*, pages 99–104. IEEE, 2010.
- 3 Mumtaz Ahmad, Ashraf Aboulnaga, Shivnath Babu, and Kamesh Munagala. QShuffler: Getting the query mix right. In *Proceedings of the IEEE International Conference on Data Engineering*, pages 1415–1417, 2008.
- 4 Mumtaz Ahmad, Ashraf Aboulnaga, Shivnath Babu, and Kamesh Munagala. Interaction-aware scheduling of report-generation workloads. *The VLDB Journal*, 20:589–615, August 2011.
- 5 Mumtaz Ahmad, Songyun Duan, Ashraf Aboulnaga, and Shivnath Babu. Predicting completion times of batch query workloads using interaction-aware models and simulation. In *Proceedings of the 14th International Conference on Extending Database Technology*, pages 449–460, New York, NY, USA, 2011. ACM.
- 6 Ron Avnur and Joseph M. Hellerstein. Eddies: Continuously adaptive query processing. In *ACM SIGMOD International Conference on Management of Data*, pages 261–272, 2000.
- 7 Pedro Bizarro, Nicolas Bruno, and David J. DeWitt. Progressive parametric query optimization. *IEEE Transactions on Knowledge and Data Engineering*, 21:582–594, 2009.
- 8 Pedro G. Bizarro. *Adaptive query processing: dealing with incomplete and uncertain statistics*. PhD thesis, University of Wisconsin at Madison, Madison, Wisconsin, 2006.
- 9 Surajit Chaudhuri, Hongrae Lee, and Vivek R. Narasayya. Variance aware optimization of parameterized queries. In *ACM SIGMOD International Conference on Management of Data*, pages 531–542, 2010.
- 10 Surajit Chaudhuri, Vivek R. Narasayya, and Ravishankar Ramamurthy. A pay-as-you-go framework for query execution feedback. *Proceedings of the VLDB Endowment*, 1(1):1141–1152, 2008.
- 11 Mitch Cherniack. *Building Query Optimizers with Combinators*. PhD thesis, Brown University, Providence, Rhode Island, May 1999.
- 12 Richard Cole, Florian Funke, Leo Giakoumakis, Wey Guy, Alfons Kemper, Stefan Krompass, Harumi Kuno, Raghunath Nambiar, Thomas Neumann, Meikel Poess, Kai-Uwe Sattler, Michael Seibold, Eric Simon, and Florian Waas. The mixed workload CH-benCHmark. In *Proceedings of the Fourth International Workshop on Testing Database Systems*, New York, NY, USA, 2011. ACM.
- 13 Amol Deshpande, Zachary G. Ives, and Vijayshankar Raman. Adaptive query processing. *Foundations and Trends in Databases*, 1(1):1–140, 2007.
- 14 Kwanchai Eurviriyankul, Norman W. Paton, Alvaro A. A. Fernandes, and Steven J. Lynden. Adaptive join processing in pipelined plans. In *13th International Conference on Extending Database Technology (EDBT)*, pages 183–194, 2010.
- 15 Parke Godfrey, Jarek Gryz, and Calisto Zuzarte. Exploiting constraint-like data characterizations in query optimization. In *ACM SIGMOD International Conference on Management of Data*, pages 582–592, Santa Barbara, California, May 2001. Association for Computing Machinery.
- 16 Goetz Graefe, Arnd Christian König, Harumi Kuno, Volker Markl, and Kai-Uwe Sattler. Robust query processing. Dagstuhl Workshop Summary 10381, Leibniz-Zentrum für Informatik, Wadern, Germany, September 2010.

- 17 Michael Greenwald. Practical algorithms for self-scaling histograms or better than average data collection. *Performance Evaluation*, 20(2):19–40, June 1996.
- 18 Jarek Gryz, Berni Schiefer, Jian Zheng, and Calisto Zuzarte. Discovery and application of check constraints in DB2. In *Proceedings, Seventeenth IEEE International Conference on Data Engineering*, pages 551–556, Heidelberg, Germany, April 2001. IEEE Computer Society Press.
- 19 Waqar Hasan and Hamid Pirahesh. Query rewrite optimization in STARBURST. Research Report RJ6367, IBM Corporation, Research Division, San Jose, California, August 1988.
- 20 Stratos Idreos, Martin L. Kersten, and Stefan Manegold. Database cracking. In *CIDR*, pages 68–78. www.cidrdb.org, 2007.
- 21 Ihab F. Ilyas, Volker Markl, Peter J. Haas, Paul Brown, and Ashraf Aboulnaga. CORDS: Automatic discovery of correlations and soft functional dependencies. In *ACM SIGMOD International Conference on Management of Data*, pages 647–658, Paris, France, June 2004.
- 22 Martin L. Kersten, Alfons Kemper, Volker Markl, Anisoara Nica, Meikel Poes, and Kai-Uwe Sattler. Tractor pulling on data warehouses. In *Proceedings of the Fourth International Workshop on Testing Database Systems*, pages 7:1–7:6, New York, NY, USA, 2011. ACM.
- 23 Jonathan J. King. QUIST—A system for semantic query optimization in relational databases. In *Proceedings of the 7th International Conference on Very Large Data Bases*, pages 510–517, Cannes, France, September 1981. IEEE Computer Society Press.
- 24 Volker Markl, Vijayshankar Raman, David E. Simmen, Guy M. Lohman, and Hamid Pirahesh. Robust query processing through progressive optimization. In *ACM SIGMOD International Conference on Management of Data*, pages 659–670, 2004.
- 25 Rimma V. Nehme, Elke A. Rundensteiner, and Elisa Bertino. Self-tuning query mesh for adaptive multi-route query processing. In *12th International Conference on Extending Database Technology (EDBT)*, pages 803–814, 2009.
- 26 G. N. Paulley and Per-Åke Larson. Exploiting uniqueness in query optimization. In *Proceedings, Tenth IEEE International Conference on Data Engineering*, pages 68–79, Houston, Texas, February 1994. IEEE Computer Society Press.
- 27 Hamid Pirahesh, Joseph M. Hellerstein, and Waqar Hasan. Extensible/rule based query rewrite optimization in STARBURST. In *ACM SIGMOD International Conference on Management of Data*, pages 39–48, San Diego, California, June 1992. Association for Computing Machinery.
- 28 H. J. A. van Kuijk. The application of constraints in query optimization. Memoranda Informatica 88–55, Universiteit Twente, Enschede, The Netherlands, 1988.

Participants

- Martina-Cezara Albutiu
TU München, DE
- Peter A. Boncz
CWI – Amsterdam, NL
- Renata Borovica
EPFL – Lausanne, CH
- Surajit Chaudhuri
Microsoft – Redmond, US
- Campbell Fraser
Microsoft – Redmond, US
- Johann Christoph Freytag
HU Berlin, DE
- Goetz Graefe
HP Labs – Madison, US
- Ralf Hartmut Güting
FernUniversität in Hagen, DE
- Wey Guy
Independent, US
- Theo Härder
TU Kaiserslautern, DE
- Fabian Hüske
TU Berlin, DE
- Stratos Idreos
CWI – Amsterdam, NL
- Ihab Francis Ilyas
University of Waterloo, CA
- Alekh Jindal
Universität des Saarlandes, DE
- Martin L. Kersten
CWI – Amsterdam, NL
- Harumi Anne Kuno
HP Labs – Palo Alto, US
- Andrew Lamb
Vertica Systems – Cambridge, US
- Allison Lee
Oracle Corporation – Redwood
Shores, US
- Stefan Manegold
CWI – Amsterdam, NL
- Anisoara Nica
Sybase – Waterloo, CA
- Glenn Paulley
Conestoga College –
Kitchener, CA
- Ilia Petrov
TU Darmstadt, DE
- Meikel Poess
Oracle Corp. –
Redwood Shores, US
- Ken Salem
University of Waterloo, CA
- Bernhard Seeger
Universität Marburg, DE
- Krzysztof Stencel
University of Warsaw, PL
- Knut Stolze
IBM Deutschland –
Böblingen, DE
- Florian M. Waas
EMC Greenplum Inc. – San
Mateo, US
- Jianliang Xu
Hong Kong Baptist Univ., CN
- Marcin Zukowski
ACTIAN – Amsterdam, NL



Report from Dagstuhl Seminar 12331

Mobility Data Mining and Privacy

Edited by

Christopher W. Clifton¹, Bart Kuijpers², Katharina Morik³, and Yucel Saygin⁴

1 Purdue University, US, clifton@cs.purdue.edu

2 Hasselt University – Diepenbeek, BE, bart.kuijpers@uhasselt.be

3 TU Dortmund, DE

4 Sabanci University – Istanbul, TR

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 12331 “Mobility Data Mining and Privacy”. Mobility data mining aims to extract knowledge from movement behaviour of people, but this data also poses novel privacy risks. This seminar gathered a multidisciplinary team for a conversation on how to balance the value in mining mobility data with privacy issues. The seminar focused on four key issues: Privacy in vehicular data, in cellular data, context-dependent privacy, and use of location uncertainty to provide privacy.

Seminar 12.–17. August, 2012 – www.dagstuhl.de/12331

1998 ACM Subject Classification K.4.1 Public Policy Issues: Privacy

Keywords and phrases Privacy, Mobility, Cellular, Vehicular Data

Digital Object Identifier 10.4230/DagRep.2.8.16

1 Executive Summary

Chris Clifton

Bart Kuijpers

License  Creative Commons BY-NC-ND 3.0 Unported license
© Chris Clifton and Bart Kuijpers

Mobility Data Mining and Privacy aimed to stimulate the emergence of a new research community to address mobility data mining together with privacy issues. Mobility data mining aims to extract knowledge from movement behaviour of people. This is an interdisciplinary research area combining a variety of disciplines such as data mining, geography, visualization, data/knowledge representation, and transforming them into a new context of mobility while considering privacy which is the social aspect of this area. The high societal impact of this topic is mainly due to the two related facets of its area of interest, i.e., people’s movement behaviour, and the associated privacy implications. Privacy is often associated with the negative impact of technology, especially with recent scandals in the US such as AOL’s data release which had a lot of media coverage. The contribution of *Mobility Data Mining and Privacy* is to turn this negative impact into positive impact by investigating how privacy technology can be integrated into mobility data mining. This is a challenging task which also imposes a high risk, since nobody knows what kinds of privacy threats exist due to mobility data and how such data can be linked to other data sources.

The seminar looked closely at two application areas: Vehicular data and cellular data. Further discussions covered two specific new general approaches to protecting location privacy: context-dependent privacy, and location uncertainty as a means to protect privacy. In each



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

Mobility Data Mining and Privacy, *Dagstuhl Reports*, Vol. 2, Issue 8, pp. 16–53

Editors: Christopher W. Clifton, Bart Kuijpers, Katharina Morik, and Yucel Saygin



DAGSTUHL
REPORTS

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

of these areas, new ideas were developed; further information is given in the working group reports.

The seminar emphasized discussion of issues and collaborative development of solutions – the majority of the time was divided between working group breakout sessions followed by report-back and general discussion sessions. While the working group reports were written by subgroups, the contents reflect discussions involving all 22 participants of the seminar.

The seminar concluded that there are numerous challenges to be addressed in mobility data mining and privacy. These challenges require investigation on both technical and policy levels. Of particular importance is educating stakeholders from various communities on the issues and potential solutions.

2 Table of Contents

Executive Summary

<i>Chris Clifton and Bart Kuijpers</i>	16
--	----

Overview of Talks

Dynamic privacy adaptation in ubiquitous computing <i>Florian Schaub</i>	19
Privacy-Aware Spatio-Temporal Queries on Unreliable Data Sources <i>Erik Buchmann</i>	20
Privacy-preserving sharing of sensitive semantic locations under road-network constraints <i>Maria-Luisa Damiani</i>	20
Methods of Analysis of Episodic Movement Data <i>Thomas Liebig</i>	20
Privacy-preserving Distributed Monitoring of Visit Quantities <i>Christine Körner</i>	21
A visual analytics framework for spatio-temporal analysis and modelling <i>Gennady Andrienko</i>	22
Tutorial: Privacy Law <i>Nilgun Basalp</i>	22
“Movie night”: Tutorial on Differential Privacy <i>Christine Task (video of talk by non-participant)</i>	22

Working Groups

Working Group: Cellular Data <i>Gennady Andrienko, Aris Gkoulalas-Divanis, Marco Gruteser, Christine Körner, Thomas Liebig, Klaus Rechert, and Michael Marhöfer</i>	23
Working Group: Vehicular Data <i>Glenn Geers, Marco Gruteser, Michael Marhoefer, Christian Wietfeld, Claudia Sánta, Olaf Spinczyk, and Ouri Wolfson</i>	32
Working Group: Context-dependent Privacy in Mobile Systems <i>Florian Schaub, Maria Luisa Damiani, and Bradley Malin</i>	36
Working Group: Privacy through Uncertainty in Location-Based Services <i>Nilgün Basalp, Joachim Biskup, Erik Buchmann, Chris Clifton, Bart Kuijpers, Walied Othman, and Erkay Savas</i>	44

Open Problems

What we learned	51
New Discoveries	51
What needs to be done	52
Future plans	52


Participants	53
-------------------------------	----

3 Overview of Talks

We tried to maximize the time spent in discussions, which limited the time available for talks. However, we did have a few talks, primarily from young researchers. We also had some short tutorial talks given as the need arose based on the direction of the ongoing discussions. Titles and brief abstracts are given below.

3.1 Dynamic privacy adaptation in ubiquitous computing

Florian Schaub, (Universität Ulm, DE)

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

Ubiquitous and pervasive computing is characterized by the integration of computing aspects into the physical environment. Physical and virtual worlds start to merge as physical artifacts gain digital sensing, processing, and communication capabilities. This development introduces a number of privacy challenges. Physical boundaries lose their meaning in terms of privacy demarcation. Furthermore, the tight integration with the physical world necessitates the consideration not only of observation and information privacy aspects but also disturbances of the user's physical environment [1].


By viewing privacy as a dynamic regulation process and developing respective privacy mechanisms, we aim to support users in ubiquitous computing environments in gaining privacy awareness, making informed privacy decisions, and controlling their privacy effectively. The presentation outlined our dynamic privacy adaptation process for ubiquitous computing that supports privacy regulation based on the user's current context [2]. Furthermore, our work on a higher level privacy context model [3] has been presented. The proposed privacy context model captures privacy-relevant context features and facilitates the detection of privacy-relevant context changes in the user's physical and virtual environment. When privacy-relevant context changes occur, an adaptive privacy system can dynamically adapt to the changed situation by reasoning about the context change and learned privacy preferences of an individual user. Individualized privacy recommendations can be offered to the user or sharing behavior can be automatically adjusted to help the user maintain a desired level of privacy.

References

- 1 Bastian Könings, Florian Schaub, "Territorial Privacy in Ubiquitous Computing", Proc. 8th Int. Conf. on Wireless On-demand Network Systems and Services (WONS '11), IEEE 2011 DOI: 10.1109/WONS.2011.5720177
- 2 Florian Schaub, Bastian Könings, Michael Weber, Frank Kargl, "Towards Context Adaptive Privacy Decisions in Ubiquitous Computing", Proc. 10th Int. Conf. on Pervasive Computing and Communications (PerCom '12), Work in Progress, IEEE 2012 DOI: 10.1109/PerComW.2012.6197521
- 3 Florian Schaub, Bastian Könings, Stefan Dietzel, Michael Weber, Frank Kargl, "Privacy Context Model for Dynamic Privacy Adaptation in Ubiquitous Computing", Proc. 6th Int. Workshop on Context-Awareness for Self-Managing Systems (Casemans '12), Ubicomp '12 workshop, ACM 2012

3.2 Privacy-Aware Spatio-Temporal Queries on Unreliable Data Sources


Erik Buchmann (KIT – Karlsruhe Institute of Technology, DE)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Erik Buchmann

A declarative spatio-temporal query processor is an important building block for many kinds of location-based applications. Such applications often apply methods to obfuscate, anonymize or delete certain spatio-temporal information for privacy reasons. However, the information that some data has been modified is privacy-relevant as well. This talk is about hiding the difference between spatio-temporal data that has been modified for privacy reasons, and unreliable information (e.g., missing values or sensors with a low precision), on the semantics level of a query processor. In particular, we evaluate spatio-temporal predicate sequences like Enter (an object was outside of a region first, then on the border, then inside) to true, false, maybe. This allows a wide range of data analyses without making restrictive assumptions on the data quality or the privacy methods used.

3.3 Privacy-preserving sharing of sensitive semantic locations under road-network constraints

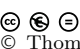
Maria-Luisa Damiani (University of Milano, IT)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Maria-Luisa Damiani

This talk illustrates recent research on the protection of sensitive positions in real time trajectories under road network constraints

3.4 Methods of Analysis of Episodic Movement Data

Thomas Liebig (Fraunhofer IAIS – St. Augustin, DE)

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

Analysis of people's movements represented by continuous sequences of spatio-temporal data tuples received lots of attention within the last years. Focus of the studies was mostly GPS data recorded on a constant sample rate. However, creation of intelligent location aware models and environments requires also reliable localization in indoor environments as well as in mixed indoor outdoor scenarios. In these cases, signal loss makes usage of GPS infeasible, therefore other recording technologies evolved.

Our approach is analysis of episodic movement data. This data contains some uncertainties among time (continuity), space (accuracy) and number of recorded objects (coverage). Prominent examples of episodic movement data are spatio-temporal activity logs, cell based tracking data and billing records. To give one detailed example, Bluetooth tracking monitors presence of mobile phones and intercoms within the sensors footprints. Usage of multiple sensors provides flows among the sensors.

Most existing data mining algorithms use interpolation and therefore are infeasible for this kind of data. For example, speed and movement direction cannot be derived from episodic data; trajectories may not be depicted as a continuous line; and densities cannot be computed.

Though this data is infeasible for individual movement or path analysis, it bares lots of information on group movement. Our approach is to aggregate movement in order to overcome the uncertainties. Deriving number of objects for spatio-temporal compartments and transitions among them gives interesting insights on spatio-temporal behavior of moving objects. As a next step to support analysts, we propose clustering of the spatio-temporal presence and flow situations. This work focuses as well on creation of a descriptive probability model for the movement based on Spatial Bayesian Networks.

We present our methods on real world data sets collected during a football game in Nîmes, France in June 2011 and another one in Dusseldorf, Germany 2012. Episodic movement data is quite frequent and more methods for its analysis are needed. To facilitate method development and exchange of ideas, we are willing to share the collected data and our findings.

3.5 Privacy-preserving Distributed Monitoring of Visit Quantities


Christine Körner (Fraunhofer IAIS – St. Augustin, DE)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Christine Körner

The organization and planning of services (e.g. shopping facilities, infrastructure) requires up-to-date knowledge about the usage behavior of customers. Especially quantitative information about the number of customers and their frequency of visiting is important. In this paper we present a framework which enables the collection of quantitative visit information for arbitrary sets of locations in a distributed and privacy-preserving way. While trajectory analysis is typically performed on a central database requiring the transmission of sensitive personal movement information, the main principle of our approach is the local processing of movement data. Only aggregated statistics are transmitted anonymously to a central coordinator, which generates the global statistics. In this presentation we introduce our approach including the methodical background that enables distributed data processing as well as the architecture of the framework. We further discuss our approach with respect to potential privacy attacks as well as its application in practice. We have implemented the local processing mechanism on an Android mobile phone in order to ensure the feasibility of our approach.

3.6 A visual analytics framework for spatio-temporal analysis and modelling

Gennady Andrienko (Fraunhofer IAIS – St. Augustin, DE)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Gennady Andrienko


Main reference N. Andrienko, G. Andrienko, “A Visual Analytics Framework for Spatio-temporal Analysis and Modelling,” *Data Mining and Knowledge Discovery*, 2013 (accepted).

URL <http://dx.doi.org/10.1007/s10618-012-0285-7>

To support analysis and modelling of large amounts of spatio-temporal data having the form of spatially referenced time series (TS) of numeric values, we combine interactive visual techniques with computational methods from machine learning and statistics. Clustering methods and interactive techniques are used to group TS by similarity. Statistical methods for TS modelling are then applied to representative TS derived from the groups of similar TS. The framework includes interactive visual interfaces to a library of modelling methods supporting the selection of a suitable method, adjustment of model parameters, and evaluation of the models obtained. The models can be externally stored, communicated, and used for prediction and in further computational analyses. From the visual analytics perspective, the framework suggests a way to externalize spatio-temporal patterns emerging in the mind of the analyst as a result of interactive visual analysis: the patterns are represented in the form of computer-processable and reusable models. From the statistical analysis perspective, the framework demonstrates how TS analysis and modelling can be supported by interactive visual interfaces, particularly, in a case of numerous TS that are hard to analyse individually. From the application perspective, the framework suggests a way to analyse large numbers of spatial TS with the use of well-established statistical methods for TS analysis.

3.7 Tutorial: Privacy Law


Nilgun Basalp (Istanbul Bilgi University, TR)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Nilgun Basalp

This session is a brief tutorial on the EC 95/46 privacy directive, with a focus on issues affecting mobile data.

3.8 “Movie night”: Tutorial on Differential Privacy

Christine Task (video of talk by non-participant)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Christine Task (video of talk by non-participant)

Differential Privacy is a relatively new approach to privacy protection, based on adding sufficient noise to query results on data to hide information of any single individual. This tutorial, given as a CERIAS seminar at Purdue in April, is an introduction to Differential Privacy targeted to a broad audience. Several of the seminar participants gathered for a “movie night” to watch a video of this presentation.




The presenter, Christine Task, is a Ph.D. student at Purdue University working with Chris Clifton. She has a B.S. in theoretical math from Ohio State and M.S. in computer science from Indiana University.

http://www.cerias.purdue.edu/news_and_events/events/security_seminar/details/index/j9cvs3as2h1qds1jrdqfdc3hu8

4 Working Groups

4.1 Working Group: Cellular Data

Gennady Andrienko, Aris Gkoulalas-Divanis, Marco Gruteser, Christine Körner, Thomas Liebig, Klaus Rechert, and Michael Marhöfer

License    Creative Commons BY-NC-ND 3.0 Unported license
© Gennady Andrienko, Aris Gkoulalas-Divanis, Marco Gruteser, Christine Körner, Thomas Liebig, Klaus Rechert, and Michael Marhöfer

4.1.1 Introduction

The phones we carry around as we go about our daily lives do not only provide a convenient way to communicate and access information, but also pose privacy risks by collecting data about our movements and habits. For example, they can record when we get up in the morning, when we leave our homes, whether we violate speed limits, how much time we spend at work, how much we exercise, whom we meet, and where we spend the night. The places we visit allow inferences about not just one, but many potentially sensitive subjects: health, sexual orientation, finances or creditworthiness, religion, and political opinions. For many, such inferences can be embarrassing, even if they are untrue and simply misinterpretations of the data. For some, this movement data can even pose a danger of physical harm, such as in stalking cases.

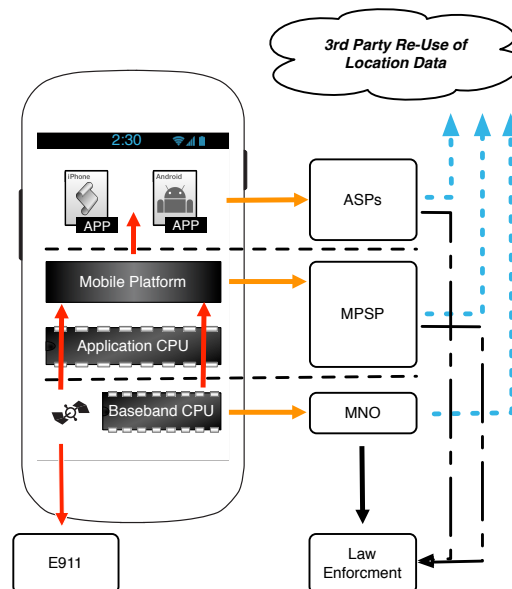
These risks have been amplified by the emergence of smartphones and the app economy over the last few years. We have witnessed a fundamental shift in mobility data collection and processing from a selected group of tightly regulated cellular operators to a complex web of app providers and Internet companies. This new ecosystem of mobility data collectors relies on a more sophisticated mix of positioning technologies to acquire increasingly precise mobility data. In addition, smartphones also carry a much richer set of sensors and input devices, which allow collection of a diverse set of other data types in combination with the mobility data. Many of these types of data were previously unavailable. While individual aspects of these changes have been highlighted in a number of articles as well as in a string of well-publicized privacy scandals, the overall structure of current mobility data streams remains confusing.

The goal of the Dagstuhl cellular data working group was to survey this new mobility data ecosystem and to discuss the implications of this broader shift. Section 4.1.2 gives an overview about the types of data collected by mobile network operators (MNO), mobile platform service providers (MPSP) and application service providers (ASP). In Section 4.1.3 we discuss privacy threats and risks that arise from the data collection. We conclude our work with a section on the manifold implications of this rapidly evolving mobility data ecosystem. We find that it is difficult to understand the data flows, apparently even for the service providers and operators themselves [5, 13, 18]. There appears to be a much greater need for transparency, perhaps supported by technical solutions that monitor and raise awareness

of such data collection. We find that location data is increasingly flowing across national borders, which raises questions about the effectiveness of current regulatory protections. We also find that applications are accessing a richer set of sensors, which allows cross-referencing and linking of data in ways that are not yet fully understood.

4.1.2 Location Data Collection

We distinguish three groups of data collectors (observers): mobile network operators (MNO), mobile platform service providers (MPSP), and application service providers (ASP). While for the first group of observers (MNOs), location data is generated and collected primarily due to technical reasons, i.e. efficient signaling, in the case of MPSP and ASPs location information is usually generated and collected to support positioning, mapping, and advertising services and to provide various kinds of location based services. Figure 1 provides a schematic overview on location data generated by mobile phones but also highlights the specific components and building blocks of mobile phones which are controlled by the different entities. Furthermore, available location data originating from the aforementioned layers may be re-used to support various new (third-party) businesses. Typically the data is then anonymized or aggregated in some way before being transferred to third parties.



■ **Figure 1** A schematic overview on a today's smartphone, its essential building blocks and their controllers illustrating the generation and information flow of location data.

Most of the data collected by MNO, MPSP and ASP can be referred as 'Episodical Movement Data': data about spatial positions of moving objects where the time intervals between the measurements may be quite large and therefore the intermediate positions cannot be reliably reconstructed by means of interpolation, map matching, or other methods. Such data can also be called 'temporally sparse'; however, this term is not very accurate since the temporal resolution of the data may greatly vary and occasionally be quite fine.

4.1.2.1 Collection and Usage of Mobile Telephony Network Data

As an example for mobile telephony networks we discuss the widely deployed GSM infrastructure, as its successors UMTS (3G) and LTE (4G) have a significantly smaller coverage and share most of its principal characteristics. A typical GSM network is structured into cells, each served by a single base transceiver station (BTS). Larger cell-compounds are called location areas. To establish a connection to the mobile station (MS) e.g. in the case of an incoming connection request, the network has to know if the MS is still available and in which location area it is currently located. To cope with subscriber mobility the location update procedure was introduced. Either periodically or when changing the location area, a location update is triggered. The time lapse between periodic location updates is defined by the network and varies between infrastructure providers.

Additionally, the infrastructure's radio subsystem measures the distance of phones to the serving cell to compensate for the signal propagation delay between the MS and BTS. The timing advance (TA) value (8-bit value) is used to split the cell radius into virtual rings. In the case of GSM these rings have a size of roughly 550 m in diameter. The TA is regularly updated and is sent by the serving infrastructure to each mobile phone.

For billing purposes so-called call data records (CDR) are generated. This datum usually consists of the cell-ID where a call has been started (either incoming or outgoing), the cell where a call has been terminated, start time, duration, ID of the caller and the phone number called. A typical GSM cell size ranges from a few hundred meters in diameter to a maximum size of 35 km. In a typical network setup a cell is further divided into three sectors. In that case also the sector ID is available, and the sector ID is also part of a call record. CDRs are usually not available in real-time. However, MNOs store CDRs for a certain time span, either because of legal requirement (e.g. EU data retention directive [9]) or accounting purposes.

Mobile telephony networks and their physical characteristics are able to help locating mobile phone users in the case of an emergency and may be a valuable tool for search and rescue (SAR) [21]. For instance, [6] analyzed post-disaster populations displacement using SIM-card movements in order to improve the allocation of relief supplies.

Furthermore, location information gathered through mobile telephony networks is now a standard tool for crime prosecution and is used by the EC Data Retention Directive with the aim of reducing the risk of terror and organized crime [9]. As an example, law enforcement officials seized CDRs over a 48 hour timespan resulting in 896,072 individual records containing 257,858 call numbers after a demonstration in Dresden, Germany, went violent [20]. Further, the police of North Rhine-Westphalia issued 225,784 active location determinations on 2,644 different subjects in 778 preliminary proceedings in 2010 [23]. While in principle law enforcement could also collect location- and movement-data from MPSP and ASPs, difficulties arise if such data is stored outside of the respective jurisdiction.

Additionally, commercial services are based on the availability of live mobility patterns of larger groups. (e.g. for traffic monitoring or location-aware advertising [19]). Thus, location information of network subscribers might be passed on to third parties. Usually, subscribers are neither aware of the extent of their information disclosure (just by carrying a switched-on mobile phone), nor of how the collected data is used and by whom. Even sporadic disclosure of location data, e.g. through periodic location updates, are able to disclose a users frequently visited places (i.e. preferences) in an accuracy similar to continuous location data after 10-14 days [25].

4.1.2.2 Collection and Usage of Data Through MPSPs and ASPs

Over the past years, researchers and journalists have started to analyze apps and mobile operating systems w.r.t. the collection of personal data [12, 16, 7, 1, 26, 27, 5]. The analyses show that sensitive information is accessed and transmitted. There are mainly three reasons for MPSP to collect location information: positioning, map services, and advertising.

Even though a mobile phone may not be equipped with GPS, a position may be obtained by approximate location determination based on mobile telephony infrastructure or WiFi. The sensor data is sent to external services and external information sources are used to improve (i.e. speed-up) the determination of the user's current location. For instance, in Spring 2011 it was found that Apple's iPhone generates and stores a user's location history, more specifically, data records correlating visible WiFi access-points or mobile telephony cell-ids with the device's GPS location on the user's phone. Moreover, the recorded data-sets are frequently synchronized with the platform provider. Presumably, this data is used by MPSPs to improve database-based, alternative location determination techniques for situations where GNSS or similar techniques are not available or not operational.

By aggregating location information of many users, such information could improve or enable new kinds of services. For instance, Google Mobile Maps makes use of user contributed data (with the user's consent) to determine and visualize the current traffic situation.

Finally, mobile advertising is one of the fastest growing advertising media, doubling its yearly revenue over the next years by a prediction of [10]. The availability of smartphones in combination with comprehensive and affordable mobile broadband communication has given rise to this new generation of advertising media, which allows to deliver up-to-date information in a context-aware and personalized manner. However, personal information as a user's current location and personal preferences are prerequisite for a tailored advertisement delivery. Consequently, MPSPs and ASPs are interested to profile users and personal data is disclosed in an unprecedented manner to various (unknown) commercial entities which poses serious privacy risks.

In order to perform dedicated tasks apps, also access other data such as the user's contacts, calendar and bookmarks as well as sensors readings (e.g. camera, microphone). If these apps have access to the Internet, they are potentially able to disclose this information and are a serious thread to user privacy [16]. Most often, advertisement libraries (e.g., as part of an app) require access to the phone information and location API [12] in order to obtain the phone's IMEI number and geographic position. For instance, Apple Siri records, stores and transmits any spoken request to Apple's cloud-based services where it is processed through speech recognition software, is analyzed to be understood, and is subsequently serviced. The computed result of each request is communicated back to the user. Additionally, to fully support inferencing from context, Siri is "expected to have knowledge of users' contact lists, relationships, messaging accounts, media (songs, playlists, etc) and more"¹, including location data to provide the context of the request, which are communicated to Apple's data center. As an example of Siri's use of location data, users are able to geo-tag familiar locations (such as their home or work) and set a reminder when they visit these locations. Moreover, user location data is used to enable Siri to support requests for finding the nearest place of interest (e.g., restaurant) or to report the local weather.

¹ <http://privacycast.com/siri-privacy-and-data-collection-retention/>, Online, Version of 9/6/2012

4.1.3 Privacy Threats and Risks

From a business perspective, mobility data with sufficiently precise location estimation are often valuable data enabling various location-based services; from the perspective of privacy advocates, such insights are often deemed a privacy threat or a privacy risk. Location privacy risks can arise if a third-party acquires a data tuple (user ID, location), which proves that an identifiable user has visited a certain location. In most cases, the datum will be a triple that also includes a time field describing when the user was present at this location. Note that in theory there are no location privacy risks if the user cannot be identified or if the location cannot be inferred from the data. In practice, however, it is difficult to determine when identification and such inferences are possible.

4.1.3.1 Collection of Location Information with Assigned User ID

Collecting location information along with a user id is the most trivial case of observing personal movement information, as long as the location of the user is estimated with sufficient accuracy for providing the intended LBS. In case the location is not yet precise enough, various techniques (e.g. fusion of several raw location data from various sensors) allow for improving the accuracy.

Additionally, ASP may have direct access to a variety of publicly available spatial and temporal data such as geographical space and inherent properties of different locations and parts of the space (e.g. street vs. park), or various objects existing or occurring in space and time: static spatial objects (having particular constant positions in space), events (having particular positions in time) and moving objects (changing their spatial positions over time). Such information either exists in explicit form in public databases like OSM, WikiMapia or in ASP's data centers, or can be extracted from publicly available data by means of event detection or situation similarity assessment [3][4]. Combining such information with positions and identities of users allow deep semantic understanding of their habits, contacts, and lifestyle.

4.1.3.2 Collection of Anonymous Location Information

When location data is collected without any obvious user identifiers, privacy risks are reduced and such seemingly anonymous data is usually exempted from privacy regulations. It is, however, often possible to re-identify the user based on quasi-identifying data that has been collected. Therefore, the aforementioned risks can apply even to such anonymous data.

The degree of difficulty in re-identifying anonymous data depends on the exact details of the data collection and anonymization scheme and the adversaries access to background information. Consider the following examples:

Re-identifying individual samples. Individual location records can be re-identified through observation re-identification [22]. The adversary knows that user Alice was the only user in location (area) l at time t , perhaps because the adversary has seen the person at this location or because records from another source prove it. If the adversary now finds an anonymous datum (l, t) in the collected mobility data, the adversary can infer that this datum could only have been collected from Alice and has re-identified the data. In this trivial example, there is actually no privacy risk from this re-identification because the adversary knew a priori that Alice was at location l at time t , so the adversary has not learned anything new. There are, however, three important variants of this trivial case that can pose privacy risks. First, the anonymous datum may contain a more precise location l' or a more precise time t' than the adversary knew about a priori. In this case, the adversary learns this more

precise information. Second, the adversary may not know that Alice was at l but simply know that Alice is the only user who has access to location l . In this latter case, also referred to as restricted space identification, the adversary would learn when Alice was actually present at this location. Third, the anonymous datum may contain additional fields with potentially sensitive information that the adversary did not know before. Note, however, that such additional information can also make the re-identification task easier.

Re-identifying time-series location data. Re-identification can also become substantially easier when location data is repeatedly collected and time-series location traces are available. We refer to time-series location traces, rather than individual location samples when it is clear which set of location samples was collected from the same user (even though the identity of the user is not known). For example, the location data may be stored in separate files for each user or a pseudonym may be used to link multiple records to the same user.

Empirical research [11] has further observed that the pair (home location, work location) is often already identifying a unique user. A recent empirical study [29] explains various approaches for re-identification of a user. Another paper has analyzed the consequences of increasingly strong re-identification methods to privacy law and its interpretation [24]. Further re-identification methods for location data rely on various inference and data mining techniques.

4.1.3.3 Collection of Data without Location

Even in absence of actual location readings provided by positioning devices, location disclosures may occur by means of other modern technologies. Recent work by Han, et al. [17] demonstrated that the complete trajectory of a user can be revealed with a 200 m accuracy by using accelerometer readings, even when no initial location information is known. What is even more alarming is that accelerometers, typically installed in modern smartphones, are usually not secured against third-party applications, which can easily obtain such readings without requiring any special privileges. Acceleration information can thus be transmitted to external servers and be used to disclose user location even if all localization mechanisms of the mobile device are disabled.

Furthermore, several privacy vulnerabilities may be exposed through the various resource types that are typically supported and communicated by modern mobile phone applications. Hornyack, et al. [16] examined several popular Android applications which require both internet access and access to sensitive data, such as location, contacts, camera, microphone, etc. for their operation. Their examination showed that almost 34% of the top 1100 popular Android applications required access to location data, while almost 10% of the applications required access to the user contacts. As can be anticipated, access of third-party applications to such sensitive data sources may lead to both user re-identification as well as sensitive information disclosure attacks, unless privacy enabling technology is in place.

4.1.4 Implications

Potentially sensitive location data from the use of smartphones is now flowing to a largely inscrutable ecosystem of international app and mobile platform providers, often without knowledge of the data subject. This represents a fundamental shift from the traditional mobile phone system, where location data was primarily stored at more tightly regulated cellular carriers that operated within national borders.

A large number of apps customize the presented information or their functionality based on user location. Examples of such apps include local weather information, location-based

reminders, maps and navigation, restaurant rating, and friend finders. Such apps often transmit the user location to a server, where it may be stored for a longer duration.

It is particularly noteworthy, however, that mobile advertisers and platform providers have emerged as an additional entity that aggregates massive sets of location records obtained from user interactions with a variety of apps. When apps request location information, the user location can also be disclosed to the mobile platform service provider as part of the wireless positioning service function. Even apps that do not need any location information to function, often reveal the user location to mobile advertisers. The information collected by these advertising and mobile providers is arguably more precise than the call data records stored by cellular carriers, since it is often obtained via WiFi positioning or the GPS. In addition, privacy notices by app providers often neglect to disclose such background data flows [1]. While the diversity of location-based apps has been foreseen by mobile privacy research to some extent—for example, research on spatial cloaking [14] has sought to provide privacy-preserving mechanisms for sharing location data with a large number of apps—this aggregation of data at mobile platform providers was less expected. In essence, this development is for economic reasons. Personal location information has become a tradable good: users provide personal information for targeted advertising in exchange for free services (quite similar to web-based advertising models). The advertising revenue generated out of such data, finances the operation of the service provider. Because of this implicit bargain between users and service providers, there is little incentive to curb data flows or adopt stronger technical privacy protections as long as it is not demanded by users or regulators.

We suspect, however, that many users are not fully aware of this implicit bargain. Therefore, we believe that it is most important from a privacy perspective to create awareness of these data flows among users, which is not incidentally the very first core principle of the fair information practice principles [28]. It is well understood that lengthy privacy disclosures, if they exist for smartphone apps, are not very effective at reaching the majority of users and even the recent media attention regarding smartphone privacy² does not appear to have found a sufficiently wide audience as our workshop discussions suggest. Raising awareness and empowering users to make informed decisions about their privacy will require novel approaches, user-interfaces, and tools.

When using smartphones, users should not only be aware of *what* data they are revealing to third-parties and how frequently it is revealed but also should be able to understand the potential risks of sharing such data. For instance, users/subscribers in the EU are currently entitled to get a full copy of their personal data stored by a commercial entity³ but such voluminous datasets can currently only be analyzed by experts. Even then, it will be difficult to judge what sensitive information can be learned from this dataset when it is linked with other data about the same person or when it is analyzed by a human expert with powerful visual analysis tools [2]. Was the precision of a location record sufficient to determine the building that a user has entered? Is it possible to reconstruct the path a users has taken between two location records? How easily can one infer habits or health of a person based on the location records collected from smartphones?

As another example, some service providers claim to collect location data only in anonymous form. The methods for re-identification, however, have evolved quickly. When can 'anonymized' time-series location data really qualify as data that is not personally identifiable information and remain outside most current privacy regulations? Finally, even

² For instance, <http://blogs.wsj.com/wtk-mobile/> Retrieved 2012/10/18.

³ For example, an Austrian student requested all personal data from Facebook and received a CD [15]

non-georeferenced data provided by the sensors embedded in a smartphone (camera, accelerometer, microphone, etc.), as well as the files stored in the internal memory (photos, music, playlists), allow extracting knowledge about a person's location and mobility. Overall, it appears necessary to investigate what associations can be established and what inferences can be made by a human when the data is considered in context, and how such information can be conveyed to users of services.

Users should also be able to learn in which countries their data is stored or processed, since this can have important implications for the applicable legal privacy framework. While the European Union has achieved some degree of harmonization of privacy standards for exported data from its citizens through the safe harbor provisions [8], differences still exist, for example, with respect to law enforcement access to user data. We believe that providing transparency of cross-border data flows would lead to a more meaningful public discussion of data protection policies. For example, when data is handled by multi-national corporations, should data subjects be given a choice where their data is processed and stored?

We hope that the research community will help address these questions and will interface with data protection authorities and policy experts to actively define privacy for this mobility data ecosystem.

References


- 1 Mobile apps for kids: Current privacy disclosures are disappointing. Technical report, Federal Trade Commission, 2012. http://www.ftc.gov/os/2012/02/120216mobile_apps_kids.pdf.
- 2 G. Andrienko and N. Andrienko. Privacy issues in geospatial visual analytics. In Georg Gartner, Felix Ortog, William Cartwright, Georg Gartner, Liqiu Meng, and Michael P. Peterson, editors, *Advances in Location-Based Services*, Lecture Notes in Geoinformation and Cartography, pages 239–246. Springer Berlin Heidelberg, 2012.
- 3 G. L. Andrienko, N. V. Andrienko, C. Hurter, S. Rinzivillo, and S. Wrobel. From movement tracks through events to places: Extracting and characterizing significant places from mobility data. In *IEEE VAST*, pages 161–170. IEEE, 2011.
- 4 G. L. Andrienko, N. V. Andrienko, M. Mladenov, M. Mock, and C. Pölitiz. Identifying place histories from activity traces with an eye to parameter impact. *IEEE Trans. Vis. Comput. Graph.*, 18(5):675–688, 2012.
- 5 C. Arthur. iphone keeps record of everywhere you go. *The Guardian*, 2011. <http://www.guardian.co.uk/technology/2011/apr/20/iphone-tracking-prompts-privacy-fears>.
- 6 L. Bengtsson, X. Lu, A. Thorson, R. Garfield, and J. von Schreeb. Improved response to disasters and outbreaks by tracking population movements with mobile phone network data: A post-earthquake geospatial study in haiti. *PLoS Med*, 8(8):e1001083, 08 2011.
- 7 W. Enck, P. Gilbert, B.-G. Chun, L. P. Cox, J. Jung, P. McDaniel, and A. N. Sheth. Taintdroid: an information-flow tracking system for realtime privacy monitoring on smartphones. In *Proceedings of the 9th USENIX conference on Operating systems design and implementation*, OSDI'10, pages 1–6, Berkeley, CA, USA, 2010. USENIX Association.
- 8 European Parliament and European Council. Directive 95/46/ec on the protection of individuals with regard to the processing of personal data and on the free movement of such data. *Official Journal of the European Communities*, (L281), 1995.
- 9 Council European Parliament. Directive 2006/24/ec of the european parliament and of the council of 15 march 2006 on the retention of data generated or processed in connection with the provision of publicly available electronic communications services or of public communications networks and amending directive 2002/58/ec. *Official Journal of the European Union*, L 105:54 – 63, 2006.

- 10 Gartner, Inc. Gartner says worldwide mobile advertising revenue forecast to reach \$3.3 billion in 2011, 2011. <http://www.gartner.com/it/page.jsp?id=1726614>.
- 11 Philippe Golle and Kurt Partridge. On the anonymity of home/work location pairs. In Hideyuki Tokuda, Michael Beigl, Adrian Friday, A. Brush, and Yoshito Tobe, editors, *Pervasive Computing*, volume 5538 of *Lecture Notes in Computer Science*, pages 390–397. Springer Berlin / Heidelberg, 2009.
- 12 M. C. Grace, W. Zhou, X. Jiang, and A.-R. Sadeghi. Unsafe exposure analysis of mobile in-app advertisements. In *Proceedings of the fifth ACM conference on Security and Privacy in Wireless and Mobile Networks*, WISEC '12, pages 101–112, New York, NY, USA, 2012. ACM.
- 13 A. Greenberg. Phone 'rootkit' maker carrier iq may have violated wiretap law in millions of cases. *Forbes*, 2011. <http://www.forbes.com/sites/andygreenberg/2011/11/30/phone-rootkit-carrier-iq-may-have-violated-wiretap-law-in-millions-of-cases/>. Retrieved 2012/10/18.
- 14 M. Gruteser and D. Grunwald. Anonymous usage of location-based services through spatial and temporal cloaking. In *Proceedings of the 1st international conference on Mobile systems, applications and services*, MobiSys '03, pages 31–42, New York, NY, USA, 2003. ACM.
- 15 K. Hill. *Forbes*, 2012. <http://www.forbes.com/sites/kashmirhill/2012/02/07/the-austrian-thorn-in-facebooks-side/>. Retrieved 2012/10/18.
- 16 P. Hornyack, S. Han, J. Jung, S. Schechter, and D. Wetherall. These aren't the droids you're looking for: retrofitting android to protect data from imperious applications. In *Proceedings of the 18th ACM conference on Computer and communications security*, CCS '11, pages 639–652, New York, NY, USA, 2011. ACM.
- 17 H. Jun, E. Owusu, L. T. Nguyen, A. Perrig, and J. Zhang. Accomplice: Location inference using accelerometers on smartphones. In *Communication Systems and Networks (COMSNETS), 2012 Fourth International Conference on*, pages 1–9, 2012.
- 18 D. Kravets. An intentional mistake: The anatomy of google's wi-fi sniffing debacle. *Wired*, 2011. <http://www.wired.com/threatlevel/2012/05/google-wifi-fcc-investigation/>. Retrieved 2012/10/18.
- 19 J. Krumm. Ubiquitous advertising: The killer application for the 21st century. *Pervasive Computing, IEEE*, 10(1):66–73, jan.-march 2011.
- 20 S. Landtag. Drucksache 5/6787. Sächsischer Landtag 5. Wahlperiode, 2011.
- 21 C. Ling, M. Loschonsky, and L. M. Reindl. Characterization of delay spread for mobile radio communications under collapsed buildings. In *IEEE 21st International Symposium on Personal Indoor and Mobile Radio Communications (PIMRC)*, pages 329–334, Sept. 2010.
- 22 C. Y. T. Ma, D. K. Y. Yau, N. K. Yip, and N. S. V. Rao. Privacy vulnerability of published anonymous mobility traces. In *Proceedings of the sixteenth annual international conference on Mobile computing and networking*, MobiCom '10, pages 185–196, New York, NY, USA, 2010. ACM.
- 23 Ministerium für Inneres und Kommunales NRW. Funkzellenauswertung (FZA) und Versenden "Stiller SMS" zur Kriminalitätsbekämpfung. MMD 15/3300, 2011.
- 24 P. Ohm. Broken promises of privacy: Responding to the surprising failure of anonymization. *UCLA Law Review*, Vol. 57, p. 1701, 2010, 2009.
- 25 K. Rechert, K. Meier, B. Greschbach, D. Wehrle, and D. von Suchodoletz. Assessing location privacy in mobile communication networks. In H. Li X. Lai, J. Zhou, editor, *ISC 11*, LNCS 7001, pages 309–324. Springer, Heidelberg, 2011.
- 26 E. Smith. iphone applications & privacy issues: An analysis of application transmission of iphone unique device identifiers (udids). Technical report, PSKL, 2010. <http://www.pskl.us/wp/wp-content/uploads/2010/09/iPhone-Applications-Privacy-Issues.pdf>.

- 27 S. Thurm and I. Yukari Kane. Your apps are watching you. *The Wall Street Journal*, 2010. <http://online.wsj.com/article/SB10001424052748704694004576020083703574602.html>.
- 28 W. H. Ware. Records, computers, and the rights of citizens. Secretary's Advisory Committee on Automated Personal Data Systems, Department of Health, Education and Welfare, Washington, D.C., 1973.
- 29 H. Zang and J. Bolot. Anonymization of location data does not work: a large-scale measurement study. In *Proceedings of the 17th annual international conference on Mobile computing and networking*, MobiCom '11, pages 145–156, New York, NY, USA, 2011. ACM.

4.2 Working Group: Vehicular Data

Glenn Geers, Marco Gruteser, Michael Marhoefer, Christian Wietfeld, Claudia Santa, Olaf Spinczyk, and Ouri Wolfson

License  Creative Commons BY-NC-ND 3.0 Unported license
 © Glenn Geers, Marco Gruteser, Michael Marhoefer, Christian Wietfeld, Claudia Santa, Olaf Spinczyk, and Ouri Wolfson

4.2.1 Vehicular Applications (An Introduction)

Currently there is much research in development of vehicular applications.

► **Example 1.** Imagine the year 2025 it is Saturday before Christmas and you want to go to a shopping mall to get some Christmas gifts for your friends. You go to your garage, enter your destination in the navigation device and start your e-car. Your current location, destination, preferable routes and maximum accepted arrival time is transferred to a central server and the fastest way is returned back by taking into account all other navigation requests. You follow the advised route through the city and the green phase of the traffic lights are optimized in your direction, because everybody else also wants to go to get the last gifts. Your car automatically adapts to the optimal speed of the green wave. To do so your position and type of car is send to the traffic light server every second. During your drive a friend is calling you and despite of the complete fusion of vehicle and mobile phone, you lose concentration for a short while and cannot see the person on the pedestrian crossing. You feel a jerky braking movements triggered by the pupil warning system. After some time you finally reach the shopping mall. Your car guides you to a free parking space next to the wine shop because it knows that your last bills showed a preference for wine. You have to cross the whole mall because this time you only need computer games for the children. After having found all gifts you walk back to your car and while leaving the parking place you are automatically charged based on the duration of your stay. The parking time is send to the shopping mall to analyze the buying behavior of customers. Late in the evening after your exhaustive shopping trip you finally arrive back home and your driving data including length of the trip, speed and driving behavior is send to your insurance to settle your account.

This vision of the future involves a lot of benefits in terms of safety, comfort and efficiency, for the driver as well as for society as a whole: accidents are prevented, parking fees are charge automatically, minimum stops at traffic lights are needed, no searching for parking spaces is needed and insurance fees are paid according to your individual risks. All these applications are currently under development in academia and/or industry.

Even though these applications offer a clear benefit, it is obvious that the involved systems gather a lot of individual data, which needs to be properly protected against privacy violations. It is therefore necessary to design the systems in a way, which protects the privacy,

not only to comply with legal regulations, but also to ensure, that users will accept these systems. The following paragraphs will show that vehicular applications may be implemented in different ways and users may decide to make a trade-off between utility, cost and privacy risks.

4.2.2 Glossary

Identifier : temporary device name

Identity : personally linked name

Data criticality : level of perceived damage due to data revelation (level of detail, amount of detail)

Identity revelation risk : describes the risk and effort to match an identifier to an identity

4.2.3 Expectations on Privacy & Privacy Risks

The expectations of users towards vehicular data analysis regarding privacy differ from other Internet applications due to several reasons:

- Driving a vehicle with a number plate takes place in public space.
- The operation of a vehicle requires to follow a set of well-established rules, which are enforced by police, with the help of electronic systems (e.g. cameras). It is a well established fact, that a driver/owner of a vehicle can be identified by officials can via the license plate, especially in case of traffic rule violations, but also in other legally relevant cases (e.g. bank robbery).


On the other hand, existing observation and enforcement is rather sparse and local. Names of persons are only revealed in case of violations and following certain regulations. Therefore anybody who does not violate the rules can move anonymously and it requires a very high technical and organizational effort to produce movement tracks of individuals.

With the introduction of ITS, many user benefits are associated, in particular fewer accidents, more comfort and less traffic jams. Nevertheless, with Intelligent Transport Systems (ITS) technical options are potentially introduced for deriving movement tracks with much lower effort. In ITS, machine-readable **identifiers** are used to address the devices (in particular for the communication devices inside vehicles), which collect vehicular data. These identifiers can be eventually matched with the **identity** of a person, i.e. a personally linked name.

This might lead for example to actual or perceived privacy risks, as the vehicular data may be used

- for traffic law enforcement (e.g. speed violations, traffic accidents)
- for general law suits (employments, divorce, etc.)
- to derive embarrassing interpretation of location tracks (actual or wrongly interpreted visits of certain doctors, places, red light districts).
- to derive transport-specific personal information about driving styles, vehicle usage profiles, etc.
- to reveal general personal information about movement pattern, relationships, etc.

Therefore during the design and implementation of vehicular-related applications, privacy risks need to be assessed. Even so the actual risks might be limited compared to other technical systems impacting the privacy (such as mobile phone usage), the user acceptance of ITS applications also strongly relies on appropriate privacy protection mechanisms.



■ **Figure 2** Assessing privacy risk.

For the systematic privacy risk assessment and risk management of applications involving vehicular data, we introduce two dimensions of risks:

First, the **identity revelation risk** describes the risk and effort to match an identifier to an identity. As a second dimension, the **criticality of the data** is assessed. The data criticality indicates the level of perceived damage due to data revelation (e.g. level of detail, amount of detail).

In the following section, the methodology of the proposed risk assessment is introduced in more detail.

4.2.4 Privacy Assessment

In order to achieve better transparency of the privacy related risks that are caused by concrete implementations of vehicular applications we propose a simple scheme that provides a privacy risk vector. A risk vector fulfills two purposes:

- It supports application developers by helping them to understand the privacy implications of the system. They can use this information to reduce the privacy risk for their future users and, thus, make their product more attractive.
- It allows users, i.e. car drivers, to assess the risk of a specific applications without having to look into the details of the implementation. This transparency helps them to manually, semi-automatically, or even fully-automatically choose which applications they want to use.

Figure 2 illustrates the meaning of the risk vector. The two-dimensional characteristic of the diagram provides a quick overview on where the privacy risk comes from: It can be the revelation of your identity, the risk of making critical data available, or a combination of both. The length of the vector can be taken a very simple indicator for the overall risk caused by using an application.

■ **Table 1** Identity Revelation Risk Quantification.

Characteristic	Contribution to Value
no identifier used	none
short/medium/long-term identifier	+ / ++ / +++
(true) identity used	+++

■ **Table 2** Data Criticality Quantification.

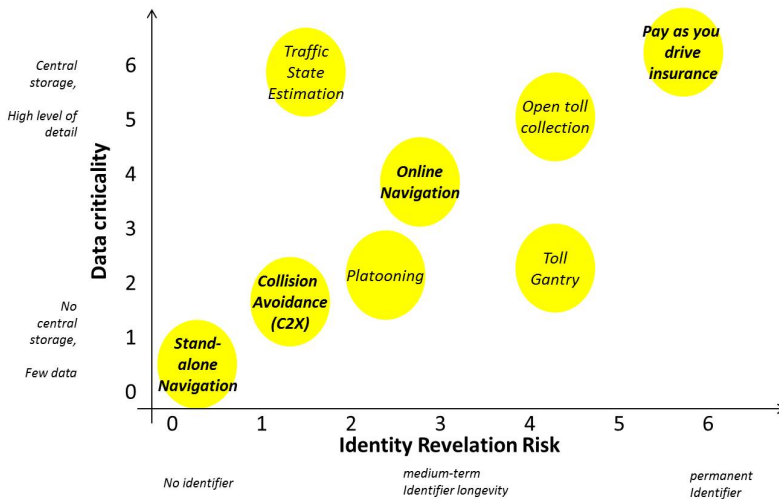
Characteristic	Contribution to Value
no storage of data / local storage / centralized storage	none / + / +++
size (level of detail) of data records	none / + / ++ / +++

The color scheme is intended to transport privacy expert knowledge to software developers and application users and intentionally kept simple: The well-known green, yellow, red known from traffic lights. Each user could, of course, define his/her own individual privacy risk boundary in this diagram, or maybe several context-dependent boundaries. Given this boundary a software solution could even automatically accept or reject the privacy risk induced by an application.

Quantification of the coordinates in the diagram is difficult but strongly desired, because it is a prerequisite for the calculation of a risk vector length. Therefore, we have analyzed a set of known vehicular applications with respect to the identity revelation risk and data criticality. Tables 1 and 2 show a proposal for mapping application characteristics to values:

Based on these tables the two coordinates can be calculated by summing up the values ('+'), resulting in a value from 0 to 6 on both axes.

4.2.5 Case Studies



4.2.6 Additional Material and Ideas

What about criminal activity? What judicial/policy controls are required?
 Is there a privacy algebra? Does the direction of the privacy vector have meaning?

If users are given full control (i.e., free version with ads or user pays for no ads) what will be the effect on app development?

Do transport apps on their own pose a privacy issue or is cross-linking with other data the worry?

4.2.7 Discussion / Conclusions

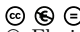
Potential title of workshop paper “Assessing and managing privacy risks during design, implementation and operation of vehicular applications”

References

- 1 Daniel Ashbrook and Thad Starner. Using gps to learn significant locations and predict movement across multiple users. *Personal and Ubiquitous Computing*, 7(5):275–286, October 2003.
- 2 Caitlin D. Cottrill. Approaches to privacy preservation in intelligent transportation systems and vehicle-infrastructure integration initiative. *Transportation Research Record: Journal of the Transportation Research Board*, 2129:9–15.
- 3 M. Gruteser and Xuan Liu. Protecting privacy, in continuous location-tracking applications. *Security Privacy, IEEE*, 2(2):28 – 34, mar-apr 2004.
- 4 B. Hoh, M. Gruteser, H. Xiong, and A. Alrabady. Enhancing security and privacy in traffic-monitoring systems. *Pervasive Computing, IEEE*, 5(4):38 –46, oct.-dec. 2006.
- 5 D. Reid. An algorithm for tracking multiple targets. *Automatic Control, IEEE Transactions on*, 24(6):843 – 854, dec 1979.
- 6 R. Shokri, G. Theodorakopoulos, J. Le Boudec, and J. Hubaux. Quantifying location privacy. In *Security and Privacy (SP), 2011 IEEE Symposium on*, pages 247 –262, may 2011.
- 7 Hairuo Xie, L. Kulik, and E. Tanin. Privacy-aware traffic monitoring. *Intelligent Transportation Systems, IEEE Transactions on*, 11(1):61 –70, march 2010.

4.3 Working Group: Context-dependent Privacy in Mobile Systems

Florian Schaub, Maria Luisa Damiani, and Bradley Malin

License  Creative Commons BY-NC-ND 3.0 Unported license
© Florian Schaub, Maria Luisa Damiani, and Bradley Malin

4.3.1 Introduction

Mobile systems, and smartphones especially, are becoming ever more present in our society. A smartphone can gather a significant quantity information about its owner’s movements and behavior. Such *mobility data* can be communicated and shared with a wide range of parties (e.g., mobile applications, location-based services, mobile operators, mobile phone vendors and providers of mobile application platforms). The accumulation of mobility data enables the mining of mobility patterns, with the goals of patterns at the level of an individual and general trends. At the same time, there are concerns that the collection, storage, and processing of such data may violate the privacy of the individuals to whom the data corresponds.

Traditionally, privacy of such personal information has been achieved through definitional and methodological approaches. From a definitional perspective, access control is commonly applied to specify policies regarding what entities are permitted access to which information under a range of specified purposes and restrictions. From a methodological perspective,

data obfuscation and perturbation manipulate the data itself prior to its flow through an information system. Definitional and methodological aspects are not mutually exclusive and can be combined when the information flow is predefined or governed by policies.

To date, such privacy protection systems have often been *one size fits all* solutions. For example, in the case of k -anonymity [31] (i.e., the size of the group to which a piece of information corresponds), k is often fixed to a constant value for everyone in all situations. This is not always the case ([14]) and, more generally there exists a variety of solutions that enable users to configure their privacy settings. Yet, the personalization of such controls to a range of settings is a cumbersome process, which can lead to misconfiguration and dissonance between privacy expectations and their actual realization [27]. However, the context (e.g., “In what setting is the individual situated?”) provide meaningful indications about the level of required protection. While there is prior research on protecting contextual information, there are limited investigations into how such contextual information can be leveraged for enhancing privacy protection.

In this report, we discuss several perspectives on context-dependent privacy, with a focus on several application domains. Subsequently, we argue for a consolidated view of context-dependent privacy to facilitate the exchange of results between computational subcommunities. When appropriate, we outline a number of research challenges and future directions. The report closes with an outlook on ongoing and planned actions.

4.3.2 Perspectives on Context-dependent Privacy

Context-dependent privacy has been considered, to a limited extent, in certain subcommunities of computer science. We give an overview of the endeavors in three representative fields to highlight the different viewpoints and approaches taken.

4.3.2.1 Context-aware Privacy in Ubiquitous Computing

Ubiquitous computing (ubicomp) centers on the vision of an interconnected world of smart devices integrated into a users’ surroundings and daily activities [33]. Some ubicomp application areas are smart homes and environments, ambient assisted living, support in work environments, life logging, and everyday computing. Context awareness is a salient characteristic of ubicomp applications. Applications use sensors to obtain information about the current situation of users, applications, and devices. Depending on the application, different context features are used to reason about the situation and detect user activities in order to adapt the application’s features and behavior accordingly. Dey broadly defines context as “any information that can be used to characterize the situation of an entity. An entity is a person, place, or object that is considered relevant to the interaction between a user and an application, including the user and applications themselves.” [11]. Commonly used context features [1] center around the *where* and *when* of the current situation and present entities (*who*). Furthermore, *what* the user is doing, as well as *why* they are doing it, are of particular interest. In addition to physically-sensed phenomena, contextual information includes information available (or inferred) about the user, such as data derived from a calendar, social network connections, and user preferences.

The connected nature and deep integration of technology with the environment in ubicomp scenarios raises a number of privacy issues [21, 33, 29]. Specifically, the collection scale, often invisible data collection, and non-transparent information flow pose issues. The physicality of ubicomp systems also extends the privacy scope beyond information collection to intrusions and disturbances in the user’s physical environment [20]. Furthermore, the complexity and

distributed nature of ubicomp systems makes it difficult for users to estimate the privacy implications of their actions and properly express their privacy expectations [27].

Context awareness has been utilized to improve the awareness of users about their current, as well as potential, privacy risks in ubicomp systems. For example, discovery protocols and detection mechanisms can be used to model what entities are present in a user's environment and what privacy-sensitive items they can access [21, 28, 20]. Context-aware privacy mechanisms, such as Confab [16], have been proposed to govern disclosure and granularity of specific information items. Privacy preferences are typically formalized in context-aware privacy policies. Moncrieff et al. [23] employ a context-aware privacy system in an ambient assisted living facility to dynamically govern what information is available to care givers. The information provided is dependent on the user's activities and potential hazards (e.g., a stove that is turned on and left unattended). User-centric privacy support systems use context awareness and heuristics in order to help users express their privacy preferences for specific situations [27]. Palen and Dourish [25] argue that, in order to ensure that the privacy provided by technical mechanisms matches a user's privacy expectations, privacy management needs to be viewed as a dynamic and continuous regulation process. Context awareness can enable ubicomp systems to support dynamic privacy adaptation either by 1) providing users with context-dependent individual recommendations or 2) automatically reconfiguring privacy mechanisms [28].

4.3.2.2 Context-based Location Privacy in LBS

Location-based services (LBS) comprise a wide spectrum of information services and applications (e.g., searching for business points of interest within a vicinity, sharing a personal location with the members of a social network). These applications typically rely on a client-server architecture in which clients convey their location to a LBS provider in exchange for spatially-contextualized information, returned in real time. Most current research on privacy in LBS is driven by the consideration that 1) location can reveal much about an individual and 2) LBS providers are potentially untrustworthy and, thus, can exploit location data for illegitimate purposes. In the literature, privacy requirements for LBS are often voiced from multiple perspectives, which calls for different classes of protection solutions. Historically, two broadly investigated privacy requirements are 1) identity privacy and 2) location privacy [17]. In the former case, the goal is to forestall the re-identification of seemingly anonymous users based on their location (e.g. home) [15]. In the latter case, the goal is to prevent the disclosure of detailed location or trace of a user (e.g., [35]). More recently, we emphasized the emergence of a third privacy requirement called behavioral privacy [8]. Behavioral information can be extracted from the traces of users by relating location with context. For example, context can reveal what the person may be doing (e.g., a person visiting a retail store is likely to be browsing or purchasing merchandise) or who they interact with (e.g., two people frequenting the same fitness club in the same period are likely to know each other). Preventing the extraction of behavioral information calls for techniques that are capable of recognizing geographical, temporal and social context.

Behavioral privacy is related to the concept of semantic location privacy and user-defined private places. Semantic location privacy refers to the capability of protecting semantic locations (i.e., the places in which users are located), such as a jewelry store [10]. In this scenario, a gentleman who never visits a jewelry store might be looking to surprise his girlfriend with an engagement ring (or some other token of affection) which he wishes to keep secret until the right time. The motivation behind this model is that places contribute more semantic information than the traditional coordinates (e.g., latitude and longitude)

commonly used to represent the user's location in traditional privacy-enhancing techniques. Moreover some places can be perceived more sensitive than others by users. For example, while standing on a crowded street may be considered innocuous, the revelation of one's presence in a jewelry store may represent a sensitive piece of knowledge. The protection of such information is challenging, particularly when users are tracked (i.e., their position is continuously reported to an LBS provider) and their movement is confined to road networks (i.e., a restricted set of trajectories) [34].

By contrast, the notion of a user-defined private place is motivated by the fact that many entities involved in the provision of location services are untrusted. These entities may be an LBS provider, but also third parties acting as location providers (LP), such as Google Location Service. Thus, a key question is, "To what extent can a user's location be protected against the party who computes it?" Placeprint [9] is a first attempt to protect the location information from both the LP and the LBS provider. In this system, users equipped with commodity devices can be geolocated in user-defined private places without revealing their presence to the LP. Additionally, users can specify context-based privacy rules to forestall the disclosure of private places also to LBS providers. The ultimate goal is to provide users with the capability of exercising flexible control over the disclosure of position information to LPs and LBS providers.

4.3.2.3 Context-Aware Access Control in E-Health Organizations

An individual's health status is considered to be among the most sensitive types of personal information. The collection and utilization of health information in the context of primary care (e.g., large hospitals) is not a new phenomenon. Electronic medical record (EMR) systems have been in place for over fifty years [32]. However, modern computing systems have enabled the integration of information into, and utilization of, EMRs at virtually anytime and anywhere. As a consequence, health information is now collected at all points of care; e.g., through hundreds of applications and IT systems in hospitals [24], home and remote health management systems (e.g., [2, 18, 19]), and through wearable devices (e.g., [5, 12]). The convergence of these systems is enabling mobility both in the collection of information and provision of treatment.

In many respects, the privacy issues associated with gathering and utilization of health information is not much different than other types of information. Yet, healthcare illustrates how context is related to more than just an individual and their personal preferences. Healthcare is a complex service-oriented business, composed of many interleaved processes. Consider, a patient may wish to specify which care providers can access their medical information where and for what purposes. Such specifications can be formalized in logical access control policies [22, 30]; however, the complexity of healthcare systems makes it difficult to define specific roles, as well as which care providers need access to what information and when [26].

For instance, healthcare is an inherently dynamic environment where teams of clinicians and support staff constantly interact. The notion of dynamic teams, while effective in supporting operations, leads to relationships that are significantly more complex pairwise than typically expected by access control frameworks [4]. The problem is further compounded by the fact that healthcare organizations are constantly evolving to update management protocols assimilate new employees and systems. Thus, the context of access to an individual's health information is critical to assess if the utilization is expected or not. As has been recently shown, this context may be inferred based on various features, such as the hospital service to which a patient is assigned or the location within a healthcare system the patient

is currently residing (which could be at home) [13, 36].

However, contextual features must also address the social dynamic of the organization [7, 6]. Healthcare professionals often collaborate in teams of employees with diverse skillsets. The likelihood that a particular team member will need (or choose to) access a patient's record is often dependent on the state of the patient (e.g., "Has the patient recently been admitted for treatment? Are they in a recovery room following a surgical procedure? Are they about to be discharged to an assisted living facility? Is the accessor of the patient's medical information a family member?"). Given that there are many organizational policies, most of which are not appropriately documented, it is critical to use audit logs and organizational knowledge to assess when access should be granted.

While this discussion focuses on healthcare environments, it is clear that this issue is more general. Similar issues clearly transpire in other domains, such as financial systems, intelligence environments, and just about anywhere where the relation between the subject (e.g., the patient) and the recipient (e.g., healthcare provider) is complex, dynamic, and driven by organizational behavior.

4.3.3 Discussion and Research Directions

Disparate subcommunities in computer science approach context-dependent privacy from very different angles. Context is used to facilitate dynamic and individualized adaptation of privacy mechanisms both 1) as a criterion for adjusting location and information granularity and 2) as domain semantics leveraged in access control systems. As a consequence of this diversity, proposed solutions, models, and mechanisms tend to be tailored for a specific perspective and have limited impact beyond the application domain in which they are initially introduced. However, in a preliminary analysis of problem characteristics, we noticed that, although framed differently, leveraging context for privacy protection raises similar research questions across domains. We believe that a consolidated view on context-dependent privacy could facilitate a common understanding of the associated research challenges to accelerate research and enhance the generality of results.

4.3.3.1 Shared Problem Characteristics

Despite diverse perspectives on context-dependent privacy all of the privacy mechanisms alluded to in Section 4.3.2 share common goals for a dual optimization: 1) minimize privacy risk for some information or some user and 2) maximize the utility of the information. Privacy risk can grossly be defined as the probability that, given the output of a privacy mechanism, a third party can infer the corresponding input (e.g., "To what extent can a user's exact location be ascertained from an obfuscated location?") Similarly, utility can be defined as the probability that the output can be utilized to support a specific purpose (e.g., "How similar are the results of a location-based service based on an obfuscated location and an exact location?").

Context can be invoked as an input parameter in the optimization to facilitate the determination of appropriate privacy mechanisms. Specifically, context can be used to determine the level of privacy necessary in a given situation or to perturb information in a manner that enhances utility. Additional inputs could be domain-specific constraints that restrict the abilities of individual users to determine privacy settings. Such constraints could be of an organizational or a regulatory nature.

4.3.3.2 Context Definition and Modeling

The benefits of considering context in privacy mechanisms and dynamic privacy adaptation are readily apparent. However it remains a challenge to identify which context information is relevant for privacy and how best to represent it. While there is extensive work on generic context modeling for application adaptation [3], the identification of privacy-relevant context features is often only performed for specific applications. Very few context models are specifically focused on privacy [28]. A general categorization and model for privacy-relevant context features would facilitate cross-utilization of privacy context across different domains.

When defining privacy-relevant context, it must be considered that context features and their values can depend on each other. Furthermore, context information and the data to be protected can be intertwined. For example, consider privacy mechanisms for applications that generate and process spatiotemporal data series, such as mobility tracking and vehicle-centric applications. These applications need to consider the context in which information has been generated (e.g., a traffic jam), but also ensure that other information types in the same series do not facilitate inference of the original data (e.g., rush hour).

A reasonable approach for representing privacy-relevant context is the definition of multiple context layers. Each layer could provide abstractions of context information on a different level. At the same time, each layer could represent semantics for specific domains (or provide granularity definitions for different information types). Modeling context on different levels is a common approach in context-aware systems, but the definition, management and structuring of context layers that appropriately capture generic and domain-specific privacy semantics, language, and knowledge is an open challenge. Context models must also be maintainable and scalable.

Context-dependent privacy requires reasoning about context information in order to adapt privacy mechanisms. This reasoning must extend over different context layers and be able to deal with uncertainty concerning the accuracy of context information. A related challenge is the development of formal models for context-dependent privacy that facilitate proofs of the provided privacy protection as well as the utility of information in the face of dynamic adaptation of privacy mechanisms.

4.3.3.3 Sensitivity of Context

An inherent challenge of context-dependent privacy is that the context information utilized for enhancing privacy adaptation is, in itself, privacy sensitive. This is because it contains potentially detailed information about the user's situation. Thus, privacy preserving context acquisition mechanisms are preferable to reduce privacy risk.

Beyond the sensitivity of actual context information, acting upon context can also potentially reveal information about the situation that caused the privacy adaptation. For example, consider a context-dependent location-based service that adapts location granularity based on context. When the user moves into a sensitive area, the location information becomes coarser, thereby revealing that the location is of higher sensitivity to the user than other locations. Thus, the sensitivity of a context-dependent privacy change must be considered, which requires respective mechanisms and models. One approach to address this issue could be to exploit historical context information to determine the probability that obfuscated information can be derived from previous information.

4.3.3.4 Making Privacy Mechanisms Context Aware

A further challenge is the integration of context awareness and context-dependent privacy into existing privacy and access control mechanisms. Although there is some work on context-aware access control, a challenge remains regarding how to account for context in privacy policies on a practical level.

4.3.4 Conclusions

Our investigation of context-dependent privacy in different domains identified a number of shared characteristics of privacy mechanisms utilizing context. However, there is only limited exchange and cross-utilization of results between sub-communities. Our preliminary analysis indicates that a generalized and formalized problem definition is potentially feasible. The benefits of such a model would be enhanced comparability and compatibility of approaches and results between different communities despite domain-specific requirements. We encourage the computer science community to explore the possibilities of such problem formalization in future collaborations.

Acknowledgements

We would like to thank all participants of the Dagstuhl Seminar on Mobility Data Mining and Privacy for stimulating discussions, valuable input, and feedback. We further extend our gratitude to the seminar organizers and the Dagstuhl staff for bringing us all together and making this seminar a very pleasant experience.

References

- 1 Gregory Abowd and Elizabeth Mynatt. Charting past, present, and future research in ubiquitous computing. *ACM TOCHI*, 7(1):29–58, 2000.
- 2 Hande Özgür Alemdar, Tim van Kasteren, and Cem Ersoy. Using active learning to allow activity recognition on a large scale. In *Proceedings of the 2nd International Joint Conference on Ambient Intelligence*, pages 105–114, 2011.
- 3 Claudio Bettini, Oliver Brdiczka, Karen Henriksen, Jadwiga Indulska, Daniela Nicklas, Anand Ranganathan, and Daniele Riboni. A survey of context modelling and reasoning techniques. *Pervasive and Mobile Computing*, 6(2):161–180, 2010.
- 4 B. Blobel, R. Nordberg, J. M. Davis, and P. Pharow. Modelling privilege management and access control. *Int J Med Inform*, 75(8):597–623, Aug 2006.
- 5 Shih-Lun Chen, Ho-Yin Lee, Chiung-An Chen, Hong-Yi Huang, and Ching-Hsing Luo. Wireless body sensor network with adaptive low-power design for biometrics and healthcare applications. *IEEE Systems Journal*, 3(4):398–409, 2009.
- 6 You Chen, Steve Nyemba, and Bradley Malin. Auditing medical record accesses via healthcare interaction networks. In *Proceedings of the American Medical Informatics Association Annual Symposium*, 2012. (in press).
- 7 You Chen, Steve Nyemba, and Bradley Malin. Detecting anomalous insiders in collaborative information systems. *IEEE Trans. Dependable and Secure Computing*, 9(3):332–344, 2012.
- 8 M. L. Damiani. Privacy enhancing techniques for the protection of mobility patterns in LBS: research issues and trends. In *European Data Protection: Coming of Age?* Springer, 2012. (to appear).
- 9 M. L. Damiani and M. Galbiati. Handling user-defined private contexts for location privacy in LBS. In *Proc. ACM GIS*, 2012. (to appear).
- 10 M.L. Damiani, C. Silvestri, and E. Bertino. Fine-grained cloaking of sensitive positions in location-sharing applications. *IEEE Pervasive Computing*, 10(4):64–72, 2011.

- 11 Anind K. Dey. Understanding and Using Context. *Pers Ubiquit Comput*, 5(1):4–7, 2001.
- 12 A. Dinh, D. Teng, L. Chen, S.B. Ko, Y. Shi, C. McCrosky, J. Basran, and V. del Bello-Hass. A wearable device for physical activity monitoring with built-in heart rate variability. In *Proc. 3rd Int. Conference on Bioinformatics and Biomedical Engineering (ICDBBE)*, 2009.
- 13 Daniel Fabbri, Kristen LeFevre, and David A. Hanauer. Explaining accesses to electronic health records. In *Proceedings of the 2011 ACM SIGKDD Workshop on Data Mining for Medicine and Healthcare*, pages 10–17, 2011.
- 14 Buğra Gedik and Ling Liu. Protecting location privacy with personalized k-anonymity: Architecture and algorithms. *IEEE Transactions on Mobile Computing*, 7(1):1–18, 2008.
- 15 M. Gruteser and D. Grunwald. Anonymous Usage of Location-Based Services Through Spatial and Temporal Cloaking. In *Proc. MobiSys '03*. ACM Press, 2003.
- 16 Jason I Hong and James A Landay. An architecture for privacy-sensitive ubiquitous computing. In *Proc. MobiSys '04*. ACM Press, 2004.
- 17 C. S. Jensen, H. Lu, and M.L. Yiu. Location Privacy Techniques in Client-Server Architectures. In *Privacy in Location-Based Applications: Research Issues and Emerging Trends*. Springer-Verlag, 2009.
- 18 Shanshan Jiang, Yanchuan Cao, Sameer Iyengar, Philip Kuryloski, Roozbeh Jafari, Yuan Xue, Ruzena Bajcsy, and Stephen Wicker. Carenet: an integrated wireless sensor networking environment for remote healthcare. In *Proceedings of the ICST 3rd international conference on Body area networks, BodyNets '08*, 2008.
- 19 Andrew D. Jurik and Alfred C. Weaver. Remote medical monitoring. *IEEE Computer*, 41(4):96–99, 2008.
- 20 Bastian Könings and Florian Schaub. Territorial privacy in ubiquitous computing. In *Proc. Int. Conf. Wireless On-Demand Network Systems and Services (WONS '11)*. IEEE, 2011.
- 21 Marc Langheinrich. Privacy in Ubiquitous Computing. In John Krumm, editor, *Ubiquitous Computing Fundamentals*, chapter 3, pages 95–160. CRC Press, 2009.
- 22 Ming Li, Shucheng Yu, Kui Ren, and Wenjing Lou. Securing personal health records in cloud computing: Patient-centric and fine-grained data access control in multi-owner settings. In *Proceedings of the 6th International ICST Conference on Security and Privacy in Communication Networks*, pages 89–106, 2010.
- 23 Simon Moncrieff, Svetha Venkatesh, and Geoff West. Dynamic privacy assessment in a smart house environment using multimodal sensing. *ACM TOMCCAP*, 5(2):10, 2008.
- 24 National Research Council (US) Committee on Engaging the Computer Science Research Community in Health Care Informatics, W. Stead and H. Lin, eds. *Computational Technology for Effective Health Care: Immediate Steps and Strategic Directions*. National Academies Press, Washington, DC, 2009.
- 25 Leysia Palen and Paul Dourish. Unpacking “privacy” for a networked world. In *Proc. Conf. on Human factors in computing systems (CHI '03)*, pages 129–136. ACM, 2003.
- 26 Lillian Røstad and Øystein Nytrø. Access control and integration of health care systems: An experience report and future challenges. In *Proceedings of the the 2nd International Conference on Availability, Reliability and Security*, pages 871–878, 2007.
- 27 Norman Sadeh, Jason Hong, Lorrie Cranor, Ian Fette, Patrick Kelley, Madhu Prabaker, and Jinghai Rao. Understanding and capturing people’s privacy policies in a mobile social networking application. *Personal and Ubiquitous Computing*, 13(6):401–412, 2009.
- 28 Florian Schaub, Bastian Könings, Stefan Dietzel, Michael Weber, and Frank Kargl. Privacy Context Model for Dynamic Privacy Adaptation in Ubiquitous Computing. In *6th Int. Workshop on Context-Awareness for Self-Managing Systems, UbiComp '12*. ACM, 2012.
- 29 Florian Schaub, Bastian Könings, Michael Weber, and Frank Kargl. Towards context adaptive privacy decisions in ubiquitous computing. In *PerCom 2012 Workshops*. IEEE, 2012.

- 30 Robert Steele and Kyongho Min. Healthpass: Fine-grained access control to portable personal health records. In *Proceedings of the 24th IEEE International Conference on Advanced Information Networking and Applications*, pages 1012–1019, 2010.
- 31 Latanya Sweeney. k-anonymity: A model for protecting privacy. *International Journal of Uncertainty Fuzziness and Knowledge Based Systems*, 10(5):557–570, 2002.
- 32 L. L. Weed. Medical records that guide and teach. *N. Engl. J. Med.*, 278(11):593–600, Mar 1968.
- 33 Mark Weiser. The computer for the 21st Century. *Sci. Am.*, 265(3):94–104, 1991.
- 34 E. Yigitoglu, M.L Damiani, O. Abul, and C. Silvestri. Privacy-preserving sharing of sensitive semantic locations under road-network constraints. In *Proc. IEEE MDM*, 2012.
- 35 Man Lung Yiu, C.S. Jensen, Xuegang Huang, and Hua Lu. SpaceTwist: Managing the Trade-Offs Among Location Privacy, Query Performance, and Query Accuracy in Mobile Services. In *proc. IEEE 24th International Conference on Data Engineering*, 2008.
- 36 Wen Zhang, Carl Gunter, David Liebovitz, Jian Tian, and Bradley Malin. Role prediction using electronic medical record system audits. In *Proceedings of the American Medical Informatics Association Annual Symposium*, pages 858–867, 2011.

4.4 Working Group: Privacy through Uncertainty in Location-Based Services

Nilgün Basalp, Joachim Biskup, Erik Buchmann, Chris Clifton, Bart Kuijpers, Walied Othman, and Erkay Savas

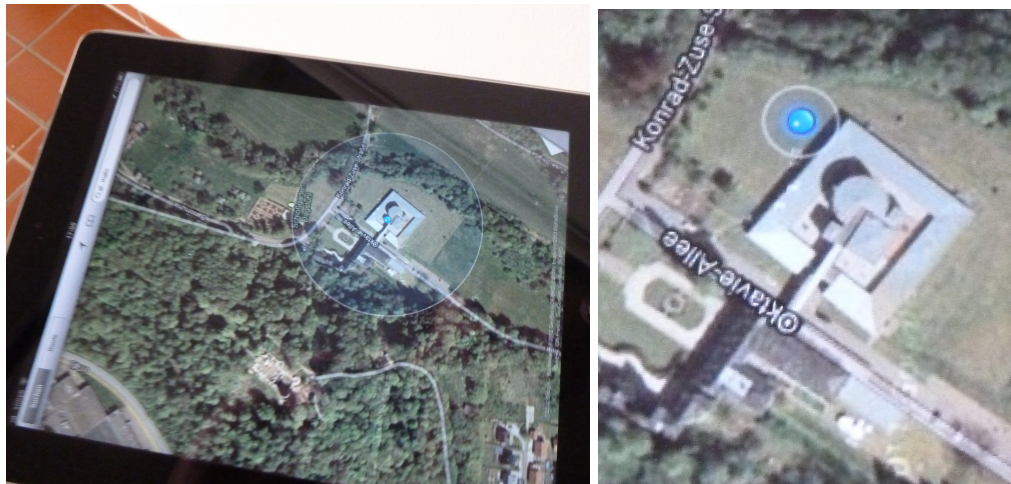
License © © © Creative Commons BY-NC-ND 3.0 Unported license
 © Nilgün Basalp, Joachim Biskup, Erik Buchmann, Chris Clifton, Bart Kuijpers, Walied Othman, and Erkay Savas

Location-Based Services (LBS) are becoming more prevalent. While there are many benefits, there are also real privacy risks. People are unwilling to give up the benefits – but can we reduce privacy risks without giving up on LBS entirely? This working group explored the possibility of introducing uncertainty into location information when using an LBS, so as to reduce privacy risk.

Uncertainty occurs naturally, so LBS likely to work in spite of uncertainty. For example, Figure 3 shows location determined by an Apple iPad at Schloss Dagstuhl. Initially, the location was reported with a high degree of uncertainty. Later, the uncertainty was reduced – but the location was not exactly as reported. A good LBS will have to accept that location may be uncertain, and give appropriate service in spite of that. Our question is, can we protect privacy by providing uncertain location, while still retaining good service?

4.4.1 Examples

In this section we explore some instances which have raised privacy issues in the past. One such case, which set off the privacy debate, is the case of Apple storing and collecting location data from its users' iPhones, unbeknownst to the user. The issue was uncovered on April 20th, 2011. Researchers discovered a file, consolidated.db, which contained longitudes and latitudes combined with a timestamp. This file contained locations that dated as far back as the release of iOS 4, which makes it contain a year's worth of location data, stored on the iPhone, synced (backed up) with iTunes and transmitted to Apple. All without the user's knowledge. A week later, Apple formally responded in a press release <http://www.apple.com/pr/library/2011/04/27Apple-Q-A-on-Location-Data.html> Apple



■ **Figure 3** Location uncertainty at Schloss Dagstuhl (actual location in hallway outside room N009).

maintained this data was anonymised, never shared with third parties, and its intent was to facilitate location look-ups by GPS in what is known as A-GPS, assisted GPS. Apple resolved the case with a software update that

- reduces the size of the crowd-sourced Wi-Fi hotspot and cell tower database cached on the iPhone,
- ceases backing up this cache, and
- deletes this cache entirely when Location Services is turned off.

In the time since, other apps have been uncovered to collect location (and other data) from mobile devices. These apps can be divided into two categories. The first category contains apps that collect location data but do not use it to provide a service, the second category uses it to provide a service. Ultimately, the goal is to provide techniques for users to protect themselves against tracking apps, while at the same time ensuring they can enjoy the same level of service from the apps. These two goals appear to be at odds with each other, but they do not need to be.

Currently, there are ways to protect yourself from being tracked. These do, however, require devices to be jailbroken (iOS) or rooted (Android). On iOS there is an app called ProtectMyPrivacy. It intercepts calls to a user's address book, playlists and location. If a user decided to hide any of these from an app that is requesting access, ProtectMyPrivacy jumbles the different fields for address book and playlist requests, and returns a fake, but preset by the user, location. This kind of protection serves well for apps that do not provide a service based on these records, but fails to ensure any service from apps in the second category.

There are a lot of research papers on such location-based social recommendation systems, and also some real systems. One particularly scary example (pointed out by Florian Schaub) is a “find a date” application that combines social media and location to find nearby women (who haven't necessarily said they want to be found):

<http://www.cultofmac.com/157641/this-creepy-app-isnt-just-stalking-women-without-their-knowledge-its-a-wake-up-call-about-facebook-privacy/>

Another examples is location-based keyword search: Modifying search results based on current location. While less seemingly privacy-invasive, this is also one where exact location is likely not needed.

4.4.2 Methods to induce uncertainty

The goal of this work is to increase the degree of uncertainty that comes from the location sensor in a way that the privacy of the individual is preserved. In particular, we strive for an amount of uncertainty that makes it impossible for an adversary to either (1) identify an individual from a particular location or a sequence of locations, (2) link a location or a sequence of locations to an individual or (3) connect a set of locations to a trajectory that belongs to the same (yet anonymous) individual. For that purpose, a wide number of anonymization approaches and obfuscation techniques exist. We subsume these techniques under the term "uncertainty methods". The applicability of uncertainty methods depends on the kind of data that needs to be modified. Thus, we distinguish between the induction of uncertainty to single locations and sequences of locations. While single locations are typical for one-shot queries sent to a location-based system, sequences of locations might be trajectories recorded by a smartphone with an activated GPS receiver or sequences of consecutive queries that have been sent from multiple positions to a location-based service.

The amount of uncertainty that is bound to a certain location or a trajectory depends on many factors. For example, if a pedestrian produces a location in the middle of a motorway or in a military exclusion area, an adversary might guess that this is implausible. Another example is a cyclist who has generated sequences of locations in distances that cannot be reached with the typical speed of a bicycle. We will address these issues in Section 4.4.6.

To ease our presentation, in the following we consider pairs of latitude, longitude only. However, all approaches described can be easily applied to more complex spatio-temporal settings by considering height and time as additional dimensions that are treated in the same way as latitude, longitude.

4.4.3 Obfuscation techniques

In general, obfuscation techniques can be applied by each user in isolation. One of the most intuitive techniques to increase the uncertainty of a location information is to add or multiply the latitude, longitude-record with a random numbers taken from a uniform distribution. The upper and lower bounds of the probability distribution function are a measure for the amount of uncertainty obtained.

A more sophisticated approach [1] takes the amount of uncertainty into account that has been already induced by the location sensor. In particular, the approach assumes that the correct position of a user is uniformly distributed over a circle that has been reported as center, radius by a GPS device. The radius specifies the accuracy of the location information. In this setting, uncertainty can be increased by shifting the center and enlarging or decreasing the radius.

Another way of obfuscation is the creation of a set of realistic dummies. With this approach, the user does not only sends a single position information to a location-based service, but a number of artificial dummy positions plus the real position. Accordingly, the service returns one result for each query. The client filters the set of results for that answer that corresponds to the real position. In this case, uncertainty is not defined as uncertainty towards a region (specified by a probability distribution), but as uncertainty towards which of the queried positions is the real one.

Finally, there exist approaches to replace latitude, longitude-pairs with the positions of prominent landmarks. For example, the exact position 49.530, 6.899 of a participant of a seminar in Dagstuhl could be reported as "Saarland", "Germany" or "Europe". The applicability of this kind of obfuscation depends on the format in which a location-based service requires a location information.

4.4.4 Anonymization approaches

Uncertainty by anonymization means to hide the identity and position of a user among a set of other users. Thus, anonymization requires the position information of multiple users. A popular approach is Spatio-Temporal Cloaking [2]. The approach adapts the concept of k-anonymity for geographic coordinates. For this purpose, the approach computes cliques of users that are close together, and releases minimum bounding rectangles that contain the positions of at least k different users each. Various variants of this concept exist, e.g. peer-to-peer-anonymization [2].

Mix zones are an approach to add uncertainty to spatiotemporal settings where the users are continuously observed by a service provider. The approach identifies each user by a pseudonym. Furthermore, it divides regions into mix zones and application zones. In predefined time intervals, all individuals within a mix zone have the option to choose a new pseudonym. Given the number of users in a mix zone is large enough, the service provider cannot link the movement of an individual in one application zone to the movement of the same individual in another application zone.

4.4.5 Legal status location privacy

The possibility to minimize violations of privacy can be achieved by creating uncertainty in location data while using location based services. An element in the definition of personal data in the EC Directive is that personal data indicates an identified or identifiable person. In other words the terms “identified or identifiable” focus on the conditions under which an individual should be considered as “identifiable”. In this regard the particular conditions of a specific case play an important role in this determination. Therefore the effect of uncertainty has to be addressed individually.

Location-based services in general process personal data in order to fulfill their contractual duties. The legal ground of using such information primarily is bound to the requirements of “informed consent” or “performance of a contractual duty” under EC Directive 95/46. Furthermore, a processing on a secondary basis requires the fulfillment of at least one of the exceptions under EC Directive such as the existence of a “legitimate interest” of the data processor or the existence of a “vital interest” of the data subject. In this context, the “legitimate interest of the data processor” criterion will be subject to further analysis. Especially the legal framework and – if any – case law will be pointed out.

4.4.6 Privacy analysis

One key challenge that must be addressed is how to analyze privacy. While there has been some research in this area (e.g., the talk by Maria Luisa Damiani 3.3), this is still a challenge with room for research.

We are considering the following agents: an application *provider* offering one or more dedicated services to a certain class of *clients*, which might be formed by subscription or even on an ad-hoc basis. Each of the clients might repeatedly request one of the services. Each request comes along with data about the requester’s location in order to enable the provider to return a location-dependent reaction to the requester. Without privacy-preserving measurements, the location would be determined in the best possible way, while however we are facing the possibility that there are technical failures leading to uncertainties regarding the actual precise position. In addition, each request is associated with a time stamp, which reflects real time by default. Introducing additional mechanisms for privacy-preservation enables a client to purposely generate further uncertainty about his or her actual location.

Though the time data might be blurred as well, for the sake of simplicity we leave this option open to future considerations. Accordingly, over the time the provider receives a sequences of requests, each consisting of at least the following components: location, time, kind of request.

Beyond supposedly honestly in each case actually providing the service requested, the provider may behave in a *curious* way, *analyzing* the collection of request data received so far for the sake of any secondary use. We identified three major basic kinds of analysis. Moreover we emphasized the following need of a client: in order to decide about the employment of uncertainty generating mechanisms, a requesting client has to evaluate the potentials for a successful analysis by the provider according to some metrics.

- Location-based reidentifiability:

In case that requests are coming without the respective requester's identification, the provider might aim to associate an identifier to each of the requests, at the best definitely the identifier of the actual requester; alternatively and less ambitiously, some assertion about the relationship between the request and the clients. As far as the provider succeeds in establishing a nontrivial and meaningful association, he would either learn precise personal data or obtain data that is somehow potentially personal and, thus, he would be able to compromise privacy to some extent.

Seen from the point of view of a client acting as requester, that client would be interested to estimate the extent of compromise achievable by the provider according to some suitable metric.

- Location identification and classification:

Whether by technical failures or by intentional blurring, a communicated location might differ from the actual one. Accordingly, the provider might aim to determine the actual location by some kind of reasoning, thereby strengthening his knowledge about the requester. As far as the requester is already identified, in this way the provider would obtain improved personal data regarding the requester. If the requester is so far not fully identified, more precise knowledge about the actual location might be helpful for the reidentification analysis or other analysis tasks.

Besides the pure geographical data about the location, the provider might also aim at determining the kind of social activities offered at the respective place and thus learning information of the requester's activities, again potentially leading to even more crucial personal data. Typically, social activities could be classified and denominated according to some ontology, e.g., distinguishing between shopping, medical care, entertainment, food services, sports, education, and so on, even possibly refined to subcategories and enhanced by further descriptive features.

Again, the client would like to evaluate the expected achievements, in particular in terms of the grade of success and the sensitivity of an identified location depending on its semantics, according to some metric.

- Subtrajectory linkage

While strictly speaking a client only communicates location-time points, she or he actually provides information about her or his movements over the time, i.e., about the resulting trajectory. Accordingly, the provider might aim at reconstructing the actual trajectory in an approximative way in order to learn more about the client. As before, besides the pure geographical data about the full trajectory, he might additionally be interested in the semantics of the curve in terms of a suitable ontology that extends the ontology for single locations.

Reconstructing an actual trajectory necessarily includes linking single locations as communicated and subtrajectories obtained before as belonging to the same client. This need

is clearly supported by already knowing the association of the requests with identifiers, but also conversely, if originally unknown learning this association might be facilitated by having established links before.

Again, the client would like to evaluate the expected achievements, in particular in terms of the grade of success and the sensitivity of a reconstructed trajectory depending on its semantics, according to some metric.

4.4.7 Parameters, how relevant, and how to estimate

The success of the analyses described above and in particular the quality of the results achieved depend on a wide range of parameters. Such parameters might refer to both the data provided by the requests and various kinds of background knowledge available to or even generated by the provider. In this subsection we briefly discuss some important examples.

- Plausible locations and their semantics:
Given a region in purely geographical terms, e.g., described by a circle, the provider is likely to possess one or more maps including this circle and its further environment as a priori background knowledge. Each such map provides some kind of semantics, in particular suggesting plausible actual positions within the region and a classification of many objects within the region.
- Density and properties of other clients:
Given data that indicates that some client, whether identified or not, stays within a region, the provider might also have gained the knowledge that other clients are staying near by. This knowledge might have various effects. For example, the existence of many unidentified clients within a region results in forming an anonymity class for the client considered, and this fact might support the client's privacy concerns. Conversely, having learned the semantics of the stay of sufficiently many other clients might suggest a semantics for the client considered.
- History-frequency of appearance/past revelation:
Analyses might not only be based on the data of the most recent requests but also refer to histories and their evaluations. For example, previously successfully identified and classified trajectories might be related to a recent trajectory under inspection, suggesting an identification or semantics based on similarities. Useful notions of similarity and the resulting suggestions might have been obtained by means of data mining applied to the data received before and stored in a repository.

4.4.8 QoS analysis

The degradation in quality of service from intentionally giving uncertain locations is a critical issue. While this is largely dependent on the particular application, and the implementation of that service provider, we can experimentally derive quality of service measures. By comparing the results from requesting a service with the actual location at the best anticipated resolution with the results from intentionally degraded uncertainty, we can establish how great the impact of a given uncertainty method is on the particular application.

While metrics are also application dependent, many location-based services return ranked lists. Examples include location-based keyword search, recommendation systems, closest point of interest, driving route finders, and public transit schedule systems. In all these cases, we can compare the impact of uncertainty by comparing the results with uncertain location against the results with actual location. There are standard ways to do this, such as KL-Divergence[3].

We plan to choose specific applications from Section 4.4.1, and evaluate against all relevant techniques described in Section 4.4.2 at various levels of uncertainty. Levels of uncertainty will be chosen to achieve interesting privacy points as determined in Section 4.4.6. Results will be presented by graphing the KL-Divergence across various parameter changes.

4.4.9 Secondary use utility analysis

Can we relate type of uncertainty to impact on classes of mining.

- What is the effect of the user-injected uncertainty in location on aggregate data collected on service provider side? Can we model it mathematically?
- The common business practice is that the service providers use whatever the user provides during the service usage. What law/regulations stipulate for the secondary use does not necessarily reflect current practice. We do not intend to change it. In fact, in certain circumstances the secondary process can create benefits.
- Following up on the previous point, the service provider may not be able to provide the adequate protection for the data submitted by the users which is subject to theft/compromise by other parties who will do the secondary (mis-)use.
- Can service provider learn more about personal information that intended after the aggregated data is obtained? To what extent? Can it be the case the uncertainty turns out to be not a worthwhile effort?
- Find a good example as a motivation for the case parties find secondary processing useful/beneficial/unharmful
- Experimental work can be useful to compare the data mining results extracted from aggregate data with uncertainty against those without.

4.4.10 Questions that remain to be answered

4.4.10.1 Other personal data

Location data isn't the only data exposed to location-based services.

- A query may contain some other personal information, which is necessary for the completion of service. What are the implications for privacy? Can a service provider get more information than intended?
- Does user provide his/her identity in the query? Does s/he query anonymously?
- What happens if the service requires tracking of user for a certain time interval?

4.4.10.2 Safety applications

Systems would need to be designed to bypass the uncertainty when safety of individual necessitates accurate location information

4.4.10.3 Acceptance/Feasibility

There are also issues concerning practical application of uncertainty. Are service providers able and willing to work with inaccurate data? Can the proposed method be implemented with current technology? (This seems feasible for Android, perhaps not for Apple.)

4.4.10.4 Adverse affects

False location declaration can lead to undesired situations; for instance putting a person in potentially problematic locations (e.g., in a crime scene)

References

- 1 Claudio Agostino Ardagna, Marco Cremonini, Ernesto Damiani, Sabrina De Capitani di Vimercati, and Pierangela Samarati. Location privacy protection through obfuscation-based techniques. In *Data and Applications Security XXI, 21st Annual IFIP WG 11.3 Working Conference on Data and Applications Security*, Redondo Beach, CA, USA, July 8-11 2007.
- 2 Marco Gruteser and Dirk Grunwald. Anonymous usage of location-based services through spatial and temporal cloaking. In *Proceedings of First ACM/USENIX International Conference on Mobile Systems, Applications, and Services (MobiSys)*, May 2003.
- 3 Chengxiang Zhai and John Lafferty. A study of smoothing methods for language models applied to information retrieval. *ACM Transactions on Information Systems*, 22(2):179–214, April 2004.

5 Open Problems

In addition to the developments and plans for future work in the working group reports, the seminar saw general lessons learned and future directions for Mobility Data Mining and Privacy.

5.1 What we learned

- Privacy issues in mobile data are real
 - Applications collect much more data than needed
 - Serious potential for harm, some evidence of actual harm
- Privacy violations may have direct societal impact
 - Not just the individual who is harmed
- Data previously highly regulated moving to lightly/unregulated businesses
- Wide variety of vocabularies used to talk about mobility, data mining, privacy
 - Poor understanding of cross-community issues by different subcommunities
- Dimensionality of uncertainty/unknowns
- Problems not clearly defined (or clearly understood even by experts)
- Age-related bias in assessment of issues

5.2 New Discoveries

- Danger of location data plus additional information gives supra-linear increase in risk
 - This linkage needs further investigation: Re-identification attacks, potential for damage, ...
- Sensors / content in smart phones lead to severely increased privacy risk
 - Lots of information
 - Easily revealed
 - Often disclosed with little awareness on part of individual
- Privacy preferences/expectations are dynamic
 - Need approaches beyond static privacy preference options
 - Location and time key components of this dynamism

- Need user studies on mobility and privacy
- Re-identification, particularly of mobile data, is too hard to prevent
 - Need broader perspectives on how to protect privacy
 - Disconnect between technology developments and law/regulation
- Technology developments are resulting in even “technology-neutral” privacy law becoming outdated

5.3 What needs to be done

- How to educate
 - Privacy technologists don’t understand law and regulation
 - Regulators don’t understand technical privacy definitions
 - Users don’t understand either
 - Application (and infrastructure) developers don’t worry about privacy
 - Curricula for privacy – not just “how to comply”
- Technology needs better guidance from law and regulation
 - And back to education – need to understand guidance
- Need to be suitably proactive
 - don’t want to let the cat out of the bag
 - Need for people to desire privacy beyond legal mandates
 - Need tools to manage (most) privacy that are general – user specifies their privacy preferences, not what they want a particular app to do.
- How do we manage cross-border data on the internet?
- Privacy research needs better understanding of economic value of data
- Venue/mechanism to better support multidisciplinary work to bring privacy into various communities

5.4 Future plans

Based on the above conclusions, the seminar participants felt that we need future study and discussion in this area. One option would be a future Dagstuhl seminar, as well as a more standard conference or workshop. More immediately, we felt that a panel would be appropriate, but this needs to be in a venue that reaches to the mobile data and location based computing community, not to the privacy community.

The study groups are pursuing further research and publication of the ideas from the seminar. However, a consolidated white paper on privacy and mobility, possibly with an accompanying tutorial, may also be a way to further raise the issues. The area also needs a study of the broader societal impact of privacy breach – social mores, economic issues, security implications – this may be a good topic to target an international grant proposal.

Participants

- Gennady Andrienko
Fraunhofer IAIS –
St. Augustin, DE
- Nilgün Basalp
Istanbul Bilgi University, TR
- Joachim Biskup
TU Dortmund, DE
- Erik Buchmann
KIT – Karlsruhe Institute of
Technology, DE
- Christopher W. Clifton
Purdue University, US
- Maria Luisa Damiani
University of Milano, IT
- Glenn Geers
NICTA – Kensington, AU
- Aris Gkoulalas-Divanis
IBM Research – Dublin, IE
- Marco Gruteser
Rutgers Univ. –
New Brunswick, US
- Christine Körner
Fraunhofer IAIS –
St. Augustin, DE
- Bart Kuijpers
Hasselt University –
Diepenbeek, BE
- Thomas Liebig
Fraunhofer IAIS –
St. Augustin, DE
- Bradley Malin
Vanderbilt University, US
- Michael Marhöfer
Oberhaching, DE
- Walied Othman
Universität Zürich, CH
- Klaus Rechert
Universität Freiburg, DE
- Claudia Santa
TU München, DE
- ErKay Savas
Sabanci University –
Istanbul, TR
- Florian Schaub
Universität Ulm, DE
- Olaf Spinczyk
TU Dortmund, DE
- Christian Wietfeld
TU Dortmund, DE
- Ouri E. Wolfson
University of Illinois –
Chicago, US



Verifying Reliability

Edited by

Görschwin Fey¹, Masahiro Fujita², Natasa Miskov-Zivanov³,
Kaushik Roy⁴, and Matteo Sonza Reorda⁵

- 1 German Aerospace Center (DLR) and University of Bremen, DE, fey@informatik.uni-bremen.de
- 2 University of Tokyo, JP, fujita@ee.t.u-tokyo.ac.jp
- 3 University of Pittsburgh, US, nam66@pitt.edu
- 4 Purdue University, US, kaushik@ecn.purdue.edu
- 5 Politecnico di Torino, IT, matteo.sonzareorda@polito.it

Abstract

Moore's law has been the driving force behind the increasing computing power of today's devices which is based on shrinking feature sizes. This shrinking process makes future devices extremely susceptible to soft errors due to, e.g., external influences like environmental radiation and internal issues like stress effects, aging and process variation. For future technology nodes "Designing reliable systems from unreliable components" ¹ will be one of the most important topics.

Seminar 19.–24. August, 2012 – <http://www.dagstuhl.de/12341>

1998 ACM Subject Classification B.6.2 Reliability and Testing [**] (B.8); B.6.3 Design Aids; B.7.3 Reliability and Testing [**] (B.8); B.7.2 Design Aids

Keywords and phrases Reliability, fault modeling, formal methods

Digital Object Identifier 10.4230/DagRep.2.8.54

Edited in cooperation with Mehdi Dehbashi (University of Bremen, DE)

1 Executive Summary


Görschwin Fey

Masahiro Fujita

Natasa Miskov-Zivanov

Kaushik Roy

Matteo Sonza Reorda

License  Creative Commons BY-NC-ND 3.0 Unported license

© Görschwin Fey, Masahiro Fujita, Natasa Miskov-Zivanov, Kaushik Roy, and Matteo Sonza Reorda

Introduction

Moore's law predicted the ever increasing computing power of the past decades from an economic perspective based on doubling the number of elements in a circuit about every two years. Moreover, Moore's law is expected to continue for another 10-20 years. On the physical level this integration is enabled by continuously shrinking feature sizes of basic components. But for future technology nodes reliability problems triggered by different

¹ Shekhar Y. Borkar, "Designing reliable systems from unreliable components: the challenges of transistor variability and degradation," IEEE Micro 25(6): 10-16 (2005)



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

Verifying Reliability, *Dagstuhl Reports*, Vol. 2, Issue 8, pp. 54–73

Editors: Görschwin Fey, Masahiro Fujita, Natasa Miskov-Zivanov, Kaushik Roy, and Matteo Sonza Reorda



DAGSTUHL
REPORTS

Dagstuhl Reports

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

sources are expected to increase rapidly. Process variations in the production process are one issue. While production processes become more accurate considering absolute measures, the relative inaccuracy compared to the component's size is increasing. One consequence are transistors with a wide range of threshold levels resulting in slightly faster or slower operating logic circuitry (both die-to die and within die). This may result for example in delay errors under certain operating conditions of a device. Increasing sensitivity to the omnipresent environmental radiation is another issue. In the past some errors induced by radiation have been observed infrequently while systems in space missions are already specified to be radiation resistant. Shrinking feature sizes result in sensitivity to radiation with lower energy causing more radiation induced events like Single Event Upsets (SEUs) even on sea level. Such effects are summarized as transient faults resulting in soft errors (as opposed to permanent faults resulting in a change of the functionality due to a modification of the physical structure). Consequently, approaches to design reliable circuits tolerating such transient faults without causing soft errors have been proposed. These design approaches to mitigate soft errors comprise all levels of design abstraction from the system specification down to the layout. Examples for these approaches are, e.g., fault tolerant algorithms and operating systems, fault tolerant processors, self-calibrating architectures, block level redundancy and error checking, synthesis approaches on the gate level, or hardening techniques on the layout level. In practical systems typically multiple mitigation techniques are implemented to guarantee reliability across the full system stack. Functional verification has been and still is a challenge in current designs containing up to hundreds millions of transistors. Mature techniques for the formal verification and the dynamic verification of large systems exist. Research in verification is ongoing to match the rapid increase of the size of the systems. The verification of reliability is an interdisciplinary topic involving at least testing technology, verification methodology, and design experience. This makes the verification of reliable implementations an even harder problem. The testing community provides underlying models for transient faults to understand the effects at the functional and eventually at the system level. Using these models, the verification community designs efficient analysis tools and verification techniques to handle large systems. As in standard verification of large circuits a concerted action of formal methods, semi-formal techniques and simulation-based validation will be required. Still knowledge from the design community is required, to further speed up the verification task. Understanding the implemented approach to reliability on the application level and the system level is required to achieve a high degree of automation in the verification task.

Organization

The seminar was organized in short slots for talks followed by extensive discussions. A panel discussion in the afternoon summarized each day and focused on further questions (Figure 1). Each day was devoted to a special topic:

- Design – Techniques to ensure reliability by design.
- Fault models – Different types of fault models are required depending on the abstraction level and the type of design considered.
- Metrics – Measuring reliability requires some kinds of metrics. These metrics can be defined with respect to the fault models. But they should also reflect potential inaccuracies.
- Engines – Different types of engines are used in Electronic Design Automation (EDA) for circuits and systems.

	Monday Design	Tuesday Fault models & Metrics I	Wednesday Engines I	Thursday Engines II + Metrics II	Friday Lessons learned
	Morning chair: Massimo Violante	Masahiro Fujita	Subhashish Mitra	Rolf Drechsler	Carsten Gebauer
9 am	Welcome, <i>Introducing everybody</i> Anand Raghunathan, Kaushik Roy: <i>Approximate Computing - Embracing Unreliability for Efficient Computing</i>	Suddhakar M. Reddy: <i>Gracefully Degradable Higher Performance Systems</i> Carsten Gebauer: <i>Issues with applying fault tolerance in safety critical automotive applications</i>	Ravishankar K. Iyer: <i>Experimental Validation of Computer Systems Dependability</i> Bernd Becker, Matthias Sauer: <i>Improving reliability by improving ATPG accuracy</i>	Cecile Braunstein: <i>A Symbolic Model-Checking Framework for Transient Fault Robustness Classification and Quantification</i> Jie Han: <i>Stochastic Computational Approaches for Accurate and Efficient Reliability Evaluation</i>	Seiji Kajihara: <i>Test Partitioning for BIST-based field test</i> Wenchao Li: <i>Requirement Analysis and Generation for Verification-Guided Error Resilience</i>
10:30 am	Coffee break	Coffee break	Coffee break	Coffee break	Coffee break
10:45 am	John P. Hayes: <i>Stochastic Computing Revisited</i> Adit Singh: <i>The Reliability Challenge from Variability Induced Timing Errors</i>	Matteo Sonza Reorda: <i>Reliability evaluation in complex systems: some cases of study and related lessons</i> Mehdi B. Tahoori: <i>Wearout Modeling and Mitigation at Higher Levels of Abstraction</i>	Bodo Hoppe: <i>Verifying architectural compliant recovery</i> Massimo Violante: <i>Validating fault tolerant designs in SRAM-based FPGAs: how to keep the route in an ocean of bits</i>	Marcela Simkova: <i>Towards Beneficial Hardware Acceleration of Functional Verification</i> Rolf Drechsler: <i>Completeness-Driven Development</i>	Robert Aitken: <i>Scaling, Errors, and Reliability: When Will the World End and Why It Hasn't So Far</i> Wrap-up-Discussion
12:15 am	Lunch	Lunch	Lunch	Lunch	Lunch
1:30 pm	Afternoon chair: Eli Arbel	Laurence Pierre	Excursion & dinner outside	Seiji Kajihara	Departure
Thursday: 1:15 pm	Michael Orshansky: <i>When Perfect is the Enemy of Efficient: Using Controlled Errors in Approximate Computing</i> Ilia Polian: <i>Towards a Cross-Layer Strategy Against Fault-based Attacks</i>	Eli Arbel: <i>Reliability closure – how can we do better?</i> Sachin Sapatnekar: <i>How does device-level reliability affect my system?</i>		Laurence Pierre: <i>On the use of semi-formal methods for reliability analysis (at RT and TLM abstraction levels)</i> Jacob A. Abraham: <i>Software-level Fault Injection: an Effective Approach to Validating System Reliability</i>	
3:00 pm	Masahiro Fujita: <i>Error Tolerance and Engineering Change with Partially Programmable Circuits and their SAT-Based Programming</i>	Ulf Schlichtmann: <i>How to efficiently analyze aging effects in large circuits - and some ideas how to use the results</i>		Shawn Blanton: <i>Improving Design, Manufacturing and Even Test through Test-Data Mining</i> Yusuke Matsunaga: <i>Probabilistic Analysis for Softerror Tolerance of Sequential Circuits</i>	
3:45 pm	Coffee break	Coffee break		Coffee break	
4:15 pm	Tomohiro Yoneda: <i>Designing a Dependable Network-on-Chip Platform for Automotive Applications</i>	Yasuo Sato: <i>Analysis of Field Test Effectiveness to LSI Reliability</i>		Yoshiki Kinoshita: <i>Validating Open Systems Dependability</i>	
5:00 pm	Panel: <i>Beyond the limitations of approximate and statistical computing</i> Chair: Suddhakar M. Reddy Panelists: J.P. Hayes, M. Orshansky, A. Raghunathan, S. Sapatnekar, A. Singh	Panel: <i>What's most urgent in industry?</i> Chair: Shawn Blanton Panelists: R. Aitken, E. Arbel, C. Gebauer, B. Hoppe, Y. Sato		Panel: <i>Fault Models, Metrics & Engines</i> Chair: Jacob A. Abraham Panelists: C. Braunstein, Y. Kinoshita, U. Schlichtmann, S. Blanton, M. Tahoori	
6 pm	Dinner	Dinner		Dinner	
8 pm	Cheese in the cafeteria	Cheese in the cafeteria		Cheese in the cafeteria	

■ Figure 1 Seminar schedule

Results

Documenting the results of intensive discussions in a compact manner is difficult. However, some results can be formulated in crisp statements. Approximate computing is a powerful technique for reliable design where the applications permit inaccuracy of operations up to a certain extent. Computing considering statistical nature of devices may be able to produce very accurate results, but providing compatible computing fabric at acceptable costs is a challenge. No single fault model will cover all aspects of reliability. In particular, fault models must be adapted to the application domain, the level of criticality and the step in the design process that is being considered. Appropriate metrics will then be applied to bridge gaps, e.g., between different levels of abstraction. An orchestration of reasoning engines ranging from formal techniques to simulation and emulation will always be required to gather data required for the different metrics. Design for Reliability will always affect all levels of abstraction. Only by concerted effort the same performance gains can be expected that we have seen in the past 50 years.

As a follow-up of the Dagstuhl Seminar, an Embedded Tutorial was successfully proposed for the DATE conference 2013. The Embedded Tutorial's title is "Reliability Analysis Reloaded: How Will We Survive?" and will include two presentations given by participants of the seminar or colleagues belonging to the research group of a participant.

2 Table of Contents

Executive Summary

<i>Görschwin Fey, Masahiro Fujita, Natasa Miskov-Zivanov, Kaushik Roy, and Matteo Sonza Reorda</i>	54
--	----

Overview of Talks



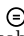

Software-Level Fault Injection: An Effective Approach to Validating System Reliability <i>Jacob A. Abraham</i>	59
Scaling, Errors, and Reliability: When Will the World End and Why It Hasn't So Far <i>Robert Aitken</i>	59
Reliability Closure – How Can We Do Better? <i>Eli Arbel</i>	59
Improving Reliability by Improving ATPG Accuracy <i>Bernd Becker, Matthias Sauer</i>	60
Improving Design, Manufacturing and Even Test through Test-Data Mining <i>Shawn Blanton</i>	60
A Symbolic Model-Checking Framework for Transient Fault Robustness Classification and Quantification <i>Cecile Braunstein</i>	61
Completeness-Driven Development <i>Rolf Drechsler</i>	61
Error Tolerance and Engineering Change with Partially Programmable Circuits and their SAT-Based Programming <i>Masahiro Fujita</i>	62
Issues with Applying Fault Tolerance in Safety Critical Automotive Applications <i>Carsten Gebauer</i>	62
Stochastic Computational Approaches for Accurate and Efficient Reliability Evaluation <i>Jie Han</i>	62
Stochastic Computing Revisited <i>John P. Hayes</i>	63
Verifying Architectural Compliant Recovery <i>Bodo Hoppe</i>	63
Test Partitioning for BIST-Based Field Test <i>Seiji Kajihara</i>	64
Validating Open Systems Dependability <i>Yoshiki Kinoshita</i>	64
Requirement Analysis and Generation for Verification-Guided Error Resilience <i>Wenchao Li</i>	64
Probabilistic Analysis for Soft-Error Tolerance of Sequential Circuits <i>Yusuke Matsunaga</i>	65

When Perfect is the Enemy of Efficient: Using Controlled Errors in Approximate Computing <i>Michael Orshansky</i>	65
On the Use of Semi-Formal Methods for Reliability Analysis (at RT and TLM Abstraction Levels) <i>Laurence Pierre</i>	65
Towards a Cross-Layer Strategy Against Fault-based Attacks <i>Iliá Polian</i>	66
Approximate Computing - Embracing Unreliability for Efficient Computing <i>Anand Raghunathan, Kaushik Roy</i>	66
Gracefully Degradable Higher Performance Systems <i>Sudhakar M. Reddy</i>	67
How does Device-Level Reliability Affect my System? <i>Sachin Sapatnekar</i>	67
Analysis of Field Test Effectiveness to LSI Reliability <i>Yasuo Sato</i>	67
How to Efficiently Analyze Aging Effects in Large Circuits - and Some Ideas How to Use the Results <i>Ulf Schlichtmann</i>	68
Towards Beneficial Hardware Acceleration of Functional Verification <i>Marcela Simkova</i>	68
The Reliability Challenge from Random Process Variability Induced Timing Errors <i>Adit Singh</i>	69
Reliability Evaluation in Complex Systems: Some Cases of Study and Related Lessons <i>Matteo Sonza Reorda</i>	69
Wearout Modeling and Mitigation at Higher Levels of Abstraction <i>Mehdi B. Tahoori</i>	69
Validating Fault Tolerant Designs in SRAM-Based FPGAs: How to Keep the Route in an Ocean of Bits <i>Massimo Violante</i>	70
Designing a Dependable Network-on-Chip Platform for Automotive Applications <i>Tomohiro Yoneda</i>	70
Panel Discussions	
Beyond the Limitations of Approximate and Statistical Computing	71
What's Most Urgent in Industry?	71
Fault Models, Metrics & Engines	72
Participants	73

3 Overview of Talks

3.1 Software-Level Fault Injection: An Effective Approach to Validating System Reliability





Jacob A. Abraham (Univ. of Texas at Austin, US)

License     Creative Commons BY-NC-ND 3.0 Unported license
© Jacob A. Abraham

Accurate evaluation of the reliability of a complex system is extremely difficult since faults at the hardware level have to be analyzed with respect to their impact on the system under varying operating conditions and workloads. Simulating a system for billions of processor cycles for different types of low-level faults is practically impossible. This talk will describe techniques which evaluate the dependability of a system by running it under normal conditions and using software routines to inject faults which emulate the effect of the hardware on system behavior. Examples of applying the ideas to a variety of systems will be described, as well as directions for exploiting virtualization and other hardware support provided by modern processors.

3.2 Scaling, Errors, and Reliability: When Will the World End and Why It Hasn't So Far





Robert Aitken (ARM Inc. - San Jose, US)

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

Technology scaling continues to follow the basics of Moore's Law, despite difficult challenges in lithography, materials, and design. Looking at what has succeeded so far can give insight into why several predicted demises of scaling have not happened, as well as showing which of the current candidates might actually succeed. Gordon Moore said "No exponential is forever, but we can delay forever" - when will forever arrive?

3.3 Reliability Closure – How Can We Do Better?


Eli Arbel (IBM - Haifa, IL)

License     Creative Commons BY-NC-ND 3.0 Unported license
© Eli Arbel

A typical design process involves dealing with multiple constraints, such as power, performance and timing. As reliability is becoming increasingly important in many applications, it can be viewed as an additional design constraint which should be met before closing the design. Many reliability features are implemented in the RT-level, thus it is of high importance to provide feedback to logic designers whether their design meets its reliability goals, and if not assist them with rectifying reliability issues. We will present some techniques for analyzing design vulnerability to soft-errors at the RT-level, how they can be used to facilitate logic implementation with respect to reliability and discuss how we can help design teams achieve their reliability goals in more accurate and faster ways.

3.4 Improving Reliability by Improving ATPG Accuracy


Bernd Becker, Matthias Sauer (Universität Freiburg, DE)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Bernd Becker, Matthias Sauer

We present SAT-based approaches - implemented in the tools Phaeton and WaveSAT- for the analysis of circuit timing and the detection of small delay defects. Phaeton enumerates all or a user-specified number of longest sensitisable paths in the whole circuit or through specific components. The algorithm encodes all aspects of the path search as an instance of the Boolean Satisfiability Problem(SAT), which allows the method not only to benefit from recent advances in SAT-solving technology, but also to avoid some of the drawbacks of previous structural approaches. The path information obtained by Phaeton can be used for several applications including design and test of circuits affected by statistical process variations, criticality analysis, and post silicon debug. The approach has been extended to sequential ATPG making use of recent advances in Bounded Model Checking. However, the computation of small-delay fault test patterns by path sensitization may result in false positives and false negatives as well. We developed WaveSAT, a SAT-based automatic test pattern generation algorithm which considers waveforms and their propagation on each relevant line of the circuit. The model incorporates individual delays for each gate and filtering of small glitches. WaveSAT generates a test if the fault is testable and is also capable of automatically generating a formal redundancy proof for undetectable small-delay faults. Experimental results for academic and industrial benchmark circuits demonstrate the methods' accuracy and scalability.

3.5 Improving Design, Manufacturing and Even Test through Test-Data Mining



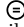
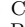
Shawn Blanton (Carnegie Mellon University - Pittsburgh, US)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Shawn Blanton

For many years now, ACTL (Advanced Chip Test Laboratory, www.ece.cmu.edu/actl) at Carnegie Mellon has been using layout information for improving manufacturing test, in particular, for changing test from a sort-only process to one that also involves learning about the design, the manufacturing process, and their interaction in producing high-yielding, high-quality parts. In this talk, I will describe METER (MEasuring Test Effectiveness Regionally), a novel approach that measures the effectiveness of arbitrary fault models and test metrics through the statistical analysis of readily-available tester data.

3.6 A Symbolic Model-Checking Framework for Transient Fault Robustness Classification and Quantification


Cecile Braunstein (UPMC - Paris, FR)

License     Creative Commons BY-NC-ND 3.0 Unported license
© Cecile Braunstein

Robustness analysis of RTL-sequential circuits impacted by transient faults is an important concern for designers. While simulation or emulation based techniques are widely used, they do not give guarantees on the robustness level of the system and are often limited to single fault models. Moreover, several robustness criterion may be adopted depending on the application being executed and the synchronisation scheme between the circuit and its environment. The use of formal methods ensures robustness level and helps in locating weak portions of a circuit to be hardened, even in case of multiple fault models. We present a framework to analyse the robustness of a RTL circuit, considering several models of faults and reparation, and show how a wideclass of robustness criteria can be mapped into our reparation model. We present an implementation of the robustness measures in the setting of BDD-based model checking and illustrate our measurements on classical benchmark circuits.

3.7 Completeness-Driven Development


Rolf Drechsler (Universität Bremen and DFKI Bremen, DE)

License     Creative Commons BY-NC-ND 3.0 Unported license
© Rolf Drechsler

Due to the steadily increasing complexity, the design of embedded systems faces serious challenges. To meet these challenges additional abstraction levels have been added to the conventional designflow resulting in Electronic System Level (ESL) design. Besides abstraction, the focus in ESL during the development of a system moves from design to verification, i.e. checking whether or not the system works as intended becomes more and more important. However, at each abstraction level only the validity of certain properties is checked. Completeness, i.e. checking whether or not the entire behavior of the design has been verified, is usually not continuously checked. As a result, bugs may be found very late causing expensive iterations across several abstraction levels. This delays the finalization of the embedded system significantly. In this work, we present the concept of Completeness-Driven Development (CDD). Based on suitable completeness measures, CDD ensures that the next step in the design process can only be entered if completeness at the current abstraction level has been achieved. This leads to an early detection of bugs and accelerates the whole design process. The application of CDD is illustrated by means of an example.

3.8 Error Tolerance and Engineering Change with Partially Programmable Circuits and their SAT-Based Programming

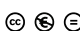
Masahiro Fujita (*University of Tokyo, JP*)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Masahiro Fujita

Introducing partial programmability in circuits by replacing some gates with *Look Up Tables* (LUTs) can be an effective way to improve post-silicon or in-field rectification and debugging. Although finding configurations of LUTs that can correct the circuits can be formulated as a QBF problem, solving it by state-of-the-art QBF solvers is still a hard problem for large circuits and many LUTs. In this paper, we present a rectification and debugging method for combinational circuits with LUTs by repeatedly applying Boolean SAT solvers. Through the experimental results, we show our proposed method can quickly find LUT configurations for large circuits with many LUTs, which cannot be solved by a QBF solver.

3.9 Issues with Applying Fault Tolerance in Safety Critical Automotive Applications


Carsten Gebauer (*Robert Bosch GmbH - Schwieberdingen, DE*)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Carsten Gebauer

Fault tolerance will be one of the key features necessary to be able to take advantage from the ever shrinking process technologies. However applying fault tolerance also has its risks. Within this talk I would like to present from an automotive point of view the issues we see regarding safety and ask for possible solutions to address these issues, in particular - latent faults of ISO 26262 - *Error Correction Codes* (ECC): Reduction of error detection due to application of correction - what is reasonable to tolerate, what not?

3.10 Stochastic Computational Approaches for Accurate and Efficient Reliability Evaluation

Jie Han (*University of Alberta, CA*)




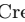
License  Creative Commons BY-NC-ND 3.0 Unported license
© Jie Han

Reliability is fast becoming a major concern due to the nanometric scaling of CMOS technology. Accurate analytical approaches for the reliability evaluation of logic circuits, however, have a computational complexity that generally increases exponentially with circuit size. This makes intractable the reliability analysis of large circuits. This talk initially presents novel computational models based on stochastic computation; using these stochastic computational models (SCMs), a simulation-based analytical approach is then proposed for the reliability evaluation of logic circuits. In this approach, signal probabilities are encoded in the statistics of random binary bit streams and non-Bernoulli sequences of random permutations of binary bits are used for initial input and gate error probabilities. By leveraging the bit-wise dependencies of random binary streams, the proposed approach

takes into account signal correlations and evaluates the joint reliability of multiple outputs. Therefore, it accurately determines the reliability of a circuit; its precision is only limited by the random fluctuations inherent in the stochastic sequences. Based on both simulation and analysis, the SCM approach takes advantages of ease in implementation and accuracy in evaluation. The use of non-Bernoulli sequences as initial inputs further increases the evaluation efficiency and accuracy compared to the conventional use of Bernoulli sequences, so the proposed stochastic approach is scalable for analyzing large circuits. It can further account for various fault models as well as calculating the soft error rate (SER). These results are supported by extensive simulations and detailed comparison with existing approaches.

3.11 Stochastic Computing Revisited




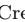
John P. Hayes (University of Michigan, US)

License     Creative Commons BY-NC-ND 3.0 Unported license
© John P. Hayes

Stochastic computing (SC) was proposed in the 1960s as a low-cost alternative to conventional computing with weighted binary numbers. It is unusual in that it represents information by means of long pseudo-random bit-streams, which can be interpreted as probabilities and processed by low-cost standard logic circuits. The SC approach is well-suited to some important new applications that require massive parallelism or extremely high tolerance of soft errors, such as ECC controllers for WiFi and flash memories. Despite some success in specialized application areas, many aspects of SC are poorly understood and a comprehensive design methodology has yet to emerge. This talk will review SC and its status from a modern perspective, focusing on design and verification issues affecting accuracy, area cost, and reliability. A novel approach to SC circuit design based on spectral transforms will also be presented.

3.12 Verifying Architectural Compliant Recovery


Bodo Hoppe (IBM Deutschland - Böblingen, DE)

License     Creative Commons BY-NC-ND 3.0 Unported license
© Bodo Hoppe

A lot of research has been done in investigating whether soft errors can be detected in hardware designs. Structural and static analysis are used to ensure any error can be detected. Formal analysis can be used to verify detectability of errors. However, today exhaustive verification has to be executed to verify that the design is reliable and able to recover properly according to the architecture. This has a lot of design complexities especially in presence of a multicore environment with coherent memory as well as high-frequency superscalarmicroprocessor designs with hardware recovery functionality. A lot of side effects may occur in the hardware and may lead to unwanted effects, for example *Potential Unexpected Loss of Data* (PULD). A method is being presented, that improves significantly the efficiency of proving that the design actually can maintain the architectural state including non-corrupted memory. But a much bigger question is can a certain design approach or rules can allow an easy proof that exhaustive fault simulation can be avoided. Or can design rulechecker in a combination with formal verification be used to ensure recoverability of the design?

3.13 Test Partitioning for BIST-Based Field Test

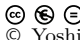
Seiji Kajihara (Kyushu Institute of Technology, JP)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Seiji Kajihara

A BIST-based field test has been used to guarantee high reliability of VLSIs. But it is not easy for field test to achieve high test quality due to the limitation of short test application time. Test partitioning and rotating test is an effective way to satisfy such a constraint. A test set of a circuit is partitioned into a number of subsets, and apply only one subset to the circuit at one test session of field test. On the other hand, because test partitioning would cause fault coverage loss at each test session, aging-induced faults will be undetected at some test sessions. The longer the detection interval is, the higher the likelihood of a system failure would be. In this talk we discuss on a metric to estimate the failure rate for test partitioning and approaches to partition a given test set into several subsets aiming to minimize the failure rate.

3.14 Validating Open Systems Dependability


Yoshiki Kinoshita (AIST - Hyogo, JP)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Yoshiki Kinoshita

I wish to convey the idea that checking the appropriateness of the claims to be verified (i.e., validation) is important in achievement of Open Systems Dependability. I also present assurance cases, which documents the result of validation. I shall use the system *European AIS [Aeronautical Information Service] Database (EAD)* as my leading example. Time permitting, I also introduce our ongoing work on D-Case/Agda, a system that supports development and checking of assurance cases using Proof Assistant technology. Verification is critical, but it should be associated with validation. Open Systems Dependability is dependability, the notion evolved from reliability, treated from Open Systems View, which our DEOS project has introduced to consider ever-changing and vague aspects of huge and complicated systems. Many systems have open systems aspects, and hardware seems no exception. Development and maintenance of the assurance case is central in achievement of Open Systems Dependability. Assurance cases are documents where all information about verification and validation of the system is available; they may be considered as "qualitative metric" of dependability.

3.15 Requirement Analysis and Generation for Verification-Guided Error Resilience

Wenchao Li (University of California - Berkeley, US)


License  Creative Commons BY-NC-ND 3.0 Unported license
© Wenchao Li

All error resilience techniques employ some form of redundancy, resulting in added cost such as area or power overhead. Formal Verification has been shown to be effective in judiciously

guiding the deployment of added error resilience. However, such analysis is sensitive to the quality and completeness of specifications. In this talk, I will discuss works on requirement analysis and generation with application to error resilience, as well as other useful applications such as error localization.

3.16 Probabilistic Analysis for Soft-Error Tolerance of Sequential Circuits


Yusuke Matsunaga (Kyushu University, JP)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Yusuke Matsunaga

This talk presents a method for estimating the error propagation probabilities in sequential circuits when one or more FFs' values are changed due to Soft Error Effects.

3.17 When Perfect is the Enemy of Efficient: Using Controlled Errors in Approximate Computing


Michael Orshansky (Univ. of Texas at Austin, US)

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

This talk describes our recent work on enabling low-level approximate computation and development of design principles for energy-optimal approximate ("sloppy") addition. We identify a fundamental trade-off between error frequency and error magnitude in a timing-starved adder and introduce a formal model to prove that for signal processing applications using a quadratic signal-to-noise ratio error measure, reducing bit-wise error frequency is sub-optimal. Instead, energy-optimal approximate addition requires limiting maximum error magnitude. The remaining approximation error can be reduced by conditional bounding logic for lower significance bits. We also show how the existence of an intrinsic notion of quality floor present in typical digital signal processing circuits can be used to reduce their energy consumption by strategically accepting some runtime errors. The basic philosophy is to prevent signal quality from severe degradation by using data statistics. The introduced innovations include techniques for carefully controlling possible errors and exploiting the specific patterns of errors for low-cost post-processing to minimize image quality degradation.

3.18 On the Use of Semi-Formal Methods for Reliability Analysis (at RT and TLM Abstraction Levels)

Laurence Pierre (TIMA - Grenoble, FR)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Laurence Pierre

Evaluating the robustness of hardware systems (ranging from digital IP blocks to complex systems on chip) with respect to soft errors has become an important part of the design flow

for many applications, at various levels of abstraction. To avoid the well-known state explosion problem that may occur when using formal approaches (static analysis), we investigate the use of semi-formal (simulation or emulation-based) techniques to improve dependability analysis. We target the analysis of the consequences of soft errors with respect to the application, no matter their origin. We propose to formalize robustness or reliability properties as PSL assertions, and to verify them at runtime using automatically derived property checkers. We discuss two illustrative examples: dependability properties of a cryptographic component for a programmable hardware device to be integrated in networking infrastructures (VHDL RTL description), and safety requirements for an avionics flight control remote module (SystemC TLM description).

3.19 Towards a Cross-Layer Strategy Against Fault-based Attacks

Ilia Polian (Universität Passau, DE)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Ilia Polian

Mobile and embedded systems process sensitive information, including personal data such as health records or financial transactions, and confidential parameters of technical systems, e.g., car engines. Protection is provided by cryptographic hardware blocks that are vulnerable to fault-based attacks. Over 700 such attacks have been published so far. Individual, attack-specific countermeasures no longer sufficient. There is strong need for a generic methodology to counter fault-based attacks with reasonable costs. The presentation will introduce fault-based attacks using the recent attack on the LED block ciphers as an example. After that, initial first results on across-layer protection strategy combining specially designed error-detecting codes with selective hardening of individual circuit elements will be presented.

3.20 Approximate Computing - Embracing Unreliability for Efficient Computing

Anand Raghunathan, Kaushik Roy (Purdue University, US)




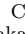
License  Creative Commons BY-NC-ND 3.0 Unported license
© Anand Raghunathan, Kaushik Roy

Designers of computing systems have been remarkably consistent in their approach to unreliability - attempt to eliminate it at all costs. While this approach has its merits and is necessary in some applications, it incurs very high costs in terms of efficiency (power, performance, or cost). With the explosion of digital data, computing platforms are increasingly being used to execute applications (such as web search, data analytics, sensor data processing, recognition, mining, and synthesis) that are inherently resilient or forgiving to errors in most of the computations. This forgiving nature is due to several factors including the redundancy and noisiness of the input data, the statistical nature of the computations themselves, and the acceptability (and often, inevitability) of less-than-perfect results. Approximate computing is an approach to designing computing platforms that are more efficient, by leveraging the forgiving nature of applications. I will outline a range of approximate computing techniques that we have developed from software to architecture to circuits, which have shown promising results. In the context of the increasing unreliability of scaled semiconductor technologies,

approximate computing suggests that embracing unreliability rather than attempting to eliminate it could be a promising approach to eschew the high costs of defect and fault tolerance. I will outline some of the challenges that need to be addressed to realize this promise, including a shift in designer mindset, and the development of systematic design methodologies to ensure that unreliability is exposed to applications in a controlled manner.

3.21 Gracefully Degradable Higher Performance Systems





Sudhakar M. Reddy (University of Iowa - Iowa City, US)

License     Creative Commons BY-NC-ND 3.0 Unported license
© Sudhakar M. Reddy

In this talk design of systems that permit trade off between performance and defect tolerance is proposed. It is suggested that additional features for defect tolerance should also facilitate enhanced performance when defects are not present or fewer than planned for number of defects are present. Enhanced performance could be higher frequency of operation or higher computational power or additional functionality.

3.22 How does Device-Level Reliability Affect my System?





Sachin Sapatnekar (University of Minnesota, US)

License     Creative Commons BY-NC-ND 3.0 Unported license
© Sachin Sapatnekar

The chip design process is inherently based on abstracting details of the design, and this is essential for complexity management. How can physics-based reliability models be used to analyze and optimize systems at higher levels? This talk will attempt to provide partial answers in this direction by discussing modeling methods for device-level failure mechanisms such as bias temperature instability, hot carrier injection, and gate oxide breakdown, where a great deal of advanced research has been carried out in the device community, but has not yet percolated far beyond.

3.23 Analysis of Field Test Effectiveness to LSI Reliability

Yasuo Sato (Kyushu Institute of Technology, JP)


License     Creative Commons BY-NC-ND 3.0 Unported license
© Yasuo Sato

Potential of field errors caused by LSI degradation such as NBTI, HCI or so on is increasing. These errors essentially look different from the conventional errors such as permanent errors or soft errors. It means that the conventional dependability theories might not well reflect their impacts on safety systems. The author analyzes their impacts on safety systems and tries to find a proper index, which shows the effectiveness of the concept of proposed field test DART. DART (Dependable Architecture with Reliability Testing) repeatedly measures the maximum delay of LSI and monitors delay margin in field. Using this approach, delay degradation can be detected before it will cause an actual system error. The technology

detail of DART is not discussed, but the relevance to the problem and what technologies should be developed by research people will be discussed.

3.24 How to Efficiently Analyze Aging Effects in Large Circuits - and Some Ideas How to Use the Results


Ulf Schlichtmann (TU München, DE)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Ulf Schlichtmann

Aging effects such as NBTI, PBTI, HCI are becoming more relevant as process technologies continue to scale. To date, commercial EDA tools only support the analysis of aging effects on transistor levels, thus severely limiting the size of designs that can be analyzed. We propose first a technique to efficiently analyze aging on gate level. We then introduce the concept of "potentially critical paths (PCPs)" which allows us to take the modeling of aging higher to the module level. We show how the concept of PCPs results in a further speedup of about 30x, without any loss in accuracy. Finally, we present some more ideas how the PCP concept can be used for further applications.

3.25 Towards Beneficial Hardware Acceleration of Functional Verification


Marcela Simkova (Brno University of Technology, CZ)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Marcela Simkova

Functional verification is a widespread technique to check whether a hardware design satisfies a given correctness specification. It is typically used in the pre-silicon phase of the design cycle to verify not only functional aspects but also reliability and safety properties. However, after the system is manufactured there are often found some previously uncovered errors. Moreover, further errors can be introduced by synthesis, mapping, place and route or fabrication processes. In order to eliminate as many remaining bugs as possible before a device is fabricated, verification is currently applied even in the post-silicon phase of the design cycle. Unfortunately, it is not possible to directly use the techniques from the pre-silicon phase (stimuli generation, assertion and coverage analysis, scoreboarding), and it is a challenging task to come up with techniques for post-silicon verification that would have strength comparable to the pre-silicon ones. In the presentation, I will talk about how to handle the gap between pre- and post-silicon verification using hardware acceleration with functional verification features. Furthermore, I will present HAVEN, an open framework for hardware acceleration of functional verification that provides means for seamless transition from pre- to post-silicon verification.

3.26 The Reliability Challenge from Random Process Variability Induced Timing Errors


Adit Singh (Auburn University, US)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Adit Singh

Current state-of-the-art timing test methods are not very effective in detecting delay faults in complex integrated circuits. Thus far, this has not been a major problem because manufacturing defects that cause subtle "delay only" failures are rare; they appear to be at least one to two orders of magnitude less frequent than defects that manifest as more easily detectable slow speed DC failures. Even if a significant fraction of delay defects remain undetected during production testing, the DPM impact on all but the lowest yielding ICs is generally quite modest. However, random transistor threshold voltage variation in aggressively scaled nanometer technologies is introducing a new source of timing variability. This variation is further amplified at the low operating voltages necessary for power minimization. Given the large number of transistors in a chip, hundreds of random "delay defects" can be statistically expected in every manufactured part in end-of-roadmap CMOS technologies. Moreover, because the increase in delay caused by a statistically slow transistor is potentially unbounded, each IC will need to be carefully tested for timing and reliably speed binned for use. This is a different and more formidable problem than the traditional speed binning of processors which is mostly aimed at handling systematic process variations –there are concerns whether delay test methodologies will be up to the task. The incorrect assignment of a higher speed to an IC because of improperly tested slow paths can result in operational failures in the field from timing errors. We discuss the significance of this emerging reliability challenge, and test and fault tolerance methods needed to address it.

3.27 Reliability Evaluation in Complex Systems: Some Cases of Study and Related Lessons


Matteo Sonza Reorda (Politecnico di Torino, IT)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Matteo Sonza Reorda

Assessing the reliability of complex systems is a challenging task. The talk will provide a couple of examples where this task has been performed, with details about the environments, results, and difficulties. Lessons will be drawn and open issues highlighted.

3.28 Wearout Modeling and Mitigation at Higher Levels of Abstraction

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


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

As CMOS technology enters in nanoscale regimes, the reliability of VLSI chips is threatened by various issues such as increased process variation, radiation-induced soft errors, as well as transistor and interconnect aging. For cost-efficient resilient system design, reliability issues

must be addressed at various design steps, together with other design objectives. In this talk, I will discuss some approaches to model and mitigate wearout, mostly due to transistor aging, at architecture level and early design stages. By considering reliability together with performance, cost, power objectives, it would be possible to balance them in a cost-effective way.

3.29 Validating Fault Tolerant Designs in SRAM-Based FPGAs: How to Keep the Route in an Ocean of Bits


Massimo Violante (Politecnico di Torino, IT)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Massimo Violante

SRAM-based FPGAs are more and more attractive for applications like avionic subsystems and satellite payloads. However, due to the lack of widely usable fault tolerant SRAM-based devices, it is up to the designer to implement suitable fault tolerant designs. The validation of such designs can be challenging, especially when considering single event upset in the device multi-million bits configuration memory. In this talk, a methodology will be illustrated to help designers to keep the proper route when validating fault tolerant circuits against the soft errors that may affect the ocean of bits in the configuration memory of SRAM-based devices.

3.30 Designing a Dependable Network-on-Chip Platform for Automotive Applications

Tomohiro Yoneda (NII - Tokyo, JP)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Tomohiro Yoneda

Current automotive electronic systems contain many *Electronic Control Units* (ECUs), and their functional safety is an important issue. We have been working to develop a dependable network-on-chip platform for implementing many ECUs on it. This talk will first introduce the designs for dependability in several different levels, such as circuit level, routing algorithm level, and processor core level, adopted in this platform. Then, several issues related to requirements and metrics specified in ISO26262 international standard on functional safety for road vehicles will be discussed, and a trial evaluation of our platform will be shown. Finally, a plan for the functional verification of the platform based on HIL (Hardware In the Loop) simulation will be introduced.

4 Panel Discussions

4.1 Beyond the Limitations of Approximate and Statistical Computing

Chair: Suddhakar M. Reddy

Panelists: J.P. Hayes, M. Orshansky, A. Raghunathan, S. Sapatnekar, A. Singh

Questions discussed during this panel where:

1. Why does approximate computing work?
2. Why does not work?
3. Where will it be applied and will there ever be appropriate devices?
4. Is there a design methodology to control the bound and to make guarantees?
5. Is it scalable?
6. What is the fine print in the trade-off approximation versus efficiency?
7. How is aging handled and predicted?

The main outcome of the panel was on the verification of approximate or probabilistic systems. The complexity increases tremendously if the whole system is considered with all the details. However, we cannot separate layers easily as this is typically done in traditional equivalence checking. Thus a neat methodology for the verification of approximate computing systems or probabilistic systems is required. This methodology must provide mechanisms for abstraction from one level to the next in order to control the complexity of the verification task.

4.2 What's Most Urgent in Industry?

Chair: Shawn Blanton

Panelists: R. Aitken, E. Arbel, C. Gebauer, B. Hoppe, Y. Sato

This panel was conducted in a round-robin fashion with one prime question targeted to each panelist: "What is the biggest issue you see?". Several issues were raised in the panel.

First of all reliability is hard to sell as it does not manifest as an obvious feature to the customer. Customers do not want to spend extra space, i.e. money, to increase reliability in consumer electronics. This has to change in the future. Appropriate metrics may be the key point here.

Once additional cost is accepted, the result typically has to be exact and the data must be coherent on the functional level. Thus, hardware needs to be able to detect errors, i.e., avoiding *Potential Unexpected Loss of Data* (PULD). These features are required within a high reliable system. Then only correct data exchange to the environment needs to be ensured. An orthogonal design issue is mixed-mode operation of reliable computing applications and non-reliable computing applications that run in a single system. Some kind of "address space separation" will be required to guarantee reliability in this case.

On the verification side, verifying reliability still requires a lot of manual work, we need to have more powerful engines which can "reverse engineer" design intent, that is understand automatically how the protection in hardware works and be able to verify that it works correctly. And can we have a "killer-algorithm" which is able to verify all sorts of error detection and correction schemes in a generic way?

4.3 Fault Models, Metrics & Engines

Chair: Jacob A. Abraham,

Panelists: C. Braunstein, Y. Kinoshita, U. Schlichtmann, S. Blanton, M. Tahoori

Questions discussed during this panel where:

1. How do we generate good fault models, especially for determining the “faulty” behavior of an application (a.k.a. How do we get rid of faulty fault models?)
2. What metrics should we use, why, and how are they validated?
3. What engines would allow the calculation of validated reliability metrics for the entire system?

Some fault model free approaches have been proposed to make fault modeling more easy, but these are typically too conservative, i.e., they allow for situations that are virtually impossible in a real system. So mostly for new technologies fault models will still be determined by empirical studies. Engines that can classify a system or determine metrics will depend on the type of system. The co-existence of formal approaches and simulation techniques will continue.

Participants

- Jacob A. Abraham
Univ. of Texas at Austin, US
- Robert Aitken
ARM Inc. - San Jose, US
- Eli Arbel
IBM - Haifa, IL
- Bernd Becker
Universität Freiburg, DE
- Shawn Blanton
Carnegie Mellon University -
Pittsburgh, US
- Cecile Braunstein
UPMC - Paris, FR
- Mehdi Dehbashi
Universität Bremen, DE
- Giuseppe Di Guglielmo
Università degli Studi di Verona,
IT
- Rolf Drechsler
Universität Bremen, DE
- Görschwin Fey
Universität Bremen, DE
- Masahiro Fujita
University of Tokyo, JP
- Carsten Gebauer
Robert Bosch GmbH -
Schwieberdingen, DE
- Jie Han
University of Alberta, CA
- John P. Hayes
University of Michigan, US
- Bodo Hoppe
IBM Deutschland - Böblingen,
DE
- Ravishankar K. Iyer
University of Illinois - Urbana,
US
- Seiji Kajihara
Kyushu Institute of Technology,
JP
- Yoshiki Kinoshita
AIST - Hyogo, JP
- Wenchao Li
University of California -
Berkeley, US
- Igor L. Markov
University of Michigan, US
- Yusuke Matsunaga
Kyushu University, JP
- Subhasish Mitra
Stanford University, US
- Michael Orshansky
Univ. of Texas at Austin, US
- Laurence Pierre
TIMA - Grenoble, FR
- Iliia Polian
Universität Passau, DE
- Anand Raghunathan
Purdue University, US
- Sudhakar M. Reddy
University of Iowa - Iowa City,
US
- Kaushik Roy
Purdue University, US
- Sachin Sapatnekar
University of Minnesota, US
- Yasuo Sato
Kyushu Institute of Technology,
JP
- Matthias Sauer
Universität Freiburg, DE
- Ulf Schlichtmann
TU München, DE
- Marcela Simková
Brno University of Technology,
CZ
- Adit Singh
Auburn University, US
- Matteo Sonza Reorda
Politecnico di Torino, IT
- Mehdi B. Tahoori
KIT - Karlsruhe Institute of
Technology, DE
- Massimo Violante
Politecnico di Torino, IT
- Tomohiro Yoneda
NII - Tokyo, JP



Engineering Multiagent Systems

Edited by

Jürgen Dix¹, Koen V. Hindriks², Brian Logan³, and
Wayne Wobcke⁴

- 1 TU Clausthal, DE, dix@tu-clausthal.de
- 2 TU Delft, NL, k.v.hindriks@tudelft.nl
- 3 University of Nottingham, GB, bsl@cs.nott.ac.uk
- 4 UNSW - Sydney, AU, wobcke@cse.unsw.edu.au

Abstract

This report documents the programme and outcomes of Dagstuhl Seminar 12342 “Engineering Multiagent Systems”. The seminar brought together researchers from both academia and industry to identify the potential for and facilitate convergence towards standards for agent technology. As such it was particularly relevant to industrial research. A key objective of the seminar, moreover, has been to establish a roadmap for engineering multiagent systems. Various research areas have been identified as important topics for a research agenda with a focus on the development of multiagent systems. Among others, these include the integration of agent technology and legacy systems, component-based agent design, standards for tooling, establishing benchmarks for agent technology, and the development of frameworks for coordination and organisation of multiagent systems. This report presents a more detailed discussion of these and other research challenges that were identified. The unique atmosphere of Dagstuhl provided the perfect environment for leading researchers from a wide variety of backgrounds to discuss future directions in programming languages, tools and platforms for multiagent systems, and the roadmap produced by the seminar will have a timely and decisive impact on the future of this whole area of research.

Seminar 19.–24. August, 2012 – <http://www.dagstuhl.de/12342>

1998 ACM Subject Classification I.2.11 Distributed Artificial Intelligence – Multiagent systems, Intelligent agents, D.2 SOFTWARE ENGINEERING

Keywords and phrases Agent-oriented programming, Multiagent systems, Software methodologies for distributed systems, Programming distributed systems, Empirical evaluation

Digital Object Identifier 10.4230/DagRep.2.8.74

Edited in cooperation with Federico Schlesinger

1 Executive Summary

Jürgen Dix

Koen V. Hindriks

Brian Logan

Wayne Wobcke

License  Creative Commons BY-NC-ND 3.0 Unported license
© Jürgen Dix, Koen V. Hindriks, Brian Logan, and Wayne Wobcke

In 1993, Yoav Shoham’s paper on agent-oriented programming was published in the Artificial Intelligence Journal. Shoham’s ideas, and the work on agent-oriented programming it inspired, has had a profound impact on the field of multiagent systems, as evidenced by Shoham’s paper receiving a 2011 IFAAMAS Influential Paper Award recognising seminal work in the



Except where otherwise noted, content of this report is licensed under a Creative Commons BY-NC-ND 3.0 Unported license
Engineering Multiagent Systems, *Dagstuhl Reports*, Vol. 2, Issue 8, pp. 74–98
Editors: Jürgen Dix, Koen V. Hindriks, Brian Logan, and Wayne Wobcke



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

field. Agent-oriented programming offers a natural approach to the development of complex systems in dynamic environments, and technology to support the development of agents and multiagent systems is beginning to play a more important role in today's software development at an industrial level.

Since Shoham's initial work, a range of platforms that support agent orientation have become available, and considerable experience has been gained with these platforms. Some key issues have also emerged from this work, however. First, given the plethora of systems and approaches that have become available in the field for developing multiagent systems, it is no longer clear which of these technologies is most appropriate for developing a particular application or what the distinctive benefits of various approaches are. It is especially important for practitioners to understand the benefits resulting from a particular choice of technology, when and how to apply it, and to develop standards that support the application of agent technology. Secondly, the very different style of agent-oriented programming potentially hampers the uptake of agent development tools and methods. To successfully apply the agent-oriented paradigm and to support the implementation and testing phases of agent-oriented development it is therefore very important to establish best practices and evaluate lessons learned from applying the technology in practice.

The aim of this seminar was to bring together researchers from both academia and industry to identify the potential for and facilitate convergence towards standards for agent technology. The seminar was very relevant for industrial research. The seminar meetings were meant to enable interaction, cross-fertilisation, and mutual feedback among researchers and practitioners from the different, but related areas, and provide the opportunity to discuss diverse views and research findings. The interaction in a Dagstuhl seminar was considered to be ideal for establishing common ground for defining standards, identifying best practices, and developing approaches to applying agent technology to the large scale, realistic scenarios found in industry. The aim of the discussions that were planned was therefore to establish a future research agenda, i.e. a *roadmap*, based on an evaluation of current state-of-the-art of agent-oriented programming languages, tools and techniques that are particularly important for large scale industrial applications.

The seminar took place August 19–24, 2012, with 37 participants from 15 countries. The programme included presentations by the participants and group discussions. Presentations were about 30 minutes long, including questions. We specifically asked participants not to present current research (their next conference paper), but rather asked for what should be considered the next step in their research area.

Participants were encouraged to use their presentations to provide input for discussion about the roadmap. They should show their perspectives and discuss what they think should be on the research agenda, try to explain why, and what it is they think this community should be aiming for. The group discussions took place in the afternoon, after the coffee break until 6pm. We put together four groups of 8-10 members, each headed by one discussion leader (see Section 4 below for more details). The results of each working group were then presented to all participants before dinner. The seminar concluded with a general discussion on Friday morning and a wrap-up.

We identified the following important outcomes of the seminar.

MAS: Understanding of the uptake of multiagent systems technology in industry is seriously hampered by the current situation concerning paper acceptance at scientific conferences and workshops: While new theoretical approaches easily find their way into these events, papers about serious implementations that scale up and put theoretical concepts to work are often considered not innovative enough and are thus not considered appropriate as

scientific papers. We need a forum to publish such papers in order to generate research on the transfer of agent technology to industry.

Merger: During the seminar, eight out of 12 steering committee members of three important workshops in the area of agent systems development (ProMAS, DALT, and AOSE) met to discuss the possibility of merging the workshops. Based on the discussions at the seminar, it was generally agreed that greater focus is needed, and a single venue to present work in the field would be desirable. The workshop steering committees therefore decided (during the seminar) to merge ProMAS, DALT and AOSE to form a new workshop *Engineering Multiagent Systems*. 2012 will therefore be the last year in which the workshops will be held separately: They will be replaced by the new EMAS workshop at next year's AAMAS.

Roadmap: The organisers agreed to start a draft on the roadmap, based on the results of the group sessions. We plan to include the group leaders to produce a first draft, discuss it with the participants and afterwards, to finalise it.

2 Table of Contents

Executive Summary

<i>Jürgen Dix, Koen V. Hindriks, Brian Logan, and Wayne Wobcke</i>	74
--	----

Overview of Talks

Challenges in MAS Verification <i>Natasha Alechina</i>	79
Agents in Unmanned Aerial Vehicle Applications <i>Jeremy Baxter</i>	79
Reflections on Multiagent Oriented Programming <i>Rafael H. Bordini</i>	79
Building Multiagent Systems for the Real World: A Company's Perspective <i>Paolo Busetta</i>	80
Experiences with Agent Factory <i>Rem Collier</i>	80
Handling High Frequency Perception / Agents and Enterprise Computing <i>Stephen Cranefield</i>	80
Engineering Multiagent Systems: Lessons and Challenges <i>Mehdi Dastani</i>	81
Timeliness Issues in Agent Based Control of Satellites, Among Other Things <i>Louise Dennis</i>	81
What We Talk About When We Talk About Agents <i>Virginia Dignum</i>	82
Agent Technology integration with Maven for an Ambient Assisted Living Case Study <i>Jorge J. Gomez-Sanz</i>	82
Perspectives and Roadmap for Engineering Multiagent Systems <i>Christian Guttman</i>	83
Multiagent Oriented Programming with JaCaMo <i>Jomi Hübner</i>	83
Lessons and Perspectives on Agent Languages <i>Yves Lespérance</i>	84
Programming Agents <i>Brian Logan</i>	84
On Engineering Emotional Agent Systems <i>John-Jules Ch. Meyer</i>	84
Observations from Current and Past Projects: 1. Shaping the Intelligent Home of the Future, 2. Settlers of Catan <i>Berndt Müller</i>	84
Application Impact of Multiagent Systems and Technologies <i>Jörg P. Müller</i>	85

Exploring Agents as a Mainstream Programming Paradigm: The simpAL Project <i>Alessandro Ricci</i>	85
MAS for Engineering Complex Systems <i>Amal El Fallah Seghrouchni</i>	86
Agents in Space for Real: Lessons Learned from Applying Agent Technology in NASAs Mission Control <i>Maarten Sierhuis</i>	86
Empirical Software Engineering for Agent Programming <i>Birna van Riemsdijk</i>	87
Engineering Multiagent Systems - Reflections <i>Jørgen Villadsen</i>	87
Challenges and Directions for Engineering Multiagent Systems <i>Michael Winikoff</i>	88
Decoupling in Industry <i>Cees Witteveen</i>	88
Engineering Multiagent Systems: Where is the Pain (and the Opportunity)? <i>Wayne Wobcke</i>	89
Working Groups	
Integration and Validation	89
Coordination and Organisation	92
Tools, Languages and Technologies	93
Component-Based Agent Design	94
Open Problems	96
Panel Discussions	96
Participants	98

3 Overview of Talks

3.1 Challenges in MAS Verification

Natasha Alechina (University of Nottingham, GB)

License © © ⊖ Creative Commons BY-NC-ND 3.0 Unported license
© Natasha Alechina

Joint work of Alechina, Natasha; Logan, Brian; Nguyen, Nga; Rakib, Abdur; Doan, Trang; Dastani, Mehdi; Meyer, John-Jules

I give a brief overview of the state of the art in verification of multiagent systems where agents are implemented in BDI agent programming languages, and list the challenges. The main challenges are: - Ability to represent MAS at a suitably high level of abstraction- Ability to formulate properties in a suitable language- Scalability improvements

3.2 Agents in Unmanned Aerial Vehicle Applications

Jeremy Baxter (QinetiQ - Malvern, GB)

License © © ⊖ Creative Commons BY-NC-ND 3.0 Unported license
© Jeremy Baxter

I discuss my background in multiagent systems and my experience with using multiagent toolkits. I have used agents to control small teams of unmanned air vehicles both in simulation and in test flight. The main focus of the work has been using agents to co-ordinate multiple vehicles and to integrate different types of planning and reasoning. When developing systems only a small part of the effort is a core agent system, the majority is interfaces. Testing is a major element of the development and is not well supported by current agent tools. There is a steep learning curve with new tools and languages which can be hard to justify in a project. I conclude that design patterns and libraries of components might gain better acceptance than complete new languages and development environments.

3.3 Reflections on Multiagent Oriented Programming

Rafael H. Bordini (PUCRS - Porto Alegre, BR)


License © © ⊖ Creative Commons BY-NC-ND 3.0 Unported license
© Rafael H. Bordini

URL <http://dx.doi.org/10.1016/j.scico.2011.10.004>

In this talk I discussed some reflections on multiagent oriented programming based on my own experiences, and in particular recent experiences with the JaCaMo platform, in joint work with Jomi Hübner, Olivier Boissier, Alessandro Ricci, and Andrea Santi. I made the point that after years of research we have not yet been able to define precisely what multiagent orientation entails as a programming paradigm. I also argued that achieving such shared view of a paradigm is essential if our work is to reach out to other research communities within computer science.

3.4 Building Multiagent Systems for the Real World: A Company's Perspective


Paolo Busetta (AOS Ltd. - Cambridge, GB)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Paolo Busetta

AOS is one of the few companies on the market whose business is focused on multiagent platforms and applications. AOS' main product, JACK, was originally released in 1997. Since then, AOS has been involved in a large number of diverse research and applicative projects, varying from cognitive simulation in virtual reality to safety-critical embedded systems. In this talk, I will present a few important technical challenges that AOS has faced in its 15 years of existence. These experiences have contributed to shape the new agent platform under development, called C-BDI. I will briefly introduce some of its novelties and how they are meant to address the needs of its expected main domains of application, in particular autonomous operational systems and serious games.

3.5 Experiences with Agent Factory

Rem Collier (University College Dublin, IE)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Rem Collier

Joint work of Collier, Rem; O'Hare, Gregory

This talk is broken into two parts. The first part presents an overview of the history of Agent Factory; a cohesive framework for the development and deployment of multiagent systems that has been under development and in constant use since 1996. It briefly reflects on the design choices made for each version of the framework and the improvements made. Where relevant a selection of applications built using the specific version of the framework are described. The second part reflects on some experiences gained from the use of Agent Factory both in terms of the development of demonstrators and in terms of its use as a teaching platform. Specific comments made in part 2 include: the lack of online community resources that promote the field; the challenge of meeting users expectations in terms of tool support; and the lack of significant work on evaluation of agent programming languages / comparison of agent-based solutions with non-agent based solutions.

3.6 Handling High Frequency Perception / Agents and Enterprise Computing

Stephen Crane field (University of Otago, NZ)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Stephen Crane field

Joint work of Crane field, Stephen; Ranathunga, Surangika; Purvis, Martin

Main reference Surangika Ranathunga, Stephen Crane field and Martin Purvis. Identifying Events Taking Place in Second Life Virtual Environments. Applied Artificial Intelligence, (26)1-2:137-181, 2012

URL <http://dx.doi.org/10.1080/08839514.2012.629559>

In the first part of this talk I briefly discussed some work at the University of Otago on connecting agents with virtual worlds such as Second Life, and used this to motivate a proposal

for enhancing agent platforms with support for handling high frequency state changes. In the second part, I considered the role that agents might play in enterprise computing and how they might be integrated into enterprise applications and business processes. I argued that agents can play a useful role as components of larger businesses processes, and that agent development tools should provide an interface between agents and the existing integration technology used in enterprise computing. In particular, I proposed that a simple interface between agents and enterprise computing infrastructure can be provided by defining agent “endpoints” for enterprise message routing and mediation engines such as Apache Camel. These configurable endpoints would translate (selectively) between internal agent entities such as beliefs and ACL messages and the message exchange abstraction used in the enterprise integration patterns (EIPs) of Hohpe and Woolf [1]. Message routing and mediation rules could then be defined outside the agent platform to interconnect the agents with any other protocols and services that have endpoints defined (such as the 130+ that are available for Apache Camel).

References

- 1 G. Hohpe and B. Woolf. *Enterprise Integration Patterns: Designing, Building, and Deploying Messaging Solutions*. Addison-Wesley, 2004.

3.7 Engineering Multiagent Systems: Lessons and Challenges



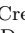
Mehdi Dastani (Utrecht University, NL)

License     Creative Commons BY-NC-ND 3.0 Unported license
© Mehdi Dastani

In my presentation I explained the aims of multiagent programming research field as formulated in this community and gave a brief presentation of the activities within this community in the last decade. A distinction is made between academic and industry perspectives. I argued that although both perspectives are valuable and challenging, their activities and aims are different. For each perspective I presented some challenges and future directions for research. I ended the presentation by emphasising the role of transfer of knowledge from the academic perspective to the industry perspective.

3.8 Timeliness Issues in Agent Based Control of Satellites, Among Other Things

Louise Dennis (University of Liverpool, GB)

License     Creative Commons BY-NC-ND 3.0 Unported license
© Louise Dennis


Joint work of Fisher, Michael; Veres, Sandor; Lincoln, Nicholas; Lisitsa, Alexei; Gao, Yang; Bordini, Rafael; Muller, Berndt

The talk consisted of three parts: 1) The Engineering Autonomous Space Software project investigated the integration of real-time control systems with a rational agent layer for decision making. The focus of the project was on the abstraction of continuous data to discrete data. The implementation ran into a number of issues related to the speed with which data or commands generated by one part of the system could be processed by another part of the system. Since this did not form the core focus of the project these issues were

worked around in an ad hoc fashion. This talk provided an overview of the problems, the “quick fixes” and throw out a couple of ideas for how the problems might be dealt with more coherently. 2) An overview of the Agent Infrastructure Layer, a Java-based toolkit for implementing the operational semantics of BDI agent systems and then model checking programs written in the systems. In particular I considered the question of whether the Agent Infrastructure Layer constituted a Virtual Machine for BDI agent languages. 3) An overview of the aims of the newly awarded Reconfigurable Autonomy project.

3.9 What We Talk About When We Talk About Agents

Virginia Dignum (TU Delft, NL)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Virginia Dignum

Joint work of Dignum, Virginia; Frank Dignum


Main reference Virginia Dignum, Frank Dignum: Designing agent systems: state of the practice; International Journal of Agent-oriented Software Engineering - IJAOSE , vol. 4, no. 3, pp. 224-243, 2010

URL <http://dx.doi.org/10.1504/IJAOSE.2010.036983>

In this presentation, I discussed different views on agent technology and its applications. Guidelines to decide on agent approaches and its consequences for (agent-oriented) software engineering lifecycle. I furthermore introduced a few extra issues to be included in the roadmap: Scaling / multi-level models; Evolution / re-design; and the role of people in the loop.

3.10 Agent Technology integration with Maven for an Ambient Assisted Living Case Study

Jorge J. Gomez-Sanz (Univ. Comp. de Madrid, ES)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Jorge J. Gomez-Sanz


URL <http://social.sf.net>, <http://ingenias.sf.net>

The talk introduces some recent advances in the INGENIAS Development Kit to deal with challenges found in an Ambient Assisted Living (AAL) oriented project and how the Maven project management tool contributed to this goal. Ambient Assisted Living systems tries to make the life of people easier by assisting them in different ways. Most of them have to do with a smart use of sensors and actuators situated all of them in the environment of the user. Literature in AAL tells a natural candidate for become the main building block in this kind of proposal is agent technology. The name of project mentioned in this talk is SociAAL, because it tries to focus on social aspects that influence this kind of systems. The project is inherently challenging because of the mixture of different technologies for which standard agent oriented development environments are not the best choice. Current tools do not strongly support integration and require installing different plugins that do not ensure two technologies can work together. A candidate solution is the Maven framework. Maven is a tool created by the Apache Foundation. Quoting them “Apache Maven is a software project management and comprehension tool”. In SociAAL project, Maven integration has meant a possibility of putting together in a straight forward way the development of OSGi objects, XML documents, XML transformations, INGENIAS agents, and other artefacts.

The introduction and justification of this framework has served to explore stages of the development which are not covered usually by AOSE methodologies. As a result, the talk introduces how well is AOSE dealing with a complete software development using as driver the software lifecycle according to the standard IEEE Glossary of Software Engineering. The conclusion of the talk is that many work is needed in AOSE to understand the role of agent technology in a long term development. In this endeavour, frameworks like Maven can help, since they are widely used mainstream software engineering tools and can be trusted to identify meaningful aspects of a development. By integrating with Maven, AOSE will have to tell what “compiling”, “generating sources”, “documenting”, “testing”, or “packaging” the multiagent system means. This will introduce better our technology to people used to work with software.

3.11 Perspectives and Roadmap for Engineering Multiagent Systems


Christian Guttman (IBM R&D Labs, Melbourne, AU)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Christian Guttman

On our agent roadmap, we have not advanced as far as we ought to. Academic and industrial agent projects still lack consistent and unified design and engineering patterns, and advantages of agent engineering over other engineering approaches are not entirely clear. Hence, it is difficult to evaluate the benefit and potential of agents as an approach and methodology for ambitious projects, and hence it is difficult to make a well informed choice of using agents. I will support this statement by revisiting how far we have come on existing agent roadmaps, and also by reporting on my recent experience on defining and leading R&D projects that extend and use agent technologies in the area of health and medicine. A few ideas are offered to extend the agent roadmaps. Our community may benefit from engaging more in the technology transfer process (showing the value of agent engineering), and engaging more with other research communities and stakeholders, where the key is to identify and define challenges together, rather than in isolated labs and research groups.

3.12 Multiagent Oriented Programming with JaCaMo


Jomi Hübner (Federal University of Santa Catarina - Brazil, BR)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Jomi Hübner

This talk brings together agent-oriented programming, organisation-oriented programming and environment-oriented programming, all of which are programming paradigms that emerged out of research in the area of multiagent systems. In putting together a programming model and concrete platform called JaCaMo which integrates important results and technologies in all those research directions, we show in this paper that with the combined paradigm, that we prefer to call “multiagent oriented programming”, the full potential of multiagent systems as a programming paradigm. JaCaMo builds upon three existing platforms: Jason for programming autonomous agents, Moise for programming agent organisations, and CArTAgO for programming shared environments.

3.13 Lessons and Perspectives on Agent Languages


Yves Lespérance (York University - Toronto, CA)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Yves Lespérance

In the talk I briefly review the main features of the Golog family of Situation Calculus-based agent programming languages and application where they have been used and put out some ideas for future research on Engineering Multiagent Systems. One topic raised for future work is modeling and reasoning about the mental states of other agents (Theory of Mind) in agent programming languages, and I briefly discuss some initial work in that area.

3.14 Programming Agents

Brian Logan (University of Nottingham, GB)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Brian Logan

Joint work of Logan, Brian; Alechina, Natasha; Bordini, Rafael; Dastani, Mehdi; Gordon, Elizabeth; Hindriks, Koen; Madden, Neil; Meyer, John-Jules; Sloman, Aaron; Vikhorev, Konstantin

In this talk, I consider agent programming from the perspective of Artificial Intelligence. I briefly outline some lessons learned from our work on developing approaches to tractable deliberation for intention scheduling in the agent programming languages ARTS, AgentSpeak(RT) and N-2APL, and highlight some unsolved problems in deliberation about deadlines and plan durations. I also sketch some future directions for agent programming.

3.15 On Engineering Emotional Agent Systems

John-Jules Ch. Meyer (Utrecht University, NL)

License  Creative Commons BY-NC-ND 3.0 Unported license
© John-Jules Ch. Meyer

In this talk I go through the following: Why emotional agent systems? The main idea. Methodology. How far we have got. Intuition of 4 basic types of emotion. Deliberation with emotions. Future work.

3.16 Observations from Current and Past Projects: 1. Shaping the Intelligent Home of the Future, 2. Settlers of Catan

Berndt Müller (University of Glamorgan, GB)

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

We discuss our experience from projects using agent-based techniques and point out some research topics that will have to be addressed if multiagent based systems become ubiquitous. These include legal and ethical issues, security, and verification. The latter needs to be more rigorous (generating reasoning engines from semantic specifications) and ideally needs to take

notions of resource and location into account. Of further importance to the acceptance of MAS as a programming paradigm, is the availability of a component-based approach and the availability of agent libraries. This is illustrated by an example of agent-based development of a turn-based game using high-level Petri nets based on the nets-within-nets paradigm.

3.17 Application Impact of Multiagent Systems and Technologies

Jörg P. Müller (TU Clausthal, DE)

License © ⓘ ⊖ Creative Commons BY-NC-ND 3.0 Unported license
© Jörg P. Müller

There appears to be a common perception among Multiagent Systems (MAS) researchers that their research field has still some room left in creating impact outside our own research community. However, there are no recent studies that allow us to back up or rebut this hypothesis. In this talk I am providing some thoughts and observations related to the application impact of MAS. I review some models of ICT impact known from the literature and discuss their applicability to MAS research. Further, I discuss previous work related to research on application impact in our research community. I come to the conclusion that it is beneficial to include research activities related to the study of application impact into a research roadmap on multiagent systems and technologies. I propose some desiderata for such research, and inform about an ongoing survey activity.

3.18 Exploring Agents as a Mainstream Programming Paradigm: The simpAL Project

Alessandro Ricci (University of Bologna, IT)

License © ⓘ ⊖ Creative Commons BY-NC-ND 3.0 Unported license
© Alessandro Ricci


Main reference Alessandro Ricci, Andrea Santi. “Designing a general-purpose programming language based on agent-oriented abstractions: the simpAL project”. Proceeding of SPLASH ’11 Workshops Proceedings of the compilation of the co-located workshops on DSM’11, TMC’11, AGERE!’11, AOOPEs’11, NEAT’11, VMIL’11.

URL <http://dx.doi.org/10.1145/2095050.2095078>

Agent-Oriented Programming has been explored so far mainly in the context of Distributed AI and Multiagent Systems, and it is almost totally unknown in the context of programming languages and software engineering. In spite of that, we argue that agent-oriented concepts and abstractions could be effective to tackle main problems that affect modern programming, beyond object-oriented programming and actor-based programming. Accordingly, our medium-term objective is to shape a new programming paradigm based on agent-oriented abstractions, as a natural evolution of the object and actor ones. This calls for devising programming languages – as well as related models and technologies – that, besides being based on agent-oriented abstractions, would provide features and mechanisms that are important when programming and software development is of concerns. In this presentation we discuss our progress in that direction, represented by the simpAL programming language and platform.

3.19 MAS for Engineering Complex Systems

Amal El Fallah Seghrouchni (UPMC - Paris, FR)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Amal El Fallah Seghrouchni


This talk will present my main experiences of MAS engineering within industrial context. It aims to show where and how MAS may bring an added-value for complex systems design.

The first part of my talk outlines some examples of simulations of complex systems and also some prospective projects we have developed in the aerospace domain. Two main projects are described: 1) SCALA is a project for mission interception based on reactive multiagent planning (collaboration with Dassault-Aviation) and 2) the coordination of fleet of UAVs where several aspects of MAS are involved such as planning, elicitation of preferences and multiagent decision (collaboration with Thales Airborne Systems).

The second part of my talk goes on to relate some lessons learnt from the development of two languages for MAS programming, namely CLAIM and its extension S-CLAIM (Smart Claim) to deploy MAS on smart devices. S-CLAIM is a declarative agent-oriented language for Ambient Intelligence (AmI) - S-CLAIM - that allows programming reactive or cognitive mobile agents in a simple, easy-to-use manner while meeting AmI requirements. Based on a hierarchical representation of the agents, the language offers a natural solution to achieve context-sensitivity. S-CLAIM is an evolution of the CLAIM language, its predecessor. It is light-weight and, being transparently underpinned by the JADE framework, allows deployment on mobile devices and easy interoperation with other components by means of web services. The usefulness of the proposed language for AmI is illustrated through a scenario and a demo featuring an AmI application in a Smart Room (see also the attached paper presented at ANT'2012). My talk concludes with a positive note concerning the transfer in the field of MAS from academia to industry.

3.20 Agents in Space for Real: Lessons Learned from Applying Agent Technology in NASA's Mission Control


Maarten Sierhuis (Ejenta Inc., US)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Maarten Sierhuis

This talk provides lessons learned from developing and implementing the first multiagent workflow system that automates the work of the OCA flight controller in NASA's Mission Control Centre for the International Space Station. OCAMS was first simulated and then developed and deployed using the Brahms agent language and NASA's Brahms environment. Ejenta, Inc. is a startup in San Francisco, CA and has as its mission to develop intelligent personal agent technology. Ejenta provides a commercial version of the Brahms agent simulation and development environment and its associated NASA applications and technology, including the OCAMS multiagent procedure execution workflow environment and the Individual Mobile Agent System. For more information, please contact the author.

3.21 Empirical Software Engineering for Agent Programming

Birna van Riemsdijk (TU Delft, NL)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Birna van Riemsdijk


Joint work of van Riemsdijk, M. Birna; Hindriks, Koen; Jonker, Catholijn M.

Main reference M. B. van Riemsdijk, K. V. Hindriks, and C. M. Jonker. An empirical study of cognitive agent programs. *Multiagent and Grid Systems (MAGS)*, 8(2):187-222, 2012.

In this talk I argue for increasing use of empirical software engineering in the development of agent programming languages and techniques. Empirical software engineering is a branch of computer science in which empirical methods are used to study how people use the technologies and to what extent certain techniques are better than others. We need to investigate how software quality characteristics as identified in mainstream software engineering apply in the context of engineering multiagent systems, and define dedicated attributes and metrics to measure to what extent these are present in the software product. In this way we can improve the techniques based on data. Also we need to develop or come to agreement concerning what 'counts' as good empirical research for engineering multiagent systems.

3.22 Engineering Multiagent Systems - Reflections

Jørgen Villadsen (Technical University of Denmark - Lyngby, DK)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Jørgen Villadsen

Joint work of Villadsen, Jørgen; Jakobsen, Troels Christian


Main reference Troels Christian Jakobsen: "I wouldn't have thought of it myself" - Emergence and unexpected intelligence in theater performances designed as self-organising critical systems. In proceedings: Algolog Multiagent Programming Seminar 2011 (AMAPS2011) - Technical University of Denmark - Lyngby

URL <http://www.imm.dtu.dk/algolog/index.php?n=Home.AMAPS>

In the first part I look at a theater performance by artistic director Troels Christian Jakobsen as a multiagent system. It is designed as a self-organising critical system using a framework where within its borders but without a script there is real interaction between the elements of the performance. In the second part I discuss the ideas behind my recent monograph on propositional attitudes and inconsistency tolerance. Natural language sentences are parsed using a categorial grammar and correctness of arguments are decided using a paraconsistent logic. In the third part I present a curriculum for the MSc in Computer Science and Engineering program at the Technical University of Denmark with a focus on multiagent systems. As the director of studies I have observed that the students are working hard and with much creativity in advanced courses and projects involving intelligent agents, in particular in the agent contest 2009-2012.

3.23 Challenges and Directions for Engineering Multiagent Systems

Michael Winikoff (University of Otago, NZ)

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

Main reference “Future directions for agent-based software engineering”, *Int. J. Agent-Oriented Software Engineering*, Vol. 3, No. 4, pp. 402-410.

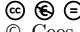
URL <http://dx.doi.org/10.1504/IJAOSE.2009.025319>

URL <http://arxiv.org/pdf/1209.1428.pdf>

In this talk I review where we stand regarding the engineering of multiagent systems. There is both good news and bad news. The good news is that over the past decade we’ve made considerable progress on techniques for engineering multiagent systems: we have good, usable methodologies, and mature tools. Furthermore, we’ve seen a wide range of demonstrated applications, and have even begun to quantify the advantages of agent technology. However, industry involvement in AAMAS appears to be declining (as measured by industry sponsorship of the conference), and industry affiliated attendants at AAMAS 2012 were few (1-2%). Furthermore, looking at the applications of agents being reported at recent AAMAS, usage of Agent Oriented Software Engineering (AOSE) and of Agent Oriented Programming Languages (AOPL) is quite limited, which is also supported by the results of a 2008 survey by Frank and Virginia Dignum (“Designing agent systems: state of the practice”, *IJAOSE* 2010, 4(3):224-243). Based on these observations, I make five recommendations: 1. Re-engage with industry 2. Stop designing AOPLs and AOSE methodologies ... and instead ... 3. Move to the “macro” level: develop techniques for designing and implementing interaction, integrate micro (single cognitive agent) and macro (MAS) design and implementation 4. Develop techniques for the Assurance of MAS 5. Re-engage with the US.

3.24 Decoupling in Industry


Cees Witteveen (TU Delft, NL)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Cees Witteveen

I discuss an application of agent technology in maintenance. The industrial partner is part of the Dutch Railway company responsible for maintenance. They are interested in flexible schedules and distribution of a global operational scheduling problem over several teams. These teams should be able to schedule their activities independently. We apply some ideas derived from temporal decoupling, but also from classical OR to solve their problems. The main lessons learned are: (1) use the language and concepts of your partner, (2) make very concrete promises and fulfil them in a verifiable way; (3) do not hesitate to consider agents only as a useful metaphor.

3.25 Engineering Multiagent Systems: Where is the Pain (and the Opportunity)?

Wayne Wobcke (UNSW - Sydney, AU)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Wayne Wobcke

I first describe two multiagent systems, one a “smart personal assistant” providing spoken dialogue interaction with a collection of personal agents in e-mail and calendar management (built using JACK), and the second an agent-based model for risk assessment of routine clinical processes, estimating the risk associated with patient misidentification and infection control (built using Brahms). I then reflect on the main difficulties in engineering these systems and the opportunities presented for further research. In summary, the main problems are not with particular programming languages or platforms, but of two types: (i) integration, where the agent aspect of the system is a small part of a much larger system, and (ii) validation of an agent model against reality. I conclude with a proposal of developing tools for semi-automatically constructing agent models using a mixture of knowledge acquisition and machine learning/data mining techniques, validating against traces of existing system behaviour, with particular application to the medical domain. This approach has recently become feasible due to the availability of “big” data sets.

4 Working Groups

There were three group-discussion sessions. The organisers separated participants into four groups (of size 8–10), which varied for the different sessions. The outcomes of the discussions in each group were presented to the other groups at the end of each session and a general discussion followed.

For the first session, we did not specify a particular topic. As this session took place on the very first day of the seminar, we just collected ideas and identified important topics.

For the second session, the four groups were assigned different topics, focussing on particular research areas: *Integration and Validation*, *Coordination and Organisation*, *Tools, Languages and Technologies* and *Component-Based Agent Design*.

The third session was again directed towards the structure of the roadmap. The motto for this session was:

What do you want to do in the next 10 years (research, applications)? Pick the top 10 topics out of a list or add new ones. Put them into clusters or state how they relate to each other.

4.1 Integration and Validation

One objective of this Dagstuhl seminar was to bring together academic researchers and industry practitioners working on a variety of applications, to compare experiences, identify common problems in deploying agent technology, and propose ways to alleviate these problems in the future. A premise of the seminar was that there is currently a large variety of deployed agent applications, but that this was not widely recognised by the agents research community. The breadth of existing work was confirmed with presentations on: (i) avionics and defence

(unmanned aerial vehicles), (ii) smart cities (ambient intelligence, crowd simulation), (iii) transportation (traffic modelling and simulation), (iv) logistics (warehouse management, maintenance scheduling), (v) workflow management, (vi) decision support (process monitoring, disaster management), (vii) healthcare (agent-based modelling, care coordination, patient monitoring), (viii) personal assistants (dialogue management), (ix) real-time control systems (mission planning), (x) power engineering (grid management), (xi) information integration (sensor networks), and (xii) virtual environments (games, training).

Most participants in this working group agreed on two major distinctive benefits of agent technology: **autonomous decision making** and **explicit problem decomposition and coordination mechanisms**. There was also clear consensus that the main problems facing deployment of agent/multiagent systems fall under three related areas: **integration**, **validation** and **software engineering**.

4.1.1 Integration

A basic problem is that an agent/multiagent system is generally only a small part of a much larger system containing any number of other complex hardware/software components. While traditionally a system can be organised “hierarchically” as a collection of agents (providing interfaces to the user and wrappers for other components), this is not generally the case for many applications. More commonly, agents need to interact with other non-agent components without the use of agent communication protocols, and need to be provided with information about the rest of the system’s behaviour in order to perform their function. This means that the agent part of the system cannot easily be isolated from the remainder of the system, both in software development and for the purpose of reasoning about agent and system behaviour.

Integration can take place at multiple layers of abstraction, e.g. the implementation level, conceptual level, business level, etc. Furthermore, the design of the non-agent part of the system is typically outside the agent developer’s control, so the agent programming language/platform needs to be highly flexible in allowing integration with interchangeable components (that may be at different abstraction layers), facilitate customisation of the agent/multiagent system to different application scenarios as needed, and support system maintenance of the whole system as it evolves. It is also desirable (for simplicity and efficiency) to be able to select only a subset of features of an agent platform needed for a given application rather than being required to use all features of a large and complicated agent platform.

4.1.2 Validation

The major potential benefits of agent technology arising from autonomous decision making present, paradoxically, a significant barrier to the adoption of the technology, since users typically require performance guarantees (preferably quantifiable and verifiable) that agent actions, though understood to be not completely predictable, are within *acceptable* bounds. Thus *validation* or *assurance* of agent/multiagent systems is particularly important. The level of system assurance required varies with the application area, but is most stringent in defence, where formal certification is required. Other sorts of applications need to be “trusted” by users, regulators and the community.

The type of properties that may need validation include safety, security, scalability, quality, maintainability, performance and interoperability. Validation is extremely difficult due to the complexity of agent behaviour and interactions, particularly as validation needs to be

not only of the agent part of the system but of the system as a whole (and as noted above, agent/multiagent systems are tightly integrated into larger systems of agent and non-agent components). Validation of the agent part of a system is difficult because this part of the system cannot (in general) be isolated from the rest of the system, and because some desired properties may need to be derived from those of the agent development platform, which in turn may need to be validated/certified. In the specific area of agent-based modelling, validation is often ignored and datasets are often insufficient to provide adequate rigorous validation of models. Verification may sometimes be possible, but this is the exception rather than the rule with today's complex applications, considering the limitations of current approaches for specification and verification.

4.1.3 Software Engineering

It was commented by one participant that agent-based software development is 10% multiagent systems engineering and 90% standard software engineering. Whatever the breakdown, a consequence of the need for integration of agents into larger systems is that standard software engineering is heavily involved in the deployment of agent/multiagent systems. To improve the ease of adoption of agent/multiagent systems, what is needed are *not* more special-purpose “agent-oriented software engineering” methodologies, which often emphasise the distinctive nature of agent-based systems or which are closely tied to particular agent languages or platforms. Similarly, special-purpose agent programming languages present a barrier to deployment if they do not provide support for integration with existing software, creating undesirable “lock-in” to particular platforms and/or duplication of effort when non-agent components need to be translated into a particular agent model/language to enable interaction with agents.

Instead, what is required is better incorporation of agent-oriented software development *within* mainstream software engineering practices, and conversely, the use of more standard software engineering methodologies and tools within agent-oriented software development. This involves: (i) providing support for agent development within the whole software development lifecycle, from requirements engineering and architectural design through to testing and maintenance, (ii) integration with mainstream software development environments and especially tools, (iii) adoption of widely used software engineering approaches such as design patterns and pattern languages, (iv) compliance with software engineering standards, (v) “reaching out” to the software engineering research community through publication in software engineering venues, and (vi) comparison of agent-based software development platforms with standard programming language environments. The overall objective is to make it easier to deploy and maintain agent/multiagent systems within mainstream applications.

The term *component* is used above loosely to refer to some part of a larger hardware/software system. The topic of a specific “component-based” software engineering paradigm for agents was a subtheme of this working group discussion, but meant a number of things: (i) treating agents as interchangeable components in a larger system (“plug-and-play” agents or agent components, perhaps taken from a component library or repository), (ii) a declarative platform-independent representation for agents to enable reuse of agents from one system to another or to make it easier to construct agent models using third party tools, and (iii) building single agents out of simpler interchangeable components (e.g. belief database, reasoning engine, etc.). Despite the unresolved ambiguity, the idea of component-based agent software engineering was felt worthy of much further research.

4.2 Coordination and Organisation

The notions of *interaction* and *organisation* are important in Multiagent Systems (MAS), but they are important in other systems as well. However, the coordination and organisation have never been studied in a homogeneous fashion. We need answers to why there are such big differences. For example it should be explained why elements which are a concern in other systems are not a concern in MAS. As an example, for a newcomer with some knowledge in distributed systems, it is surprising that classical problems in concurrent systems, like deadlock or starvation, are hardly mentioned today. While we cannot define the problems with deadlock and starvation away, we claim that it is the agent paradigm, with its levels of abstraction and decentralised solutions, that helps to define these notions appropriately. We feel agent technology helps to better understand the problem, provides tools to deal with them and, in the end, verification of these and other properties is possible. One community dedicated to these topics is the COIN (COIN: Coordination, Organisation, Institutions and Norms). A big part of COIN is more concerned with abstractions and much less with implementations. There are important connections to the multiagent oriented paradigm that are not yet fully explored. Different technologies to enable coordination are not able to represent concepts needed in MAOP (Multi-Agent Oriented Programming) beyond messages. We believe that a semantical underpinning (as opposed to classical tuple spaces without any messages) makes message interchange easily possible and also helps to implement it. In a solution for *interaction* and *organisation* applied to a multiagent system, we would look for an *organisation model* subsuming both aspects. In an organisational model, we need to determine (1) the elements required to define a coordination, and (2) what goals should be pursued. Agents then acquire or are given these goals and commit to them in ways compliant with the organisation model.

Ideally, we are looking for an organisation specification language that could be translated into explicit organisations at the *execution level*. We can look for optimisations at two levels: (1) looking for first-class-citizen representation of concepts belonging to the organisation model; or (2) focusing on protocols/algorithms that implement the coordination/organisation which have certain properties (e.g., being deadlock-free) we might even verify. Also, we assume such MAS can change its coordination behavior (the proper (local or global) strategy) each time. In any case, the organisational approach makes explicit the strategy of the system in those cases. These observations are also important for our planned roadmap. Coordination in industry is often solved with dirty hacks. There is no general methodology. Techniques and concepts are needed. We feel we really need an agreement about the kind of high level concepts that define the coordination. In academia we often program just single agents instead of defining the MAS from the very beginning. Once the language is chosen, the coordination problem needs to be worked out and different coordination solutions can be compared. We need some kind of *coordination engineering*.

To sum up, we considered the following tasks to be particularly important: (1) to develop coordination mechanisms for large distributed open systems, (2) to take *runtime organisation* seriously, (3) to develop platforms that incorporate coordination/organisation support, (4) to make an organisation live as a distributed system (this is not just the design of the organisation), (5) to develop both top-down as well as bottom-up methods (from agents to organisations and back).

4.3 Tools, Languages and Technologies

Tooling and programming languages for multiagent systems are very important for the uptake of the multiagent programming paradigm. In particular, the need for more sophisticated approaches and tools for testing and debugging was clear to the participants of the seminar. Multiagent systems pose many new challenges in this area. First of all, the behavior of agents is often dynamic and may change over time. Multiagent systems are also typically distributed systems which introduces additional challenges. Multiagent systems, moreover, are used to control complex, dynamic and non-deterministic environments. Such environments do not support, for example, easy replay of one and the same test case. A key challenge therefore is to establish an approach for testing such complex systems. In order to manage this complexity tests are needed at different levels of a multiagent system similar to unit and integration tests in more traditional object-oriented approaches. There is a need to identify the equivalents of these tests for agent technology. Different techniques may need to be introduced such as mock agents for protocol testing. For example, it was felt that a language for expressing test cases may advance the state of the art significantly. Such a language would need to provide support for defining the state of the agents' environment and for the state of the agents themselves.

A need for test beds that are made widely available for collecting data on and for comparing various platforms was also identified. The variety of platforms available for engineering multiagent systems for developers raises the issue of which to choose. Various benchmarks should be developed to identify the benefits and weaknesses of platforms. As a community, we should agree on a list of common benchmarks that relate to specific aspects of a multiagent system. Relevant aspects that are specific for multiagent systems include, for example, components of agents such as percept processing, belief revision, and intention reconsideration, as well as more general aspects such as reactivity and scalability. An important issue is how to ensure that similar things are measured in different platforms. One solution would be to use standardised interfaces for e.g. connecting to environments. It was suggested to create and use a web portal to publish and discuss benchmarks to make progress in this area.

A related but different topic concerns the usability of different agent platforms. The learning curve associated with one platform may differ greatly from that of another, while usability from the perspective of the capabilities offered by the latter platform may be rated higher by expert users than that of the former platform. Methodologies are needed to be able to perform systematic studies into usability aspects. Here both qualitative as well as quantitative techniques will need to be used, in particular at this stage of the research where only few studies are available that look at usability issues. As in the case for benchmarks, a range of different tasks will need to be designed that can be used as test cases in usability studies. Moreover, the skill and experience level of users will need to be taken into account. In order to be able to study differences between platforms we also need to develop techniques for comparing different solutions programmed in different agent programming languages (e.g. a simple comparison of lines of codes will not do). Usability studies may also be used to enhance teaching and improve courses on how to engineer multiagent systems. This could even lead to an increase of the number of universities that adopt agent technology and programming languages in courses related to engineering intelligent and multiagent systems.

One of the main contributions of the multiagent programming paradigm is to introduce a new set of abstractions for programming software systems. Apart from the notion of an agent, the focus of the multiagent paradigm on concurrent and event-driven programming may provide the proper level of abstraction for programming distributed systems by abstracting

away low-level concerns related to, for example, threads. If research in our community would focus more on these aspects this could possibly lead to (re)connecting multiagent programming to mainstream computing science research on related issues. Also, at the conceptual level, multiagent systems raise important new issues such as how to incorporate norms and program the organisational structure of a multiagent system. Finally, various more technical challenges need to be faced relating to the scalability of multiagent systems. The management of huge number of agents in, for example, large-scale (cognitive) agent-based simulations remains an important challenge that needs to be addressed.

4.4 Component-Based Agent Design

The working group explored the hypothesis that the monolithic nature of many current agent programming languages and platforms is both a barrier to the adoption of agent technology in “mainstream” software development and industry and an impediment to research, and results in a dilution of effort in the development and maintenance of agent platforms, with useful new features or capabilities being reimplemented for different platforms rather than being improved.¹

Feedback from industrial participants and academics working on large scale deployed applications stressed both the relatively small size of the “agent component” in many systems employing agent technology, and the need for the agent components to integrate with existing software engineering methodologies and tools (see Section 4.1.3). In this context, both the overarching agent-centric nature of many agent development methodologies and the monolithic nature of agent platforms are an issue. In some cases, ideas prototyped in an agent programming language/platform have been re-implemented using “traditional” software development methodologies and languages when the system is deployed, either to facilitate integration and maintenance of the agent components by traditional software developers or because the overheads of a complex agent platform could not be justified when only a subset of its capabilities is required. A more modular approach would address these concerns, by allowing developers to directly incorporate only those features that are required for a particular application (“allowing agent language complexity to be application specific”). In addition, the overhead of learning new agent technologies is also reduced for mainstream developers and in undergraduate teaching (seen as a barrier for many). Only those components and APIs relevant to the current project must be mastered, facilitating a piecemeal, demand-driven integration of agent technology, starting with simple applications of agents, e.g., simple decision making, and progressively expanding outwards to more complex capabilities, e.g., negotiation, as developers gain experience with, and confidence in, agent technology.

A more component-based approach would facilitate research, particularly at the single-agent level, where considerable work remains to be done. Currently, extending a feature of an agent language or platform, such as extending deliberation to incorporate reasoning under uncertainty, or adding a new feature, such as learning, involves mastering the details of the agent platform, and often requires a project of PhD length. This hinders innovation and makes it difficult to create ad-hoc prototypes, e.g., to demonstrate the benefits of agent technology to other communities. A more modular approach with standardised interfaces would address these concerns by allowing researchers to target a single component or small

¹ Component-based agent software engineering was also discussed in the Integration and Validation group.

number of components, rather than the agent platform as a whole. The loose coupling inherent in a component-based approach may also facilitate the development of novel agent architectures incorporating, e.g., multiple asynchronous deliberation cycles, or concurrent reactive and deliberative cycles. This may in turn help “bridge the gap” between architectures for software agents and those found in autonomous robots such as UAVs and spacecraft. Much can be learned from the experiences of the robotics community, which is coalescing around component-based platforms such as ROS that provide libraries and tools to help researchers and software developers create robot applications.² A more modular approach also raises novel short-term research challenges at the component integration level. Key issues include interfaces between components (should these be based on standard languages for representing beliefs, goals and plans, or on syntax neutral approaches based on queries), how coordination between components can be coordinated, and how such coordination rules can best be expressed. In the medium to longer term, componentisation of agent designs should foster the development of a reference model for agent technology. Such a reference model would be a powerful tool both for structuring agents research, and in clarifying to the mainstream software development community that “agents” represents a suite of technologies that can be adopted as needed in applications.

Lastly, it was argued that component-based approaches would also help agent technologies achieve critical mass, both within the agent programming community and, more generally, in the mainstream software development community. The relatively small agent programming community is currently structured around several competing agent programming languages and platforms. While this has been very successful in driving innovation, useful innovations must be re-implemented for each platform (at considerable cost) rather than effort being concentrated on expanding and improving innovative features and their documentation. Focussing on common components would lead to more rapid advances in “whole platform” capabilities (since features no longer have to be re-implemented), promote standardisation regarding key concepts and technologies, and should result in higher quality implementations more likely to be adopted by mainstream developers. Again, there is much that can be learned from recent developments in the robotics community and in other related communities, such as computer vision and the re-use of high-quality BDD libraries in model checking.

A key challenge in adopting a component-based approach is to identify and develop components and their APIs. Fortunately, there is a pool of existing agent platform implementations that can serve as a basis for components, and the agent development community has already made some initial steps in the direction of common interfaces, such as the environment interface standard,³ and modularisation (within a single platform) is now common. However much remains to be done. Another key challenge is in the development of middleware to support the interaction of components, and how this interaction can best be specified. This area is less explored, but even here there is preliminary work on which the community can build.

² www.ros.org

³ <http://sourceforge.net/projects/apeis>

5 Open Problems

- While much research will continue to be devoted to foundational work, there needs to be an increased appreciation within the community of the challenges involved in engineering large-scale multiagent systems. Agents conferences and workshops should encourage submission and acceptance of papers that address these concerns, and should work towards setting and maintaining standards to ensure that work of this kind is of high quality.
- From the industry perspective, although it is very clear what steps can be taken to facilitate the deployment of agent technology, it is uncertain (a) which organisations are best capable of doing this work, (b) how this work will be funded, and (c) whether the needs of end users are sufficient to provide the impetus for the work to be done commercially.
- There remain significant technological barriers to the deployment of multiagent systems which requires research into new techniques, lessons learned from applications, and more generally software engineering type of papers that use existing agent technologies (and not common languages such as Java to build multiagent systems).
- Another challenge that remains is to identify the application areas and types of applications where agent technology provides a critical advantage (such as “autonomous decision making” or explicit multiagent coordination mechanisms), and if possible, to quantify the benefits of using the technology.

6 Panel Discussions

A plenary session was organised on Friday, the last day of the seminar, in which summary reports of the four groups were presented and discussed. The purpose of this session was to identify key challenges and ideas for future research based on discussions during the seminar.

One of the ideas that repeatedly came up during discussions relates to the development of a *modular* or *component-based* agent architecture. The idea is that engineering multiagent systems may be facilitated by a set of components that can be relatively easily exchanged and reused between agent-based applications. Developers within industry may be interested in using some instead of all components of existing agent architectures. Moreover, developing such components may also give rise to some degree of standardisation. It may also give rise to a reference model for agent technology. The main challenge here remains to identify and develop these components.

A related topic concerns the need to continue *research at the single agent level*. The notion of BDI+ was coined to refer to the need to integrate, for example, emotions in a more systematic way into agent architectures. Another example in this area concerns learning. Integrating learning into the agent architecture raises new and interesting challenges that are different from the typical issues studied in the machine learning community. Another challenge is to design new components that extend the capabilities of agents in order to support, for example, reasoning under uncertainty. Finally, more research is needed on the capability of agents to explain their behavior which not only may provide a selling point for the technology but also may be used in debugging tools to identify the reasons for observed behavior.

Another topic that has been quite extensively discussed during the seminar and obtained quite a lot of support as a topic for the research agenda concerns *metrics* and the development of *benchmarks* for agent technology. Some agent programming languages and frameworks

may, for example, facilitate the design of scalable systems. But how do we identify these languages and the features that support the engineering of scalable multiagent systems? Are there specific metrics that apply to multiagent systems. How can we measure, for example, concepts that are often mentioned in the literature such as believability and flexibility of agents?

Tooling has been identified as a main topic for future research as it is very important for the uptake of any technology. The application of agent technology in commercial and business applications requires integration of this technology into the full software life cycle. However, we should not reinvent but rather reuse techniques and tools wherever possible. More research is needed to more clearly identify where tooling developed within the more broader software engineering community can be used to provide this support and where agent-specific tools are needed. Generally, a need was felt to focus on debugging support initially as providing assurance for a multiagent system seems to be of key importance. Moreover, it may be useful to connect to and integrate the work from the Agent-Oriented Programming and Software Engineering communities better.

Participants

- Natasha Alechina
University of Nottingham, GB
- Jeremy Baxter
QinetiQ - Malvern, GB
- Michal Bida
Charles University - Prague, CZ
- Olivier Boissier
Ecole des Mines - St. Etienne, FR
- Rafael H. Bordini
PUCRS - Porto Alegre, BR
- Lars Braubach
Universität Hamburg, DE
- Paolo Busetta
AOS Ltd. - Cambridge, GB
- Rem Collier
University College Dublin, IE
- Stephen Crane
University of Otago, NZ
- Mehdi Dastani
Utrecht University, NL
- Louise Dennis
University of Liverpool, GB
- Virginia Dignum
TU Delft, NL
- Jürgen Dix
TU Clausthal, DE
- Jorge J. Gomez-Sanz
Univ. Comp. de Madrid, ES
- Christian Guttmann
IBM R&D Labs; AU; EBTIC
Abu Dhabi, AE; Monash
University, AU
- Axel Hefler
TU Berlin, DE
- Koen V. Hindriks
TU Delft, NL
- Tom Holvoet
KU Leuven, BE
- Jomi Hübner
Federal University of Santa
Catarina - Brazil, BR
- Yves Lespérance
York University - Toronto, CA
- Brian Logan
University of Nottingham, GB
- John-Jules Ch. Meyer
Utrecht University, NL
- Berndt Müller
University of Glamorgan, GB
- Jörg P. Müller
TU Clausthal, DE
- Alexander Pokahr
Universität Hamburg, DE
- Alessandro Ricci
University of Bologna, IT
- Andrea Santi
University of Bologna, IT
- Federico Schlesinger
TU Clausthal, DE
- Amal El Fallah Seghrouchni
UPMC - Paris, FR
- Maarten Sierhuis
Ejenta Inc., US
- Marija Slavkovic
University of Liverpool, GB
- Bas J. G. Testerink
Utrecht University, NL
- Birna van Riemsdijk
TU Delft, NL
- Jørgen Villadsen
Technical University of Denmark
- Lyngby, DK
- Michael Winikoff
University of Otago, NZ
- Cees Witteveen
TU Delft, NL
- Wayne Wobcke
UNSW - Sydney, AU



Information Flow and Its Applications

Edited by

Samson Abramsky¹, Jean Krivine², and Michael W. Mislove³

¹ University of Oxford, GB, samson.abramsky@comlab.ox.ac.uk

² University of Paris Diderot, FR, jean.krivine@pps.univ-paris-diderot.fr

³ Tulane University, US, mwm@math.tulane.edu

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 12352 “Information Flow and Its Applications”. This seminar brought together mathematicians, computer scientists, physicists and researchers from related disciplines such as computational biology who are working on problems concerning information and information flow.

Seminar 26.–31. August, 2012 – www.dagstuhl.de/12352

1998 ACM Subject Classification H.1.1 Systems and Information Theory

Keywords and phrases Information flow, semantics of computation, quantum computing, systems biology, information theory, informatics

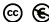
Digital Object Identifier 10.4230/DagRep.2.8.99

1 Executive Summary

Samson Abramsky

Jean Krivine

Michael W. Mislove

License  Creative Commons BY-NC-ND 3.0 Unported license
© Samson Abramsky, Jean Krivine, and Michael W. Mislove

The seminar “Information Flow and Its Applications” that took place in Schloss Dagstuhl in August 2012, has been the latest in a series of meetings concerning information flow that began with the 2008 Clifford Lectures by Samson Abramsky at Tulane University, and continued with two further meetings on informatic phenomena at Tulane, as well as a previous Dagstuhl seminar on “The Semantics of Information”¹. The seminar “Information Flow and Its Applications” brought together mathematicians, computer scientists, physicists and researchers from related disciplines such as computational biology who are working on problems concerning information and information flow.

The seminar gathered 21 participants in addition to the 3 organizers, in the studios but cosy atmosphere of Schloss Dagstuhl. Armed with slides and chalks, each speaker described in terms as simple as can be, the questions and problems they were trying to solve, which, as the title of the seminar suggests, had all in common the issue of the representation and analysis of information flows.

The hypothesis underlying the organization of the seminar was the following: information flows leave on substrates which transport and transform data along time and space. From the modeling, analysis or simulation of these substrates will emerge unifying techniques or

¹ <http://www.dagstuhl.de/10232>



concepts. It is understood that such substrate can be *artificial*, for instance in the case of an electrical circuit, or *natural*, as in the complex signaling pathways that govern cellular fate. Moreover, information may be treated by systems in a *designed* manner, for instance a computer that processes its inputs according to a determined program, or be the result of *evolution*, like the internet which is a perfect example of a system that carries and processes information in spite of the absence of a pre-existing specification.

Although traditionally information processing is studied by distinct communities, scattered along the *Artificial-Natural* and *Designed-Evolved* axes, it is noteworthy that this separation is, to some extent, a historical artifact in the sense that artificial systems may be the fruit of evolution (as the internet) while natural ones may be used in a purely specified manner (as in synthetic biology). It is therefore natural to expect that tools and techniques developed in one field may be also relevant to others.

Another unifying scheme of the seminar was the emphasis on the use of formal languages in the representation of information flows. Indeed once "a real world" computing system, such as the cell or a quantum circuit, is abstracted as a formal programming language, one may then start to apply techniques imported from theoretical computer science. In the study of evolved systems, these techniques may be used to *extract* a specification of what is being observed, while in the context of systems where a specification is *a priori* at disposal, one may use these techniques to *verify* that the way information is processed conforms to the expectation.

Over the 4 days of talks, which gave rise to feedback that went beyond the expectation of the organizers, the participants of the seminar "Information Flow and its Applications" have had the opportunity to listen to talks ranging from Systems Biology to Theoretical Physics, from Quantum Computing to the study of Ecological systems. As organizers, we believe that the original guess that Information Flow should be a topic of its own was largely a good one.

2 Table of Contents

Executive Summary

Samson Abramsky, Jean Krivine, and Michael W. Mislove 99

Overview of Talks

Galois group of a symmetric measurement <i>Marcus Appleby</i>	103
Information and distributed computation <i>David Balduzzi</i>	103
A Hypothesis Test For Bell's Inequality <i>Peter Bierhorst</i>	103
Complex Information Systems <i>Robert J. Bonneau</i>	104
Structured Data Analysis <i>Gunnar Carlson</i>	104
Ensemble signaling in a MAP kinase cascade <i>Eric Deeds</i>	104
Information Flow in Quantum Computing: Circuits, Entanglement, and MBQC <i>Ross Duncan</i>	105
Formal model reduction <i>Jerome Feret</i>	105
Hidden Bayesian networks <i>Tobias Fritz</i>	106
Statistics, causality and Bell's theorem <i>Richard Gill</i>	106
Graphs, Rewriting and Pathway Reconstruction for Rule-Based Models <i>Jonathan Hayman</i>	107
Coherence in Hilbert's hotel <i>Peter Hines</i>	107
Minimal glueings and unambiguous stoichiometry in Kappa <i>Ricardo Honorato-Zimmer</i>	108
From Contextuality to Nonlocality <i>Shane Mansfield</i>	108
Analyzing continuous channels <i>Michael W. Mislove</i>	108
Differential Privacy: An Overview <i>Catuscia Palamidessi</i>	109
Compact Closed Categories and Frobenius Algebras for Computing Natural Language Meaning <i>Mehrnoosh Sadrzadeh</i>	109
Rigid geometric constraints for Kappa models <i>Sandro Stucki</i>	110

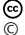

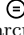
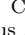
102 12352 – Information Flow and Its Applications

Computer and Information Science – Future Paradigm for Complex System Models? <i>Baltasar Tranco y Widemann</i>	110
Coalgebraic Infinite Games without Discounting – Towards Reflexive Economics <i>Viktor Winschel</i>	111
Probabilistic event structures <i>Glynn Winskel</i>	111
Participants	112

3 Overview of Talks

3.1 Galois group of a symmetric measurement





Marcus Appleby (Perimeter Institute – Waterloo, CA)

License     Creative Commons BY-NC-ND 3.0 Unported license
© Marcus Appleby

The problem of proving (or disproving) the existence of symmetric informationally complete positive operator valued measures (SICs) has been the focus of much effort in the quantum information community during the last 12 years. In this talk we describe the Galois invariances of Weyl-Heisenberg covariant SICs (the class which has been most intensively studied). It is a striking fact that the published exact solutions (in dimensions 2–16, 19, 24, 35 and 48) are all expressible in terms of radicals, implying that the associated Galois groups must be solvable. Building on the work of Scott and Grassl (J. Math. Phys. 51 042203 (2010)) we investigate the Galois group in more detail. We show that there is an intriguing interplay between the Galois and Clifford group symmetries. We also show that there are a number of interesting regularities in the Galois group structure for the cases we have examined. We conclude with some speculations about the bearing this may have on the SIC existence problem.

3.2 Information and distributed computation

David Balduzzi (MPI für Intelligente Systeme – Tübingen, DE)

License     Creative Commons BY-NC-ND 3.0 Unported license
© David Balduzzi




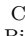
Main reference D. Balduzzi, “On the information-theoretic structure of distributed measurements,” in Proc. of 7th Int’l Workshop on Developments of Computational Methods (DCM’11), EPTCS, vol. 88, pp. 28–42, 2012.

URL <http://dx.doi.org/10.4204/EPTCS.88.3>

Computations implemented in physical systems can be described at many different spatio-temporal granularities. For example, the work performed by the brain can be described at the level of atoms, molecules, neurons and potentially higher-order structures. I present an information-theoretic approach to coarse-graining distributed computations. The first step is to introduce effective information, which can be shown to incorporate Kolmogorov complexity, mutual information and VC-entropy (an important measure of complexity in statistical learning theory) as special cases. A second measure, excess information, provides a geometric characterization of indecomposable computations. The two measures are then applied to study coarse-grainings in Conway’s Game of Life and Hopfield networks.

3.3 A Hypothesis Test For Bell’s Inequality

Peter Bierhorst (Tulane University, US)

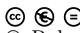
License     Creative Commons BY-NC-ND 3.0 Unported license
© Peter Bierhorst

Experimental tests of Bell inequalities require statistical analysis. Usually, successive trials are taken to be independent and identically distributed. Thus the expectation quantities in Bell’s inequality can be estimated by appealing to the Law of Large Numbers. Though

the i.i.d. assumption is a natural one, it need not be obeyed by a local hidden variable theory. Luckily, statistical methods can still distinguish Quantum Mechanics from local hidden variable theories over large numbers of trials.

3.4 Complex Information Systems

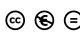
Robert J. Bonneau (AFOSR – Arlington)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Robert J. Bonneau

The talk will provide an overview of complex information systems including quantifying, managing, and designing heterogeneous networked systems. Methods of measuring and assessing the performance of networked, software, and hardware integrated systems such as cloud architectures will be discussed including techniques of sparse approximation in systems measurements, and algebraic and topological statistical metrics for performance. Strategies of quantifying risk over different geometric and statistical classes of distributed systems will be examined as well as methods of tracking and coding dynamic information flows.

3.5 Structured Data Analysis

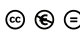
Gunnar Carlson (Stanford University, US)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Gunnar Carlson

We discuss methods for representing and "measuring" the shape of data sets. They are represented by simplicial complexes, and the shape is measured using homological signatures as extended to the world of point cloud data via the persistent homology methodology. Examples were given in both cases, and suggestions were made about how to represent even more complicated data types, involving dynamical systems and control systems.

3.6 Ensemble signaling in a MAP kinase cascade

Eric Deeds (University of Kansas, US)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Eric Deeds





Joint work of Deeds, Eric; Suderman, Ryan

A cell's ability to adapt to constantly changing environmental conditions is derived from its signaling networks. Despite their importance, there is currently no consensus regarding the nature of the protein complexes such networks employ. One prominent view involves signaling machines, while the inherent combinatorial complexity of such networks has led to the more recent proposal of pleiomorphic ensembles. In this work, we use rule-based modeling techniques to explore this question in the case of the yeast pheromone MAPK cascade. We constructed a model of this cascade based on current understanding of the interactions in the pathway. We found that, despite exhibiting considerable ensemble character, this model can replicate existing experimental data for the cascade. We also considered a model designed to

exhibit more machine-like character. This model could not replicate the behavioral changes observed in cascade when Ste5, the signaling scaffold, is overexpressed. These findings indicate that ensemble signaling can indeed produce "realistic" cellular behavior, and that machine and ensemble systems can exhibit distinctly different phenotypes.

3.7 Information Flow in Quantum Computing: Circuits, Entanglement, and MBQC



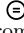

Ross Duncan (Université Libre de Bruxelles, BE)

License     Creative Commons BY-NC-ND 3.0 Unported license
© Ross Duncan

The measurement based quantum quantum computer (MBQC) works by pushing quantum information through a network of entangled quantum states. The structure of these states, and the paths that the information takes within the network are easily studied using a high-level presentation of quantum theory based on symmetric monoidal categories and its graphical language. In this talk I'll show how to apply these techniques to derive a quantum circuit equivalent to a given MBQC program, thus verifying the correctness of the original program.

3.8 Formal model reduction

Jérôme Feret (ENS – Paris, FR)

License     Creative Commons BY-NC-ND 3.0 Unported license
© Jerome Feret

Joint work of Camporesi, Ferdinanda; Danos, Vincent; Feret, Jérôme; Fontana, Walter; Harmer, Russell; Krivine, Jean


Modelers of molecular signaling networks must cope with the combinatorial explosion of protein states generated by post-translational modifications and complex formation. Rule-based models provide a powerful alternative to approaches that require an explicit enumeration of all possible molecular species of a system. Such models consist of formal rules stipulating the (partial) contexts for specific protein-protein interactions to occur. These contexts specify molecular patterns that are usually less detailed than molecular species. Yet, the execution of rule-based dynamics requires stochastic simulation, which can be very costly. It thus appears desirable to convert a rule-based model into a reduced system of differential equations by exploiting the lower resolution at which rules specify interactions.

In this talk, we present a formal framework for constructing coarse-grained systems. We track the flow of information between different regions of chemical species, so as to detect and abstract away some useless correlations between the state of sites of molecular species.

The result of our abstraction is a set of molecular patterns, called fragments, and a system which describes exactly the concentration (or population) evolution of these fragments. The method never requires the execution of the concrete rule-based model and the soundness of the approach is described and proved by abstract interpretation.

3.9 Hidden Bayesian networks

Tobias Fritz (Institute of Photonic Sciences – Castelldefels, ES)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Tobias Fritz

Main reference T. Fritz, “Beyond Bell’s theorem: correlation scenarios,” *New Journal of Physics*, vol. 14, issue 10, Oct. 2012.

URL <http://dx.doi.org/10.1088/1367-2630/14/10/103001>


This will be an outline of a research program aiming at generalizing Bell’s Theorem to arbitrary causal structures. It will probably be joint work with Rob Spekkens.

An arbitrary causal structure is given by a (finite) poset represented by a directed acyclic graph. The vertices in the graph represent spacetime events at which measurements are conducted. The edges are the worldlines in spacetime of physical systems being transmitted from one vertex to another; these messages are what allow the creation of correlations between measurements at different vertices. Such correlations are called classical if the correlations can be modeled in terms of classical messages, or, equivalently, in terms of a “hidden Bayesian network”. (A hidden Bayesian network is like a hidden Markov model, just on an arbitrary causal structure.) The correlations are called quantum if they can be modeled by quantum systems being sent along the edges and quantum measurements being conducted on the vertices.

Standard Bell scenarios as well as the scenarios studied in arXiv:1206.5115 are subclasses of this formalism; interesting examples beyond these remain to be found.

3.10 Statistics, causality and Bell’s theorem

Richard Gill (Leiden University, NL)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Richard Gill

Main reference R.D. Gill, “Statistics, causality and Bell’s theorem,” submitted, arXiv:1207.5103 [stat.AP].


URL <http://arxiv.org/abs/1207.5103v1>

Bell’s (1964) theorem is popularly supposed to establish the non-locality of quantum physics as a mathematical-physical theory. Building from this, observed violation of Bell’s inequality in experiments such as that of Aspect and coworkers (1982) is popularly supposed to provide empirical proof of non-locality in the real world. My talk reviews recent work on Bell’s theorem, linking it to issues in causality as understood by statisticians. The talk starts with a new proof of a strong (finite sample) version of Bell’s theorem which relies only on elementary arithmetic and (counting) probability. This proof underscores the fact that Bell’s theorem tells us that quantum theory is incompatible with the conjunction of three cherished and formerly uncontroversial physical principles, nicknamed here locality, realism, and freedom. I will argue that (accepting quantum theory) Bell’s theorem should lead us to seriously consider relinquishing not locality, but realism, as a fundamental physical principle. In the talk I hope also to discuss statistical issues, in the interpretation of state-of-the-art Bell type experiments, related to post-selection in observational studies. Finally I state an open problem concerning the design of a quantum Randi challenge: a computer challenge to Bell-deniers. Can we prove useful probabilistic Bell inequalities for the situation that Alice and Bob’s measurement stations receive all their measurement settings at once, and all signals from the source at once, and then (disconnected from the one another and the source) generate all their outputs at once. A “useful” Bell inequality in this context is one where the

probability of violation of the local realism expectation bound by some small amount δ decreases with N , the number of settings to be processed at a time.

3.11 Graphs, Rewriting and Pathway Reconstruction for Rule-Based Models

Jonathan Hayman (ENS – Paris, FR)

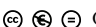
License  Creative Commons BY-NC-ND 3.0 Unported license
© Jonathan Hayman

Joint work of Hayman, Jonathan; Danos, Vincent; Feret, Jérôme; Fontana, Walter; Harmer, Russell; Krivine, Jean; Thompson-Walsh, Chris; Winskel, Glynn

We introduce a novel way of constructing concise causal histories (pathways) to represent how specified structures are formed during simulation of systems represented by rule-based models. This is founded on a new, clean, graph-based semantics introduced in the first part of this paper for Kappa, a rule-based modeling language that has emerged as a natural description of protein-protein interactions in molecular biology. The semantics is capable of capturing the whole of Kappa, including subtle side-effects on deletion of structure, and its structured presentation provides the basis for the translation of techniques to other models. In particular, we give a notion of trajectory compression, which restricts a trace culminating in the production of a given structure to the actions necessary for the structure to occur. This is central to the reconstruction of biochemical pathways due to the failure of traditional techniques to provide adequately concise causal histories, and we expect it to be applicable in a range of other modeling situations.

3.12 Coherence in Hilbert’s hotel

Peter Hines (University of York, GB)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Peter Hines

Main reference A series of papers based on this talk is in progress. Please contact the author for drafts.


This talk describes the interaction of MacLane’s categorical coherence with self-similarity $S \cong S \otimes S$ and untyped (i.e. single-object) monoidal structures.

A coherence result is presented, giving a decision procedure for commutativity of diagrams built up from typed and untyped monoidal tensors and structural isomorphisms, and canonical isomorphisms between the two settings.

Applications are discussed, including word problems in Thompson groups, and deciding equality of arithmetic expressions based on modular arithmetic and related operations.

3.13 Minimal glueings and unambiguous stoichiometry in Kappa

Ricardo Honorato-Zimmer (University of Edinburgh, GB)


License  Creative Commons BY-NC-ND 3.0 Unported license
© Ricardo Honorato-Zimmer

Joint work of Honorato-Zimmer, Ricardo; Stucki, Sandro; Danos, Vincent

Rule-based modeling languages have been developed to represent biomolecular interactions in a concise way. They achieve this by using patterns that omit unnecessary details or context from the description of those interactions. Hence, combinatorial systems with a large or even infinite number of different species can then be studied and analyzed. However, in these languages it is usually not immediately clear when a rule creates or destroys an instance of a pattern or observable, that is, the stoichiometry of the rule for that observable is often ambiguous. In this work, we formally define the concept of minimal glueings in the category-theoretical framework developed by Harmer et al (2011) that allow us to check if a set of observables has an unambiguous stoichiometry with respect to a rule.

3.14 From Contextuality to Nonlocality


Shane Mansfield (University of Oxford, GB)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Shane Mansfield

We outline a route to constructing n-partite Bell-type models from more general measurement configurations, e.g. Kochen-Specker configurations, using a sheaf theoretic formulation of measurement scenarios. The construction has the property that the resultant model is no-signaling and that it is nonlocal if and only if the original model is contextual. This could provide a new route to Bell tests for contextuality. It also raises an interesting and novel connection between the simplest possible contextual model, the triangle, and PR boxes, the maximally nonlocal (2,2,2) correlations.

3.15 Analyzing continuous channels

Michael W. Mislove (Tulane University, US)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Michael W. Mislove

Main reference Slides of the talk available at <http://www.entcs.org/mislove/dagstuhl12.pdf>




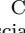
URL <http://www.entcs.org/mislove/dagstuhl12.pdf>

In this talk we describe how the Cantor Fan can be used as a basis for approximating continuous channels with discrete ones. The Cantor Fan is our term for the order-ideal completion of the rooted full binary tree, with the path order, and is so named because the set of maximal elements is homeomorphic to the middle-third Cantor set, when the completion is regarded as an algebraic domain. We describe the effect of applying the Probability monad to this structure, and how it then gives a setting to describe how a channel with the Cantor set as input/output set can be approximated by a sequence of channels, C_n , in a hierarchy, where C_n has 2^n inputs/outputs for each n . We discuss how entropy and capacity of the

individual approximations can be viewed in this setting. Since this is very preliminary work, no definitive results are presented.

3.16 Differential Privacy: An Overview

Catuscia Palamidessi (Ecole Polytechnique – Palaiseau, FR)

License     Creative Commons BY-NC-ND 3.0 Unported license
© Catuscia Palamidessi

Joint work of Palamidessi, Catuscia; Chatzikokolakis, Konstantinos; Alvim, Mario S.; Andrés, Miguel E.



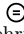

Main reference M.S. Alvim, M.E. Andrés, K. Chatzikokolakis, and C. Palamidessi, “On the relation between Differential Privacy and Quantitative Information Flow,” in Proc. of ICALP’11, LNCS 6756, Springer, pp. 60–76, 2011.

URL <http://hal.inria.fr/inria-00627937/en>

We discuss the general problem of protecting private information and we present differential privacy, a framework which has been recently – and quite successfully - introduced in the area of statistical databases. We discuss the trade-off between privacy and utility, and present some fundamental result in the area. Then, we generalize the notion of differential privacy so to make it applicable to domains other than databases. We start from the observation that the standard notion of differential privacy relies on the notion of Hamming distance on the set of databases, and we extend it to arbitrary metric spaces. We show various examples, and we revise some of the fundamental results of differential privacy in this extended setting. As a particular case study, we consider location-based applications, and the resulting notion of geo-indistinguishability.

3.17 Compact Closed Categories and Frobenius Algebras for Computing Natural Language Meaning

Mehrnoosh Sadrzadeh (University of Oxford, GB)

License     Creative Commons BY-NC-ND 3.0 Unported license
© Mehrnoosh Sadrzadeh

Joint work of Sadrzadeh, Mehrnoosh; Clark, Stephen; Coecke, Bob; Grefenstette, Edward; Kartsaklis, Dimitri; Pulman, Stephen

Main reference (1) B. Coecke, M. Sadrzadeh, S. Clark, “Mathematical Foundations for a Compositional Distributional Model of Meaning,” *Linguistic Analysis*, volume dedicated to Lambek’s Festschrift, 2010.

(2) E. Grefenstette, M. Sadrzadeh, “Experimental Support for a Categorical Compositional Distributional Model of Meaning,” in Proc. of the Conf. on Empirical Methods in Natural Language Processing (EMNLP’11), pp. 1394–1404, ACL, 2011.

URL <http://arxiv.org/abs/1003.4394>


URL <http://www.aclweb.org/anthology/D11-1129>

Compact closed categories have found applications in modeling quantum information protocols by Abramsky-Coecke. They also provide semantics for Lambek’s pregroup algebras, applied to formalizing the grammatical structure of natural language, and are implicit in a distributional model of word meaning based on vector spaces. In particular, in previous work, Coecke-Clark-Sadrzadeh used the product category of pregroups with vector spaces and provided a distributional model of meaning for sentences. We recast this theory in terms of strongly monoidal functors and advance it via Frobenius algebras over vector spaces. The former are used to formalize topological quantum field theories by Atiyah and Baez-Dolan, and the latter are used to model classical data in quantum protocols by Coecke-Pavlovic-Vicary. The Frobenius algebras enable us to work in a single space in which lives meanings of words,

phrases, and sentences of any structure. Hence we can compare meanings of different language constructs and enhance the applicability of the theory. We report on experimental results on a number of language tasks such as word sense disambiguation and term/definition extraction and show how our theoretical predictions are verified on real large scale data from British National Corpus.

3.18 Rigid geometric constraints for Kappa models

Sandro Stucki (University of Edinburgh, GB)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Sandro Stucki

Joint work of Danos, Vincent; Honorato-Zimmer, Ricardo; Jaramillo-Riveri, Sebastian; Stucki, Sandro

Rule-based modeling languages such as Kappa and BNGL allow for a concise description of combinatorially complex biochemical processes. However, these languages do not provide means to directly express the three-dimensional geometry of chemical species. We propose an extension to the Kappa modeling language allowing the annotation of the structure of chemical species with three-dimensional geometric information. This naturally introduces rigidity constraints on the species and reduces the state space of the resulting model by excluding species that are not geometrically sound. We show that geometrically enhanced Kappa models can still be simulated efficiently, albeit at the cost of a greater number of null-events occurring during the simulation.

3.19 Computer and Information Science – Future Paradigm for Complex System Models?

Baltasar Trancon y Widemann (Universität Bayreuth, DE)


License  Creative Commons BY-NC-ND 3.0 Unported license
© Baltasar Trancon y Widemann

Joint work of Trancon y Widemann, Baltasar; Hauhs, Michael

We review the meta-theoretical situation of the scientific discipline of ecology, taken as a typical example of the family of ‘complex system sciences’. We argue that its current relation to classical physics as paradigmatic supplier of concepts, and to computer science as operational supplier of tools is outdated. We propose to regard computer science on the same level as physics, and to explore the scientific potential of concepts imported from theoretical computer science. In particular, we predict a major impact of the richer CS concept of behavior, seen as primary target of empirics and formalization rather than a mechanistic consequence of state. We exemplify this claim by putting the concepts of safety and liveness into ecological context. We show how they can be used as both metaphors and mathematical entities, informing scientific discourse, methodology and modeling.

3.20 Coalgebraic Infinite Games without Discounting – Towards Reflexive Economics


Viktor Winschel (Universität Mannheim, DE)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Viktor Winschel

After a short introduction into the economic field of optimal currency areas and the current problems of the European Monetary Union we sketch the many reflexive issues in economic theory that result from the fact that economic modeling takes place within the modeled system. As a first example towards reflexive structures in economics we present coalgebraic infinite games. The technical part introduces the predicate coinduction proof principle that allows to proof subgame perfect equilibria defined as predicates on the carrier of coalgebras. The economic important result shows that this proof principle does not depend on the usual discounting that essentially transforms infinite structures into finite ones where backwards induction can be applied.

3.21 Probabilistic event structures

Glynn Winskel (University of Cambridge, GB)

License  Creative Commons BY-NC-ND 3.0 Unported license
© Glynn Winskel

This talk presented a new definition of probabilistic event structures, extending existing definitions, and characterized as event structures together with a continuous valuation on their domain of configurations. Probabilistic event structures possess a probabilistic measure on their domain of configurations. This prepares the ground for a very general definition of a probabilistic strategies, which are shown to compose, with probabilistic copy-cat strategies as identities. The result of the play-off of a probabilistic strategy and counter- strategy in a game is a probabilistic event structure so that a measurable pay- off function from the configurations of a game is a random variable, for which the expectation (the expected pay-off) is obtained as the standard Lebesgue integral.

Participants

- Samson Abramsky
University of Oxford, GB
- Marcus Appleby
Perimeter Institute –
Waterloo, CA
- David Balduzzi
MPI für Intelligente Systeme –
Tübingen, DE
- Peter Bierhorst
Tulane University, US
- Robert J. Bonneau
AFOSR – Arlington, US
- Gunnar Carlson
Stanford University, US
- Eric Deeds
University of Kansas, US
- Ross Duncan
Université Libre de Bruxelles, BE
- Jerome Feret
ENS – Paris, FR
- Tobias Fritz
Institute of Photonic Sciences –
Castelldefels, ES
- Richard Gill
Leiden University, NL
- Jonathan Hayman
ENS – Paris, FR
- Peter Hines
University of York, GB
- Ricardo Honorato-Zimmer
University of Edinburgh, GB
- Jean Krivine
University Paris-Diderot, FR
- Shane Mansfield
University of Oxford, GB
- Michael W. Mislove
Tulane University, US
- Catuscia Palamidessi
Ecole Polytechnique –
Palaiseau, FR
- Prakash Panangaden
McGill University –
Montreal, CA
- Mehrnoosh Sadrzadeh
University of Oxford, GB
- Sandro Stucki
University of Edinburgh, GB
- Baltasar Trancón y Widemann
Universität Bayreuth, DE
- Viktor Winschel
Universität Mannheim, DE
- Glynn Winskel
University of Cambridge, GB

