DAGSTUHL
REPORTS

Volume 4, Issue 8, August 2014

ISSN 2192-5283

Published online and open access by
Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, Dagstuhl Publishing, Saarbrücken/Wadern, Germany. Online available at http://www.dagstuhl.de/dagpub/2192-5283

Publication date
January, 2015

Bibliographic information published by the Deutsche Nationalbibliothek
The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available in the Internet at `http://dnb.d-nb.de`.

Digital Object Identifier: 10.4230/DagRep.4.8.i

Aims and Scope
The periodical *Dagstuhl Reports* documents the program and the results of Dagstuhl Seminars and Dagstuhl Perspectives Workshops.
In principal, for each Dagstuhl Seminar or Dagstuhl Perspectives Workshop a report is published that contains the following:

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

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

Contact
Schloss Dagstuhl – Leibniz-Zentrum für Informatik
Dagstuhl Reports, Editorial Office
Oktavie-Allee, 66687 Wadern, Germany
reports@dagstuhl.de
http://www.dagstuhl.de/dagrep

Report from Dagstuhl Seminar 14331

Querying and Reasoning Under Expressive Constraints

Edited by

Michael Benedikt[1], Carsten Lutz[2], and Balder Ten Cate[3]

1 **University of Oxford, GB,** michael.benedikt@gmail.com
2 **Universität Bremen, DE,** clu@informatik.uni-bremen.de
3 **LogicBlox Inc. and University of California – Santa Cruz, US,** balder.tencate@gmail.com

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 14331 "Querying and Reasoning Under Expressive Constraints" which took place from August 10th to August 14th, 2014. The seminar aimed to bring together researchers in databases, knowledge representation, decidable fragments of first-order logic, and constraint satisfaction to identify and discuss common themes and technique as well as complementary ones, identify future research issues, and foster cooperation and cross-fertilization between the communities.

Seminar August 10–14, 2014 – http://www.dagstuhl.de/14331
1998 ACM Subject Classification H.2.1 Database Management: Logical Design, I.2.4 Artificial Intelligence: Knowledge Representation, Formalisms and Methods, F.4.1 Mathematical Logic and Formal Languages: Mathematical Logic
Keywords and phrases Integrity constraints, Open-World Query Answering, Ontology-Based Data Access, Knowledge Representation, Automated Reasoning, Decidable Fragments of First-Order Logic
Digital Object Identifier 10.4230/DagRep.4.8.1
Edited in cooperation with Cristina Civili (University of Rome "La Sapienza", IT)

1 Executive Summary

Michael Benedikt
Carsten Lutz
Balder ten Cate

Motivation

Query answering in the presence of expressive constraints and logical rules is a topic that has drawn attention from several different research communities. In databases, the interaction of constraints and queries arises in the context of query optimization – for example, how to make use of integrity constraints such as inclusion dependencies and functional dependencies in running a query more efficiently. The topic is also central to the more recent database topics of data integration and data exchange, where constraints are used in the specification of schema mappings. In the area of knowledge representation, the interaction of constraints and queries plays a great role as well – particularly in ontology-based query answering.

The work in these areas is closely related also to another fundamental topic in theoretical computer science, namely decidable fragments of first-order logic. In particular, many of the

query answering and query analysis techniques used in recent work within databases and knowledge representation have close links to static analysis of guarded logics, a family of logics that arose out of work by the modal logic and finite model theory communities.

The seminar focused on the convergence of interest of the databases, knowledge representation, and computational logic communities. Its goal was to make visible the connections between these distinct communities, to look at tools and algorithms in one community that can be applied within others, to understand which formalisms and techniques are most promising from the perspective of practical applications, and to propose new ways to combine techniques across communities.

Overview and Outcome

The week started with three overview lectures from well-known authorities in databases, description logics, and decidable fragments of first-order logic. These talks introduced the necessary background for participants and raised research themes that would be explored in later talks. The week then proceeded with a wide-ranging series of talks by participants. In addition to finite model theory, description logics, and databases, there were also talks concerning the interaction of querying problems with constraint satisfaction. The presentations included theoretical work as well as system demonstrations and discussion of practical obstacles to efficient querying with constraints. There were two presentations by participants from industry (IBM and LogicBlox), describing products that implement integrity constraint-based approaches to entity resolution and data analytics, respectively. There was also a presentation on the status of constraint-based reasoning within the W3C endorsed query language SPARQL. In addition to the formal talks, the seminar had an open discussion session, which included a mention of some major open problems and directions to be explored for the communities, as well an attempt at mapping the distinct vocabularies of the different communities.

A main outcome of the discussion was a desire for further interaction between the communities. There were a number of proposals put forward for how to achieve this, including co-location of a KR-related conference with a database conference like VLDB or SIGMOD/PODS. Another outcome was a collection of topics that were particularly worth pursuing by all communities. The handling of inconsistency in databases was one of these – both further investigation of the most widely-used approach for inconsistency-handling, based on repair and consistent query answering, and the examination of alternative approaches. The notion of repair tied into the question of investigating the relationship of data uncertainty and constraints. Markov logic networks (MLNs) are likely to play a role in reconciling "hard" integrity constraints with probabilities, although the interplay of probabilistic data and classical approaches to integrity constraints will involve a more general revision of the major computational problems with uncertainty in mind. Another topic identified for future work was the notion of incremental checking of constraints. Incremental computation was alluded to in several talks, but there appears to be a need to take a more holistic look at models for incremental computation and their application in constraint maintenance. The recent activity within dynamic complexity makes the topic of incremental computation within constraint handling particularly ripe for revisiting.

Conclusion

We believe that the seminar was very successful in bringing together the involved communities and in promoting interaction and exchange between them. Similarities as well as differences

between the communities' research efforts became clearly visible and the participants conceived the seminar as a significant step forwards in bridging the gap and raising mutual awareness. Many participants expressed interest in a followup event.

2 Table of Contents

3 Overview of Talks

3.1 Open-World Finite Query Answering Under Number Restrictions

Antoine Amarilli (Telecom ParisTech, FR)

Open-world finite query answering (QA) is the problem of deciding, given a database instance, a set of constraints and a query, whether the query holds over all possible finite completions of the instance satisfying the constraints. It is used to reason over incomplete information and find out if a query is entailed by constraints given non-exhaustive data. Though finite QA is in general undecidable under expressive constraint languages, decidable cases are known: the guarded fragment, which cannot express number restrictions such as functional dependencies, or the guarded fragment with number restrictions but on a signature of arity only two.

We show that finite QA is decidable under unary inclusion dependencies and functional dependencies. More specifically, we prove that, up to an existing finite closure operation on the dependencies, finite controllability holds: namely, finite QA is equivalent to query answering for arbitrary models (finite and infinite), for which efficient techniques are known. This provides, to our knowledge, the first decidability result for finite QA on arbitrary arity signatures under tuple-generating and equality-generating dependencies with complex interaction.

3.2 The Use of Integrity Constraints at LogicBlox

Molham Aref (LogicBlox – Atlanta, US)

LogicBlox provides an integrated platform for high-performance data management and analytics, based on a declarative language called LogiQL. LogiQL is, at its core, an extension of Datalog that offers native language support for expressing data-intensive tasks such as machine learning and combinatorial optimization. The presentation gives an overview of the platform and the language. In particular, it focuses on the important role of integrity constraints in LogiQL, which are used not only for maintaining data integrity, but also, for example, for the specification of complex optimization problems and probabilistic programming.

3.3 Expressive languages for querying the semantic web

Marcelo Arenas (Pontificia Universidad Catolica de Chile, CL)

Joint work of Arenas, Marcelo; Gottlob, Georg; Pieris, Andreas

The problem of querying RDF data is a central issue for the development of the Semantic Web. The query language SPARQL has become the standard language for querying RDF, since its standardisation in 2008. However, the 2008 version of this language missed some

important functionalities: reasoning capabilities to deal with RDFS and OWL vocabularies, navigational capabilities to exploit the graph structure of RDF data, and a general form of recursion much needed to express some natural queries. To overcome these limitations, a new version of SPARQL, called SPARQL 1.1, was recently released, which includes entailment regimes for RDFS and OWL vocabularies, and a mechanism to express navigation patterns through regular expressions. Unfortunately, there are still some useful navigation patterns that cannot be expressed in SPARQL 1.1, and the language lacks of a general mechanism to express recursive queries.

To the best of our knowledge, there is no RDF query language that combines the above functionalities, and which can also be evaluated efficiently. It is the aim of this work to fill this gap. Towards this direction, we focus on the OWL 2 QL profile of OWL 2, and we show that every SPARQL query enriched with the above features can be naturally translated into a query expressed in a language which is based on an extension of Datalog which allows for value invention and stratified negation. However, the query evaluation problem for this language is highly intractable, which is not surprising since it is expressive enough to encode some inherently hard queries. We identify a natural fragment of it, and we show it to be tractable and powerful enough to define SPARQL queries enhanced with the desired functionalities.

3.4 Inconsistency-tolerant query answering in ontology-based data access

Meghyn Bienvenu (University Paris South, FR)

In this talk, I will give an overview of a recent line of work on inconsistency-tolerant query answering in the setting of ontology-based data access. After reviewing some basic notions related to querying data in the presence of ontologies, I will present various inconsistency-tolerant semantics, discuss their computational properties, and describe a practical method based upon the use of incomplete methods and calls to a SAT solver. At the end of the talk, I will mention some open questions and directions for future research.

References for results mentioned in the talk: [1, 2, 3, 4, 5, 6, 7, 8, 9].

References

1 Lembo et al. Inconsistency-tolerant semantics for description logics. In Proc. of RR (2010).
2 Lembo et al. Query rewriting for inconsistent DL-Lite ontologies. In Proc. of RR (2011).
3 Rosati. On the complexity of dealing with inconsistency in description logic ontologies. In Proc. of IJCAI (2011).
4 Bienvenu. First-order expressibility results for queries over inconsistent DL-Lite knowledge bases. In Proc. of DL (2011).
5 Bienvenu. On the complexity of consistent query answering in the presence of simple ontologies. In Proc. of AAAI (2012).
6 Lukasiewicz, Martinez, and Simari. Inconsistency handling in datalog+/- ontologies. In Proc. of ECAI (2012).
7 Lukasiewicz, Martinez, and Simari. Complexity of inconsistency-tolerant query answering in datalog+/-. In Proc. of OTM (2013).

8 Bienvenu and Rosati. Tractable approximations of consistent query answering for robust ontology-based data access. Proc. of IJCAI (2013).
9 Bienvenu, Bourgaux, and Goasdoué;. Querying inconsistent description logic knowledge bases under preferred repair semantics. In Proc. of AAAI (2014).

3.5 Complexity of Constraints: a short introduction to the universal-algebraic approach

Manuel Bodirsky (Ecole Polytechnique – Palaiseau, FR)

The complexity of Constraint Satisfaction Problems (CSPs) is in multiple ways connected to the theory of databases. It is still an open problem whether the class of all CSPs over a fixed finite domain exhibits a complexity dichotomy: Feder and Vardi conjectured that all such CSPs are in P or NP-complete. In this talk I will give a short introduction to a universal-algebraic approach to this conjecture, including a description of a conjecture about the boarder betwen NP-complete and polynomial-time tractable CSPs. We also present a universal-algebra description of those CSPs that can be solved by a Datalog program. Finally, we give an outlook on how these techniques can be used to study the complexity of (well-behaved) classes of CSPs where the domain is infinite, which creates a link to the subsequent talk of Florent Madelaine on the logic of MMSNP.

3.6 Towards Efficient Reasoning Under Guarded-based Disjunctive Existential Rules

Pierre Bourhis (ENS – Cachan, FR)

Joint work of Bourhis, Pierre; Morak, Michael, Pieris Andreas
Main reference P. Bourhis, M. Morak, A. Pieris, "The Impact of Disjunction on Query Answering Under Guarded-Based Existential Rules," in Proc. of the 23rd Int'l Joint Conf. on Artificial Intelligence (IJCAI'13), IJCAI/AAAI, 2013.
URL http://www.aaai.org/ocs/index.php/IJCAI/IJCAI13/paper/view/6504

The complete picture of the complexity of answering (unions of) conjunctive queries under the main guarded-based classes of disjunctive existential rules has been recently settled. It has been shown that the problem is very hard, namely 2 Exptime-complete, even for fixed sets of rules expressed in lightweight formalisms. This gives rise to the question whether its complexity can be reduced by restricting the query language.

Several subclasses of conjunctive queries have been proposed with the aim of reducing the complexity of classical database problems such as query evaluation and query containment. Three of the most prominent subclasses of this kind are queries of bounded hypertree-width, queries of bounded treewidth and acyclic queries.

The central objective of the talk is to understand what whether the above query languages have a positive impact on the complexity of query answering under the main guarded-based classes of disjunctive existential rules.

We show that (unions of) conjunctive queries of bounded hypertree-width and of bounded treewidth do not reduce the complexity of our problem,

Regarding acyclic queries, although our problem remains 2Exptime-complete in general, in some relevant settings the complexity reduces to Exptime-complete; in fact, this requires to bound the arity of the predicates, and for some expressive guarded-based formalisms, to fix the set of rules.

3.7 On the Data Complexity of Consistent Query Answering over Graph Databases

Gaelle Fontaine (University of Chile, CL)

Joint work of Barcelo, Pablo; Fontaine, Gaelle

Graph database applications such as RDF, social networks and scientific databases are prone to inconsistency, mainly due to interoperability issues. This raises the need for understanding query answering over inconsistent graph databases in a framework that is simple yet general enough to accommodate many of its applications. We follow the well-known approach of consistent query answering (CQA), and study the data complexity of CQA over graph databases for the commonly used regular path queries (RPQs) and regular path constraints (RPCs). In this talk we will present the main complexity results and compare those results to the ones obtained in the setting of relational databases.

3.8 SPARQL 1.1 Entailment Regimes

Birte Glimm (Universität Ulm, DE)

Main reference B. Glimm, C. Ogbuji, (eds.), "SPARQL 1.1 Entailment Regimes – W3C Recommendation 21 March 2013," W3C, 2013.
URL http://www.w3.org/TR/sparql11-entailment/

SPARQL is a semantic web query language and protocol for RDF data standardised by the World Wide Web Consortium in 2008. Since 2013, SPARQL 1.1 extends the original standard by several features such as update capabilities, new query language features, new result formats or the ability to also query for implicit knowledge that can be inferred under RDF(S) or OWL semantics. The talk introduces the SPARQL entailment regimes, which allow for retrieving inferred knowledge. The entailment regimes are specified in a very general way by extending SPARQL's standard query evaluation mechanism (simple entailment/subgraph matching). As a consequence, the standard allows for using different semantics and implementation techniques. While the resulting query language is very expressive, it also lacks some commonly used features such as existentially quantified variables as known from conjunctive queries. The main focus of the talk is on the OWL Direct Semantics entailment regime (based on Description Logics) and some optimisations for query evaluation.

3.9 Non-monotonicity in Data Exchange and Ontological Reasoning

André Hernich, (University of Liverpool, GB)

The semantics of a set E of existential rules (a.k.a. tuple-generating dependencies) is typically defined as the set of all its models. The answer to a query over E is then defined as the set of all tuples that are answers to the query in all models of E. While this notion of query answer is quite robust for positive queries (e. g., it has several natural alternative characterizations, including a procedural one in terms of the chase), it is not so clear whether this semantics is the "right" one for queries that involve negation, or if we allow existential rules extended with negative body literals. Indeed, in the context of data exchange where existential rules serve as a specification of how to translate a given source database into a target database, it has been argued that the model-theoretic, or open world based semantics leads to counter-intuitive answers for non-monotone queries. This triggered several proposals of alternative semantics based on variants of the closed world assumption. Furthermore, existential rules extended with negative body-literals, and corresponding semantics have been studied recently in the context of ontology-based data access.

In this talk, I'll discuss ways of how to deal with negation in queries and/or existential rules. The first part of the talk focuses on closed-world approaches from data exchange, whereas the second part will deal with existential rules extended with negative body literals.

3.10 Beyond DL-Lite: Pay-as-you-go query answering

Ian Horrocks (University of Oxford, GB)

The attractive features of DL-Lite are well known: data can be left in legacy/scalable data stores, and data complexity is low (AC0). However, these benefits come at the cost of a severely constrained ontology language. Many applications seem to require more expressive ontologies – certainly many existing ontologies do not satisfy the relevant constraints. In this talk I will survey techniques for (empirically) scalable query answering in cases where more expressive ontologies are used, focusing on recent work on enhanced materialisation-based techniques.

3.11 Ontology-based data access with Ontop and databases: rewriting and optimisations

Roman Kontchakov (University of London, GB)

Joint work of Kontchakov, Roman; Zakharyaschev, Michael

In the ontology-based data access (OBDA) paradigm, an ontology defines a high- level global schema and provides a vocabulary for user queries, thus isolating the user from the details of the structure of data sources (which can be relational databases, triple stores, datalog

engines, etc.). The OBDA system transforms user queries into the vocabulary of the data and then delegates the actual query evaluation to the data sources.

In this talk, we focus on the rewriting algorithm and optimisation techniques implemented in the OBDA system Ontop in the context of relational databases [8, 6]. In particular, we discuss the tree- witness rewriting, which considers all possible ways of splitting the query into fragments that are mapped onto the elements present in the database and the labelled nulls derived by the axioms of the ontology; the latter fragments are called *tree- witnesses* [5, 7].

From the theoretical point of view, tree witnesses over *OWL 2 QL* ontologies give rise to exponential UCQ rewritings and can be used to obtain exponential lower bounds even for non-recursive datalog rewritings [1, 2, 3, 4]. We also remark that over *OWL 2 EL* ontologies, the same approach results in (recursive) datalog rewritings of polynomial size [7].

In practical scenarios, it appears that the tree witnesses do not produce too many choices and the size of the rewriting can be significantly reduced by taking account of the integrity constraints (in particular, inclusion dependencies) from datasources. These observations confirm that the more general Semantic Query Optimisation methods are important for making OBDA systems efficient [8].

References

1 Bienvenu, M., Kikot, S., and Podolskii, V. V. Succinctness of query rewriting in OWL 2 QL: the case of tree-like queries. In *Informal Proc. of the 27th Int. Workshop on Description Logics (DL 2014)*, volume 1193 of CEUR-WS, pages 45–57, 2014.

2 Gottlob, G., Kikot, S., Kontchakov, R., Podolskii, V. V., Schwentick, T., and Zakharyaschev, M. The price of query rewriting in ontology-based data access. *Artif. Intell.*, 213:42–59, 2014.

3 Kikot, S., Kontchakov, R., Podolskii, V., and Zakharyaschev, M. Exponential lower bounds and separation for query rewriting. In *Proc. of the 39th Int. Colloquium on Automata, Languages, and Programming (ICALP 2012), Part II*, volume 7392 of *LNCS*, pages 263–274. Springer, 2012.

4 Kikot, S., Kontchakov, R., Podolskii, V., and Zakharyaschev, M. On the succinctness of query rewriting over shallow ontologies. In *Proc. of the 23rd EACSL Conf. on Computer Science Logic and the 29th ACM/IEEE Symposium on Logic in Computer Science (CSL-LICS 2014)*, pages 57:1–57:10. ACM, 2014.

5 Kikot, S., Kontchakov, R., and Zakharyaschev, M. Conjunctive query answering with OWL 2 QL. In *Proc. of the 13th Int. Conf. on Principles of Knowledge Representation and Reasoning (KR 2012)*, pages 275–285. AAAI Press, 2012.

6 Kontchakov, R., Rezk, M., Rodríguez-Muro, M., Xiao G., and Zakharyaschev, M. Answering SPARQL queries under the OWL 2 QL entailment regime with databases. In *Proc. of the 13th Int. Semantic Web Conf. (ISWC 2014), Part I*, volume 8796 of *LNCS*, pages 552–567. Springer, 2014.

7 Kontchakov, R., and Zakharyaschev, M. An Introduction to description logics and query rewriting. In *The 10th Reasoning Web Summer School*, volume 8714 of *LNCS*, pages 195–244. Springer, 2014.

8 Rodríguez-Muro, M., Kontchakov, R., and Zakharyaschev, M. (2013). Ontology-based data access: Ontop of databases. In *Proc. of the 12th Int. Semantic Web Conf. (ISWC 2013)*, volume 8218 of *LNCS*, pages 558–573. Springer, 2013.

3.12 Schema-Agnostic Query Rewriting in SPARQL 1.1

Markus Krötzsch (TU Dresden, DE)

Joint work of Bischof, Stefan; Krötzsch, Markus; Polleres, Axel; Rudolph, Sebastian
Main reference S. Bischof, M. Krötzsch, A. Polleres, S. Rudolph, "Schema-Agnostic Query Rewriting in SPARQL 1.1," in Proc. of the 13th Int'l Semantic Web Conf. (ISWC'14), LNCS, Vol. 8796, pp. 585–600, Springer, 2014; pre-print available from author's webpage.
URL http://dx.doi.org/10.1007/978-3-319-11964-9_37
URL http://korrekt.org/page/Schema-Agnostic_Query_Rewriting_in_SPARQL_1.1

SPARQL 1.1 supports the use of ontologies to enrich query results with logical entailments, and OWL 2 provides a dedicated fragment OWL QL for this purpose. Typical implementations use the OWL QL schema to rewrite a conjunctive query into an equivalent set of queries, to be answered against the non-schema part of the data. With the adoption of the recent SPARQL 1.1 standard, however, RDF databases are capable of answering much more expressive queries directly, and we ask how this can be exploited in query rewriting. We find that SPARQL 1.1 is powerful enough to "implement" a full-fledged OWL QL reasoner in a single query. Using additional SPARQL 1.1 features, we develop a new method of schema-agnostic query rewriting, where arbitrary conjunctive queries over OWL QL are rewritten into equivalent SPARQL 1.1 queries in a way that is fully independent of the actual schema. This allows us to query RDF data under OWL QL entailment without extracting or preprocessing OWL axioms.

3.13 Infinite CSP and dichotomy : a quest for nice logics with dichotomy

Florent R. Madelaine (Clermont University, FR)

In this talk, we recall 2 logics MMSNP and MMSNP2 that have been studied in the context of finite domain Constraint Satisfaction Problems (CSP), that correspond naturally to 2 Ontology Based Data Access query languages (cf. talk of Frank Wolter).

The first one, MMSNP, was introduced by Feder and Vardi in their influential 93 paper in which they motivated the CSP dichotomy conjecture. Though too expressive to capture only finite domain CSP, it is known to be Ptime equivalent to the class of finite domain CSP. Consequently, any MMSNP problem would also be either in Ptime or NP-complete, if the conjecture holds. The logic MMSNP enjoys a simple combinatorial definition in terms of forbidden patterns problems via vertex coloured obtructions, and via some suitable normal form it is possible to decide effectively interesting properties such as containment.

The second one, MMSNP2, extends MMSNP. The original definition was combinatorial and involved an extension of forbidden patterns problems via edge coloured obstructions. An equivalent and more logically flavoured definition is to define MMSNP2 as GMSNP where guarded second order quantification is used in lieu of monadic quantification. It is open whether this logic enjoys a dichotomy.

Both logics fall within the framework of CSP with infinite domains, given by a template that is omega-categorical, a framework extensively studied by Manuel Bodirsky et al. (cf. his talk). This could be used to attack the above open question.

3.14 All–instances termination of chase is undecidable

Jerzy Marcinkowski (University of Wroclaw, PL)

Joint work of Gogacz, Tomasz; Jerzy Marcinkowski, Jerzy
Main reference T. Gogacz, J. Marcinkowski, "All-instances Termination of Chase is Undecidable," in Proc. of the 41st Int'l Colloquium on Automata, Languages, and Programming (ICALP'14), LNCS, Vol. 8573, pp. 293–304, Springer, 2014.
URL http://dx.doi.org/10.1007/978-3-662-43951-7_25

We show that all–instances termination of chase is undecidable. More precisely, there is no algorithm deciding, for a given set $\mathcal{T}$ consisting of Tuple Generating Dependencies (a.k.a. Datalog$^{\exists}$ program), whether the $\mathcal{T}$-chase on D will terminate for every finite database instance D. Our method applies to Oblivious Chase, Semi-Oblivious Chase and – after a slight modification – also for Standard Chase. This means that we give a (negative) solution to the all–instances termination problem for all version of chase that are usually considered.

3.15 Efficient Separability of Regular Languages by Subsequences and Suffixes

Wim Martens (Universität Bayreuth, DE)

Main reference W. Czerwinski, W. Martens, T. Masopust, "Efficient Separability of Regular Languages by Subsequences and Suffixes," in Proc. of the 40th Int'l Colloquium on Automata, Languages, and Programming (ICALP'13), LNCS, Vol. 7966, pp. 150–161, Springer, 2013.
URL http://dx.doi.org/10.1007/978-3-642-39212-2_16

When can two regular word languages K and L be separated by a simple language? We investigate this question and consider separation by piecewise- and suffix-testable languages and variants thereof. We give characterizations of when two languages can be separated and present an overview of when these problems can be decided in polynomial time if K and L are given by nondeterministic automata.

3.16 Query Rewriting and Optimization for Ontological Databases

Giorgio Orsi (University of Oxford, GB)

Joint work of Gottlob, Georg; Orsi, Giorgio; Pieris, Andreas
Main reference G. Gottlob, G. Orsi, A. Pieris, "Query Rewriting and Optimization for Ontological Databases," ACM Transactions on Database Systems (TODS), 39(3):25:1–25:46, 2014.
URL http://dx.doi.org/10.1145/2638546

Ontological queries are evaluated against a knowledge base consisting of an extensional database and an ontology (i. e., a set of logical assertions and constraints which derive new intensional knowledge from the extensional database), rather than directly on the extensional database. The evaluation and optimization of such queries is an intriguing new problem for database research.

In this talk, we discuss two important aspects of this problem: query rewriting and optimisation [1]. Query rewriting consists of the compilation of an ontological query into

an equivalent first-order query against the underlying extensional database. We present a set of optimization techniques for query rewriting under rather general classes of ontological constraints which is well-suited for practical implementations. In particular, we show how a conjunctive query against a knowledge base, expressed using linear and sticky existential rules, that is, members of the recently introduced Datalog+/- family of ontology languages, can be compiled into a union of conjunctive queries (UCQ) against the underlying database. Ontological query optimisation, in this context, attempts to improve this rewriting process so to produce possibly small and cost-effective UCQ rewritings for an input query.

This is a joint work with Georg Gottlob and Andreas Pieris. The prototype IRIS+/- engine is available at: https://bitbucket.org/giorsi/nyaya.

References

1 Georg Gottlob, Giorgio Orsi, and Andreas Pieris. *Query Rewriting and Optimization for Ontological Databases*. ACM Transactions on Database Systems (TODS), 39(3):25:1–25:46, 2014.

3.17 Decidable Logics for Managing Change in Graph Databases

Magdalena Ortiz (TU Wien, AT)

Joint work of Ahmetaj, Shqiponja; Calvanese, Diego; Ortiz, Magdalena; Simkus, Mantas
Main reference S. Ahmetaj, D. Calvanese, M. Ortiz, M. Simkus, "Managing Change in Graph-Structured Data Using Description Logics," in Proc. of the 28th AAAI Conf. on Artificial Intelligence (AAAI'14), pp. 966–973, AAAI, 2014.
URL http://www.aaai.org/ocs/index.php/AAAI/AAAI14/paper/view/8238

We study different reasoning problems relevant for managing data that evolves as a result of operations (carried out by users or applications). The problems range from ensuring the satisfaction of a given set of integrity constraints after a given sequence of updates, to deciding the (non-)existence of a sequence of actions that would take the data to an (un)desirable state, starting either from a specific data instance or from an incomplete description of it. Our goal is to identify settings that are expressive enough for realistic application domains, yet are computationally manageable and allow for decidable reasoning. To this aim, we focus on graph-databases, understood broadly as relational databases restricted to unary and binary predicates, and rely on decidable fragments of first order predicate logic whose computational properties are well understood, such as description logics and the two variable fragment of first order predicate logic. We consider an elegant action language in which actions are finite sequences of insertions and deletions of nodes and labels, and use the mentioned decidable logics for describing integrity constraints and (partial) states of the data. We then formalize the data management problems mentioned above as a static verification problem and several planning problems, and we provide algorithms and tight complexity bounds for the formalized problems.

3.18 Generic finite realisations of binary overlap specifications

Martin Otto (TU Darmstadt, DE)

Main reference M. Otto, "Finite Groupoids, Finite Coverings and Symmetries in Finite Structures," arXiv:1404.4599v2 [math.CO], 2014.
URL http://arxiv.org/abs/1404.4599v2

We present a generic construction of finite hypergraphs and relational structures based on reduced products with Cayley graphs of suitable groupoids. To this end we need to obtain groupoids whose Cayley graphs have large girth not just in the usual graph-theoretic sense, but with respect to the length of cycles formed by cosets w.r.t. subsets of generators. Reduced products with such Cayley graphs are sufficiently generic to provide highly symmetric and locally acyclic finite realisations of specifications given in terms of binary local overlap patterns (like tgd-type extension requirements or guarded bisimulation types). These in turn yield corresponding hypergraph coverings and finite model constructions as well as a new proof of Herwig-Lascar style extension properties for partial isomorphisms that lift local to global symmetries. Suitable finite coverings of controlled acyclicity and Herwig-Lascar extension properties have previously found several applications in the finite model theory of guarded logics. Of special interest here are: finite model properties for guarded logics, finite controllability of conjunctive queries w.r.t. guarded constraints, and characterisation theorems for guarded logics in the style of classical preservation properties that also work in restriction to just finite models.

3.19 Certain Answers to Well-Designed SPARQL Queries

Reinhard Pichler (TU Wien, AT)

The need to answer queries over many possible worlds arises in many different settings such as data integration/exchange or in ontology based query answering. The generally agreed approach to query answering in such settings is to compute the so-called certain answers, i. e., answers that one gets by answering the query over any possible world.

In this talk, we study the problem of computing certain answers for the semantic web query language SPARQL. More precisely, we focus on an important fragment of SPARQL – the so-called *well-designed* SPARQL – under OWL2-QL entailment. The main challenge comes from the non-monotonicity of the OPTIONAL operator. We thus first define an intuitive semantics of certain answers based on the subsumption relation – building upon the work of Arenas and Pérez [1]. To actually compute the certain answers of well-designed SPARQL queries under QOL2-QL entailment, we follow the DL-Lite approach of Calvanese et al. [2], i. e.: we start by showing that the certain answers can be obtained from the (in general infinite) canonical model of the RDF graph and the ontology. This fact is then used to show the correctness of a rewriting-based query evaluation procedure.

Acknowledgement. The results presented here stem from unpublished joint work with Shqiponja Ahmetaj, Wolfgang Fischl, Mantas Simkus, and Sebastian Skritek. This research has been funded by the Vienna Science and Technology Fund (WWTF) through project ICT12-015.

References

1 Marcelo Arenas and Jorge Pérez. Querying semantic web data with SPARQL. In *Proc. PODS 2011*, pages 305–316. ACM, 2011.
2 Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, and Riccardo Rosati. Tractable reasoning and efficient query answering in description logics: The *DL-Lite* family. *J. Autom. Reasoning*, 39(3):385–429, 2007.

3.20 Polynomial Combined Rewritings for Existential Rules

Andreas Pieris (University of Oxford, GB)

We consider the scenario of ontology-based query answering where a conjunctive query is evaluated against a database enriched with intensional knowledge via an ontology. It is generally accepted that true scalability of query answering in this setting can only be achieved by using standard relational database management systems (RDBMSs). An approach to query answering that enables the use of RDBMSs is the so-called polynomial combined approach. We investigate this approach for the main guarded- and sticky-based classes of existential rules, and we highlight the assumptions on the underlying schema that are sufficient for the polynomial combined first-order rewritability of those classes. To the best of our knowledge, this is the first work which explicitly studies the polynomial combined approach for existential rules.

3.21 A Declarative Constraint-Based Framework for Linking Entities

Lucian Popa (IBM Almaden Center, US)

Joint work of Douglas Burdick; Ronald Fagin; Phokion Kolaitis; Popa, Lucian; Wang-Chiew Tan

We investigate the theoretical foundations towards a declarative framework for entity linking and, in particular, for entity resolution, which is the problem of identifying whether two records represent the same real-world entity. The framework that we adopt is based on link-to-source constraints, unlike in some earlier approaches where source-to-link constraints were used to dictate how to generate links. Our approach makes it possible to focus entirely on the intended properties of the outcome of entity linking, thus separating the specification from any procedure that implements it. The link-to-source constraints specify the desired properties of a link relation in terms of source relations and built-in predicates such as similarity measures. A key feature of the link-to-source constraints is that they employ disjunction, which enables the declarative listing of all the reasons as to why two entities should be linked. We also consider extensions that capture collective entity resolution, by allowing inter-dependence between links. We study the semantics and the computational complexity of this declarative framework, and also make connections to some well-known probabilistic approaches for entity resolution, including ones based on Markov Logic Networks.

This is joint work [1] with Douglas Burdick, Ronald Fagin, Phokion Kolaitis, and Wang-Chiew Tan. Much of this work is motivated from Midas [2, 3], a broader research project at

IBM Research – Almaden, focused on developing the high-level languages and tools needed for the large-scale integration of data from public unstructured sources. Midas has been applied to various scenarios, including a data integration scenario in the financial domain, based on regulatory filings archived by the US Securities and Exchange Commission (SEC). A live demonstration of this application shows the importance of entity linking in such integration.

References

1 D. Burdick, R. Fagin, P. G. Kolaitis, L. Popa and W.-C. Tan. *A Declarative Framework for Linking Entities*. Submitted for publication. 2014.

2 D. Burdick, M. A. Hernández, H. Ho, G. Koutrika, R. Krishnamurthy, L. Popa, I. Stanoi, S. Vaithyanathan and S. R. Das. *Extracting, Linking and Integrating Data from Public Sources: A Financial Case Study*. IEEE Data Eng. Bull., 34(3), pages 60–67, 2011.

3 M. A. Hernández, G. Koutrika, R. Krishnamurthy, L. Popa, and R. Wisnesky. *HIL: A High-Level Scripting Language for Entity Integration*. In EDBT, pages 549–560, 2013.

3.22 Linear programming and integer linear programming for logical satisfiability problems

Ian Pratt-Hartmann (University of Manchester, GB)

In this talk, I will give a brief overview of some of the ways in which satisfiability and finite satisfiability problems for logical fragments have been reduced to linear programming and integer linear programming problems. I will argue that such reductions are not merely a useful source of upper complexity bounds, but also a way of thinking about these logics that provides real insight into their structure.

3.23 Query answering with key constraints: certainty, counting, and probabilities

Dan Suciu (University of Washington – Seattle, US)

Main reference P. Koutris, D. Suciu, "A Dichotomy on the Complexity of Consistent Query Answering for Atoms with Simple Keys," in Proc. of the 17th Int'l Conf. on Database Theory (ICDT'14), pp. 165–176, 2014.
URL http://dx.doi.org/10.5441/002/icdt.2014.19

In order to answer a query on a database that violates key constraints one has to considers all possible ways to repair the database. I will discuss three semantics to query answering: certainty (checking if the query is true in all repairs), counting the number of repairs that satisfy the query, and computing the marginal probability that the query is true. In each case there exists a dichotomy theorem for answering conjunctive queries without self-joins: for every query the data complexity is either in PTIME, or it is NP-complete / #P- complete. (For "certainty" the dichotomy holds only for single-attribute keys.)

References

1 Nilesh Dalvi, Dan Suciu. *Management of probabilistic data: foundations and challenges.* PODS, 2007
2 Maslowski, Wijsen. *A dichotomy in the complexity of counting database repairs.* J. Comput. Syst. Sci. 2013
3 Paris Koutris, Dan Suciu. *A Dichotomy on the Complexity of Consistent Query Answering for Atoms with Simple Keys.* ICDT 2014

3.24 Decidable Fragments of First-Order Logic

Lidia Tendera (University of Opole, PL)

In this talk we give an outline of recent work in the quest for expressive fragments of first-order logics with good algorithmic properties. We concentrate on two related problems: the satisfiability problem and the query answering problem (under the open world assumption). We are equally interested in reasoning in unrestricted and in finite models, as in some applications we want to model systems and computations that are essentially finite.

We focus mainly on fragments defined by restricting the number of variables, usage of quantifiers or usage of negation, and on most popular classes of queries.

Presenting the material we highlight most effective techniques used in this context, their advantages and limitations. We also point out some future directions of study.

3.25 The Automata/Logic Connection for Expressive Guarded Logics

Michael Vanden Boom (University of Oxford, GB)

Joint work of Vanden Boom, Michael; Benedikt, Michael; ten Cate, Balder; Colcombet, Thomas

The guarded fragment (GF) and guarded negation fragment (GNF) are expressive logics, capturing many query classes and integrity constraints of interest in databases and knowledge representation. In this talk, we consider the fixpoint extensions GFP and GNFP of these logics, which are known to have decidable satisfiability and other nice computational and model-theoretic properties.

In particular, these logics have tree-like models, which make them amenable to tree automata techniques. We describe automata that can be used to show optimal complexity bounds for satisfiability. We also describe recent unpublished work demonstrating how these automata can be adapted to decide the following boundedness problem: given phi(X,x) in GFP or GNFP and positive in X, is there a natural number n such that the least fixpoint of the operator defined by phi is reached within n iterations over all structures?

3.26 On Reasoning about Duplicates over Range Restricted Queries with Bag Semantics

Grant Weddell, (University of Waterloo, CA)

We review earlier work on reducing reasoning about duplicate elimination in queries over an object relational data source to reasoning about knowledge base consistency in a dialect of the CFD family of description logics, logics in which unary function symbols replace binary predicate symbols in underlying signatures. We also give an overview of how such capabilities are used in a more general framework for query compilation based on Craig interpolation.

3.27 Presentations, Invariance, and Definability

Scott Weinstein (University of Pennsylvania, US)

Joint work of Weinstein, Scott; Lindell, Steven

We develop the notion of a presentation for abstract relational structures and investigate presentation invariant elementary definability on classes of finite structures. In particular, we focus on linear orderings induced by traversals of simple graphs, and show how they allow us to strictly extend the power of first- order logic in a natural fashion (defining reachability for example). We show how this leads to a new descriptive characterization of logspace in terms of traversal- invariant elementary definability. Continuing, we demonstrate how an elementary partial order presentation can be used to give a very simple definition of tree-width, providing a new normal form for tree decompositions which is especially compelling. We explore algorithmic aspects of this presentation.

3.28 Two-Variable Logic on 2-Dimensional Structures

Thomas Zeume (TU Dortmund, DE)

Joint work of Manuel, Amaldev; Schwentick, Thomas; Zeume, Thomas

In this talk results for the finite satisfiability problem of the two-variable fragment of first-order logic (FO2) over two-dimensional structures have been discussed. Here, two-dimensional structures are structures with two orders, their induced successor relations and arbitrarily many unary relations. Two types of orders have been considered: linear orders and preorders (i. e. equivalence relations whose equivalence relations are linearly ordered).

Participants

- Isolde Adler
Goethe-Universität Frankfurt am Main, DE
- Antoine Amarilli
Telecom ParisTech, FR
- Molham Aref
LogicBlox – Atlanta, US
- Marcelo Arenas
Pontificia Universidad Catolica de Chile, CL
- Michael Benedikt
University of Oxford, GB
- Meghyn Bienvenu
University Paris South, FR
- Manuel Bodirsky
Ecole Polytechnique – Palaiseau, FR
- Pierre Bourhis
ENS – Cachan, FR
- Cristina Civili
University of Rome "La Sapienza", IT
- Gaelle Fontaine
University of Chile, CL
- Birte Glimm
Universität Ulm, DE
- Tomasz Gogacz
University of Wroclaw, PL
- André Hernich
University of Liverpool, GB
- Ian Horrocks
University of Oxford, GB
- Yazmin Angélica Ibáñez-García
Free University of Bozen-Bolzano, IT
- Emanuel Kieronski
University of Wroclaw, PL
- Phokion G. Kolaitis
University of California – Santa Cruz, US
- Roman Kontchakov
Birkbeck, Univ. of London, GB
- Tomer Kotek
TU Wien, AT
- Markus Krötzsch
TU Dresden, DE
- Katja Losemann
Universität Bayreuth, DE
- Carsten Lutz
Universität Bremen, DE
- Florent R. Madelaine
Clermont University, FR
- Jerzy Marcinkowski
University of Wroclaw, PL
- Wim Martens
Universität Bayreuth, DE
- Jakub Michaliszyn
Imperial College London, GB
- Giorgio Orsi
University of Oxford, GB
- Magdalena Ortiz
TU Wien, AT
- Martin Otto
TU Darmstadt, DE
- Reinhard Pichler
TU Wien, AT
- Andreas Pieris
University of Oxford, GB
- Lucian Popa
IBM Almaden Center, US
- Ian Pratt-Hartmann
University of Manchester, GB
- Sebastian Rudolph
TU Dresden, DE
- Vadim Savenkov
TU Wien, AT
- Inanc Seylan
Universität Bremen, DE
- Dan Suciu
University of Washington – Seattle, US
- Balder Ten Cate
University of California – Santa Cruz, US
- Lidia Tendera
University of Opole, PL
- Michaël Thomazo
TU Dresden, DE
- David Toman
University of Waterloo, CA
- Michael Vanden Boom
University of Oxford, GB
- Grant Weddell
University of Waterloo, CA
- Scott Weinstein
University of Pennsylvania, US
- Frank Wolter
University of Liverpool, GB
- Michael Zakharyaschev
Birkbeck, Univ. of London, GB
- Thomas Zeume
TU Dortmund, DE

Report from Dagstuhl Seminar 14332

Formal Methods for Coordinating Multi-Agent Systems

Edited by

Thomas Ågotnes[1] and Nils Bulling[2]

1 **University of Bergen, NO,** thomas.agotnes@infomedia.uib.no
2 **TU Clausthal, DE,** bulling@in.tu-clausthal.de

Abstract

This report documents the programme and outcomes of the Dagstuhl Serminar 14332 "Formal Methods for Coordinating Multi-Agent Systems", that took place from 10 to 14 August, 2014. This seminar brought together researchers from the following subfields of multi-agent systems: logic, game theory, and agreement technologies. It is set up at the intersection of these active fields of research and aimed at fostering collaborations between them. A key objective of the seminar has been to shed light on formal methods for coordinating multi-agent systems, in particular, how to combine research and tools from the different areas to obtain new techniques for coordinating the behavior of agents. The *coordination problem* is a key problem in multi-agent systems: how can we coordinate the individual behaviour of the agents such that the global behaviour of the system as a whole satisfies our needs? Dagstuhl was an excellent venue to bring together leading researchers from logics, game theory, and agreement technologies to learn about their research activities, to discuss as well as to work on timely problems, and to establish new collaborations between researchers. The outcome of the working groups and discussions provides promising avenues and open questions for future research in the field.

Seminar August 10–14, 2014 – http://www.dagstuhl.de/14332
1998 ACM Subject Classification I.2.11 Distributed Artificial Intelligence – Multiagent systems, Intelligent agents
Keywords and phrases multi-agent systems, coordination, formal methods, game theory, logic, agreement technologies
Digital Object Identifier 10.4230/DagRep.4.8.21
Edited in cooperation with Marija Slavkovic

1 Executive Summary

Thomas Ågotnes
Nils Bulling

Formal methods form an active and broad field of research in multi-agent systems, ranging from bottom-up to top-down approaches. Properties of individual agents, e. g., aspects related to decision making and knowledge representation, are rather low-level, while the specification and verification of multi-agent systems are higher-level. In particular, *logic-based approaches* have been successfully used for the modeling of intelligent agents and for reasoning about them: epistemic logics allow to talk about knowledge; temporal logics to reason about the evolution of actions; and strategic logics have been proposed to reason about abilities of agents and coalitions. Alternating-time temporal logics and STIT logics are prominent members of the latter type, and more expressive logics like Strategy Logic have recently

been proposed. What they all have in common is their descriptive flavor. Typically, they are not used to actively change the state of the agent system but to talk and to reason about the system. Multi-agent logics are particularly relevant for the coordination problem. The latter is concerned with global properties of a system. Since the global behaviour of a system emerges from the individual behaviour of agents, it is not obvious what the global properties are. By specifying global properties using multi-agent logics, verification techniques can be employed to verify what of these properties are met by the system; thus, to find out what the global properties are. Interaction between rational decision makers in general, and coordination problems in particular, have been studied in game theory for decades. However, game theory is not concerned with *computational* or *logical* aspects of coordination: how we can represent and reason about coordination in computers. In contrast, many *agreement technologies* are used in an interactive way, e. g., for arriving at agreements about joint actions or coalition structures. Techniques like norms and social laws coordinate the agent's behavior and often require less interaction of agents with their peers. Agents have to decide whether to comply with the rules or not. A difficult problem is the synthesis of appropriate norms and social laws. Related issues important for appropriate control techniques include the detection of norm violations and sanctioning mechanisms.

The seminar aimed at opening up new directions of research into the coordination problem, by bringing together researchers working in different areas of multi-agent systems as well as related fields, and in particular, to combine insights from research in the following fields:

- formal methods and verification, and multi-agent logics in particular,
- game theory in multi-agent systems, and
- agreement technologies.

The seminar took place between 10 and 14 August, 2014. This medium-size, four day seminar was highly international: the 27 participants came from 12 different countries. The scientific program consisted of presentations, discussions and working groups. We scheduled presentations of three different types: overview, medium, and short. The aim of the four one hour overview talks was to give a broad introduction of the main fields relevant to the seminar – to provide a common ground. They covered Argumentation Theory, Normative Systems, Judgement Aggregation, and Computational Social Choice. Then, we scheduled ten medium (20 minutes long) and ten short (15 minutes long) presentations. We encouraged the speakers to give rather informal, non conference-style talks focussing on high-level ideas in order to provide input for the discussion groups.

From the discussions, two working groups emerged which focused on one of the following topics (cf. Sections 4.1 and 4.2):

- Concepts: conceptual definition and classification – what is coordination, coordination problems, and solutions?
- Formalisation of coordination

We organized three meetings for the working groups and two joint discussion sessions for presenting and discussing the results of the working groups.

In addition to the scientific program, we enjoyed a hike which was followed a Barbecue, and the unique atmosphere of Dagstuhl.

2 Table of Contents

3 Overview of Talks

3.1 Strategic Voting and Strategic Candidacy

Markus Brill (Duke University, US)

Joint work of Brill, Markus; Conitzer, Vincent
Main reference M. Brill, V. Conitzer, "Strategic Voting and Strategic Candidacy," in Proc. of the 29th AAAI Conf. on Artificial Intelligence (AAAI'15), to appear; pre-print available from author's webpage.
URL http://www.cs.duke.edu/~brill/papers/candidacy.pdf

Models of strategic candidacy analyze the incentives of candidates to run in an election. Most work on this topic assumes that strategizing only takes place among candidates, whereas voters vote truthfully. In this paper, we extend the analysis to also include strategic behavior on the part of the voters. (We also study cases where only candidates or only voters are strategic.) We consider two settings in which strategic voting is well-defined and has a natural interpretation: majority-consistent voting with single-peaked preferences and voting by successive elimination. In the former setting, we analyze the type of strategic behavior required in order to guarantee desirable voting outcomes. In the latter setting, we determine the complexity of computing the set of potential outcomes if both candidates and voters act strategically.

3.2 Many questions on the semantics of AT(E)L(*)-type logics and some answers

Jan M. Broersen (Utrecht University, NL)

I propose to analyse the semantics of AT(E)L(*) type logics using standard tree-based semantics with epistemic indistinguishability relations. The aim is to come to a unified view on (1) memory, (2) uncertainty and (3) uniformity of strategies in the semantics for these logics. The new setting should also enable one to develop a theory of strategic ability that includes strategies with observations.

3.3 Agents with Perfect and Truly Perfect Recall

Nils Bulling (TU Clausthal, DE)

Joint work of Bulling, Nils; Jamroga, Wojciech; Popovici, Matei
Main reference N. Bulling, W. Jamroga, M. Popovici, "ATL* with truly perfect recall: Expressivity and validities," in Proc. of the 21st European Conf. on Artificial Intelligence (ECAI'14), Frontiers in Artificial Intelligence and Applications, Vol. 263, pp. 177–182, IOS Press, 2014.
URL http://dx.doi.org/10.3233/978-1-61499-419-0-177

In alternating-time temporal logic ATL* [1], agents with perfect recall assign choices to sequences of states, i. e., to possible finite histories of the game. However, when a nested strategic modality is interpreted, the new strategy does not take into account the previous sequence of events. It is as if agents collect their observations in the nested game again from

scratch, thus effectively forgetting what they observed before. Intuitively, it does not fit the assumption of agents having perfect recall of the past.

In this talk I shall review ATL and its semantic variants. Then, I present a new semantics for ATL* where the past is not forgotten in nested games [3, 4]. I give a formal treatment and show that the standard semantics of ATL* coincides with the new semantics in case of agents with perfect information. On the other hand, both semantics differ if agents have imperfect information about the state of the game. I compare the expressivity of the logics and their sets of validities; the latter characterizes general properties of the underlying classes of games (cf. [2]).

References

1 R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time Temporal Logic. *Journal of the ACM*, 49:672–713, 2002.
2 Nils Bulling and Wojciech Jamroga. Comparing variants of strategic ability: how uncertainty and memory influence general properties of games. *Journal of Autonomous Agents and Multi-Agent Systems*, 28(3):474–518, 2014.
3 Nils Bulling, Wojciech Jamroga, and Matei Popovici. Agents with truly perfect recall in alternating-time temporal logic (extended abstract). In *Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2014)*, pages 1561–1562, Paris, France, May 2014. ACM Press.
4 Nils Bulling, Wojciech Jamroga, and Matei Popovici. ATL* with truly perfect recall: Expressivity and validities. In *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI 2014)*, pages 177–182, Prague, Czech Republic, August 2014.

3.4 Reasoning About Norms under Uncertainty in Dynamic Environments

Natalia Criado (John Moores University – Liverpool, GB)

One of the main goals of the agent community is to provide a trustworthy technology that allows humans to delegate some specific tasks to software agents. Frequently, laws and social norms regulate these tasks. As a consequence, agents need mechanisms for reasoning about these norms. Up until now the existing proposals on normative agents assume that agents interact within a deterministic environment that is certainly perceived. In practice, agents interact by means of sensors and actuators under uncertainty with non-deterministic and dynamic environments. In this talk, I presented my work on normative agents that are able to deal with uncertainty in dynamic environments

3.5 Norm-based Coordination in Multi-Agent Systems

Mehdi Dastani (Utrecht University, NL)

Norms have been widely proposed as a means of coordinating the behaviours of agents. This presentation discusses some application areas where norms can be used for coordination

purposes. A programming view on norms will be provided and it will be explained how programmed norms can be enforced and regimented by means of sanctions. In order to study the behaviours of norm programs, a logical analysis will be provided that allows reasoning about the behaviours of multi-agent systems under norm enforcement. We also discuss the enforcement of norms by multiple sanctions.

References

1 Mehdi Dastani and Davide Grossi and John-Jules Meyer. A Logic for Normative Multi-Agent Programs. Journal of Logic and Computation, 2011.
2 Max Knobbout and Mehdi Dastani. Reasoning under Compliance Assumptions in Normative Multiagent Systems. Eleventh International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), 2012.
3 Nils Bulling and Mehdi Dastani and Max Knobbout. Monitoring Norm Violations in Multi-Agent Systems, Twelfth International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), 2013.
4 Natasha Alechina and Mehdi Dastani ad Brian Logan. Norm Approximation for Imperfect Monitors. Thirteenth International Conference Autonomous Agents and Multi-Agent Systems (AAMAS), 2014.

3.6 Verifying Agents that Plan

Louise A. Dennis (University of Liverpool, GB)

Joint work of Dennis, Louise A.; Fisher, Michael

We are interested in the issue of "graceful degradation" of BDI based autonomous systems. If such as system gracefully degrades then it can adapt to the failure of some of its components, for instance by modifying its goals, adapting its plans, or requesting assistance from other agents. A key part of such a capability is likely to be access to planning sub-systems.

This raises some interesting issues for the verification of such systems. This will be a largely speculative short talk, exploring some of those issues in a preliminary way.

3.7 Automata techniques for temporal epistemic logics

Catalin Dima (University Paris-Est – Créteil, FR)

Joint work of Bozianu, Rodica; Dima, Catalin; Filiot, Emmanuel; Maubert, Bastien; Pinchinat, Sophie
Main reference R. Bozianu, C. Dima, E. Filiot, "Safraless Synthesis for Epistemic Temporal Specifications," in Proc. of the 26th Int'l Conf. on Computer Aided Verification (CAV'14), LNCS, Vol. 8559, pp. 441-456, Springer, 2014.
URL http://dx.doi.org/10.1007/978-3-319-08867-9_29
Main reference C. Dima, B. Maubert, S. Pinchinat, "The Expressive Power of Epistemic mu-Calculus," arXiv:1407.5166v1 [cs.LO], 2014.
URL http://arxiv.org/abs/1407.5166v1

The classical duality between automata and logics lies at the foundations of many results on decidability and expressivity of modal and temporal logics, with applications in model-checking and synthesis. It's therefore natural to investigate the possibility of extending this duality to the case of temporal epistemic logics. We present two recent results in which

automata techniques and their relation with temporal epistemic logics play an essential role: a non-expressivity result concerning ATL with imperfect information and the epistemic mu-calculus, and an implemented technique for controller synthesis from temporal epistemic goals.

3.8 Sharing information in teams: what, when and with whom?

Maaike Harbers (TU Delft, NL)

A multi-agent system with a joint goal, a team, can increase its performance by sharing information between team members. For instance, sharing information helps to anticipate each other's actions, plan own activities efficiently, and help each other when possible and needed. However, sharing more information is not always better. Agents may have a limited capacity for sending and processing information, and the exchange of certain information can lead to privacy loss. In this talk, I will discuss a formal model of teamwork that captures positive and negative effects of information sharing in teams. The model allows for the exploration of different information sharing policies, thus benefiting the design of effective teams.

3.9 Exploiting Speculative Computation with Defeasible Reasoning in Multi-Agent System

Ho-Pun Lam (NICTA – Brisbane, AU)

Joint work of Lam, Ho-Pun; Governatori, Guido; Satoh, Ken; Hosobe, Hiroshi
Main reference H.-P. Lam, G. Governatori, K. Satoh, H. Hosobe, "Distributed Defeasible Speculative Reasoning in Ambient Environment," in Proc. of the 13th Int'l Workshop on Computational Logic in Multi-Agent Systems (CLIMA'12), LNCS, Vol. 7486, pp. 43–60, Springer, 2012.
URL http://dx.doi.org/10.1007/978-3-642-32897-8_5

Speculative Computation is an effective means for solving problems with incomplete information in an open and distributed environment, such as peer- to-peer environment. It allows such a system to compute tentative (and possibly final) solutions using default knowledge about the current environment, or the agent's perception, even if the communications between peers are delayed or broken. However, previous work in speculative reasoning assumed that agents are hierarchically structured, which may not be the case in reality. We propose a more general multi-agents system with no centralized control. Agents in the framework have equivalent functionalities and can collaborate with each other to achieve their common goals. We characterize the framework using the argumentation semantics of defeasible logic, which provides support of speculative reasoning in the presence of conflicting information, and provide an operational model for the framework.

3.10 Formal Argumentation and Its Roles in Multi-Agent Systems

Beishui Liao (Zhejiang University, CN)

Formal argumentation is an increasingly active research area in artificial intelligence. In this talk, after discussing some properties of various kinds of reasoning in multi-agent systems and the limitations of first-order logic and traditional nonmonotonic formalisms, I briefly introduce formal argumentation, and some research directions of formal argumentation in multi-agent systems.

3.11 Fair allocation of group tasks according to social norms

Brian Logan (University of Nottingham, GB)

Joint work of Alechina, Natasha; van der Hoek, Wiebe; Logan, Brian
Main reference N. Alechina, W. van der Hoek, B. Logan, "Fair Allocation of Group Tasks According to Social Norms," in Proc. of the 15th Int'l Workshop on Computational Logic in Multi-Agent Systems (CLIMA'14), LNCS, Vol. 8624, pp. 19–34, Springer, 2014.
URL http://dx.doi.org/10.1007/978-3-319-09764-0_2

This talk considers the problem of decomposing a group norm into a set of individual obligations for the agents comprising the group, such that if the individual obligations are fulfilled, the group obligation is fulfilled. The group norms we consider may be non-repeating or repeating (e. g., a group obligation that should be discharged each week). We assume that the assignment of tasks in a group norm to agents is subject to additional social or organisational norms that specify permissible ways in which tasks can be assigned. An important type of social norms are 'fairness constraints', that seek to distribute individual responsibility for discharging the group norm in a 'fair' or 'equitable' way, e. g., an agent may be required to perform a particular task no more than once a week. I briefly present our initial attempts to formalise both group norms and social norms/fairness constraints, and highlight some open problems.

3.12 Towards future road networks: considering traffic system's fairness trap

Marin Lujak (University Rey Juan Carlos, ES)

Joint work of Lujak, Marin; Giordani, Stefano; Ossowski, Sascha
Main reference M. Lujak, S. Giordani, S. Ossowski, "Fair route guidance: bridging system and user optimization," in Proc. of the 17th Int'l IEEE Conf. on Intelligent Transportation Systems (ITSC'14), IEEE, to appear.

In this talk we study the problem of the assignment of road paths to vehicles. Due to the assumption that a low percentage of vehicles follow the routes proposed by route guidance systems (RGS) and the increase of the use of the same, the conventional RGS might shortly result obsolete. Assuming a complete road network information at the disposal of RGSs,

their proposed paths are related with user optimization which in general can be arbitrarily more costly than the system optimum. However, the user optimum is fair for the drivers on the same Origin-Destination (O-D) but it doesn't guarantee fairness for different O-D pairs. Contrary, the system optimum can produce unfair assignments both for the vehicles of the same as of different O-D pairs. This is the reason why, in this talk, we present and discuss an optimization model which bridges this gap between the user and system optimum, and propose a new mathematical programming formulation based on Nash Welfare optimization which results in a good egalitarian and utilitarian welfare for all O-D pairs. Furthermore, we discuss issues related with the proposed distributed approach based on multi-agent system principles and paradigms.

3.13 Practical Reasoning, Norms and Argument

Nir Oren (University of Aberdeen, GB)

Joint work of Oren, Nir; van Riemsdijk, Birna

This talk describes a model of normative system based on a transition system wherein agents act based on goals and directed norms. These impose a partial preference relation over paths through the system, from which an equilibrium can be computed. I then map the preference relations for an agent to argument schemes, from which an argumentation framework for the system can be derived; evaluating this system aims to again identify system equilibria, and the arguments themselves can also be used to explain the functioning of the system.

3.14 An action language approach to normative specification, analysis and revision

Julian Padget (University of Bath, GB)

Joint work of Li, Tingting; Balke, Tina; De Vos, Marina; Padget, Julian; Satoh, Ken
Main reference T. Li, T. Balke, M. De Vos, J. Padget, K. Satoh, "Legal conflict detection in interacting legal systems," in Kevin D. Ashley, editor, "Legal Knowledge and Information Systems," Frontiers in Artificial Intelligence and Applications, Vol. 259, pp. 107–116. IOS Press, 2013.
URL http://dx.doi.org/10.3233/978-1-61499-359-9-107

InstAL (Institutional Action Language) is both a declarative domain-specific language for the specification of collections of interacting normative systems and a framework for a set of associated tools. The computational model is realized by translating the specification language to AnsProlog, a logic programming language under the answer set semantics (ASP), and is underpinned by a set-theoretic formal model and a formalized translation process. Among its notable features are: (i) non-inertial fluents, which allow the creation of institutional fluents that hold for as long as some condition over existing institutional facts is true, (ii) interacting institutions, which allow for events in one institution to affect another and support the modularization of institutional specifications, and (iii) durations, which allow a fluent to hold at some number of time steps in the future. A recent substantive extension is the means to carry out conflict detection and resolution for coordinated institutions. This enables the specification of normative positions through use cases comprising a (partial)

sequence of events and a (partial) state description and the consequent synthesis using inductive logic programming of a minimal self-consistent rule set.

3.15 Anonymity in ATL – Strategic Homogeneity

Truls Pedersen (University of Bergen, NO)

We discuss anonymous alternating-time temporal logic (ATL) which is ATL interpreted over symmetric concurrent game structures (CGSs). By symmetry in the model, we mean that all permutations of any action profile leads to bisimilar states. Symmetry in the model permit representation for which the model checking problem is no longer exponential in the number of agents. We also give a complete axiomatization of anonymous ATL over symmetric models. As the symmetry assumption is very strong, we indicate how we can reintroduce heterogeneity in the model by roles without seriously affecting the complexity of the model checking problem.

3.16 Modal Logic for Mixed Strategies in Games

Joshua Sack (University of Amsterdam, NL)

Joint work of Sack, Joshua; van der Hoek, Wiebe
Main reference J. Sack, W. van der Hoek, "A Modal Logic for Mixed Strategies," Studia Logica, 102(2):339–360, 2014.
URL http://dx.doi.org/10.1007/s11225-014-9548-1

Mixed strategies are useful for reasoning about equilibria in games. Although it is true that every finite strategic form game has a Nash equilibrium, this is not true if one restricts oneself to pure strategies, as can be exemplified by the Matching Pennies Game. This talk introduces Modal Logic for Mixed Strategies, a modal logic that can reason about mixed strategies and mixed Nash equilibria in games. A sound and strongly complete proof system for it is given that makes use of a number of non-standard infinitary rules.

3.17 Abstract Formal Basis for Digital Crowds

Marija Slavkovik (University of Bergen, NO)

Joint work of Slavkovik, Marija; Dennis, Louise; Fisher, Michael
Main reference M. Slavkovik, L. A. Dennis, M. Fisher, "An Abstract Formal Basis for Digital Crowds," arXiv:1408.1592v1 [cs.LO], 2014.
URL http://arxiv.org/abs/1408.1592v1

Crowdsourcing, together with its related approaches, has become very popular in recent years. All crowdsourcing processes involve the participation of a digital crowd, a large number of people that access a single Internet platform or shared service. In this paper we explore the

possibility of applying formal methods, typically used for the verification of software and hardware systems, in analysing the behaviour of a digital crowd. More precisely, we provide a formal description language for specifying digital crowds. We represent digital crowds in which the agents do not directly communicate with each other. We further show how this specification can provide the basis for sophisticated formal methods, in particular formal verification.

3.18 Judgment Aggregation – an overview

Marija Slavkovik (University of Bergen, NO)

Judgment aggregation is a social choice theory concerned with aggregation of collection of judgments, or truth-value assignments, made for a set of logically related issues. This talk was a general overview of the field. It presumed that the participants are entirely unfamiliar with judgment aggregation and somewhat familiar with social choice theory. The talk focused on answering the following questions: what is judgment aggregation, how does it relate to the better known voting theory, what are the points of interest from a MAS and AI perspective, and what are the interesting open problems in this area.

3.19 Resource-sensitive interactions

Nicolas Troquard (National Research Council – Povo (Trento), IT)

I insist on the importance of considering resources in multi-agent interactions. I present some motivations with examples in socio-technical systems. I introduce a resource-sensitive logic of agency in a more technical part.

3.20 Pre-vote negotiations and voting games

Paolo Turrini (Imperial College London, GB)

Joint work of Grandi, Umberto; Grossi, Davide; Turrini, Paolo
Main reference U. Grandi, D. Grossi, P. Turrini, "Pre-vote negotiations and binary voting with constraints," arXiv:1404.5433v1 [cs.GT], 2014.
URL http://arxiv.org/abs/1404.5433v1

I have talked about voting games on possibly interconnected issues, where voters might hold a principled opinion about a subset of the issues at stake while willing to strike deals on the remaining ones, and can influence one another before casting their ballots in order to obtain an individually more favourable outcome. We analyse voters' rational behaviour in a two-phase game, allowing players to undergo a negotiation phase before their vote, and showing under what conditions undesirable equilibria can be removed as an effect of the pre-

vote phase. What I presented is joint work with Davide Grossi (Liverpool) and Umberto Grandi (Toulouse).

3.21 Computational Reasoning for Socially Adaptive Electronic Partners

Birna Van Riemsdijk (TU Delft, NL)

Main reference M. B. van Riemsdijk, L. Dennis, M. Fisher, K. V. Hindriks, "Agent reasoning for norm compliance: a semantic approach," in Proc. of the 12th Int'l Conf. on Autonomous Agents and Multi-Agent Systems (AAMAS'13), pp. 499–506, IFAAMAS, 2013; pre-print available from author's webpage.
URL http://dl.acm.org/citation.cfm?id=2485000
URL http://ii.tudelft.nl/~birna/publications/2013/riemsdijk13aamas.pdf

Technology is becoming an integral part of our daily lives. To ensure that this process unfolds with sufficient respect for important values such as privacy and freedom, I develop software that adapts to norms and values of people. In this talk I introduce a new computational reasoning framework for Socially Adaptive Electronic Partners (SAEPs) that support people in their daily lives without people having to adapt their way of living to the software. The computational reasoning techniques are aimed at determining when and to what extent norm-compliance can be guaranteed, and deciding what to do if in exceptional situations a norm cannot be complied with. The reasoning framework is based on executable temporal logic, integrating the agent's execution semantics with adopted norms for guaranteeing compliance.

3.22 Collective Intention Revision from a Database Perspective

Marc Van Zee (University of Luxembourg, LU)

Joint work of Van Zee, Marc; Dastani, Mehdi; Van der Torre, Leon; Shoham, Yoav
Main reference M. van Zee, M. Dastani, Y. Shoham, L. van der Torre, "Collective Intention Revision from a Database Perspective," in Proc. of the Collective Intentionality Conference 2014; pre-print available from author's webpage.
URL http://icr.uni.lu/marc/publications/ci2014.pdf

Icard et al. recently formalized Shoham's "database perspective" with a logical model to capture action, belief and intention. We extend this model to a multi-agent setting by introducing a collective intention base that captures dependencies between intentions of different agents. We provide AGM-like postulates for multi-agent revision of beliefs, individual intentions, and collective intentions, and conjecture a representation theorem relating our postulates to the formal model.

3.23 Adaptation of social control and trust mechanisms

Laurent Vercouter (INSA – St-Etienne-du-Rouvray, FR)

Multiagent systems offer an interesting approach to deal with large scale distributed systems as it proposes to focus on local design and development of agents on on their interaction rather than trying to consider the system as a whole. However, it is not an easy task to define control mechanisms ensuring that the global behavior comply to what is expected. Agreement technologies and more specifically multiagent trust and reputation models brings solution for controlling such networks when it is not possible to get a global and complete observation of the agents' behaviors. Following the early proposition of Castelfranchi & Falcone, social control mechanisms have been developed these last years in the multiagent community as a way to provide to artificial agent societies the means to control itself using trust and reputation concepts in order to allow each agent to observe, evaluate and sanction its neighborhood. This talk will present the general issues involved in the development of social control mechanisms. We will present some specific concrete examples of their applications to decentralized systems (sensor networks, peer-to-peer systems, social networks). We will consider especially situations where uncertainty (in the agents' identity or in the definition of expectations) prevents the use of classical control mechanisms and for which adapted solutions have been developped. Finally, we will present our view of open issues and perspectives for this topic.

4 Working Groups

There were four working group discussion sessions of about 75 minutes each, including an initial meeting where the organizers proposed general topics, centered around coordination, for the discussions groups. They were then concretized by the seminar participants. After some general discussion among all seminar participants the organizers separated the seminar into two groups of about equal size. The outcomes of the working groups were presented to the whole seminar at two wrap-up sessions, in the middle and at the very end of the seminar.

For the first meeting of the plenary discussion group, the seminar participants identified interesting sub-topics of the general topic. These were then, after some discussion, divided into two general topics:

- Concepts: conceptual definition and classification – what is coordination, coordination problems, and solutions?
- Formalisation of coordination

The following two subsections are reports that sums up the outcomes of the two working groups, edited by the respective working group coordinator(s).

4.1 Working group: Concepts – What is coordination?

Joint work of working group members; edited by: Louise Dennis (group coordinator)

4.1.1 Introduction

The working group began by considering the question "What is a Coordination Problem" and then worked on understanding the space of such problems and their solutions.

4.1.2 What is a coordination problem?

▶ **Definition 1** (Coordination Problem)**.** Given a purpose/requirement and n agents (where $n > 1$), where one of those agents alone may not achieve the purpose then the complete solution to the coordination problem is a system which guarantees the achievement of the purpose.

Partial solutions

- achieve compromises,
- increase the chance of achieving goals

4.1.3 Purpose

We assume that a coordination is designed by some person for the purpose of achieving some thing.

Requirements Ultimately the system designed wants to realise the requirements. They may be informal – in fact they will probably be informal.

Motivations Motivations are a formal set of expressions that explicitly state a purpose of the coordination. If the requirements were formal then the motivations can equal the requirements. Motivations may be exogenous to the system (i. e., nothing in the system has any explicit awareness of them) or they may be endogenous to the system (i. e., at least one entity in the system is aware of the motivations).

We note that motivations can be engineered, particularly if they are to be endogenous to the system. i. e., finding a suitable formal expression for the requirements may be a part of the solution.

Soft Constraints Additionally there may be a number of properties which are desirable in the system but which are not required for a solution. "Optimal" solutions (and compromises) will optimise these soft constraints.

Goals

Note: Both motivations and soft constraints are referred to as *goals*. A *goal* is therefore an explicit formal property that we wish the system to have.

We categorise goals into three types.

A Type A goals are discrete yes/no properties. The system either has them or it does not.

B Type B goals are metrics with a threshold. The system is required to exceed the threshold in order for the goal to be achieved. Optimal systems will more than exceed the threshold

C Type C goals are metrics with no threshold. The system is asked to optimise these but there is no minimum requirement.

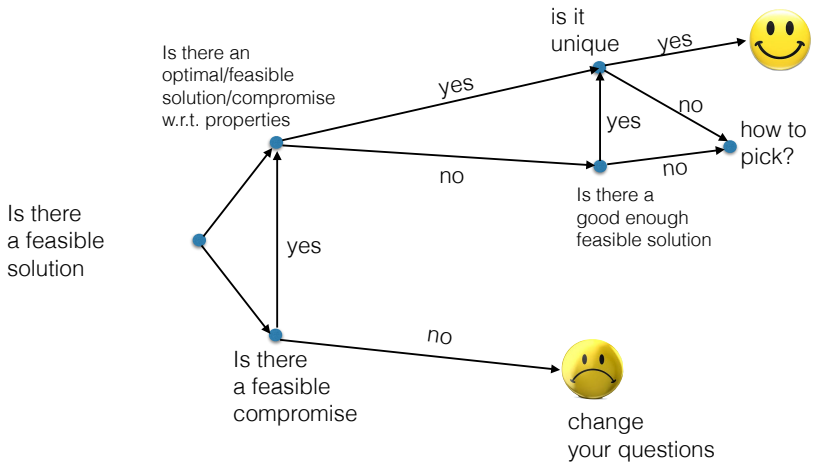

Figure 1 Idealised Coordination Design Flow.

Solutions and Compromises

This terminology let us further refine our definitions of solutions and compromises. We note that:

- A solution achieves all type A and type B goals.
- Given some evaluation function over goals of type B and type C, and optimal solution achieves the maximal value for that function.
- A compromise fails to satisfies all type A and type B goals.
- Givens some evaluation function over goals of type B, type C and over the number-/desirability of goals of type A and type B, an optimal compromise maximises the value of that function.

Coordination Measures

We will refer to goals of type B and type C as *coordination measures*. Typical coordination measures include: Fairness, How much free riding is tolerated, Envy-freeness, Efficiency, Social welfare, Robustness, Scalability, Security, Safety, Computation time.

Coordination Design Flow

Ideally we should have a clear formal framework for assessing coordination problems and answering questions about whether solutions exist and are optimal. This is unlikely to be the case in general but it seems reasonable to suppose that classes of problems could be identified which were amenable to such formal analysis. This would give a design flow something like that shown in Figure 1.

4.1.4 Patterns and Mechanisms

We identify a number of features of existing coordinated multi-agent systems. We separate these into coordination patterns and coordination mechanisms.

Coordination Patterns

Orchestration There is some global entity in the system that computes solutions and directs actions.

Synchronisation/Choreography Groups of entities with the system come together to coordinate their actions.

Cooperation Entities in the system may have different endogenous goals but these facilitate each other – e. g., if one entity wishes a table to be in room A and another wishes it removed from room B then they may cooperate to move the table from room B to room A

Collaboration Entities in the system share the same goals.

Competition Entities in the system compete for resources in order to maximise resource usage/utility/or something.

Coordination Mechanisms

There are a wide range of techniques that have been used to implement coordination in multi-agent systems. We identified the following: Incentives, Negotiation, Argumentation, Norms, Social Choice, Auctions, Operations Research Methods, Roles.

4.1.5 Features

In some cases a designer may have no initial system to work with and will be able to design both the entities within the system as well as the coordination mechanisms. In most cases, however, a coordination problem will include some pre-existing entities and environments which the problem needs to include. These may have certain features.

Features of Entities

When considering the individual entities or agents within a system we identified the following features each of which exists on a spectrum: Compatability of Goals, Reliability, Altruism, Malice.

Features of the Final System

We believe the final system representing a solution to a coordination problem can be placed in a two-dimensional space governed by how much centralisation there is, and how strict the control of the individual entities is. We illustrate that in figure 2.

We have mapped out some common coordination designs in this space. Regimentation is a coordination in which agents are strictly controlled by some centralised authority. Normative systems exist somewhere in the middle of the Strict-Lax axis since agents may ignore norms but will incur sanctions if they do so. These systems may be centralised with some organisation or institution making judgements about violations and imposing sanctions or de-centralised in which individual agents make judgements about the behaviour of other agents and then apply sanctions individually.

While it is difficult to imagine systems in which there is strict de-centralised control, there are systems in which the centralisation may be both static and not strictly-speaking an entity in the system. This would include situations where the built environment passively controls the agent behaviour – e. g., by physically restricting access to certain locations.

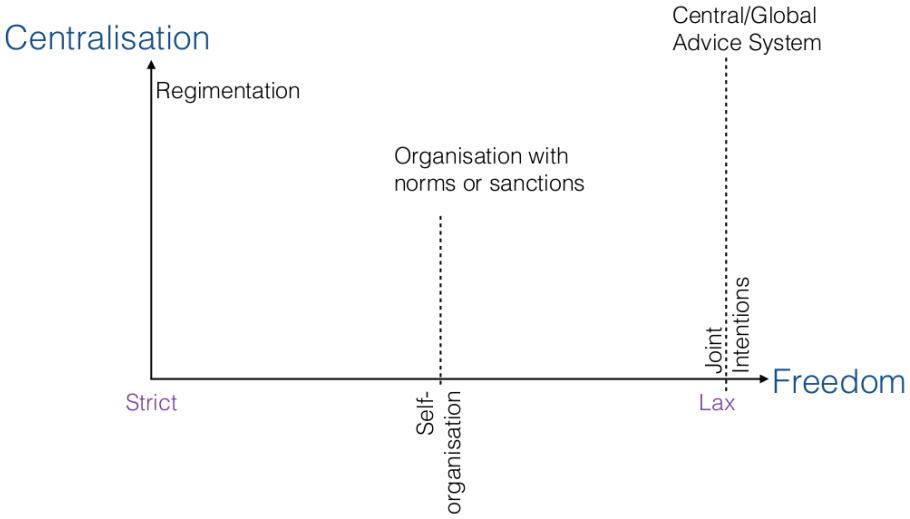

Figure 2 Features of a Coordination.

4.1.6 Examples of Coordination Problems

We conidered two high level examples of significant coordination problems and discuss how they relate to the engineering concepts outlined above.

Traffic Control

Driverless cars have a high profile and strong commercial support eager to bring them to market. The presence of autonomous vehicles, as well as new opportunities to network drivers and traffic systems allows us to consider more sophisticated systems for controlling traffic in place of the current systems which rely on norms, regulation, and minimally networked signalling.

Smart Homes and Smart Grids

Our other example considered the interaction of two coordinations. A *smart grid* manages the flow of electricity around a large network. A *smart home* forms a node in a grid and consists of a number of devices which aim to make the life of the home owner easier and safer.

Comparing and Contrasting the Two Systems

We note

- Humans in the traffic control example accept a far greater loss of autonomy than they do in the smart home. It makes more sense to talk about using norms to regulate human behaviour in one than in the other.
- Although legacy coordination systems exist in both cases, it will be harder to get rid of them entirely in the traffic control system where people have strong beliefs that they are required for safety.
- In the traffic control example there are a large number of different activities that require coordinations to achieve.

- In the traffic control example all entities have similar capabilities (they are all vehicles) while in the smart home the entities all have very different capabilities.
- In the traffic control example we must account for a much higher degree of self-interest (or even malice) from the agents than in the home where all computational entities exist to serve the home owner.

4.1.7 Open Research Questions

Given the framework we have outlined above, we can identify some open research questions.

- What are the techiques for answering the following question?
  Given X pre-existing components for a system, and Y requirements (or more realistically motivations – a formal expression of the requirements) is there a feasible coordination that achieves the requirements and, if so, is it optimal? What happens if we restrict ourselves only to certain patterns or mechanisms for building the solution.
- What are coordination patterns and how do they differ from mechanisms?
- How does monitoring fit in with the scheme?
- We note that coordinations do not exist in isolation. There may be external entities that influence the coordination and which the coordination influences in turn. In some cases these external entities may be additional coordinations. Coordinating a set of coordinations is also a design problem.
  Examples of these are, for instance, the interactions between parliaments and peoples. The people choose (design) the parliament which in turn regulates the behaviour of the people.

4.2 Working group: Formalization of Coordination

Joint work of working group members; written and edited by: Jan Broersen, Marc van Zee, Joshua Sack, Paolo Turrini

4.2.1 Definitions of coordination

The working group came up with several views on the nature of coordination, resulting in different definitions. There was a separation between two camps, roughly described by the following definitions of coordination.

> **Definition 1**: Coordination is the merging of plans of individual agents into a single group plan, independent of any goals on the individual or group level.
> **Definition 2**: Coordination is the aggregation / merging of goals of individual agents into a group goal and the ensuing process of coming to, and implementing, a joint plan to reach the aggregated group goal.

To the surprise of the coordinator of this group (and co-author of this report), support for the first definition was strong. This led some to claim that under that definition just any joint behaviour would count as a coordination. That there is this dichotomy is an interesting phenomenon. We see this in many areas. See the discussion on 'capability theorists' versus 'utility theorists' in economics. Or the dichotomy between 'mentalists' and 'physicalists' in theories of action (e. g. in stit theory).

The working group arrived at the conclusion that the notion of 'protocol' is crucial, because the result of a coordination is best thought of as a protocol. We came, roughly, to the following definition.

Definition 3: A protocol is a restriction on the maximal game tree containing all combinations of free choices of the group of coordinating agents.

Note that the definition of 'protocol' does not go into the issue of the 'character' of the restriction. The restriction may be either such that particular actions are made impossible (regimentation) or they may lead to changes in the utility or goal structures. Again, the difference between these options is a subject of heated debate.

An issue that came up several times is that it is important to distinguish between coordination 'upfront' and coordination 'on the fly'. The definitions are roughly as follows.

Definition 4: Coordination 'upfront' separates the deliberation phase and the execution phase of a coordination: first there is the group deliberation phase and then there is the execution phase where the interaction game is played with the protocol in place.

Definition 5: Coordination 'on the fly' does not separate the deliberation phase and the execution phase of a coordination: during execution, agents make new observations, get new gaols, stumble into impossibilities, etc, etc.

When looking at concrete examples of coordination, it is clear that 'coordination on the fly' is the more realistic conceptualisation. But, as always, to come to a reasonable formalisation, it makes sense to simplify reality and focus on idealised scenarios where the execution phase and deliberation phase are treated separately.

4.2.2 A concrete example

In order to understand and illustrate the many ways in which a multi-agent system is able to coordinate we will consider a football team (for the American reader: a soccer team) that is engaged in playing a football match. We will explain the coordination process within the football team as a layered process consisting of the following three levels:

3. Normative powers (obligations, permissions, counts as)
2. Collective mental attitudes (collective beliefs, intentions, goals)
1. Individual mental attitudes (individual beliefs, intentions, goals),

where the individual agents each contain their own doxastic and motivational attitudes, i. e. their beliefs, goals, intentions, preferences, etc. These attitudes may depend on each other in different ways, and changing one of them may cause a change in others. For instance, if an agent believes that it is impossible to reach his goal to receive a ball that has been played by another team member, the agent may decide to pursue another goal such as defending.

We can view coordination in a multiagent system as a coordination process of these different mental attitudes, giving rise to higher-level group attitudes such as collective beliefs, collective intentions, and collective goals. There is still much debate whether these higher-level attitudes can in some way be reduced to the individual attitudes[1].

Going up another layer in the coordination hierarchy, we can say that agents do not just form collective intentions and beliefs, they also decide to follow the rules of the game. For instance, they agree as a group that the man in blak counts as a referee, and that playing the ball to a player of the own team, when there is no more than one player of the other team between him and the enemy goal, counts as offside, which is penalized with a free kick. Thus, speaking in John Searle's terms, a group of agent is able to designate "status functions" to objects, giving them a certain "deontic power". Note that these deontic powers are only

[1] There is a good Stanford encyclopedia article about this called "Collective Intentionality".

there because the group has chosen to accept them. Another example of a concept that we as humans have given a special status is money. It is interesting to note that most of the money does not even exist. It is merely a concept that exist because we want it, but there is no physical representation of it.

How does this affect choice of action of an agent? We can say that there are two kinds of reasons for actions, namely desire-dependent reasons for actions and desire independent reasons for action. The desire dependent reasons for action come from within the agent. For instance, a football player aims to score a goal because of his desire to be celebrated by his team. Desire independent reasons come from the deontic powers, and explain why we sometimes do things we don't want. For instance, a defender might not even like to play left behind and would rather play as a striker, but since he counts as a good defender, he has to play there anyway.

4.2.3 Towards formalization

Most contributors within the working group thought of 'Alternating Time Temporal Logic [1] / Coalition Logic [10] / Strategic *stit* Logic [3]' as the typical formal setting in which to study coordination. However, a drawback of the ontological commitments made by these logic is maybe that they do not easily connect to logics of programs or logics of action that depart from an event-based perspective (Dynamic Logic [6], Situation Calculus [8], etc.). To accommodate this criticism, we may look at the following alternatives:

- dynamic logic for strategies and concurrent processes [2, 11, 12]
- strategy logic (making strategies explicit in the object language) [7]
- interpreted systems (the aim of interpreted systems is to ground abstract logical models as runs of concrete systems) [4, 5]
- protocols as sets of strategies [9]

We can extract several important ingredients necessary to model coordination from the following example. Consider two agents, Ann and Bob, who are choosing a dress for a party. They can only choose between a black suit and a white suit. They both prefer to wear a suit of the same color, but they have no way to know what the other has chosen. This situation is common knowledge among the two.

Necessary ingredients

- There must be agents that can perform different actions independently and concurrently. Independently means that their choices of action cannot interfere in the choice of actions of other agents. Concurrently means that agents don't know, before making their choice, what choices the others have made. Note that concurrently is not the same as simultaneously, which means that actions must happen at the same time.
- There must be desired outcomes (i. e., A& B choose the same suit) and undesired ones. In particular, each agent must have an action that is not a sure success, but only a success conditional to an action of another agent. All actions by Ann and Bob are not a sure success (success being choosing the same suit).

Unneccesary ingredients. These ingredients are not necessary to model coordination and we can abstract away from them.

- It is not necessary to have a strict linear order on the outcomes, nor a numerical utility function, but only a dichotomous (success failure) relation.

- It is not necessary to have paths, time lines etc. but only a dichotomous (now, later) one. Clearly time adds notions such as interleaving, but every extensive structure can be translated into a normal form one (plus information sets), therefore we can do with the latter.
- It is not necessary to have more than two agents. Each problem involving more agents can be simulated.
- It is not necessary to have more than two actions. Also because each problem involving more agents can be simulated.
- It is not necessary to have probabilistic beliefs. They don't add anything to knowledge relations in modelling the problem.

The big claim is that each coordination problem can be ultimately described with the ingredients used above and no more.

4.2.4 Conclusion

The above insights are the result of 4 days of discussion and later consideration and rationalisations by the four authors. We think they form a sound basis for further study.

The discussions during the workshop tended to focus on rather basic differences of a philosophical nature, which hampered progress. This was somewhat unexpected. The joint goal of the workshop contributors is to come to formalisation and conceptualisations of coordination in Multi–Agent Systems. To reach this goal it would be good if for the next edition of the workshop we find a format (indeed: a coordination) that prevents discussions on preliminaries that are too long.

References

1 R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, 2002.
2 J. F. A. K van Benthem. In praise of strategies. Unpublished manuscript, 2007.
3 Jan Broersen. CTL.STIT: enhancing ATL to express important multi-agent system verification properties. In *Proceedings of the ninth international joint conference on Autonomous agents and multiagent systems (AAMAS '10)*, New York, NY, USA, 2010. ACM.
4 R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning about Knowledge.* The MIT Press, 1995.
5 J. Y. Halpern and R. Fagin. Modelling knowledge and action in distributed systems. *Distributed Computing*, 3(3):159–177, 1989.
6 D. Harel, D. Kozen, and J. Tiuryn. *Dynamic Logic.* The MIT Press, 2000.
7 Thomas A. Henzinger Krishnendu Chatterjee and Nir Piterman. Strategy logic. In *Proceedings of the 18th International Conference on Concurrency Theory (CONCUR)*, Lecture Notes in Computer Science 4703, pages 59–73. Springer, 2007.
8 J. McCarthy and P. Hayes. Some philosophical problems from the standpoint of artificial intelligence. In B. Meltzer and D. Michie, editors, *Machine Intelligence*, volume 4, pages 463–502. Edinburgh University Press, 1969.
9 Rohit Parikh and Ramaswamy Ramanujam. A knowledge based semantics of messages. *Journal of Logic, Language and Information*, 12(4):453–467, September 2003.
10 Marc Pauly. A modal logic for coalitional power in games. *Journal of Logic and Computation*, 12(1):149–166, 2002.
11 D. Peleg. Communication in concurrent dynamic logic. *Journal of Computer and System Sciences*, 35:23–58, 1987.
12 Jan van Eijck. PDL as a multi-agent strategy logic. In *Proceedings TARK 2013*, 2013.

5 Open Problems

The seminar aimed at bringing together researchers from (at least) three disciplines – logic, game theory and agreement technologies – to discuss and solve (formal) problems related to coordination. As expected not all problems could have been solved and new open problems have emerged. Some are listed below.

- What are the techniques for answering the following question?
  Given X pre-existing components for a system, and Y requirements (or more realistically motivations – a formal expression of the requirements) is there a feasible coordination that achieves the requirements and, if so, is it optimal? What happens if we restrict ourselves only to certain patterns or mechanisms for building the solution.
- What are coordination patterns and how do they differ from mechanisms?
- How does monitoring fit in with the scheme?
- We note that coordinations do not exist in isolation. There may be external entities that influence the coordination and which the coordination influences in turn. In some cases these external entities may be additional coordinations. Coordinating a set of coordinations is also a design problem.
- What are interesting coordination properties that it would be useful to express in a formal logic?
- How suitable are different existing logical formalisms (e. g., coalition logic or PDL) for expressing interesting coordination properties?
- How should a general and flexible model of coordination look like, if there is one at all?
- Should coordination always be conscious or can it just emerge incidentally?

Participants

- Thomas Ågotnes
University of Bergen, NO
- Markus Brill
Duke University, US
- Jan M. Broersen
Utrecht University, NL
- Nils Bulling
TU Clausthal, DE
- Natalia Criado
John Moores University – Liverpool, GB
- Mehdi Dastani
Utrecht University, NL
- Louise A. Dennis
University of Liverpool, GB
- Catalin Dima
University Paris-Est – Créteil, FR
- Valentin Goranko
Technical Univ. of Denmark – Lyngby, DK
- Maaike Harbers
TU Delft, NL
- Andreas Herzig
Paul Sabatier University – Toulouse, FR
- Max Knobbout
Utrecht University, NL
- Ho-Pun Lam
NICTA – Brisbane, AU
- Beishui Liao
Zhejiang University, CN
- Brian Logan
University of Nottingham, GB
- Marin Lujak
University Rey Juan Carlos, ES
- Nir Oren
University of Aberdeen, GB
- Julian Padget
University of Bath, GB
- Truls Pedersen
University of Bergen, NO
- Joshua Sack
University of Amsterdam, NL
- Marija Slavkovik
University of Bergen, NO
- Bas J. G. Testerink
Utrecht University, NL
- Nicolas Troquard
National Research Council – Trento, IT
- Paolo Turrini
Imperial College London, GB
- Birna Van Riemsdijk
TU Delft, NL
- Marc Van Zee
University of Luxembourg, LU
- Laurent Vercouter
INSA – St-Etienne-du-Rouvray, FR

Report from Dagstuhl Seminar 14341

Resource-bounded Problem Solving

Edited by

Yll Haxhimusa[1], Iris van Rooij[2], Sashank Varma[3], and H. Todd Wareham[4]

1 **Vienna University of Technology, AT,** yll@prip.tuwien.ac.at
2 **Radboud University Nijmegen, NL,** i.vanrooij@donders.ru.nl
3 **University of Minnesota, US,** sashank@umn.edu
4 **Memorial Univiversity of Newfoundland, CA,** harold@mun.ca

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 14341 'Resource-bounded Problem Solving'. This seminar is a successor to Dagstuhl Seminar 11351: 'Computer Science & Problem Solving: New Foundations', held in August 2011, which was the first Dagstuhl event to bring together computer scientists and psychologists to exchange perspectives on problem solving in general. The current seminar built on the previous seminar by (1) narrowing the focus to issues related to resource-bounded problem solving and (2) broadening the range of perspectives on the specific topic by including a balanced number of attendees with expertise in computer science, psychology, artificial intelligence, and cognitive neuroscience.

Seminar August 17–22, 2014 – http://www.dagstuhl.de/14341
1998 ACM Subject Classification F.1.1 Models of Computation, F.1.3 Complexity Measures and Classes, F.2.0 [Analysis of Algorithms and Problem Complexity] General, F.2.3 Tradeoffs Between Complexity Measures, I.2.0 [Artificial Intelligence] General, I.2.6 Learning, I.2.4 Knowledge Representation Formalisms and Methods, I.2.8 Problem Solving, Control Methods, and Search, J.4 Social and Behavioral Sciences
Keywords and phrases complexity theory, problem solving, cognitive psychology, computational trade-offs
Digital Object Identifier 10.4230/DagRep.4.8.45

1 Executive Summary

Yll Haxhimusa
Iris van Rooij
Sashank Varma
H. Todd Wareham

This Dagstuhl Seminar on 'Resource-bounded Problem Solving' was a successor to Dagstuhl Seminar 11351: 'Computer Science & Problem Solving: New Foundations', held in August 2011, which was the first Dagstuhl event to bring together computer scientists and psychologists to exchange perspectives on problem solving in general. Before summarizing the content of the seminar itself, we describe the theoretical motivations for the topic of 'Resource-bounded Problem Solving', and the choice for the interdisciplinary composition of participants, ranging complexity theory, cognitive psychology, artificial intelligence and cognitive neuroscience.

Resource-bounded Problem Solving, *Dagstuhl Reports*, Vol. 4, Issue 8, pp. 45–72
Editors: Yll Haxhimusa, Iris van Rooij, Sashank Varma, and H. Todd Wareham

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

Background and Motivation

Problem solving, whether by humans or machines, is bounded by the resources at hand. For machines, these resources fundamentally include hardware and processing speed. For humans, important resources also include mental representations, memory capacities, inferential capacities, and time to name a few. All these resources can be severely limited, by constraints imposed by both the implementing hardware (current technology in the case of machines, the organization of our brains in the case of humans) and the physical and social operating environments.

The study of resource-bounded problem solving has a long and productive history within computer science, which has resulted in a number of efficient exact-solution algorithms and algorithm design techniques as well as, where such algorithms are not possible, various widely-applicable approximate-solution heuristics. Given that brains have evolved to solve problems under severe resource constraints, resource-bounded problem solving may also provide one of the best windows on the organization of cognitive brain function. Trying to model exactly how humans solve really hard—that is, resource demanding—problems efficiently seems a good strategy from the perspective of deriving predictive and explanatory cognitive models. After all, many cognitive models may be able to match human performance on really easy problems, but only a few can for really hard problems. Hence, if one finds even one cognitive model that can solve a hard problem as well as humans one can be much more confident that it captures fundamental principles of human cognition.

What makes tasks resource demanding? Here computational complexity theory is a key source of information for the study of human and machine problem solving. Computational complexity theory studies the intrinsic resource demands of various computational problems. It allows us to assess when resource demands are low, reasonable, high, or impractically high. Though the importance of computational complexity has been recognized in computer science for decades, it has been underutilized to date by cognitive scientists. This is not for want of opportunities, as cases are known where cognitive scientists have studied principles of resource-bounded problem solving in apparent ignorance of relevant computational complexity results. The following example illustrates such a situation.

- Solving constrained Traveling Salesman problems: MacGregor and Ormerod (1996,Attention, Perception & Psychophysics) hypothesized that the difficulty of solving Euclidean versions of the Traveling Salesperson problem (E-TSP) may be more a function of the number of inner points (i. e., the points not lying on the convex hull of the point set) than of the total number of points. This hypothesis was tested and confirmed for human subjects solving pen-and-paper instances of E-TSP. Independently, it was shown that E-TSP is fixed-parameter tractable when the number of inner points is the parameter (Deineko et al.COCOON 2004). In other words, it is possible to solve E-TSP in time $f(k)n^c$, where fs is a non-polynomial function of the number k of inner points, n is the total number of points, and c is a constant. This result is relevant for explaining the findings of MacGregor and Ormerod (1996) as it gives a computational formalization of their hypothesis.

There may be many other such opportunities waiting to be noticed. There is also evidence that when such opportunities are exploited and cognitive scientists and complexity theorists establish collaborations, these collaborations can yield novel perspectives and approaches. Below are two examples of such ongoing collaborations and their products.

- Analyzing resource demands of cognitive models: Using computational complexity concepts and techniques, psychologists can systematically study how human problem solving

proceeds under various resource demands. Also, complexity theory can predict how resource demands scale as a function of a problem's parameters. Psychologists can then in turn use these predictions to test the models being studied. This approach has been successfully implemented by Iris van Rooij, Todd Wareham, and others in a wide variety of domains, including decision-making (2005, Journal of Mathematical Psychology) and analogical problem solving (2011, Journal of Problem Solving). This program has also led to the development of a theory of structure approximability which has produced new results within both computer science (2007, Proceedings of Dagstuhl seminar07281) and psychology (2012, Journal of Mathematical Psychology).

- Pyramid data structures and efficient search: Humans and animals are able to delineate, detect and recognize objects in complex scenes at a blink of an eye. Tsotsos (1990, Behavioral and Brain Sciences) performed a complexity analysis and showed that hierarchical representation of visual information and hierarchical processing of this information is one of the best, if not the best, way for brains to solve visual problems. Pyramid data structures provide an effective model for efficient hierarchical search of the problem space. This perspective has led to fruitful collaborations between Yll Haxhimusa, Zygmunt Pizlo, Walter Kropatsch and others, yielding new algorithmic techniques in computer vision (2005, Pattern Recognition Letters; 2009, Vision & Computing),as well as inspiring new cognitive models of visual problem solving in psychology (2006 and 2011, Journal of Problem Solving).

With this seminar we aimed to actively stimulate the exchange of ideas and results between computational complexity theorists, psychologists, cognitive neuroscientists, and AI-researchers studying problem solving. In particular, we wanted to ensure that this exchange would be of use to each (and not just one) of these communities. We believe that such n-way productivity is crucial to fruitful long-term interdisciplinary collaboration, in that it encourages the continued interest of members of all communities in collaborating.

Organization of seminar

On Day 1 of the seminar, Iris van Rooij opened the seminar by explaining its history, motivation and aims. This was following by a round of introductions, in which each participant introduced themselves, who they are, what their home disciplines were, what their relevant research interests were, and what they hoped to both bring to the seminar and get out off it.

Given the interdisciplinary nature of the questions of interest and the wide range of expertise of the seminar participants, it was crucial that a common understanding of the different goals and assumptions of the disciplines involved at the meeting be established early in the meeting. To this end, the first full day of the seminar was devoted to primers on basic terminology, assumptions, and goals of four major sub disciplines represented at the seminar (namely, computational complexity theory, artificial intelligence, psychology, and cognitive neuroscience).

The four introductory keynote speakers, Zygmunt Pizlo (Cognitive Psychology), Todd Wareham (Computational Complexity Theory), Rineke Verbrugge (Artificial Intelligence), and David Noelle (Cognitive Neuroscience), all addressed each of the following questions for their own respective fields.

- What are the goals of that discipline?
- What are the techniques used in that discipline?

- What do the terms “problem solving” and “resource bounds” mean in that discipline?
- What does that discipline have to offer to other disciplines in the context of this seminar?
- What does that discipline want from the other disciplines in the context of this seminar?
- What are some tentative research questions and collaboration opportunities?

Following these introductory key notes there was a panel discussion. Day 1 closed with a Town Hall meeting, in which the set-up and organization for the next days was discussed with all participants and a preliminary schedule was established (later on, as needed, this schedule was updated). At the Town Hall meeting the concept of Birds-of-a-Feather (B.O.F.) sessions was also explained (see Section 5), which turned out to be a very successful format for self-organized working groups.

Days 2, 3 and 4 involved a mix of talks, posters, B.O.F. sessions, and Town Hall sessions. Pairs or triplets of talks were followed by panel discussions to stimulate cross talk connections. Poster sessions allowed for more in-depth discussion in an informal setting, and B.O.F. sessions allowed people to gather and discuss more specific topics of common interest. At Town Hall sessions, a plenary report on the insights gained from each B.O.F. session was given, so that all participants were kept up to date of the outcome of such events. All in all, this set-up worked very well, stimulating active exchange and discussion between participants that crossed disciplinary boundaries.

On the morning of Day 5 we closed with reflections on the overall organization and content of the seminar. It was a shared perspective among participants that the seminar had been exceptionally successful in bringing together the different fields involved in the seminar, initiating many first-time theoretical exchanges and conceptualizations of possible common research questions. Many participants indicated that following this seminar, they would be interested in more focused seminars specializing in subtopics within the domain of problem solving or specializing in specific modeling techniques.

2 Table of Contents

3 Overview of Talks

3.1 Where is the chocolate?: Modeling Development of Reasoning about False Beliefs of Others

Burcu Arslan (University of Groningen, NL)

Reasoning about false beliefs of others develops with age. We constructed a computational cognitive model in order to show the developmental transitions. These start from a child's reasoning from his/her own point of view (zero-order) to taking into consideration another agent's beliefs (first-order), and later to taking into consideration another agent's beliefs about again other agents' beliefs (second-order). The model is based on a combination of rule-based and simulation approaches. We modeled the gradual development of reasoning about false beliefs of others by using activation of declarative knowledge instead of utility learning. Initially, in addition to the story facts, there is only one strategy chunk, namely a zero-order reasoning chunk, in declarative memory. The model retrieves this chunk each time it has to solve a problem. Based on the feedback, the model will strengthen a successful strategy chunk, or it will add or strengthen an alternative strategy if the current one failed.

3.2 When Thinking Never Comes to a Halt: Tractability and Approximability in AI

Tarek R. Besold (Universität Osnabrück, DE)

Joint work of Besold, Tarek R.; Robere, Robert
Main reference T. R. Besold, R. Robere, "When Thinking Never Comes to a Halt: Using Formal Methods in Making Sure Your AI Gets the Job Done Good Enough," in V. C. Mueller, editor, Fundamental Issues of Artificial Intelligence (Synthese Library), Springer, to appear.

A growing number of researchers in Cognitive Science advocate the thesis that human cognitive capacities are constrained by computational tractability. If right, this thesis also can be expected to have far-reaching consequences for work in Artificial General Intelligence: Models and systems considered as basis for the development of general cognitive architectures with human-like performance would also have to comply with tractability constraints, making in- depth complexity theoretic analysis a necessary and important part of the standard research and development cycle already from a rather early stage. We present an application case study for such an analysis based on results from a parameterized complexity and approximation theoretic analysis of the Heuristic Driven Theory Projection (HDTP) analogy-making framework.

3.3 The Challenge of Optimization

Sarah Carruthers (University of Victoria, CA)

How do humans cope with complexity? One approach that can assist our understanding in this area is to study human performance on hard computational problems. Studies to date have focused on the optimization version of a handful of problems, including the Traveling Salesperson Problem, Minimum Vertex Cover, and n-Puzzle. In this work, we show that for these hard optimization problems, participants' internal representations cannot be consistent with the given problem. Identifying how internal representations are generated under these conditions presents an opportunity for new interpretation of results from previous studies. With this in mind, we propose alternative ways of presenting hard computational problems such that participants' internal representations can be consistent with the intended task. Finally we present study results, which support our hypothesis that the internal representation generation differs between optimization and other versions of two hard computational problems.

3.4 Brain-Wiring Optimization: Lessons for Cognitive Modeling

Christopher Cherniak (University of Maryland – College Park, US)

A (a) bounded-resource rationality framework links together independent research programs in (b) cognitive psychology and (c) computer science. Neural wiring minimization sometimes appears either virtually perfect, or as good as can be currently detected (vs only moderately well-engineered). – A predictive success story. Quick & dirty heuristics for such neuroanatomy optimization may be germane to problem solving procedures:

1. Neural network minimization 'for free, directly from physics'. Via exploiting anomalies of the mathematical / computational order / the optimization landscape.
   a. Evading exhaustive search of alternative topologies, when 'topology does not matter'.
   b. Avoiding local minima traps in energy-minimization layout processes.
   c. Finetuning a fast heuristic to match only best cases of wirecost minimization.
2. Detecting large-scale system optimization indirectly, via a 'Size Law': The larger the proportion of a complete optimal system that an evaluated sub-system is, the better the optimization of the subset.

3.5 Ifs and Ands and Ors

Nicole Cruz de Echeverria Loebell(University of London, GB)

Joint work of Loebell, Nicole Cruz de Echeverria ; Baratgin, Jean; Feeney, Aidan; Oaksford, Mike; Over, David

The probabilistic approach in the study of reasoning generalizes binary logic to cover the correctness of deductive inferences from uncertain premises. In this new framework, the

concept of binary validity is extended to that of probabilistic validity, p-validity, and the concept of binary consistency is extended to that of coherence. Many results in judgment and decision making show that people are not always coherent in their probability judgments. Our question was whether they are more coherent, and conform more to p-validity, when they make explicit deductive inferences. This question was investigated in an experiment using conditionals and disjunctions. Participants gave confidence judgments about a list of inferences (inference group) or about the statements these inferences were composed of (statements group) based on information about a character from a short context story. Overall participants' probability judgments conformed to p-validity and coherence at above chance levels. They conformed less often to p-validity for binary valid, but p-invalid inferences, and they conformed more often to p-validity and coherence in the inference group than in the statements group. The results provide support for the psychological relevance of the normative standards of p-validity and coherence. And they suggest that conformance to these standards is not automatic, but increases in the context of an explicit inference task. The possibility of reducing the complexity of computing coherence when engaged in an inference task through the manipulation of graphic representations of logical partitions is discussed.

3.6 Emergence of Complex, Problem-Solving Cognitive Structure

Liane Gabora (University of British Columbia – Vancouver, CA)

How does a problem solving cognitive architecture come about? Humans are unique in extent to which we take constraints and resources at hand into account by adapting existing problem solving methods to new situations to meet own goals, preferences. This requires more than heuristics. It requires more than assimilation (reframe information in terms of what you already know) and accommodation (revise what you know to account for new information). It requires that representations be iteratively restructured by re-viewing them from different, relevant perspectives until they achieve a form that solves problems. This in turn requires integrated, self-modifying conceptual structure. In the talk I discuss how such a problem-solving structure comes about.

3.7 An Analytic Tableaux Model for Deductive Mastermind Empirically Tested with a Massively Used Online Learning System

Nina Gierasimczuk (University of Amsterdam, NL)

The paper is concerned with the psychological relevance of a logical model for deductive reasoning. We propose a new way to analyze logical reasoning in a deductive version of the Mastermind game implemented within a popular Dutch online educational learning system (Math Garden). Our main goal is to derive predictions about the difficulty of Deductive Mastermind tasks. By means of a logical analysis we derive the number of steps needed for solving these tasks (a proxy for working memory load). Our model is based on the analytic tableaux method, known from proof theory. We associate the difficulty of Deductive

Mastermind game-items with the size of the corresponding logical trees obtained by the tableaux method. We derive empirical hypotheses from this model. A large group of students (over 37 thousand children, 5–12 years of age) played the Deductive Mastermind game, which gave empirical difficulty ratings of all 321 game-items. The results show that our logical approach predicts these item ratings well, which supports the psychological relevance of our model.

3.8 Real-World Problem-Solving: Back to the Future

Vinod Goel (York University – Toronto, CA)

The machinery of Newell and Simon's information processing theory has been the basis cognitive theory for the past 50 years. It has been nurtured on toy or well structured problems and plagued by doubt about its ability to scale up to explain real-world problems. I introduce some data from the neuropsychology literature that reinforces the distinction between toy problems and real world problems and further sheds doubt on the ability of information processing theory to explain significant aspects of human cognition.

3.9 Coordinated Reasoning

Justine Jacot (Lund University, SE)

Logical reasoning tasks have focused on testing the way lay people use some logical skills – be it an innate ability or an acquired competence – often without questioning the nature of the reasoning itself, taking for granted that the 'logic' used is classical logic and that reasoning equates with being rational. A failure to complete reasoning tasks is thus either interpreted as a failure of logical skills or a failure of rationality. However, if one agrees with Stenning and Van Lambalgen (2008) that people must reason 'to' an interpretation before reasoning 'from' an interpretation when performing a reasoning task, coordination on the meaning of instructions between experimenters and subjects is not only a precondition to drawing conclusions about reasoning (logical or otherwise), but also a requirement all along the performing of the tasks. Moreover, when those tasks are formulated in natural language, pragmatic constraints apply, identified as 'intuitions' of relevance by Relevance Theory (Sperber & Wilson 1995), which takes them as an argument that experimental tasks do not say much about the theory behind the reasoning competence of subjects. Although I agree with Relevance Theory on the fact that (un)successful performance in reasoning tasks is compatible with different theories of reasoning, the explanation they provide in terms of intuitions of relevance does not help much in understanding how people reason, nor does it allow to understand how experimenters assess the (lack of) success of subjects. I propose a simple model of experimental reasoning tasks in terms of signaling games, which highlights the need for coordination between experimenter and subject throughout the task, as well as the difficulty for experimenter to evaluate properly the performance of subject.

3.10 Abductive Reasoning Using Random Examples

Brendan Juba (Harvard University, US)

In 'abductive reasoning,' a known or desired condition is given, and one seeks to find plausible (but not necessarily entailed) premises that imply the given condition. Several key cognitive processes, notably including (sub)goal formation, can be viewed as instances of abductive reasoning problems. I propose a new model of abductive reasoning as searching for a conditional distribution using random examples, in the same spirit as Valiant's PAC-learning model. In contrast to previous probabilistic models of abductive reasoning, this model does not rely on a prior distribution over the possible conditions to indicate which conditions are more plausible.

Much like PAC-learning, this simple model is well suited to the analysis of algorithms. As a consequence, the model suggests an interesting picture of which representations can be found efficiently, and which are computationally intractable. In particular, k-DNF representations (for small k) can be found efficiently in this model, but recent results suggest that an efficient algorithm for finding conjunctive representations (or any stronger representation) may not exist.

3.11 Prefrontal Function, Reasoning and Adaptive Behavior

Etienne Koechlin (ENP – Paris, FR)

The prefrontal cortex subserves executive control and decision-making, i. e. the coordination and selection of thoughts and actions in the service of adaptive behavior. We present a computational theory describing the evolution of the prefrontal cortex in humans as gradually adding new inferential Bayesian capabilities for dealing with a computationally intractable decision problem: exploring and learning new behavioral strategies vs. exploiting and adjusting previously learned ones through reinforcement learning. The theory clarifies the integration of model-free and model-based reinforcement learning through the notion of strategy creation. The theory also shows that counterfactual inferences in humans yield to the notion of hypothesis testing, a critical reasoning ability for approximating optimal adaptive processes and presumably endowing humans with a qualitative evolutionary advantage in adaptive behavior. We also present recent empirical data from behavioral and neuroimaging experiments conducted in our lab supporting the proposed theory.

3.12 Computing Probabilities and Probabilistic Computations – A Primer

Johan H. P. Kwisthout (Radboud University Nijmegen, NL)

In this presentation, I elaborate on the difference (complexity-wise) between 'computing probabilities' and 'probabilistic computations' and show how and why they are distinct. I show that NP-hard problems cannot be approximated both efficiently (in polynomial time) and effectively (with small probability of error) using such a probabilistic computation unless $NP \subseteq BPP$.

3.13 Complexity of Computations in Spiking Neural Networks

Wolfgang Maass (TU Graz, AT))

Spike-based computations promise to be substantially more power-efficient than traditional clocked processing schemes. However it turned out to be surprisingly difficult to design networks of spiking neurons that are able to carry out demanding computations. We present a new method for organizing such networks out of simple stereotypical network motifs in such a way that they can solve hard constraint satisfaction problems from the domains of planning / optimization and verification / logical inference. We use here noise as a computational resource, in spite of the fact that the timing of spikes (rather than just spike rates) are essential for the resulting computations. This new organization scheme is supported by a new theoretical understanding of spike-based stochastic computations. Surprisingly, one can identify in this context also a concrete computational advantage of spiking networks compared with traditional non-spiking stochastic neural networks (Boltzmann machines).

3.14 Cognitive Neuroscience: Foundations and Insights for Problem Solving

David Noelle (University of California – Merced, US)

This presentation provides a general introduction to the field of cognitive neuroscience, including its goals, its primary methods, and its major findings. The domain of 'problem solving' is discussed from a cognitive neuroscience perspective, focusing on cognitive control and the prefrontal cortex. A variety of 'resource bounds' on neural computation are introduced. Finally, opportunities for cross-fertilization between the fields of computational complexity, artificial intelligence, cognitive psychology, and cognitive neuroscience are suggested, using as an example some recent work involving how anatomical connection patterns in the prefrontal cortex naturally give rise to a mechanism that resembles 'pointers' in computer science, supporting novel compositional knowledge structures in the brain.

3.15 From Channel Capacity to Semantic Relevance: Accurate Retrieval as the Ultimate Boundary on Complex Cognition

Stellan Ohlsson (University of Illinois at Chicago, US)

Inspired by work on channel capacity in statistical information theory, the first wave of cognitive psychology emphasized limits on transient storage (short-term memory, working memory). Digital computers focused computer scientists on boundaries on long-term memory storage and cpu processing speed. Processing limits of these sorts put boundaries on the amount of heuristic search an intelligent agent could engage in, and hence which problems it can solve. To overcome this boundary, intelligent agents need large amounts of knowledge to focus search on a few promising search paths. This focus in turn implies that learning rate is the limiting factor: It takes 10 years or more to acquire the required knowledge base. This, in turn, focus attention on the retrieval of the knowledge that is needed at each moment in time. J. R. Anderson has argued that because the world is unpredictable, retrieval cannot, even, in principle, be 100% accurate. The boundary on cognition is thus the unpredictable nature of reality. In this talk, I'll take the focus on retrieval one step further, and argue that the ultimate boundary is the need for a massively parallel relevance calculation that is independent of experience. This limit might be ultimate in the sense that there is no way to overcome it.

3.16 Psychology of Problem Solving

Zygmunt Pizlo (Purdue University – West Lafayette, US)

The talk begins with a brief discussion of problem solving as a goal vs. a task to achieve some goal. Visual integration of a contour is an example of the latter. In the presence of noise in the retinal image, this problem is computationally intractable. But, if natural constraints are used, such as smoothness, proximity, convexity and closure, this problem becomes tractable. Specifically, it can be solved as a shortest path problem in $n \log n$ time using the cortical representation known as a log-polar mapping. The second part of the talk discusses a recent model of how humans produce near-optimal TSP tours using working memory that can hold only a few pieces of information at a time. This is an illustration of problem solving as a goal in itself. The new model is an elaboration of a multiresolution-multiscale pyramid algorithm whose computational architecture is based on the known anatomical and physiological characteristics of the human visual system. The computational complexity of the model is linear and its working memory stores at any given time information about 2-5 clusters. The model has visual attention that is moved the way humans move their eyes. Performance of this model matches closely performance of human subjects.

3.17 Problem Solving: Representing the (right) Computational Problem

Ulrike Stege (University of Victoria, CA)

When problem solving computational optimization problems with the goal to understand how good human performance is when solving instances of the problem or understanding their problem solving strategy in general, it can be crucial to describe the task in a way that is easy to understand and that does not hinder human performance.We give examples of different computation problems that are – in a sense – equivalent. We further discuss different ways of parameterized problems using equivalent problem definitions. This can be particularly informative for NP-hard problems as different parameterization can result in different parameterized complexity (FPT versus W[1]). We then discuss kernelization, a powerful technique for fixed-parameter tractable problems. In the second part of the talk (joint work with Balasubramanian, Muller and Srinivasan) we discuss the problem of reverse engineering definition of computational problems when the problem solving strategy is known. Using the example of the greedy technique, we discuss how tweaking the constraints or objective function of an optimization problem can yield quality or optimum solution by introducing structure to the problem. We further explain the analogy of algorithm w.r.t. to the quality of their solution with policies employed, e. g., in autonomic systems.

3.18 Towards a Theory of Learning Problem-solving Skills

Niels A. Taatgen (University of Groningen, NL)

A general theory of human problem solving that everyone starts out with a general set of problem solving algorithms that are specialized for specific tasks when necessary. This idea originated from Newell and Simon, and has also been instantiated by Anderson in ACT* and later ACT-R. However, there is a problem with this theory: it cannot deal with the fact that people are generally able to deal with problems that are computationally intractable. Therefore, the general problem solving algorithms cannot 'solve' them, and therefore cannot lead to appropriate specialized skills. The solution is to assume that humans generally build up general problem solving skills while solving specific problems. These general skills are adapted to the specific needs of an individual, and can therefore, after longer periods of time, lead to the ability to solve increasingly complex instances of a particular intractable problem, with possible transfer to other intractable problems. The new PRIMs cognitive architecture (Taatgen, 2013) has been developed to learn new general skills, including problem-solving skills. Although it has not been applied explicitly to problem solving yet, it has been used for mathematics, working memory, control and theory of mind.

3.19 Complexity, Attention, Learning and Goldilocks

John K. Tsotsos (York University – Toronto, CA)

This presentation begins by overviewing results aimed at illuminating the complexity complexity of visual processing. These results are then used to steer development of a theory of visual attention, Selective Tuning, to ensure a tractable formulation. The theory has made many predictions of how the human brain attends to visual stimuli, which many years after, are now strongly supported by experiment. One remaining problem, of many, was highlighted, namely, how to learn the parameters of the model. A hypothesis was presented on this, specifically that we consider how each mechanism within the theory might be improved via training, thus leading to overall model improvements. A pilot experiment was shown to confirm that human visual attention may be trained in such a manner. The tie to the complexity reduction of each mechanism via training was presented. Finally, the overall optimality of the process was discussed, with the conclusion that the brain is not optimal. Rather, it exhibits Goldilocks intelligence – it behaves in 'just the right manner' for organisms to be successful.

3.20 Higher-order theory of mind in mixed-motive settings: An agent-based simulation study

Harmen de Weerd (University of Groningen, NL)

In daily life, people often encounter situations in which they have little to no data on how others will react. To solve this problem, people use theory of mind by reasoning about the beliefs, goals and intentions of others. This allows them to predict the behaviour of others in new situations, as well as interpret observed behaviour. People can also use this ability recursively: they use higher-order theory of mind to reason about the theory of mind abilities of others, as in 'he thinks that I don't know that he sent me an anonymous letter'. However, humans appear to be the only species to do so. In this poster, we investigate whether exposure to mixed-motive situations, which involve both cooperative and competitive elements, can explain the emergence of human-like theory of mind abilities. We focus on repeated one-shot negotiations within the setting of the Colored Trails game. Simulation results show that in this setting, both first-order and second-order theory of mind agents outperform agents that are more limited in their theory of mind abilities. Moreover, the presence of such first-order and second-order theory of mind agents also benefits other agents. Agents did not obtain any additional benefit from the ability to make use of third-order theory of mind. However, in contrast to findings in strictly competitive settings, we find that in negotiations, agents do benefit from fourth-order theory of mind.

4 Overview of Posters

In order to provide an alternative forum to talks for presentation of research results and interests, posters were solicited prior to the workshop. These posters were on display from the second through fourth full days of the meeting, and several periods of time were allocated in the schedule for presenters to stand by their posters and answer questions.

In addition, a cash prize of $200 was donated by Purdue University Press, publisher of the *Journal of Problem Solving*, for the best poster with a student as first author. A team of four judges, one from each of the four disciplines represented at the workshop (John Tsotsos (CS), Stellan Ohlsson (PY), Reineke Verbrugge (AI), Vinod Goel (CN)), made their assessments over the course of the poster display. The prize was awarded to Harmen de Weerd on the final evening of the workshop by Zygmunt Pizlo.

Several of the posters described above were also presented as talks. To avoid duplication of material, though the author and title are given below, a reference is made to the appropriate subsection of Section 3 of this report for the associated abstract.

4.1 Where is the chocolate? Modeling Development of Reasoning about False Beliefs of Others

Burcu Arslan (University of Groningen, NL)

(See abstract in Section 3.1.)

4.2 Formally Making Sure Your AI Gets the Job Done (Good Enough)

Tarek R. Besold (Universität Osnabrück, DE)

The recognition that human minds/brains are finite systems with limited resources for computation has led researchers in cognitive science to advance the Tractable Cognition thesis: Human cognitive capacities are constrained by computational tractability. As also human-level AI in its attempt to recreate intelligence and capacities inspired by the human mind is dealing with finite systems, transferring this thesis and adapting it accordingly may give rise to insights that can help in progressing towards meeting the classical goal of AI in creating machines equipped with capacities rivaling human intelligence. Therefore, we develop the "Tractable Artificial and General Intelligence Thesis" and corresponding formal models usable for guiding the development of cognitive systems and models by applying notions from parameterized complexity theory and hardness of approximation to a general AI framework. This poster an overview of work putting special emphasis on connections and correspondences to the heuristics framework as recent development within cognitive science and cognitive psychology.

4.3 The Challenge of Optimization

Sarah Carruthers (University of Victoria, CA)

(See abstract in Section 3.3.)

4.4 Ifs and Ands and Ors

Nicole Cruz de Echeverria Loebell (University of London, GB)

(See abstract in Section 3.5.)

4.5 Computational Evidence that Self-regulation of Creativity is Good for Society

Liane Gabora (University of British Columbia – Vancouver, CA)

Excess individual creativity can be detrimental to society because creators invest in unproven ideas at the expense of propagating proven ones. Moreover, a proportion of individuals can benefit from creativity without being creative themselves by copying creators. We hypothesized that (1) societies increase their rate of cultural evolution by tempering the novelty-generating effects of creativity with the novelty-preserving effects of imitation, and (2) this is carried out by selectively rewarding and punishing creativity according to the value of the individuals' creative outputs. We tested this using an agent-based model of cultural evolution in which each agent self-regulated its invention-to-imitation ratio as a function of the fitness of its cultural outputs. In self-regulating societies, agents segregated into creators and imitators. The mean fitness of cultural outputs was higher than in non-self-regulating societies, and changes in diversity were rapider and more pronounced. We discuss limitations and possible social implications of our findings.

4.6 Fast and Loose: A Bounded-rational Analysis of Wason's Selection Task

Emmanuel Genot (Lund University, SE)

This paper proposes a bounded-rational version of the Bayesian rational analysis of Wason's Selection Task (ST), which was called for by its proponents—who acknowledged that it presupposes computations with unrealistic computational cost—but never realized. We model ST as a sequential decision problem of information-seeking by questioning, based on a non-classical understanding of the conditional. Our assumptions generalize those of the rational analysis, and are supported by independent empirical arguments, pertaining to both ST and other reasoning tasks. Our analysis generalizes the Bayesian rational analysis: it

converges to the same selection under its special assumptions, but vindicates the selection under weaker assumptions, and without the need for Bayesian rationality. In conclusion, we discuss the extension of our analysis to the deontic case.

4.7 Pearl Diving – Data for TSP Research from an Online Game

Alexandra Kirsch (Universität Tübingen, DE)

For a better understanding of how humans solve instances of the Traveling Salesperson Problem (TSP) we have implemented an online game, in which player shave to solve different variants of TSPs. The game contains features of social games that are now common on social network platforms, with an appealing colorful design, an introduction level to introduce the goal of the game, and purchasable game advantages. The levels are structured in blocks of eight levels, each containing a slightly different variant of the TSP. In previous laboratory studies of human TSP solving, our players get feedback and can repeat levels until they find the optimal solution. Up to August 2014, more than 50people have produced data of more than 2000 tours. The regularly updated dataset is available online for all interested researchers:http://hci.uni-tuebingen.de/DatasetsThe game itself can be played at www.perlentaucher.medieninformatik.uni-tuebingen.de

4.8 Resource-Rational Decision-Making

Falk Lieder (University of California – Berkeley, US)

Extreme events come to mind very easily and people overestimate their probability and overweight them in decision-making.In this paper we show that rational use of limited cognitive resources can generate these 'availability biases'. We hypothesize that availability helps people to quickly make good decisions in very risky situations. Our analysis shows that agents who decide by simulating a finite number of possible outcomes (sampling) should over-sample outcomes with extreme utility. We derive a cognitive strategy with connections to fast-and-frugal heuristics, and we experimentally confirm its prediction that an event's extremity increases the factor by which people overestimate its frequency. Our model also explains three context effects in decision-making under risk: framing effects, the Allai's paradox, and preference reversals.

4.9 Indirection and Symbol-like Processing in the Prefrontal Cortex and Basal Ganglia

David Noelle (University of California – Merced, US)

The ability to flexibly, rapidly, and accurately perform novel tasks is a hallmark of human behavior. In our everyday lives we are often faced with arbitrary instructions that we must understand and follow, and we are able to do so with remarkable ease. It has frequently been argued that this ability relies on symbol processing, which depends critically on the ability to represent variables and bind them to arbitrary values. Whereas symbol processing is a fundamental feature of all computer systems, it remains a mystery whether and how this ability is carried out by the brain. Here, we provide an example of how the structure and functioning of the pre-frontal cortex/basal ganglia working memory system can support variable binding, through a form of indirection(akin to a pointer in computer science). We show how indirection enables the system to flexibly generalize its behavior substantially beyond its direct experience (i. e., systematicity). We argue that this provides a biologically plausible mechanism that approximates a key component of symbol processing,exhibiting both the flexibility, but also some of the limitations, that area associated with this ability in humans.

4.10 Predictive Processing: A Standard Deviation from the Real World

Maria Otworowska (Radboud University Nijmegen, NL)

The question how a brain may give rise to cognition, such as the ability to form beliefs about the world and act upon them, has been a subject of decades of interdisciplinary research. Currently, a leading theory about how a brain may just do this is Predictive Processing theory. However in its current version, Predictive Processing is based on a Laplace assumption – it means that causes of sensory input are assumed to be represented probabilistically (so-called: recognition density) as an unimodal probability density function. We have shown that Laplace assumption violates expectations about typical human search behavior. Predictive Processing theory could be adjusted to explain such behavior, but this requires multimodal distributions, violating the Laplace assumption. This raises important questions about the possible (neural) representation and approximation of such multimodal distributions within a Predictive Processing framework.

4.11 Assessing the Computational Adequacy of the General Problem Solver Model

Zahra Sajedinia (Memorial University of Newfoundland, CA)

Problem solving is a core cognitive ability. Human performance varies widely in solving different types of problems. Ideally, cognitive models of problem solving should explain these variations in two ways: (1) the model should reproduce the sequences of actions applied by humans during problem solving (empirical adequacy), and (2) the time required by the model should match that required by humans, i. e., the model should be fast (slow) when humans are fast (slow) (computational adequacy). The former can be assessed by traditional psychological experiments; however, the latter requires the application of techniques from computational complexity theory. In this poster, we describe the first formal assessment of the computational adequacy of Newell and Simon's General Problem Solver model. We also discuss how our results can be used in both designing new psychological experiments and refining models of human problem solving.

4.12 Higher-order Theory of Mind in Mixed-motive Settings: An Agent-based Simulation Study

Harmen de Weerd (University of Groningen, NL)

(See abstract in Section 3.20.)

5 Working Groups and Discussions

Early in the planning for this workshop, the co-organizers realized that it would be difficult to specify topics for working groups given the diversity of disciplines represented. To mitigate this, we adopted the Bird-of-a-Feather (B.O.F.) mechanism employed at large conferences to allow people with similar interests to find each other in an unsupervised manner. A bulletin board was designated for the duration of the meeting as the B.O.F. 'clearing house'. Anyone could post a suggested B.O.F. topic along with a meeting time and venue on the board, and anyone with an interest in that topic was free to attend and contribute. Our only requirement was that the organizer of each B.O.F. provide a brief oral report summarizing the discussion and prepare a written report for inclusion in the seminar report. This mechanism was very successful, yielding 9 B.O.F. sessions. A subset of the written reports from these sessions is included below.

5.1 New Brain-Resource Inspired Computational Formalisms

Organizers: I. van Rooij, J. Kwisthout.

Participants: S. Carruthers, Ch. Cherniak, M. Fellows, N. Gierasimczuk, F. Jaekel, B. Juba, A. Kolokolova, F. Lieder, W. Maass, D. Noelle, F. Rosamond, C. Rothkopf, Z. Sajedinia, U. Stege, J. Tsotsos, S. Varma, T. Wareham, S. Watson.

One of the expected outcomes of this workshop was 'suggestions for new theories of computational complexity that better address the needs of researchers in psychology and cognitive science.' A sub-group of participants discussed this expected outcome. It was agreed that in order to be useful for the above-mentioned community of researchers, the abstraction level of such new theories should be 'just right.' It should be sufficiently abstract to generalize over concrete experiments or implementations, yet specific enough to be able to capture the key brain resources and constraints on them without generalizing them away. From the complexity-theoretic point of view, it was emphasized that the formalisms must discriminative, i. e., that important resources should not be 'lost in the abstraction.'

An important outcome of the discussion was agreement that 'networks of spiking neurons', as proposed by W. Maass and colleagues, are a promising starting point for future investigations. One of the foundational problems to be solved is how to translate the defining aspects of traditional computation devices into this neurally plausible model. For instance, in a Turing Machine the defining aspects are the input (the tape with symbols), the state transitions, the resources (e. g., time, space, non-determinism, randomness), the acceptance criteria, and the notion of reductions. These aspects must be mapped to their relevant counterparts in networks of spiking neurons, where the resources are defined as the convergence speed to the stationary distribution of the network, the number of samples from the stationary distribution that are needed to give the target output with a certain accuracy, the number or maximum frequency of spikes, the number of neurons, the number of connections, and so on.

We believe that researchers from various disciplines and with various backgrounds will be interested in this research program. Complexity theorists in computer science and artificial neural network researchers in cognitive and neuroscience can join forces to investigate the foundational problem identified above, with neuroscientists involved to make sure the formal models capture the relevant aspects of neuronal computations. The complexity theorists might then continue, for example, to prove class inclusion or separation, establish hierarchies, identify relations with more traditional complexity classes, and so on. Computational

modelers in cognitive science and neuroscience might use the formalism and proven results to show the tractability and intractability of competing theories, compare empirical data with model predictions, and generate testable hypotheses about particular brain computations.

5.2 Cognitive Architectures and Problem Solving

Organizer: T. Buwalda.

Participants: B. Arslan, T. Besold, M. Blokpoel, V. Goel, F. Jaekel, N. Loebell, D. Noelle, S. Ohlsson, N. Taatgen, J. Tsotsos, S. Varma, R. Verbrugge.

This discussion focused on the relation between problem solving and cognitive architectures. We collectively articulated the history of problem solving and cognitive architectures, inventoried the role problem solving is currently playing in cognitive architectures, and speculated about the role problem solving will play in the future development of cognitive architectures. The discussion of the historical relation focused on George Miller, who coined the definition of cognitive architectures (but not the term itself) and on Newell and Simon, who created the General Problem Solver and imported production rules from the theory of computation into cognitive science. The discussion of the contemporary relation focused on production system architectures such as Soar and ACT-R, which put problem solving first, and on Bayesian approaches, which emphasize statistical mechanisms. These architecture provide not just different views of the fundamental computational mechanisms of cognition, but also – when are applied to different domains – yield different theories and models of cognition. One existing architecture that is strongly tied to problem solving and that continues to generate new insights is Soar. A promising new architecture that was created with problem solving as a focus area is Icarus.

A number of questions for future research were identified. One was whether it is more scientifically productive to have an architecture that is relatively specialized for one domain, or to have an architecture that is more general, but that underspecifies (i. e., leaves more degrees of freedom to the builder of) models of specific domains. Although specialized architectures often support 'better' models of the domains for which they are specialized (e. g., problem solving), they make less sense if the value of architecture is in explaining all of cognition. On the other hand, general architectures and the underspecified models they support are often criticized as unfalsifiable, and thus outside of the science. We agreed that the underspecificity and unfalsifiability critiques do not hold when comparing models of the same domain (i. e., problem solving) that are implemented in the same architecture. In that case, model evaluation is possible, and the best model can be retained and the other models rejected.

We discussed the more general question of how to choose the 'best' architecture among competing alternatives based on empirical support. One viewpoint, inspired by Lakatos' philosophy of science and endorsed by scientists such as Newell, Anderson, and Cooper is that architectures can be assessed by their productivity – how many models of how many domains they support, and whether these models 'progress' by generating new empirical predictions and growing to account for new data (However, we noted that major updates or changes to architectures can make their past productivity irrelevant). Another viewpoint is that it might not be necessary to pick the 'best' architecture. Instead, all architectures contribute to progress in cognitive science, with newer architectures 'standing on the shoulders' of older architectures.

With respect to the future, it was observed that there is currently little appetite in psychology for constructing computational models; experimental research reigns supreme. By contrast, in neuroscience, modeling is more appreciated, perhaps because neuroscientists deal with more complex phenomena that generate orders of magnitude more data. This leads to a better understanding of the utility of computational models for making scientific progress. As evidence of this fact, cognitive architecture is flourishing in the interstices between artificial intelligence, cognitive science, and cognitive neuroscience, as evidenced by emerging fields and conferences such as Biologically Inspired Cognitive Architectures (BICA), Artificial General Intelligence (AGI), and Neural-Symbolic Learning and Reasoning (NeSy).

The discussion ended with the mention of two more challenges for the future: (1) building models of problem solving that are constrained by, and can explain, neuroscientific data, and (2) having a reference set of problem solving data against which to evaluate architecture-based models of problem solving.

5.3 Should 'as if' Explanations of Inferential Behavior be Resource-Bounded?

Organizers: E. Genot, J. Jacot.

Participants: N. Fleischhut, N. Loebell, M. Otworowska, S. Varma.

Theoretical frameworks such as Bayesian learning theory, Bayesian decision theory, and game theory can be fruitfully applied to describe and understand inferential behavior. It is common practice in philosophy and economics – and sometimes in cognitive science and psychology – to do so for functional-level theories (e. g., with Bayesian rational analysis models). However, people – the agents whose behavior is to be modeled – are not rational in the strong normative sense that these frameworks assume. Philosophers and economists generally ignore this problem. Cognitive scientists and psychologists do so as well by justifying their models as 'as if' explanations, shifting the burden of accommodating bounds on cognition to lower algorithmic-level theories (in the sense of Marr). This shifting strategy is usually accompanied by hand waving and the invocation of 'approximations' and 'heuristics' as cure-alls for intractability. We agreed that this strategy does not stand close scrutiny, and that an alternative must be found.

The main question we discussed was whether functional-level theories should incorporate from the start constraints on the cognitive resources that are assumed to be available, and not defer such constraints to algorithmic-level theories. We discussed as an example the rational analysis approach, which assumes the 'principle of rationality': that organisms respond optimally to their environment (and thus behave as optimal statistical learners and expected-utility maximizers). If instead organisms can be assumed to respond merely 'well enough,' and if a suitable interpretation of 'well enough' can be given that assumes only bounded resources, then 'as if' functional-level explanations can share the load with algorithmic-level theories. However, this imposes a trade-off and raises an issue of its own. The trade-off is that what constitutes a 'well enough' response to a task is highly context-dependent, and for this reason broad-spectrum theories cannot serve as unifying frameworks without local adjustments. This issue is that what constitutes an acceptable adjustment must be made on a context-by-context basis, and this raises the risk of over-fitting.

We agreed both on the usefulness of 'as if' explanations, and on the necessity to adjust them. Although this topic is broad by nature (as is the spectrum of the theoretical frameworks

we discussed), we also discussed concrete cases and applications, such as applications of rational analysis to data on reasoning tasks (e. g., variants of the Wason Selection Task). We agree that 'as if' explanations should incorporate linguistic pragmatics for understanding subjects' responses. In particular, Gricean pragmatics and Relevance Theory incorporate constraints on cognitive resources (such as attention and memory) in communication contexts while taking into account linguistic competence in the interpretation of instructions. Provided that such models can be captured formally, they can yield empirical predictions, by suggesting which parameters should be manipulated experimentally. The discussion concluded on which tools would be appropriate for this formalization.

5.4 Formal Representations of Real World Problem Solving / Complexity of Representing Real World Problems

Organizers: J. Kwisthout, F. Lieder.

Participants: [not recorded]

This meeting addressed two closely related questions: (1) what are suitable formal representations for real world problems, and (2) how can we analyze and characterize the complexity of representing real world problems by human problem solvers? A number of interesting ideas from cognitive science, artificial intelligence, and complexity analysis were brought to bear to understand these questions.

To take one example, interesting links were established between the observation that human problem solvers faced with unfamiliar problems tend to 'probe' the environment (i. e., act to get information), and with the complexity-theoretic aspect of interactive proofs. In such proofs, the traditional notion of 'verification' in NP is extended to be an interactive protocol where the verifier and an oracle exchange information.

To take another example, an interesting observation was made that problem solving is rarely done in complete isolation: we interact both with the environment and with the problem itself. Thus, there is an inherent multi-agent aspect in problem solving. This raises the possibility that communication complexity is relevant for understanding problem solving.

6 Panel Discussions and New Research Directions

The importance of choosing a proper representation in problem solving was discussed extensively. Some new insight were (1) when humans solve problem(s) the problem(s) they are solving may not be what they think they are solving and (2) representation format affects human problem solving. It was realized that research on meta-mathematics (the formal study of problem solving strategies) can benefit psychology and neuroscience knowledge of human problem solving in various ways, e. g.,

- knowledge compilation is the cognitive science equivalent of proposing a lemma in Extended Frege, and
- change detection in evolving data streams in AI can be informed by cognitive neuroscience on learning and switching strategies in changing environment.

New links between memory retrieval and Markov chains on semantic networks were also discussed.

In the B.O.F. sessions (Section 5.2) some notable common issues were (1) there is no agreement upon criterion for success (and how to measure it); (2) making falsifiable claims/models of cognitive / brain computation is important; and (3) computation is the appropriate language for modeling natural systems. Moreover, when comparing 'different' cognitive architectures, one has to know what are the key assumptions and commitments of the different architectures.

Several insights emerged during the panel discussion on alternative resource-bounded computation in spiking neural networks – namely, computational level models need to be at the right level of abstraction: low enough for neuroscience (to relate to concrete data of brain and behavior; should have features that neuroscience consider critical); but also general enough to be suitable for formal analytical methods.

It was also frequently noted that off-the-shelf complexity theory is limited in how much it can directly inform cognitive (neuro)science. Many of the sub-disciplines represented at the seminar could benefit from such a more finely-tuned computational complexity. To address this, a proposal was made for a study on alternative resource based computation. Such a study should incorporate the observation that the structure of the environment is exploited during human problem solving: this came out both in psychological talks as well as in computer science talks on parameterized complexity.

Future possible direction for further study has been considered, just to mention some:

- Machine Learning and Problem Solving,
- Problem solving by nature,
- Problem solving in cognitive architectures,
- Problem representation / interpretation,
- Learning representations and algorithms for problem solving (machine learning; how do we learn the representations and algorithms).

7 Dissemination of Results

All participants have been invited to submit their research presented at this seminar or inspired by this seminar for consideration for publication in The Journal of Problem Solving (JPS). JPS is an open access journal with an interdisciplinary readership. Considering the fact that papers in JPS can be accessed by everyone (no subscription is required), the proceedings from this workshop are expected to be read widely and have large impact. JPS (ISSN 1932-6246) is a multidisciplinary journal that publishes empirical and theoretical papers on mental mechanisms involved in problem solving. The journal welcomes original and rigorous research in all areas of human problem solving, with special interest in solving difficult problems (e. g., problems in which human beings outperform artificial systems). Examples of topics include (but are not limited to) optimization and combinatorial problems, mathematics and physics problems, theorem proving, games and puzzles, knowledge discovery problems, insight problems and problems arising in applied settings. JPS encourages submissions from psychology, computer science, mathematics, operations research and neuroscience. More information on the journal web site:

- http://docs.lib.purdue.edu/jps/
- Editor-in-Chief: Zygmunt Pizlo, Department of Psychological Sciences, Purdue University.

8 Seminar Program

See Section 3 for more information about the talks and the posters.

Monday, August the 18th

9:00	General Introduction	I. van Rooij
10:15	Intro: Psychology	Z. Pizlo
11:15	Intro: Computational Complexity	T. Wareham
13:30	Intro: AI	R. Verbrugge
14:30	Intro: Cognitive Neuroscience	D. Noelle
16:00	Panel and General Discussion	
19:15	Town Hall: Plans for the Week	

Tuesday, August the 19th

9:00	Talk	J. Tsotsos
9:30	Talk	W. Maass
10:15	Discussion:	Computational Complexity and Brain Computation
11:15	Poster session	
13:30	Talk	S. Ohlsson
14:00	Talk	M. Fellows
14:30	Talk	J. Kwisthout
15:00	Talk	T. Besold
16:00	Discussions	
19:15	Reporting + Panel I	

Wednesday, August the 20th

9:00	Talk	U. Stege
9:30	Talk	Ch. Cherniak
10:15	Talk	N. Taatgen
10:35	Talk	N. Loebell
10:55	Open problem session	
19:15	Panel II	

Thursday, August the 21st

9:00	Talk	N. Gierasimczuk
9:30	Talk	C. Rothkopf
10:30	Talk	B. Juba
10:55	Talk	E. Koechlin
11:20	Talk	V. Goel
11:45	Panel Discussion	
13:30	Talk	L. Gabora
13:55	Poster Session II	
14:50	Break out Session Talk:	B. Arslan / H. de Weerd / S. Carruthers
16:00	Discussions	
19:15	Poster Prize	
19:20	B.O.F Reports	
20:00	Town Hall Discussion	

Friday, August the 22nd

9:00	Seminar Summary

Participants

- Burcu Arslan
University of Groningen, NL
- Tarek R. Besold
Universität Osnabrück, DE
- Mark Blokpoel
Radboud Univ. Nijmegen, NL
- Sarah Carruthers
University of Victoria, CA
- Christopher Cherniak
University of Maryland – College Park, US
- Nicole Cruz de Echeverria Loebell
Birkbeck, Univ. of London, GB
- Harmen de Weerd
University of Groningen, NL
- Michael R. Fellows
Charles Darwin University – Darwin, AU
- Nadine Fleischhut
MPI für Bildungsforschung, DE
- Liane Gabora
University of British Columbia – Vancouver, CA
- Emmanuel Genot
Lund University, SE
- Nina Gierasimczuk
University of Amsterdam, NL
- Vinod Goel
York University – Toronto, CA
- Yll Haxhimusa
TU Wien, AT
- Justine Jacot
Lund University, SE
- Frank Jäkel
Universität Osnabrück, DE
- Brendan Juba
Harvard University, US
- Alexandra Kirsch
Universität Tübingen, DE
- Etienne Koechlin
ENP – Paris, FR
- Antonina Kolokolova
Memorial University of Newfoundland, CA
- Johan H. P. Kwisthout
Radboud Univ. Nijmegen, NL
- Falk Lieder
University of California – Berkeley, US
- Wolfgang Maass
TU Graz, AT
- Matthias Mnich
Universität Bonn, DE
- Martin Möhrmann
Universität Osnabrück, DE
- David Noelle
Univ. of California – Merced, US
- Stellan Ohlsson
Univ. of Illinois – Chicago, US
- Maria Otworowska
Radboud Univ. Nijmegen, NL
- Zygmunt Pizlo
Purdue University – West Lafayette, US
- Frances A. Rosamond
Charles Darwin University – Darwin, AU
- Constantin Rothkopf
TU Darmstadt, DE
- Zahra Sajedinia
Memorial University of Newfoundland, CA
- Ulrike Stege
University of Victoria, CA
- Marieke Sweers
Radboud Univ. Nijmegen, NL
- Niels A. Taatgen
University of Groningen, NL
- John K. Tsotsos
York University – Toronto, CA
- Iris van Rooij
Radboud Univ. Nijmegen, NL
- Sashank Varma
University of Minnesota – Minneapolis, US
- Rineke Verbrugge
University of Groningen, NL
- Todd Wareham
Memorial University of Newfoundland, CA
- Scott Watson
Memorial University of Newfoundland, CA

Report from Dagstuhl Seminar 14342

Equilibrium Computation

Edited by

Nimrod Megiddo[1], Kurt Mehlhorn[2], Rahul Savani[3], and Vijay V. Vazirani[4]

1 **IBM Almaden Center – San José, US,** megiddo@theory.stanford.edu
2 **MPI für Informatik – Saarbrücken, DE,** mehlhorn@mpi-inf.mpg.de
3 **University of Liverpool, GB,** rahul.savani@liverpool.ac.uk
4 **Georgia Institute of Technology, US,** vazirani@cc.gatech.edu

Abstract

This report documents the program and outcomes of Dagstuhl Seminar 14342 "Equilibrium Computation". The seminar was at the leading edge of current topics related to equilibrium computation for games and markets. We summarize these topics, give the talk abstracts, and give brief summaries of the problems that were discussed in the open problem sessions.

Seminar August 17–22, 2014 – http://www.dagstuhl.de/14342
1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems, J.4 Social and Behavioral Science: Economics
Keywords and phrases Algorithms, Computational Complexity, Equilibrium Computation, Game Theory, Market Equilibrium, Nash Equilibrium
Digital Object Identifier 10.4230/DagRep.4.8.73

1 Executive Summary

Nimrod Megiddo
Kurt Mehlhorn
Rahul Savani
Vijay V. Vazirani

The aim of this seminar was to study research issues related to algorithms and complexity for computation of equilibria in games and markets. The majority of participants were academics from computer science departments; some were from other disciplines; and several participants were from the corporate research departments of eBay, IBM, and Microsoft. All participants have strong interdisciplinary interests that typically span Economics, Game Theory, and Theoretical Computer Science.

The seminar started with a session of lightening talks, in which participants had two minutes and one slide to introduce themselves. This session was extremely well received, and it was worth the effort to ensure that everyone submitted a slide in advance. It is an effective and efficient way for everyone to get to know a little bit about each other, and thus to have things to talk about outside of talks right from the start of the seminar.

Three tutorials were given on topics chosen by the organizers. Bernhard von Stengel gave a tutorial on complementary pivoting algorithms for the Linear Complementarity Problem (LCP). The tutorial focussed on geometric aspects of LCPs and complementary pivoting algorithms, and in particular Lemke's algorithm. The LCP captures many game and market

Equilibrium Computation, *Dagstuhl Reports*, Vol. 4, Issue 8, pp. 73–88
Editors: Nimrod Megiddo, Kurt Mehlhorn, Rahul Savani, and Vijay V. Vazirani
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

problems, and it came up throughout the seminar, most directly in the final talk by Adler on reductions to bimatrix games from PPAD Lemke-verified LCPs.

Complementary pivoting algorithms inspired the complexity class PPAD, which, together with FIXP, capture the problems of finding fixed points and equilibria of games and markets. The second tutorial, given by Kousha Etessami, was about the complexity of equilibria and fixed points. It covered PPAD (= linear-FIXP), FIXP, and FIXP-a, and discussed some associated open problems. Related contributed talks included the following. Etessami, in a separate talk, showed that the complexity of computing a (perfect) equilibrium for an n-player extensive form game of perfect recall is hard for FIXP-a. Gairing showed that the problem of finding an equilibrium of a weighted congestion game is FIXP-hard. Garg presented several results on market equilibria, including the result that it is FIXP-hard to compute an equilibrium of an Arrow-Debreu exchange market with Leontief utility functions. Chen presented a PPAD-hardness result for the problem of finding an approximate equilibrium in an anonymous game with seven actions per player. Mehta showed that it is PPAD-hard to find an equilibrium of a rank-3 bimatrix game. Paparas presented PPAD-hardness results for several market settings with non-monotone utilities. The number of talks related to these complexity classes shows their ongoing importance for the field of equilibrium computation.

The third tutorial was on game dynamics and was given by Sergiu Hart. He showed that "uncoupledness" severely limits the possibilities to converge to Nash equilibria, but on the other hand, there are simple adaptive heuristics, such as "regret matching", that lead to correlated equilibria. At the end of his tutorial, Hart also presented an exponential lower bound on the query complexity of correlated equilibria. In a closely related contributed talk, Goldberg gave bounds for the query complexity of approximate equilibria of various types, including for the relatively new concept of ϵ-well-supported correlated equilibrium.

A large number of contributed talks presented algorithms for computing equilibria of games and markets. On market equilibria we had the following algorithmic talks: Cole presented an asynchronous gradient descent method that implements asynchronous tâtonnement; Mehlhorn presented a combinatorial polynomial-time algorithm for the linear exchange model; Vazirani introduced Leontief-Free Utility Functions and presented a complementary pivoting algorithm for computing an equilibrium in markets with these utilities; and Vegh presented new convex programmes for linear Arrow-Debreu markets. On other game models, we had the following algorithmic talks: Cummings presented an efficient differentially private algorithm for computing an equilibrium in aggregative games; Savani presented a gradient descent algorithm for finding an approximate equilibrium of a polymatrix game; and Skopalik presented algorithms for finding approximate pure equilibria of congestion games.

There were other contributed talks on a range of topics: Harks talked about resource competition on integral polymatroids; Hoefer talked about decentralized secretary algorithms; Jain presented an analysis of several business models and pricing schemes; and Schäfer presented results about coordination games on graphs.

Apart from the topics of the tutorials, all other talk topics were chosen by the presenters, not by the organizers. Generally talks were informal, and were very interactive, often with lengthy discussions taking place during them. All talks were well received. Open problems were discussed in two sessions, the first during a normal seminar room session, and the second with cheese and wine in the evening. Below we give abstracts for the talks and brief summaries of the open problems that were discussed.

2 Table of Contents

3 Overview of Talks

3.1 A direct reduction of PPAD Lemke-verified linear complementarity problems to bimatrix games

Ilan Adler (University of California – Berkeley, US)

Main reference I. Adler, S. Verma, "A direct reduction of PPAD Lemke-verified linear complementarity problems to bimatrix games," arXiv:1302.0067v1 [cs.CC], 2013.
URL http://arxiv.org/abs/1302.0067v1

The linear complementarity problem, LCP(q, M), is defined as follows. For given M, q find z such that $q + Mz \geq 0, z >= 0, z(q + Mz) = 0$, or certify that there is no such z. It is well known that the problem of finding a Nash equilibrium for a bimatrix game (2-NASH) can be formulated as a linear complementarity problem (LCP). In addition, 2-NASH is known to be complete in the complexity class PPAD (Polynomial-time Parity Argument Directed). However, the ingeniously constructed reduction (which is designed for any PPAD problem) is very complicated, so while of great theoretical significance, it is not practical for actually solving an LCP via 2-NASH, and it may not provide the potential insight that can be gained from studying the game obtained from a problem formulated as an LCP (e.g. market equilibrium). The main goal of this paper is the construction of a simple explicit reduction of any LCP(q, M) that can be verified as belonging to PPAD via the graph induced by the generic Lemke algorithm with some positive covering vector d, to a symmetric 2-NASH. In particular, any endpoint of this graph (with the exception of the initial point of the algorithm) corresponds to either a solution or to a so-called secondary ray. Thus, an LCP problem is verified as belonging to PPAD if any secondary ray can be used to construct, in polynomial time, a certificate that there is no solution to the problem. We achieve our goal by showing that for any M, q and a positive d satisfying a certain nondegeneracy assumption with respect to M, we can simply and directly construct a symmetric 2-NASH whose Nash equilibria correspond one-to-one to the end points of the graph induced by LCP(q, M) and the Lemke algorithm with a covering vector d. We note that for a given M the reduction works for all positive d with the exception of a subset of measure 0.

3.2 Complexity of Nash Equilibria in Anonymous Games

Xi Chen (Columbia University – New York, US)

We show that finding a $(\frac{1}{2}^n)$-approximate Nash equilibrium in an anonymous game with seven actions is PPAD-complete.

3.3 From Asynchronous Gradient Descent to Asynchronous Tatonnement

Richard Cole (New York University, US)

Gradient descent is an important class of iterative algorithms for minimizing convex functions. Classically, gradient descent is a sequential and synchronous process. Distributed and asynchronous variants of gradient descent have been studied since the 1980s, and they have been experiencing a resurgence due to demand from large-scale machine learning problems running on multi-core processors. We provide a version of asynchronous gradient descent (AGD) in which communication between the processors is minimal and there is little synchronization overhead. We also propose a new timing model for its analysis. With this model, we give the first amortized analysis of AGD on convex functions. The amortization allows for bad updates (updates that increase the value of the convex function); in contrast, most prior work makes the strong assumption that every update must be significantly improving. Typically, the step sizes used in AGD are smaller than those used in its synchronous counterpart. We provide a method to determine the step sizes in AGD based on the Hessian entries of the convex function. In certain circumstances, the resulting step sizes are a constant fraction of those used in the corresponding synchronous algorithm, enabling the overall performance of AGD to improve linearly with the number of processors. Our amortized analysis of AGD can be applied to show that tatonnement, a simple distributed price update dynamic, converges toward the market equilibrium in a number of economic markets. We use the Ongoing market model due to Cole and Fleischer [STOC'08], a fairly recent market model that supports distributed and asynchronous price updates. We show that asynchronous tatonnement converges toward the market equilibrium in Ongoing Fisher markets in which the buyers have CES utility functions; our analysis of AGD can be applied to the market problem due to the fact that tatonnement is equivalent to gradient descent for this class of markets [Cheung, Cole, Devanur STOC'13].

3.4 Privacy and Truthful Equilibrium Selection for Aggregative Games

Rachel Cummings (Northwestern University – Evanston, US)

Joint work of Cummings, Rachel; Kearns, Michael; Roth, Aaron; Wu, Zhiwei Steven
Main reference R. Cummings, M. Kearns, A. Roth, Z. S. Wu, "Privacy and Truthful Equilibrium Selection for Aggregative Games," arXiv:1407.7740v2 [cs.DS], 2014.
URL http://arxiv.org/abs/1407.7740v2

We study a very general class of games – multi-dimensional aggregative games – which in particular generalize both anonymous games and weighted congestion games. For any such game that is also large (meaning that the influence that any single player's action has on the utility of others is diminishing with the number of players in the game), we solve the equilibrium selection problem in a strong sense. In particular, we give an efficient weak mediator: an algorithm or mechanism which has only the power to listen to reported types and provide non-binding suggested actions, such that (a) it is an asymptotic Nash equilibrium for every player to truthfully report their type to the mediator, and then follow its suggested

action; and (b) that when players do so, they end up coordinating on a particular asymptotic pure strategy Nash equilibrium of the induced complete information game. In fact, truthful reporting is an ex-post Nash equilibrium of the mediated game, so our solution applies even in settings of incomplete information, and even when player types are arbitrary or worst-case (i.e. not drawn from a common prior). We achieve this by giving an efficient differentially private algorithm for computing a Nash equilibrium in such games. The rates of convergence to equilibrium in all of our results are inverse polynomial in the number of players n. We also give similar results for a related class of one- dimensional games with weaker conditions on the aggregation function, and apply our main results to a multi-dimensional market game. Our results can be viewed as giving, for a rich class of games, a more robust version of the Revelation Principle, in that we work with weaker informational assumptions (no common prior), yet provide a stronger solution concept (Nash versus Bayes Nash equilibrium). Previously, similar results were only known for the special case of unweighted congestion games. In the process, we derive several algorithmic results that are of independent interest, and that further the connections between tools in differential privacy and truthfulness in game-theoretic settings. We give the first algorithm for efficiently computing Nash equilibria in aggregative games of constant dimension $d > 1$. We also give the first method for solving a particular class of linear programs under the constraint of joint differential privacy.

3.5 The complexity of computing a (perfect) equilibrium for an n-player extensive form game of perfect recall

Kousha Etessami (University of Edinburgh, GB)

We study the complexity of computing or approximating an equilibrium for a given finite n-player extensive form game of perfect recall (EFGPR), where $n \geq 3$. Our results apply not only to Nash equilibrium (NE), but also to various important refinements of NE for EFGPRs, including: subgame-perfect equilibrium in behavior strategies, sequential equilibrium (SE), and extensive form trembling-hand perfect equilibrium (PE). Of these, the most refined notion is PE. By a classic result of Selten, a PE exists for any EFGPR. We show that, for all these notions of equilibrium, approximating an equilibrium for a given EFGPR, to within a given desired precision, is `FIXP-a`-complete. We also show that computing a "delta-almost subgame-perfect equilibrium" in behavior strategies for a given EFGPR and given $\delta > 0$, is `PPAD`-complete. In doing so, we also define the more refined notion of a "δ-almost ϵ-perfect" equilibrium, and show that computing one is `PPAD`-complete. Thus, approximating one such (delta-almost) equilibrium for n-player EFGPRs, $n \geq 3$, is `P`-time equivalent to approximating a (δ-almost) NE for a normal form game (NFG) with 3 or more players. NFGs are trivially encodable as EFGPRs without blowup in size. Thus our results extend the celebrated complexity results for NFGs to the considerably more general setting of EFGPRs. For 2-player EFGPRs, analogous complexity results follow from the algorithms of Koller, Megiddo, and von Stengel (1996), and von Stengel, van den Elzen, and Talman (2002). However, prior to the present paper, no analogous results were known for EFGPRs with 3 or more players.

3.6 Tutorial: Complexity of Equilibria and Fixed Points

Kousha Etessami (University of Edinburgh, GB)

This tutorial will discuss the complexity of equilibria and fixed points. It focusses on the complexity classes FIXP, FIXP-a, and linear-FIXP (= PPAD), and some associated open problems.

3.7 Weighted Congestion Games are FIXP-hard

Martin Gairing (University of Liverpool, GB)

In this talk, I will discuss the complexity class FIXP and show that weighted congestion games are FIXP-hard. The proof builds on a recent result of Milchtaich (https://faculty.biu.ac.il/~milchti/papers/representation.pdf).

3.8 Leontief Exchange Markets Can Solve Multivariate Polynomial Equations, Yielding FIXP and ETR Hardness

Jugal Garg (MPI für Informatik – Saarbrücken, DE)

We show FIXP-hardness of computing equilibria in Arrow-Debreu exchange markets under Leontief utility functions, and Arrow-Debreu markets under linear utility functions and Leontief production sets, thereby settling the open question of [Vazirani-Yannakakis'11]. In both cases, as required under FIXP, the set of instances mapped onto will admit equilibria, i.e., will be "yes" instances. If all instances are under consideration, then in both cases we prove that the problem of deciding if a given instance admits an equilibrium is ETR-complete, where ETR is the class defined by the Existential Theory of Reals. The main technical part of our result is the following reduction: Given a set S of simultaneous multivariate polynomial equations in which the variables are constrained to be in a closed bounded region in the positive orthant, we construct a Leontief exchange market M which has one good corresponding to each variable in S. We prove that the equilibria of M, when projected onto prices of these latter goods, are in one-to-one correspondence with the set of solutions of the polynomials.

3.9 Bounds for the query complexity of approximate equilibria

Paul W. Goldberg (University of Oxford, GB)

Joint work of Goldberg, Paul W.; Roth, Aaron
Main reference P. W. Goldberg, A. Roth, "Bounds for the query complexity of approximate equilibria," in Proc. of the 2014 ACM Conf. on Economics and Computation (EC'14), pp. 639–656, ACM, 2014.
URL http://dx.doi.org/10.1145/2600057.2602845

We analyze the number of payoff queries needed to compute approximate equilibria of multi-player games. We find that query complexity is an effective tool for distinguishing the computational difficulty of alternative solution concepts, and we develop new techniques for upper- and lower bounding the query complexity. For binary-choice games, we show logarithmic upper and lower bounds on the query complexity of approximate correlated equilibrium. For well-supported approximate correlated equilibrium (a restriction where a player's behavior must always be approximately optimal, in the worst case over draws from the distribution) we show a linear lower bound, thus separating the query complexity of well supported approximate correlated equilibrium from the standard notion of approximate correlated equilibrium. Finally, we give a query-efficient reduction from the problem of computing an approximate well-supported Nash equilibrium to the problem of verifying a well supported Nash equilibrium, where the additional query overhead is proportional to the description length of the game. This gives a polynomial-query algorithm for computing well supported approximate Nash equilibria (and hence correlated equilibria) in concisely represented games. We identify a class of games (which includes congestion games) in which the reduction can be made not only query efficient, but also computationally efficient.

3.10 Resource Competition on Integral Polymatroids

Tobias Harks (Maastricht University, NL)

Joint work of Harks, Tobias; Klimm, Max; Peis, Britta
Main reference T. Harks, M. Klimm, P. Peis, "Resource Competition on Integral Polymatroids," arXiv:1407.7650v1 [cs.GT], 2014.
URL http://arxiv.org/abs/1407.7650v1

We derive a new existence result for integer-splittable congestion games on integral polymatroids.

3.11 The Query Complexity of Correlated Equilibria

Sergiu Hart (The Hebrew Univ. of Jerusalem, IL)

Joint work of Hart, Sergiu; Nisan, Noam
Main reference S. Hart, N. Nisan, "The Query Complexity of Correlated Equilibria," arXiv:1305.4874v1 [cs.GT], 2013.
URL http://arxiv.org/abs/1305.4874v1

We consider the complexity of finding a Correlated Equilibrium in an n-player game in a model that allows the algorithm to make queries for players' utilities at pure strategy profiles.

Many randomized regret-matching dynamics are known to yield an approximate correlated equilibrium quickly: in time that is polynomial in the number of players, n, the number of strategies of each player, m, and the approximation error, $1/\epsilon$;. Here we show that both randomization and approximation are necessary: no efficient deterministic algorithm can reach even an approximate equilibrium and no efficient randomized algorithm can reach an exact equilibrium.

3.12 Tutorial: Game Dynamics

Sergiu Hart (The Hebrew Univ. of Jerusalem, IL)

An overview of work on dynamical systems in multi-player environments. On the one hand, the natural informational restriction that each participant does not know the payoff functions of the other participants – "uncoupledness" – severely limits the possibilities to converge to Nash equilibria. On the other hand, there are simple adaptive heuristics – such as "regret matching" – that lead in the long run to correlated equilibria, a concept that embodies full rationality.

3.13 Decentral Secretary Algorithms

Martin Hoefer (Universität des Saarlandes, DE)

The secretary model is a popular framework for the analysis of online admission problems beyond the worst case. In many markets, however, decisions about acceptance or rejection of applicants have to be made in a decentralized fashion and under competition. In this paper, we cope with this problem and design algorithms for decentralized secretary problems with competition among firms. In the basic model, there are m firms and each has a job to offer. n applicants arrive iteratively in random order. Upon arrival of an applicant, a value for each job is revealed. Each firm has to decide whether or not to offer its job to the current applicant without knowing the actions or values of other firms. Applicants then decide to accept their most preferred offer. We consider the overall social welfare of the matching, as well as the value of the match for each single firm. We design a decentralized randomized thresholding-based algorithm with ratio $O(\log n)$ that works in a very general sampling model. In addition, it can be used by firms hiring several applicants based on a local matroid. On the other hand, even in the basic model we show a lower bound of $\Omega(\log n/(\log \log n))$ for all thresholding-based algorithms. Moreover, we provide secretary algorithms with constant competitive ratios, e.g., when values of applicants for different firms are stochastically independent. In this case, we can show a constant ratio even when each firm offers several different jobs, and even with respect to its individually optimal assignment. We also analyze several variants with stochastic correlation among applicant values.

3.14 Business Model Analysis

Kamal Jain (eBay Research Labs, US)

The talk presented how to mathematically analyze business models to study some of their advantages and disadvantages.

3.15 A Combinatorial Algorithm for the Linear Exchange Model

Kurt Mehlhorn (MPI für Informatik – Saarbrücken, DE)

Joint work of Ran Duan; Mehlhorn, Kurt
Main reference R. Duan, K. Mehlhorn, "A Combinatorial Polynomial Algorithm for the Linear Arrow-Debreu Market," arXiv:1212.0979v3 [cs.DS], 2014.
URL http://arxiv.org/abs/1212.0979v3

We present the first combinatorial polynomial time algorithm for computing the equilibrium of the Arrow-Debreu market model with linear utilities. Our algorithm views the allocation of money as flows and iteratively improves the balanced flow as in [Devanur et al. 2008] for Fisher's model. We develop new methods to carefully deal with the flows and surpluses during price adjustments. Our algorithm performs $O(n^6 \log(nU))$ maximum flow computations, where n is the number of agents and U is the maximum integer utility. The flows have to be presented as numbers of bitlength $O(n \log(nU))$ to guarantee an exact solution. Previously, [Jain 2007, Ye 2007] have given polynomial time algorithms for this problem, which are based on solving convex programs using the ellipsoid algorithm and the interior-point method, respectively.

3.16 Resolving the Complexity of Constant-Rank Bimatrix Games

Ruta Mehta (Georgia Institute of Technology, US)

The rank of a bimatrix game (A, B) is defined as the rank of $(A + B)$, e.g., rank-0 is zero-sum games. In 2005, Kannan and Theobald asked if there exists a polynomial time algorithm for constant rank games. We answer this question affirmatively for rank-1 games, and negatively for games with rank three or more (unless PPAD =P); the status of rank-2 games remains unresolved. In the process we obtain a number of other results, including a simpler proof of PPAD-hardness for 2-Nash.

3.17 The Complexity of Non-Monotone Markets

Dimitris Paparas (Columbia University – New York, US)

Joint work of Chen, Xi; Paparas, Dimitris; Yannakakis, Mihalis
Main reference X. Chen, D. Paparas, M. Yannakakis, "The Complexity of Non-Monotone Markets," arXiv:1211.4918v1 [cs.CC], 2012.
URL http://arxiv.org/abs/1211.4918v1

We introduce the notion of non-monotone utilities, which covers a wide variety of utility functions in economic theory. We then prove that it is PPAD-hard to compute an approximate Arrow-Debreu market equilibrium in markets with linear and non- monotone utilities. Building on this result, we settle the long-standing open problem regarding the computation of an approximate Arrow-Debreu market equilibrium in markets with CES utility functions, by proving that it is PPAD-complete when the Constant Elasticity of Substitution parameter ρ is any constant less than -1.

3.18 Computing Approximate Nash Equilibria in Polymatrix Games

Rahul Savani (University of Liverpool, GB)

Joint work of Deligkas, Argyrios; Fearnley, John; Savani, Rahul; Spirakis, Paul
Main reference A. Deligkas, J. Fearnley, R. Savani, P. Spirakis, "Computing Approximate Nash Equilibria in Polymatrix Games," in Proc. of the 10th Int'l Conf. on Web and Internet Economics (WINE'14), to appear; pre-print available as arXiv:1409.3741v2 [cs.GT].
URL arxiv.org/abs/1409.3741v2

In an ϵ-Nash equilibrium, a player can gain at most ϵ by unilaterally changing his behaviour. For two-player (bimatrix) games with payoffs in $[0, 1]$, the best-known ϵ achievable in polynomial time is 0.3393 (Tsaknakis and Spirakis). In general, for n-player games an ϵ-Nash equilibrium can be computed in polynomial time for an ϵ that is an increasing function of n but does not depend on the number of strategies of the players. For three-player and four-player games the corresponding values of ϵ are 0.6022 and 0.7153, respectively. Polymatrix games are a restriction of general n-player games where a player's payoff is the sum of payoffs from a number of bimatrix games. There exists a very small but constant ϵ such that computing an ϵ-Nash equilibrium of a polymatrix game is PPAD-hard. Our main result is that a $(0.5 + \delta)$-Nash equilibrium of an n-player polymatrix game can be computed in time polynomial in the input size and $\frac{1}{\delta}$. Inspired by the algorithm of Tsaknakis and Spirakis, our algorithm uses gradient descent on the maximum regret of the players. We also show that this algorithm can be applied to efficiently find a $(0.5 + \delta)$-Nash equilibrium in a two-player Bayesian game.

3.19 Coordination Games on Graphs

Guido Schäfer (CWI – Amsterdam, NL)

We introduce natural strategic games on graphs, which capture the idea of coordination in a local setting. We show that these games have an exact potential and have strong equilibria when the graph is a pseudoforest. We also exhibit some other classes of games for which a strong equilibrium exists. However, in general strong equilibria do not need to exist. Further, we study the (strong) price of stability and anarchy. Finally, we consider the problems of computing strong equilibria and of determining whether a joint strategy is a strong equilibrium.

3.20 Approximate pure Nash equilibria

Alexander Skopalik (Universität Paderborn, DE)

Among other solution concepts, the notion of the pure Nash equilibrium plays a central role in Game Theory. Pure Nash equilibria in a game characterize situations with non-cooperative deterministic players in which no player has any incentive to unilaterally deviate from the current situation in order to achieve a higher payoff. Unfortunately, it is well known that there are games that do not have pure Nash equilibria. Furthermore, even in games where the existence of equilibria is guaranteed, their computation can be a computationally hard task. Such negative results significantly question the importance of pure Nash equilibria as solution concepts that characterize the behavior of rational players. Approximate pure Nash equilibria, which characterize situations where no player can significantly improve her payoff by unilaterally deviating from her current strategy, could serve as alternative solution concepts provided that they exist and can be computed efficiently. We discuss recent positive algorithmic and positive existence results for approximate pure Nash equilibria in unweighted and weighted congestion games.

3.21 Leontief-Free Utility Functions

Vijay V. Vazirani (Georgia Tech, US)

Joint work of Garg, Jugal; Mehta, Ruta; Vazirani Vijay V.

Leontief utility functions capture the joint utility of a bundle if the goods in it are complements. We give an analogous notion for the case of substitutable goods. Even though our utility function is non-separable, we show that it always admits an equilibrium using rational numbers and we give a complementary pivot algorithm for finding one.

3.22 Convex programmes for linear Arrow-Debreu markets

László A. Végh (London School of Economics, GB)

We give a new, flow-type convex programme for linear Arrow-Debreu markets, along with a simple proof of the existence and rationality of equilibria and some further properties. We also survey previous convex programs for the problem and investigate connections between them.

3.23 Tutorial: Geometric Views of Linear Complementarity Algorithms and Their Complexity

Bernhard von Stengel (London School of Economics, GB)

The linear complementarity problem (LCP) generalizes linear programming (via the complementary slackness conditions of a pair of optimal primal and dual solutions) and finding Nash equilibria of bimatrix games. It suffices to look only for symmetric equilibria of symmetric games, which simplifies the problem setup. Lemke's classical complementary pivoting algorithm finds a solution to an LCP in many cases.

This tutorial reviews complementary pivoting, and two main geometric views of the algorithms by Lemke-Howson and Lemke: the polyhedral view, which describes the nonnegativity (feasilibity) constraints, and the "complementary cones" view, which maintains complementarity. Using complementary cones, Lemke's algorithm is seen as inverting a piecewise linear map along a line segment. One new result is that the Lemke-Howson algorithm is a special case of this description, which allows its implementation as a special case of Lemke's algorithm.

Based on a joint talk with Rahul Savani.

4 Open Problems

Open problems were discussed in two sessions, the first in seminar room, and the second with cheese and wine in the evening. We give brief summaries of the open problems.

4.1 Session 1

- Kamal Jain described an open problem related to a computational variant of Minkowski's Linear Forms Theorem that lies in the complexity class PPP.
- Bernhard von Stengel posed the open problem of lower bounding the maximal number of extreme Nash equilibria in 5×5 bimatrix games, where for $k \times k$ games with $k \leq 4$ we know that there can be at most $2^k - 1$ extreme equilibria, and for $k \geq 6$, von Stengel has constructed games with strictly more than $2^k - 1$ extreme equilibria.

4.2 Session 2

Thanks to Yuval Rabani for making the notes on the open problems from the second session.

- Rahul Savani asked: What is the complexity of finding a Nash equilibrium in an asymmetric 2-player network congestion game? Is there a polynomial time algorithm or is the problem PLS-hard?
- Rahul Savani asked about the problem "Network coordination": What is the complexity of finding a mixed Nash equilibrium in a polymatrix game where for every edge of the graph defining the game, the 2-player game has the property that in each pair of payoffs, the two values are equal? The problem is in $\texttt{PLS} \cap \texttt{PPAD}$.
- Dimitris Paparas asked about classifying utility functions into monotone and non-monotone. In particular, WGS utilities and also CES utilities with $\rho \in [-1, 0)$ are monotone, and CES utilities with $\rho < -1$ and also Leontief utilities are non-monotone. We know that there is a polynomial time algorithm for computing a competitive equilibrium in exchange markets that use specific subsets of monotone utilities. Does this extend to all monotone utilities? We know that there is a class of exchange markets using non-monotone and linear utilities for which the problem of computing a competitive equilibrium is PPAD-hard. Does the result hold if we don't use also linear utilities?
- Bernhard von Stengel posed a problem about Nash codes. We have a noisy channel with k possible inputs and k possible outputs. A stochastic matrix defines the distribution of outputs given an input. The two sides want to convey a specific signal as accurately as possible. A strategy for the sender is a mapping of the possible signals to the possible inputs. A strategy of the receiver is a reconstruction of a signal from the received output (w.l.o.g. this is the best possible such reconstruction given the sender's mapping). Given a matrix, some mappings may lead to a Nash equilibrium and others might not. What is the complexity of deciding whether or not every mapping from signals to inputs leads to a Nash equilibrium?
- Paul Goldberg posed a problem about the communication complexity of finding approximate equilibria.: In a 2-player game, the row and column players get their own payoff matrix (R and C, respectively). Beforehand, they can agree on two functions $f = f(R)$ and $f' = f'(C)$ that are mixed strategies that depends on the partial information they have. Then, an adversary chooses R and C and the game is played by using the strategies $f(R)$ and $f'(C)$. For which values of ϵ can you get an ϵ-Nash equilibrium this way? Known: and upper bound of 3/4 and a lower bound of 0.501 (i.e., above 1/2). The upper bound is simply playing with probability 1/2 a uniform distribution on your strategies and with probability 1/2 a best response to the other player's uniform distribution on strategies. Perhaps the same question is interesting with respect to an ϵ-correlated equilibrium, using f,f' chosen at random from a joint distribution.
- Ruta Mehta asked about the computational complexity of decision and counting related to k-player games. What is the complexity of various counting problems in (symmetric) k-player games? Various such problems in 2-player games are known to be #P-hard. Also, not all decision problems concerning 2-player Nash equilibrium that are known to be NP-hard, are also known to be ETR-hard in k-player games. (The cases of superset and subset in 3-player games are known to be ETR-complete; what remains open is maximal and minimal supports as far as the list of Gilboa and Zemel goes.)
- Kousha Etessami asked an open problem related to "Solvency Games", which are a special case of one-counter MDPs.

Participants

- Ilan Adler
University of California – Berkeley, US
- Susanne Albers
TU München, DE
- Xi Chen
Columbia Univ. – New York, US
- Richard Cole
New York University, US
- Rachel Cummings
Northwestern University – Evanston, US
- Kousha Etessami
University of Edinburgh, GB
- Martin Gairing
University of Liverpool, GB
- Jugal Garg
MPI für Informatik – Saarbrücken, DE
- Paul W. Goldberg
University of Oxford, GB
- Tobias Harks
Maastricht University, NL
- Sergiu Hart
The Hebrew University of Jerusalem, IL
- Penny Haxell
University of Waterloo, CA
- Martin Hoefer
Universität des Saarlandes, DE
- Kamal Jain
eBay Research Labs, US
- Nimrod Megiddo
IBM Almaden Center, US
- Kurt Mehlhorn
MPI für Informatik – Saarbrücken, DE
- Ruta Mehta
Georgia Inst. of Technology, US
- Peter Bro Miltersen
Aarhus University, DK
- Dimitris Paparas
Columbia Univ. – New York, US
- Britta Peis
RWTH Aachen, DE
- Yuval Rabani
The Hebrew University of Jerusalem, IL
- Rahul Savani
University of Liverpool, GB
- Guido Schäfer
CWI – Amsterdam, NL
- Leonard J. Schulman
CalTech, US
- Alexander Skopalik
Universität Paderborn, DE
- Madhu Sudan
Microsoft New England R&D Center – Cambridge, US
- László A. Végh
London School of Economics, GB
- Bernhard von Stengel
London School of Economics, GB

Report from Dagstuhl Seminar 14351

Decision Procedures and Abstract Interpretation

Edited by

Daniel Kroening[1], Thomas W. Reps[2], Sanjit A. Seshia[3], and Aditya V. Thakur[4]

1 **University of Oxford, GB,** kroening@cs.ox.ac.uk
2 **University of Wisconsin-Madison, US,** reps@cs.wisc.edu
3 **University of California, Berkeley, US,** sseshia@eecs.berkeley.edu
4 **University of California, Berkeley, US,** thakur@berkeley.edu

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 14351 "Decision Procedures and Abstract Interpretation". The seminar brought together practitioners and reseachers in abstract interpretation and decision procedures. The meeting highlighted the connections between the two disciplines, and created new links between the two research communities. Joint activities were also conducted with the participants of Dagstuhl Seminar 14352 "Next Generation Static Software Analysis Tools", which was held concurrently.

Seminar August 24–29, 2014 – http://www.dagstuhl.de/14351
1998 ACM Subject Classification D.2.4 Software/Program Verification. F.3.1 Specifying and Verifying and Reasoning about Programs. F.3.2 Semantics of Programming Languages. F.4.1 Mathematical Logic
Keywords and phrases Program analysis, Abstract interpretation, Abstract domain, Fixed-point finding, Satisfiability checking, Satisfiability modulo theories, Decision procedures, Constraint programming
Digital Object Identifier 10.4230/DagRep.4.8.89

1 Executive Summary

Aditya Thakur

The seminar was successful in bringing the following two communities together:

- designers and implementors of abstract interpreters, and
- designers and implementors of decision procedures.

The abstract interpretation (AI) and decsion procedure (DP) communities have several interests in common. Tools created by each of these communities can be viewed as using symbolic techniques to explore the state state of a transition system. However, the respective repertoires of techniques used in the two disciplines are quite different, and each community has its own mindset and outlook. The seminar sought to capitalize on recent ideas that demonstrated new connections between the two disciplines, and, consequently, promote the cross-fertilization between the areas at a deep technical level.

The seminar had 27 participants from both the AI and DP communities. To keep pariticipants from both areas engaged during a session, the organizers refrained from filling a session only with talks focusing on a particular community. Instead, each session consisted of

talks by participants of both research communities. Furthermore, talks by young researchers were scheduled earlier in the week, which enabled them to get better feedback on their research.

The Dagstuhl Seminar 14352 "Next Generation Static Software Analysis Tools" was held concurrently with Dagstuhl Seminar 14351. There were a number of joint activities organized to foster interaction among participants of the two seminars:

- The first session on Monday was a joint session for participants of both seminars. In this session, all participants introduced themselves and briefly described their research interests. Furthermore, Patrick Cousot, an organizer for Seminar 14352, and Thomas Reps, an organizer for Seminar 14351, each gave a "scene-setting" talk.
- The Wednesday excursion to the steel mill and Egyptian exhibit was organized as a joint activity.
- A joint session was organized on Thursday afternoon. The talks in this session were given by participants of both seminars.
- The seating arrangement for the Friday dinner was organized so that participants from both seminars sat together.
- The schedule of talks for both seminars was shared with all participants. Hence, participants of one seminar were able to attend a specific talk in the other seminar, if they felt the talk was especially relevant.

Apart from the planned activities listed above, the week saw a lot of informal discussions among participants of these two seminars in the evenings.

The seminar also featured talks about two other research areas: constraint programming (CP) and machine learning (ML). The talks by Mine, Rueher, and Truchet highlighted the use of abstract interpretation in CP. The talks by Reps, Seshia, Sharma, and Thakur discussed the application of ML techniques, such as inductive learning, to problems in AI and DP. Both these sets of talks garnered interesting discussions about the connections among all these various reseach areas. Furthermore, this discussion indicates that future seminars should include even more researchers and practitioners from not just the AI and DP communities, but also the CP and ML communities.

2 Table of Contents

3 Overview of Talks

3.1 Symbolic Optimization with SMT Solvers

Aws Albarghouthi (University of Toronto, CA)

Joint work of Li, Yi; Albarghouthi, Aws; Kincaid, Zachary; Gurfinkel, Arie; Chechik, Marsha
Main reference Y. Li, A. Albarghouthi, Z. Kincaid, A. Gurfinkel, M. Chechik, "Symbolic Optimization with SMT Solvers," in Proc. of the 41st ACM SIGPLAN-SIGACT Symp. on Principles of Programming Languages (POPL'14), pp. 607–618, ACM, 2014.
URL http://dx.doi.org/10.1145/2535838.2535857

The rise in efficiency of Satisfiability Modulo Theories (SMT) solvers has created numerous uses for them in software verification, program synthesis, functional programming, refinement types, etc. In all of these applications, SMT solvers are used for generating satisfying assignments (e. g., a witness for a bug) or proving unsatisfiability/validity(e. g., proving that a subtyping relation holds). We are often interested in finding not just an arbitrary satisfying assignment, but one that optimizes (minimizes/maximizes) certain criteria. For example, we might be interested in detecting program executions that maximize energy usage (performance bugs), or synthesizing short programs that do not make expensive API calls. Unfortunately, none of the available SMT solvers offer such optimization capabilities.

In this talk, I describe SYMBA, an efficient SMT-based optimization algorithm for objective functions in the theory of linear real arithmetic (LRA). Given a formula Phi and an objective function t, SYMBA finds a satisfying assignment of Phi; that maximizes the value of t. SYMBA utilizes efficient SMT solvers as black boxes. As a result, it is easy to implement and it directly benefits from future advances in SMT solvers. Moreover, SYMBA can optimize a set of objective functions, reusing information between them to speed up the analysis. We have implemented SYMBA and evaluated it on a large number of optimization benchmarks drawn from program analysis tasks, namely, symbolic abstraction for a large family of numerical abstract domains. Our results indicate the power and efficiency of SYMBA in comparison with competing approaches, and highlight the importance of its multi-objective-function feature.

3.2 Spatial Interpolants

Joshua Berdine (Microsoft Research UK – Cambridge, GB)

Joint work of Albargouthi, Aws; Berdine, Josh; Cook, Byron; Kincaid Zachary

We propose SplInter, a new technique for proving safety properties of heap manipulating programs that marries (1) a new separation logic–based analysis for heap reasoning with (2) an interpolation-based technique for refining and strengthening heap shape invariants with data invariants. SplInter is property-directed, precise, and produces counterexample traces in case a property does not hold. Using the novel notion of spatial interpolants modulo theories, SplInter can infer complex invariants over general recursive predicates, e. g., of the form all data elements in a linked list are even or a binary tree is sorted. Furthermore, we treat interpolation as a black box, which gives us the freedom to encode data manipulation in whatever theory is suitable for the program at hand (e. g., bitvectors, arrays, or linear

arithmetic), so that our technique immediately benefits from any future advances in SMT solving and interpolation.

3.3 The Transformer Refinement Prover – A Work in (Lack of) Progress Talk

Martin Brain (University of Oxford, GB)

Over the past few years there have been a number of interesting papers showing that algorithms used for SAT solving (Stalmarck's, DPLL, CDCL) can be lifted to work over abstract domains. This raises the question of whether the architecture of a SAT solver can be lifted to produce a generic solver which can be parameterised with different abstract domains to produce a range of concrete decision procedures. The Transformer Refinement Prover (TRP) is an attempt to build such a tool. This talk discusses some of the issues engineering, architectural and conceptual that have been encountered during the construction.

3.4 A formal approach to the Analysis of Reliability Architectures

Alessandro Cimatti (Bruno Kessler Foundation – Trento, IT)

Joint work of Bozzano, Marco; Cimatti, Alessandro; Mattarei, Cristian
Main reference M. Bozzano, A. Cimatti, C. Mattarei, "Efficient Analysis of Reliability Architectures via Predicate Abstraction," in Proc. of the 9th Int'l Haifa Verification Conf. (HVC'13), LNCS, Vol. 8244, pp. 279–294, Springer, 2013.
URL http://dx.doi.org/10.1007/978-3-319-03077-7_19

The development of complex and critical systems calls for a rigorous and thorough evaluation of reliability aspects. Over the years, several methodologies have been introduced in order to aid the verification and analysis of such systems. Despite this fact, current technologies are still limited to specific architectures, without providing a generic evaluation of redundant system definitions.

In this talk, we present a novel approach able to assess the reliability of an arbitrary combinatorial redundant system. We rely on an expressive modeling language to represent a wide class of architectural solutions to be assessed. On such models, we provide a portfolio of automatic analysis techniques: we can produce a fault tree, that represents the conditions under which the system fails to produce a correct output; from it, we can extract a function over the components reliability, which represents the failure probability of the system. At its core, the approach relies on the logical formalism of equality and uninterpreted functions. Advanced automated reasoning techniques, in particular Satisfiability Modulo Theories decision procedures, and Predicate Abstraction, are suitably combined to achieve efficiency.

We carried out an extensive experimental evaluation of the proposed approach on a wide class of multi-stage redundant systems. We are able to obtain, in a fully automated manner, all the results that are manually obtained in previous works, and we cover a much wider class of architectures, demonstrating scalability for a large number of components, thus enabling the analysis of complex architectures of realistic size.

3.5 Lifting Satisfiability Procedures to Reachability Analysis

Vijay DSilva (Google – San Francisco, US)

Joint work of D'Silva, Vijay; Haller, Leopold; Kroening, Daniel
Main reference V. D'Silva, L. Haller, D. Kroening, "Abstract Conflict Driven Learning," in Proc. of the 40th Annual ACM SIGPLAN-SIGACT Symp. on Principles of Programming Languages (POPL'13), pp. 143–154, ACM, 2013.
URL http://dx.doi.org/10.1145/2429069.2429087

There have been many attempts at lifting ideas from the satisfiability literature to program analysis. The intuition behind the abstract satisfaction approach is that the objects manipulated by satisfiability solvers satisfy similar axioms to abstract domains used in program analysis. An immediate consequence of identifying these axioms is that satisfiability procedures can be lifted to all structures satisfying these axioms. In this work, we show that the DPLL and Conflict Driven Clause Learning (CDCL) algorithms lift to certain families of lattices and transformers used for reachability analysis. The resulting reachability analyzers automatically refine fixed points using the notions of decisions and learning. This approach has been applied successfully to bound the error of floating point computations in embedded software.

References

1 Vijay D'Silva, Leopold Haller and Daniel Kroening. *Satisfiability Solvers are Abstract Interpreters*. Static Analysis Symposium, 2012
2 Vijay D'Silva, Leopold Haller and Daniel Kroening. *Abstract Conflict Driven Learning*. Principles of Programming Languages, 2013

3.6 Solving Exists/Forall Problems With SMT

Bruno Dutertre (SRI – Menlo Park, US)

Joint work of Dutertre, Bruno; Jovanovic, Dejan

We describe an algorithm for solving problems of the form exists x. for all y. P(x, y) by relying of two SMT or SAT solvers. One solver searches for candidate x while the other attempts to refute x by exhibiting a y for which P(x, y) is false. A key component of this algorithm is generalizing from a counterexample y. We describe generalization methods that work for different quantification domains, including a method based on model-guided virtual term substitution.

3.7 Predicate Abstraction in IC3

Alberto Griggio (Bruno Kessler Foundation – Trento, IT)

Joint work of Cimatti, Alessandro; Griggio, Alberto; Mover, Sergio; Tonetta, Stefano
Main reference A. Cimatti, A. Griggio, S. Mover, S. Tonetta, "IC3 Modulo Theories via Implicit Predicate Abstraction," in Proc. of the 20th Int'l Conf. on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'14), LNCS, Vol. 8413, pp. 46–61, Springer, 2014.
URL http://dx.doi.org/10.1007/978-3-642-54862-8_4

We present a novel approach for generalizing the IC3 algorithm for invariant checking from finite-state to infinite-state transition systems, expressed over some background theories. The procedure is based on a tight integration of IC3 with Implicit (predicate) Abstraction, a technique that expresses abstract transitions without computing explicitly the abstract system and is incremental with respect to the addition of predicates. In this scenario, IC3 operates only at the Boolean level of the abstract state space, discovering inductive clauses over the abstraction predicates. Theory reasoning is confined within the underlying SMT solver, and applied transparently when performing satisfiability checks. When the current abstraction allows for a spurious counterexample, it is refined by discovering and adding a sufficient set of new predicates. Importantly, this can be done in a completely incremental manner, without discarding the clauses found in the previous search.

3.8 Property Directed Polyhedral Abstraction

Arie Gurfinkel (Carnegie Mellon University – Pittsburgh, US)

Joint work of Gurfinkel, Arie; Bjørner, Nikolaj

We show how to combine the benefits of Polyhedral Abstract Interpretation (poly-AI) with the flexibility of Property Directed Reachability (PDR) algorithms for computing safe inductive convex polyhedral invariants. We develop two algorithms that integrate Poly-AI with PDR and show their benefits on a prototype in Z3 using a preliminary evaluation. The algorithms mimic the traditional forward Kleene and a chaotic backward iterations, respectively. Our main contribution is to show how to replace the expensive convex hull and quantifier elimination computations, a major bottleneck in poly-AI, with a lazy property-directed algorithms based on interpolation and model-based projection. Our approach integrates seamlessly within the framework of PDR adapted to Linear Real Arithmetic, and allows to dynamically decide between computing convex and non-convex invariants as directed by the property.

3.9 The PAGAI static analyzer

Julien Henry (VERIMAG – Gières, FR)

Joint work of Henry, Julien; Monniaux, David; Moy, Matthieu

Pagai is a static analyzer based on combinations of abstract interpretation and SMT, that computes numerical invariants for LLVM bitcode. Abstract interpretation can be made more precise by distinguishing every paths inside loops for delaying least upper bounds, at the expense of an exponential blowup. SMT allows symbolic and sparse representation of sets of paths, and the fixpoint computation is guided by SMT queries. We present early but promising experimental results of PAGAI on the SV-Comp benchmarks. In a second part, we present a new approach to the estimation of Worst-Case execution time, by defining the problem as an instance of optimization modulo theory. Naive encodings of the problem into SMT lead to formulas intractable for any production- grade solver based on DPLL(T). We show that simple static pre-analysis of program fragments provide invariants that dramatically improve the efficiency of the SMTsolver on these examples.

3.10 Projection using Parametric Objectives

Jacob Howe (City University – London, GB)

Joint work of Howe, Jacob; King, Andy
Main reference J. M. Howe, A. King, "Polyhedral Analysis using Parametric Objectives," in Proc. of the 19th Int'l Symp. on Static Analysis (SAS'12), LNCS, Vol. 7460, pp. 41–57, Springer, 2012.
URL http://dx.doi.org/10.1007/978-3-642-33125-1_6

The abstract domain of polyhedra lies at the heart of many program analysis techniques. However, its operations can be expensive, precluding their application to polyhedra that involve many variables. This talk describes a new approach to computing polyhedral domain operations. The core of this approach is an algorithm to calculate variable elimination (projection) based on parametric linear programming. The algorithm enumerates only non-redundant inequalities of the projection space, hence permits anytime approximation of the output. Some preliminary data from experiments are included.

3.11 Abstract Conflict-Driven Learning

Daniel Kroening (University of Oxford, GB)

Joint work of D'Silva, Vijay; Haller, Leopold; Kroening, Daniel
Main reference V. D'Silva, L. Haller, D. Kroening, "Abstract Conflict Driven Learning," in Proc. of the 40th Annual ACM SIGPLAN-SIGACT Symp. on Principles of Programming Languages (POPL'13), pp. 143–154, ACM, 2013.)
URL http://dx.doi.org/10.1145/2429069.2429087

Modern satisfiability solvers implement an algorithm, called Conflict Driven Clause Learning, which combines search for a model with analysis of conflicts. We show that this algorithm can be generalised to solve the lattice-theoretic problem of determining if an additive

transformer on a Boolean lattice is always bottom. Our generalised procedure combines overapproximations of greatest fixed points with underapproximations of least fixed points to obtain more precise results than computing fixed points in isolation. We generalise implication graphs used in satisfiability solvers to derive underapproximate transformers from overapproximate ones. Our generalisation provides a new method for static analyzers that operate over non-distributive lattices to reason about properties that require disjunction.

3.12 A method to infer inductive numeric invariants inspired by Constraint Programming

Antoine Mine (ENS – Paris, FR)

Joint work of Mine, Antoine; Truchet, Charlotte; Sankaranarayanan, Sriram

In this talk, we suggest the idea of using algorithms inspired by Constraint Programming in order to infer inductive invariants on numeric programs. Similarly to Constraint Programming solvers on continuous domains, our algorithm approximates the problem from above, using decreasing iterations that may split, discard, and tighten axis-aligned boxes. If successful, the algorithm outputs a set of boxes that includes the initial states and is a post-fixpoint of the abstract semantic function of interest. Our work is very preliminary; many improvements still need to be performed to determine if the method is usable in practice, and in which contexts. Nevertheless, we show that a naive proof-of-concept implementation of our algorithm is already capable of inferring non-trivial inductive invariants that would otherwise require the use of relational or even non-linear abstract domains when using more traditional abstract interpretation iteration methods.

3.13 Automating Separation Logic with Trees and Data Using SMT Solvers

Ruzica Piskac (Yale University, US)

Joint work of Piskac, Ruzica; Wies, Thomas; Zufferey, Damien
Main reference R. Piskac, T. Wies, D. Zufferey, "Automating Separation Logic with Trees and Data," in Proc. of the 26th Int'l Conf. on Computer Aided Verification (CAV'14), LNCS, Vol. 8559, pp. 711–728, Springer, 2014.
URL http://dx.doi.org/10.1007/978-3-319-08867-9_47

Separation logic (SL) follows a discipline of local reasoning that mimics human intuition about how to prove the correctness of heap-manipulating programs. Central to this discipline is the frame rule, a Hoare logic proof rule that decomposes the global heap into a footprint, the region on which a program fragment operates, and a frame, the region that remains untouched by the program fragment. Automation of the frame rule involves the actual inference of the frame from SL assertions expressing the global heap and the footprint. In this talk, I present Grasshopper, a tool for compositional verification of heap-manipulating programs. What makes our tool unique is its decidable specification language, which supports mixing of assertions expressed in separation logic and first-order logic. We achieve this combination of specification languages through a translation to programs whose specifications are expressed in a decidable fragment of first-order logic. This logic is well-suited for automation using SMT solvers.

3.14 Verification with Recursive Functions

Régis Blanc (EPFL – Lausanne, CH)

Joint work of Blanc, Régis; Kuncak, Viktor; Suter, Philippe; Kneuss, Etienne
Main reference R. Blanc, V. Kuncak, E. Kneuss, P. Suter, "An overview of the Leon verification system: verification by translation to recursive functions," in Proc. of the 4th Workshop on Scala (SCALA'13), Article No. 1, ACM, 2013.
URL http://dx.doi.org/10.1145/2489837.2489838

We present the Leon system, a verifier for a subset of the Scala programming language. Along with several functional features of Scala, Leon supports imperative constructs such as mutations and loops, using a translation into recursive functional form. Both properties and programs in Leon are expressed in terms of user-defined functions. We discuss several techniques that led to an efficient semi-decision procedure for first-order constraints with recursive functions, which is the core solving engine of Leon. We illustrate the current capabilities of Leon on an interactive web interface.

3.15 Setting the Scene for "Decision Procedures and Abstract Interpretation"

Thomas W. Reps (University of Wisconsin – Madison, US)

Joint work of Reps, Thomas W.; Thakur, Aditya V.
Main reference T. W. Reps, A. V. Thakur, "Through the lens of abstraction," presentation at the 2014 High Confidence Software and Systems Conf. (HCSS'14); manuscript available from author's webpage.
URL http://www.cs.wisc.edu/wpis/papers/hcss14-final-version.pdf

This talk is intended as a "scene-setting" talk about Seminar 14351 for the benefit of the participants of both Seminars 14351 and 14352. It presents a somewhat personal view of the opportunity offered by the seminar, and concentrates mainly on two topics. The first is the use of logic to support abstract interpretation (i. e., for performing alpha on a set of states described by a formula, for applying the best transformer, for creating a representation of the best transformer, and for creating the reduced product of two or more values). The second is the use of abstract interpretation to support decision procedures better (e. g., by reverse-engineering existing decision procedures to identify uses of abstract domains, which allows them to be generalized by using more expressive abstract domains; and by using logic-based abstraction methods directly for unsatisfiability checking and validity checking).

The seminar is intended to expose members of two communities to each other, namely, (i) designers/implementers of abstract interpreters, and (ii) designers/implementers of decision procedures. One connection between the two communities is that the tools that are created by their respective members can be viewed as using symbolic techniques to explore a state space. However, the repertoires of techniques used in the two disciplines are quite different, and each has its own mindset and outlook. The ideas and methods presented in the talk demonstrate new connections between the two disciplines, and suggest that the time is ripe for cross-fertilization between them to occur.

3.16 On Suspicious Intervals for Floating-Point Number Programs

Michel Rueher (University of Nice, FR)

Joint work of Ponsini, Olivier ; Michel, Claude; Rueher, Michel
Main reference O. Ponsini, C. Michel, M. Rueher, "Verifying floating-point programs with constraint programming and abstract interpretation techniques," Automated Software Engineering, published online, 2014.
URL http://dx.doi.org/10.1007/s10515-014-0154-2

Programs with floating-point computations are often derived from mathematical models or designed with the semantics of the real numbers in mind. However, for a given input, the computed path with floating-point numbers may differ from the path corresponding to the same computation with real numbers. State-of-the-art tools compute an over- approximation of the error introduced by floating-point operations with respect to the same sequence of operations in an idealized semantics of real numbers. Thus, totally inappropriate behaviours of a program may be dreaded but the developer does not know whether these behaviours will actually occur, or not. That is why it is very important to estimate the accuracy of floating-point computations with respect to the same sequence of operations in an idealized semantics of real numbers. To tackle this problem, we will present some capabilities of CP techniques for: a) Computing tight approximations for value analysis; b) Identify suspicious values. The crux of the matter is the accuracy of the estimation of floating-point computations because a rough approximation may generate numerous false alarms. We show that a hybrid approach for value analysis of floating-point programs that combines abstract interpretation and CP techniques is more effective than static analyser and CP solvers, when used separately. Interestingly, the refutation capabilities of CP solvers over floating-point numbers and over real numbers can significantly improve the precision of the domains computed by abstract interpretation. When the approximation remains nevertheless too rough, CP techniques can also help to identify suspicious values, that is values for which the behaviour of the program over the floating-point numbers is significantly different from the behaviour one could expect over the real numbers. In other words, for verifying whether a program can actually produce values inside the part of the approximation over the floats that intersect with a forbidden interval.

3.17 Inferring Invariants by Strategy Iteration

Peter Schrammel (University of Oxford, GB)

Joint work of Daniel Monniaux; Schrammel, Peter
Main reference D. Monniaux, P. Schrammel, "Speeding Up Logico-Numerical Strategy Iteration," in Proc. of the 21st Int'l Symp. on Static Analysis (SAS'14), LNCS, Vol. 8723, pp. 253–267, Springer, 2014.
URL http://dx.doi.org/10.1007/978-3-319-10936-7_16

Template polyhedral analysis abstracts numerical variables inside a program by one polyhedron per control location, with a priori fixed directions for the faces. The strongest inductive invariant in such an abstract domain may be computed by a combination of strategy iteration and SMT solving. Unfortunately, the above approaches lead to unacceptable space and time costs if applied to a program whose control states have been partitioned according to predicates. We therefore propose a modification of the strategy iteration algorithm where the strategies are stored succinctly, and the linear programs to be solved at each iteration step are simplified according to an equivalence relation.

3.18 Solvers, Abstraction, and Inductive Learning

Sanjit A. Seshia (University of California – Berkeley, US)

Main reference S. A. Seshia, "Sciduction: Combining Induction, Deduction, and Structure for Verification and Synthesis," in Proc. of the 49th Annual Design Automation Conf. (DAC'12), pp. 356–365, ACM, 2012; pre-print available from author's webpage.
URL http://dx.doi.org/10.1145/2228360.2228425
URL http://www.eecs.berkeley.edu/~sseshia/pubs/b2hd-seshia-dac12.html

This talk seeks to make connections between three topics: decision procedures (SMT solvers), abstraction and abstract interpretation, and inductive learning (machine learning). There are three main messages in the talk. First, we make the point that many verification tasks are effectively solved through "reduction to synthesis". Examples include the generation of inductive invariants for proofs by induction, and the generation of abstract models for abstraction-based verification. Second, the resulting synthesis problems can be tackled through a combination of induction (learning from examples), deduction, and a structure hypothesis. An example is the counterexample-guided inductive synthesis (CEGIS) paradigm. We compare CEGIS with "traditional" machine learning algorithms. Finally, we pose some fundamental questions about the efficiency and convergence (termination) of CEGIS. Initial results are presented that draw from results in machine learning theory.

3.19 Data-Driven Invariant Inference

Rahul Sharma (Stanford University, US)

Joint work of Sharma, Rahul;Alex Aiken; Saurabh Gupta; Bharath Hariharan; Aditya Nori
Main reference R. Sharma, A. Aiken, "From Invariant Checking to Invariant Inference Using Randomized Search," in Proc. of the 26th Int'l Conf. on Computer Aided Verification (CAV'14), LNCS, Vol. 8559, pp. 88–105, Springer, 2014.
URL http://dx.doi.org/10.1007/978-3-319-08867-9_6

We discuss two applications that leverage concrete states to improve invariant inference. First, for many abstract interpretations, concrete states can help reduce the number of fixpoint iterations required to reach convergence. In the second application, concrete states guide a search based invariant inference engine. The main advantage of a search based procedure is the generality and we show how to retarget our procedure to infer invariants for many different domains.

3.20 Verification using Small and Short Worlds

Rohit Sinha (University of California – Berkeley, US)

Joint work of Sinha, Rohit; Sturton, Cynthia; Maniatis, Petros; Seshia, Sanjit; Wagner, David
Main reference R. Sinha, C. Sturton, P. Maniatis, S. A. Seshia, D. Wagner, "Verification with Small and Short Worlds," in Proc. of the 2012 IEEE Int'l Conf. on Formal Methods in Computer-Aided Design (FMCAD'12), pp. 68–77, IEEE, 2012.
URL http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=6462557

We consider the verification of safety properties in systems with large arrays and data structures. Such systems are common at the low levels of software stacks; examples are

hypervisors and CPU emulators. The very large data structures in such systems (e. g., address-translation tables and other caches) make automated verification based on straightforward state-space exploration infeasible. We present S2W, a new abstraction-based model-checking methodology to facilitate automated verification of such systems. As a first step, inductive invariant checking is performed. If that fails, we compute an abstraction of the original system by precisely modeling only a subset of state variables while allowing the rest of the state to evolve arbitrarily at each step. This subset of the state constitutes a "small world" hypothesis, and is extracted from the property. Finally, we verify the safety property on the abstract model using bounded model checking. We ensure the verification is sound by first computing a bound on the reachability diameter of the abstract model. For this computation, we developed a set of heuristics that we term the "short world" approach.

3.21 More Algorithms for Symbolic Abstraction

Aditya Thakur (University of Wisconsin – Madison, US)

Joint work of Thakur, Aditya; Reps, Thomas; Breck, Jason

This talk describes two algorithms for performing *symbolic abstraction* [1]: Given a formula φ in a logic $\mathcal{L}$ and an abstract domain $\mathcal{A}$, the symbolic abstraction of φ is the strongest consequence of φ that is expressible in $\mathcal{A}$. Symbolic abstraction has a dual use: it can be used to compute abstract transformers in abstract interpretation, and to check unsatisfiability of a formula.

The talk presents the bilateral framework [2] for performing symbolic abstraction that maintains an over-approximation and under-approximation of the final answer. The algorithm performs symbolic abstraction by intelligently querying an SMT solver. The framework was instantiated to synthesize abstract transformers for machine- code analysis.

The next algorithm for symbolic abstraction is applicable to a new fragment of separation logic (SL) [3]. The algorithm works by performing a bottom-up evaluation of the formula using an abstract domain of shape graphs. This algorithm can be used to check unsatisfiability of an SL formula.

References

1 Reps, T., Sagiv, M., Yorsh, G. (2004). *Symbolic implementation of the best transformer.* In Verification, Model Checking, and Abstract Interpretation (pp. 252-266). Springer Berlin Heidelberg.
2 Thakur, A., Elder, M., Reps, T. (2012). *Bilateral algorithms for symbolic abstraction.* In Static Analysis (pp. 111-128). Springer Berlin Heidelberg.
3 Thakur, A., Breck, J., Reps, T. (2014). *Satisfiability modulo abstraction for separation logic with linked lists* . In Symposium on Model Checking of Software (SPIN) (pp. 58–67). ACM.

3.22 Abstract Domains for Constraint Programming

Charlotte Truchet (University of Nantes, FR)

Joint work of Truchet, Charlotte; Marie Pelleau; Antoine Miné; Frédéric Benhamou
Main reference M. Pelleau, A. Miné, C. Truchet, F. Benhamou, "A Constraint Solver Based on Abstract Domains," in Proc. of the 14th Int'l Conf. on Verification, Model Checking, and Abstract Interpretation (VMCAI'13), LNCS, Vol. 7737, pp. 434–454, Springer, 2013.
URL http://dx.doi.org/10.1007/978-3-642-35873-9_26

Constraint Programming (CP) in a domain of Artificial Intelligence that offers a variety of tools to model and solve hard combinatorial problems. Abstract Interpretation (AI) is a domain of Semantics that studies approximations of program traces in order to prove correctness properties. Although they are distinct scientific areas, these two domains have a lot in common. In both cases, we are interested in some set that is hard to compute or intractable: solution set in CP, concrete domain in AI. In both cases, instead of computing this set, we study some over-approximations of it: abstract domains in AI, consistent domains in CP. But the methods differ when the over-approximations are not good enough. CP has developed sophisticated algorithmic mechanisms to exactly solve the problem if the variables are discrete, or reach a given precision if they are continuous. In AI, the abstract domains themselves are refined, either by adding operators, or by increasing their expressivity. In the end, CP provides with solving methods that are very efficient on many NP problems, but are rather monolithic. For instance, they are unable to solve mixed problems with integer and real variables. AI analyzes huge programs using a lot of expressive abstract domains, but does not feature a notion of precision.

In this talk, we showed how to introduce the notion of abstract domain in CP using the example of the Octagons, which offer a good trade between efficiency and expressivity. Then we presented Absolute, a prototype constraint solver without constraints: it is built upon a library of abstract domains called Apron, by Miné and Jeannet. It, thus, naturally copes with mixed problems.

3.23 Ideal Abstractions

Thomas Wies (New York University, US)

Joint work of Damien Zufferey, Tom Henzinger
Main reference D. Zufferey, T. Wies, T. A. Henzinger, "Ideal Abstractions for Well-Structured Transition Systems," in Proc. of the 13th Int'l Conf. on Verification, Model Checking, and Abstract Interpretation (VMCAI'12), LNCS, Vol. 7148, pp. 445–460, Springer, 2102.
URL http://dx.doi.org/10.1007/978-3-642-27940-9_29

Many concurrent infinite state systems can be seen as well-structured transition systems (WSTS). Examples include concurrent programs with shared-memory and dynamic thread creation, as well as distributed message passing systems in the actors framework. WSTS are an attractive class of systems for formal analysis because they admit generic decision procedures for important verification problems such as coverability. Unfortunately, these decision procedures often have very high complexity or provide termination guarantees only in special cases that are not of practical relevance. To obtain a practical analysis with more general termination guarantees, we propose an abstract interpretation that is inspired by decision procedures for the covering problem of WSTS. The abstract domain of our

analysis builds on the ideal completion of the well-quasi-ordered state space to obtain an efficient symbolic representation of infinite sets of states. A widening operator that mimics acceleration-based forward algorithms for computing covering sets ensures termination while controlling the loss of precision of the analysis. I will resent an instance of our analysis framework for the class of depth-bounded WSTS and its application to verifying progress properties of concurrent data structure implementations.

3.24 Parametric Program Analysis

Hongseok Yang (University of Oxford, GB)

Joint work of Hongseok Yang, Mayur Naik, Xin Zhang, Ravi Mangal, Radu Grigore, Ghila Castelnuovo, Mooly Sagiv

Recent years have seen the development of successful commercial programming tools based on static analysis technologies, which automatically verify intended properties of programs or find tricky bugs that are difficult to detect by testing techniques. One of the key reasons for this success is that these tools use clever strategies for abstracting programs—most details about a given program are abstracted away by these strategies, unless they are predicted to be crucial for proving a given property about the program or detecting a type of program errors of interest. Developing such a strategy is, however, nontrivial, and is currently done by a large amount of manual engineering efforts in most tool projects. Finding a good abstraction strategy automatically or even reducing these manual efforts involved in the development of such a strategy is considered one of the main open challenges in the area of program analysis.

In this talk, I will explain how I tried to address this challenge with colleagues in the past few years. During this time, we worked on parametric program analyses, where parameters for controlling the degree of program abstraction are expressed explicitly in the specification of the analyses. For those analyses, we developed algorithms for searching for a desired parameter value with respect to a given program and a given property, which use ideas from the neighbouring areas of program analysis such as testing, searching and optimisation. In my talk, I will describe the main ideas behind these algorithms without going into technical details. I will focus on intuitions about why and when these algorithms work. I will also talk briefly about a few lessons that I learnt while working on this problem.

3.25 A Simple and Scalable Static Analysis for Bound Analysis and Amortized Complexity Analysis

Florian Zuleger (TU Wien, AT)

Joint work of Sinn, Moritz; Zuleger, Florian; Veith, Helmut
Main reference M. Sinn, F. Zuleger, H. Veith, "A Simple and Scalable Static Analysis for Bound Analysis and Amortized Complexity Analysis," in Proc. of the 26th Int'l Conf. on Computer Aided Verification (CAV'14), LNCS, Vol. 8559, pp. 745–761, Springer, 2014.
URL http://dx.doi.org/10.1007/978-3-319-08867-9_50

We present the first scalable bound analysis that achieves amortized complexity analysis. In contrast to earlier work, our bound analysis is not based on general purpose reasoners such as abstract interpreters, software model checkers or computer algebra tools. Rather, we

derive bounds directly from abstract program models, which we obtain from programs by comparatively simple invariant generation and symbolic execution techniques. As a result, we obtain an analysis that is more predictable and more scalable than earlier approaches. We demonstrate by a thorough experimental evaluation that our analysis is fast and at the same time able to compute bounds for challenging loops in a large real-world benchmark.

Technically, our approach is based on lossy vector addition systems (VASS). Our bound analysis first computes a lexicographic ranking function that proves the termination of a VASS, and then derives a bound from this ranking function. Our methodology achieves amortized analysis based on a new insight how lexicographic ranking functions can be used for bound analysis.

Participants

- Aws Albarghouthi
University of Toronto, CA
- Joshua Berdine
Microsoft Research UK – Cambridge, GB
- Régis Blanc
EPFL – Lausanne, CH
- Martin Brain
University of Oxford, GB
- Jörg Brauer
Verified Systems International GmbH – Bremen, DE
- Alessandro Cimatti
Bruno Kessler Foundation – Trento, IT
- Vijay D'Silva
Google – San Francisco, US
- Bruno Dutertre
SRI – Menlo Park, US
- Alberto Griggio
Bruno Kessler Foundation – Trento, IT
- Arie Gurfinkel
Carnegie Mellon University, US
- Julien Henry
VERIMAG – Gières, FR
- Jacob Howe
City University – London, GB
- Daniel Kroening
University of Oxford, GB
- Akash Lal
Microsoft Research India – Bangalore, IN
- Antoine Miné
ENS – Paris, FR
- Ruzica Piskac
Yale University, US
- Thomas W. Reps
University of Wisconsin – Madison, US
- Michel Rueher
University of Nice, FR
- Peter Schrammel
University of Oxford, GB
- Sanjit A. Seshia
University of California – Berkeley, US
- Rahul Sharma
Stanford University, US
- Rohit Sinha
University of California – Berkeley, US
- Aditya Thakur
University of Wisconsin – Madison, US
- Charlotte Truchet
University of Nantes, FR
- Thomas Wies
New York University, US
- Hongseok Yang
University of Oxford, GB
- Florian Zuleger
TU Wien, AT

Report from Dagstuhl Seminar 14352

Next Generation Static Software Analysis Tools

Edited by

Patrick Cousot[1], Daniel Kroening[2], and Carsten Sinz[3]

1 **ENS – Paris, FR,** pcousot@cims.nyu.edu
2 **University of Oxford, GB,** kroening@cs.ox.ac.uk
3 **KIT – Karlsruhe Institute of Technology, DE,** carsten.sinz@kit.edu

Abstract

There has been tremendous progress in static software analysis over the last years with, for example, refined abstract interpretation methods, the advent of fast decision procedures like SAT and SMT solvers, new approaches like software (bounded) model checking or CEGAR, or new problem encodings. We are now close to integrating these techniques into every programmer's toolbox.

The aim of the seminar was to bring together developers of software analysis tools and algorithms, including researchers working on the underlying decision procedures (e. g., SMT solvers), and people who are interested in applying these techniques (e. g. in the automotive or avionics industry).

The seminar offered the unique chance, by assembling the leading experts in these areas, to make a big step ahead towards new, more powerful tools for static software analysis.

Current (academic) tools still suffer from some shortcomings:

- Tools are not yet robust enough or support only a subset of a programming language's features.
- Scalability to large software packages is not yet sufficient.
- There is a lack of standardized property specification and environment modeling constructs, which makes exchange of analysis results more complicated than necessary.
- Differing interpretations of programming language semantics by different tools lead to limited trust in analysis results.
- Moreover, a comprehensive benchmark collection to compare and evaluate tools is missing.

Besides these application-oriented questions, further, more fundamental questions have also been topics of the seminar:

- What are the right logics for program verification, bug finding and software analysis? How can we handle universal quantification? And how to model main memory and complex data structures?
- Which decision procedures are most suitable for static software analysis? How can different procedures be combined? Which optimizations to general-purpose decision procedures (SAT/SMT/QBF) are possible in the context of software analysis?

Seminar August 24–29, 2014 – http://www.dagstuhl.de/14352
1998 ACM Subject Classification D.2.4 Software/Program Verification, F.3.1 Specifying and Verifying and Reasoning about Programs
Keywords and phrases Software quality, Bug finding, Verification, Decision procedures, SMT/SAT solvers
Digital Object Identifier 10.4230/DagRep.4.8.107
Edited in cooperation with Christoph Gladisch

Next Generation Static Software Analysis Tools, *Dagstuhl Reports*, Vol. 4, Issue 8, pp. 107–125
Editors: Patrick Cousot, Daniel Kroening, and Carsten Sinz

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

1 Executive Summary

Patrick Cousot
Klaus Havelund
Daniel Kroening
Carsten Sinz
Christoph Gladisch

Software errors are still a widespread plague. They manifest themselves, e. g., in program crashes, malfunction, incorrect behavior, or security vulnerabilities. Even software that has been in use for decades and has been deployed to millions of users (e. g., the compression library zlib) still contains flaws that are revealed only now and have to be fixed. Both in academia and industry considerable effort has been undertaken to develop tools and methodologies to obtain fault-free software. Nowadays, static analysis tools, which search for program errors without running the software, have reached a state where they are, in some industries (e. g., the automotive or avionics industry), already part of the standard software development and quality assurance process (with tools and companies like, e. g., Polyspace, Coverity, KlocWork, AbsInt, or Astrée). And although these tools can help finding residual errors more quickly, they still suffer from some shortcomings:

- Lack in precision. For a certain fraction of program locations in the source code it cannot be decided whether there is an error or not. Such "undecided cases" require (often time-consuming) manual rework, limiting the value of such tools.
- Due to the manual effort required, static software analysis tools have not yet made their way to mainstream software development (besides industries, where software reliability is indispensable and considerable amounts of time and money are spent on quality assurance).

Over the last years, software analysis tools based on abstract interpretation have been refined and tools based on new core formalisms, such as model checking, have gained traction, mainly in the form of two key methods: counterexample-guided abstraction refinement (CEGAR), and bounded model checking (BMC). The success of these new tools was, to a substantial part, enabled by the enormous progress that was made on the underlying logical decision procedures (SAT and SMT solvers). New software analysis tools based on these techniques come with considerably improved precision (less false positives), but they are still not competitive with tools based on abstract interpretation with respect to scalability. Also, they are rarely used in industrial software development projects so far.

With this seminar we believe that we were able to stimulate further progress in this field by intensifying the collaboration between (a) researchers on new static software analysis tools, (b) scientists working on improved high-performance decision procedures, and (c) practitioners, who know what is needed in industry and which kind of software analysis tools are accepted by developers and which are not.

The Dagstuhl Seminar was attended by participants from both industry and academia. It included presentations on a wide range of topics such as:

- Recent trends in static analysis, consisting of new algorithms and implementation techniques.
- New decision procedures for software analysis, for example, to analyze programs with complex data structures.

- Industrial case studies: What are the problems industrial users of static analysis tools are facing?
- Experience reports and statements on current challenges.

The first day of the seminar started with an introduction round, in which each participant shortly presented his research interests. As the seminar was held concurrently with a second, closely related Dagstuhl Seminar on "Decision Procedures and Abstract Interpretation" (14351), the introductory session was held jointly by both seminars. Four overview talks were also organized jointly by both seminars, and were given by Thomas Reps, Patrick Cousot, Vijay Ganesh, and Francesco Logozzo.

There was also a tool demonstration session on Thursday afternoon, in which seven tools were presented (15 minutes each).

In further talks of the seminar young as well as senior researchers presented on-going and completed work. Tool developers and participants from industry reflected on current challenges in the realm of software analysis.

The seminar was concluded with a panel discussion about the current challenges of static software analysis for industrial application (see Sec. 5 for an extended exposition of the panel discussion).

We expect that with this Dagstuhl Seminar we were able to make a step forward towards bringing static software analysis tools to every programmer's workbench, and therefore, ultimately, improve software quality in general.

2 Table of Contents

3 Overview of Talks

3.1 Most Overlooked Static Analysis Pitfalls

Roberto Bagnara (BUGSENG & University of Parma, IT)

Quality software requires complex verification activities. Such activities cannot be practically and reliably performed without extensive use of tools. Poor-quality static analysis tools either result into higher costs of the verification process or fail their goal altogether, delivering a false sense of security instead of the promised quality for the developed software. There are well-engineered static analysis tools that are based on obsolete technology, to the point of ignoring 30+ years of research in software verification. There also are theoretically-sophisticated tools that fall short of their objectives due to poor engineering. In this talk I will illustrate what I believe are important pitfalls of the design of static analysis tools, drawing on the experience of the BUGSENG's team that developed the ECLAIR software verification platform.

3.2 CPAchecker: A Flexible Framework for Software Verification

Dirk Beyer (Universität Passau, DE)

Main reference D. Beyer, M. E. Keremoglu, "CPAchecker: A Tool for Configurable Software Verification," in Proc. of the 23rd Int'l Conf. on Computer Aided Verification (CAV'11), LNCS, Vol. 6806, pp. 184–190, Springer, 2011; pre-print available from author's webpage.
URL http://dx.doi.org/10.1007/978-3-642-22110-1_16
URL http://www.sosy-lab.org/~dbeyer/Publications/2011-CAV.CPAchecker_A_Tool_for_Configurable_Software_Verification.pdf

CPAchecker is a tool and framework that aims at easy integration of new verification components. It is based on configurable program analysis, a concept for implementing different approaches from data-flow analysis, abstract interpretation, and software model-checking in one uniform software framework. Every abstract domain, together with the corresponding operations, implements the interface of configurable program analysis (CPA). The main algorithm is configurable to perform a fixed-point analysis on arbitrary combinations of existing CPAs.

In software verification, it takes a considerable amount of effort to convert a verification idea into actual experimental results– we aim at accelerating this process. We hope that researchers and practitioners find it convenient and productive to implement new verification ideas and algorithms using this flexible and easy-to-extend platform, and that it advances the field by making it easier to perform practical experiments.

The tool is implemented in Java and runs as command-line tool or as Eclipse plug-in. CPAchecker has existing CPAs for several abstract domains already, including predicates, explicit values, octagons, and BDDs. The tool integrates CEGAR, lazy abstraction refinement, interpolation, boolean predicate abstraction, large-block encoding, bounded model checking, generation of error witnesses with test values, and several SMT solvers in a modular and flexible design. CPAchecker is publicly available under the open-source license Apache 2. The tool won several medals in competitions on software verification.

http://cpachecker.sosy-lab.org/

3.3 Abstract Interpretation: "Scene-Setting Talk"

Patrick Cousot (ENS – Paris, FR)

We unify static analysis by extrapolation (widening) with static analysis by interpolation to prove a given program specification. This unification is done in the theory of abstract interpretation using dual-narrowing. We show that narrowing and dual-narrowing are equivalent up to the exchange of their parameters. This yields new ideas for narrowing based on Craig interpolation. This unification is also possible by understanding that interpolation can be done in arbitrary abstract domains, not only logical ones. We show that an increasing iterative static analysis using extrapolation of successive iterates by widening followed by a decreasing iterative static analysis using interpolation of successive iterates by narrowing (both bounded by the specification) can be further improved by a increasing iterative static analysis using interpolation of iterates with the specification by dual-narrowing until reaching a fixpoint and checking whether it is inductive for the specification.

3.4 Abstracting Induction by Extrapolation and Interpolation

Patrick Cousot (ENS – Paris, FR)

We unify static analysis by extrapolation (widening) with static analysis by interpolation to prove a given program specification. This unification is done in the theory of abstract interpretation using dual-narrowing. We show that narrowing and dual-narrowing are equivalent up to the exchange of their parameters. This yields new ideas for narrowing based on Craig interpolation. This unification is also possible by understanding that interpolation can be done in arbitrary abstract domains, not only logical ones. We show that an increasing iterative static analysis using extrapolation of successive iterates by widening followed by a decreasing iterative static analysis using interpolation of successive iterates by narrowing (both bounded by the specification) can be further improved by a increasing iterative static analysis using interpolation of iterates with the specification by dual-narrowing until reaching a fixpoint and checking whether it is inductive for the specification.

3.5 Path-sensitive static analysis using trace hashing

Tomasz Dudziak (Universität des Saarlandes, DE)

Path-sensitivity can significantly improve precision of static program analysis but due to the inherently exponential number of possible control flow histories it often requires manual tweaking of parameters and annotations. I propose an approach based on hashing of control flow paths that can dynamically adapt to available resources. By careful construction of the hash function it can exploit additional assumptions about the nature of path-sensitive properties. Additionally, due to its randomized nature it introduces a new trade-off between cost of the analysis and probability of proving properties of interest.

3.6 An algebraic approach for inferring and using symmetries in rule-based models

Jerôme Feret (ENS – Paris, FR)

Symmetries arise naturally in rule-based models, and under various forms. Besides automorphisms between site graphs, which are usually built within the semantics, symmetries can take the form of pairs of sites having the same capabilities of interactions, of some protein variants behaving exactly the same way, or of some linear, planar, or 3D molecular complexes which could be see modulo permutations of their axis and/or mirror-image symmetries.

In this paper, we propose a unifying handling of symmetries in Kappa. We follow an algebraic approach, that is based on the single pushout semantics of Kappa. We model classes of symmetries as finite groups of transformations between site graphs, which are compatible with the notion of embedding (that is to say that it is always possible to restrict a symmetry that is applied with the co- domain of an embedding to the domain of this embedding) and we provide some assumptions that ensure that symmetries are compatible with pushouts. Then, we characterize when a set of rules is symmetric with respect to a group of symmetries and, in such a case, we give sufficient conditions so that this group of symmetries induces a forward bisimulation and/or a backward bisimulation over the population semantics.

3.7 Bounded Verification with TACO: Symmetry-breaking + tight field bounds

Marcelo Frias (University of Buenos Aires, AR)

Main reference J. P. Galeotti, N.Rosner, C. G. Lopez Pombo, M. F. Frias, "TACO: Efficient SAT-Based Bounded Verification Using Symmetry Breaking and Tight Bounds ," IEEE Trans. on Software Engineering, 39(9):1283–1307, 2013.
URL http://dx.doi.org/10.1109/TSE.2013.15

I will discuss work developed in my group on bounded verification of JML-annotated Java code, and how appropriate symmetry breaking predicates and bounds on the semantics of class fields allow us to improve the performance of TACO.

3.8 Impact of Community Structure on SAT Solver Performance

Vijay Ganesh (University of Waterloo, CA)

Joint work of Newsham, Zack; Ganesh, Vijay; Fischmeister, Sebastian; Audemard, Gilles; Simon, Laurent
Main reference Z. Newsham, V. Ganesh, S. Fischmeister, G. Audemard, L. Simon, "Impact of Community Structure on SAT Solver Performance," in Proc. of the 17th Int'l Conf. on Theory and Applications of Satisfiability Testing (SAT'14), LNCS, Vol. 8561, pp. 252–268, Springer, 2014.
URL http://dx.doi.org/10.1007/978-3-319-09284-3_20

Modern CDCL SAT solvers routinely solve very large industrial SAT instances in relatively short periods of time. It is clear that these solvers somehow exploit the structure of real-world instances. However, to-date there have been few results that precisely characterize this structure. In this paper, we provide evidence that the community structure of real-world

SAT instances is correlated with the running time of CDCL SAT solvers. It has been known for some time that real-world SAT instances, viewed as graphs, have natural communities in them. A community is a sub-graph of the graph of a SAT instance, such that this sub-graph has more internal edges than outgoing to the rest of the graph. The community structure of a graph is often characterized by a quality metric called Q. Intuitively, a graph with high-quality community structure (high Q) is easily separable into smaller communities, while the one with low Q is not. We provide three results based on empirical data which show that community structure of real-world industrial instances is a better predictor of the running time of CDCL solvers than other commonly considered factors such as variables and clauses. First, we show that there is a strong correlation between the Q value and Literal Block Distance metric of quality of conflict clauses used in clause-deletion policies in Glucose-like solvers. Second, using regression analysis, we show that the the number of communities and the Q value of the graph of real-world SAT instances is more predictive of the running time of CDCL solvers than traditional metrics like number of variables or clauses. Finally, we show that randomly-generated SAT instances with 0.05 <= Q <= 0.13 are dramatically harder to solve for CDCL solvers than otherwise.

3.9 Using a deductive verification environment for verification, bug finding, specification, and all that

Christoph Gladisch (KIT – Karlsruher Institut für Technologie, DE)

Joint work of Gladisch, Christoph; Shmuel Tyszberowicz; Bernhard Beckert; Ferruccio Damiani; Mana Taghdiri; Mattias Ulbrich; Tianhai Liu; Aboubakr Achraf El Ghazi; Daniel Grunwald

The boundaries between program analysis techniques such as static analysis, deductive verification, abstract interpretation, and model checking are overlapping. Deductive verification technology can, for instance, be used as a framework for the other techniques. In this talk deductive verification technology is taken as a basis and a set of techniques that were developed on the KeY platform are presented such as fault detection, model generation for quantified formulas, test generation, and specification techniques.

References

1 Daniel Grunwald and Christoph Gladisch and Tianhai Liu and Mana Taghdiri and Shmuel Tyszberowicz. *Generating JML Specifications from Alloy Expressions*. 10th Haifa Verification Conference (HVC), Israel, 2014
2 Aboubakr Achraf El Ghazi and Mattias Ulbrich and Christoph Gladisch and Shmuel Tyszberowicz and Mana Taghdiri. *JKelloy: A Proof Assistant for Relational Specifications of Java Programs*. NASA Formal Methods – 6th International Symposium, NFM 2014, Houston, TX, USA, 2014.
3 Wolfgang Ahrendt and Bernhard Beckert and Daniel Bruns and Richard Bubel and Christoph Gladisch and Sarah Grebing and Reiner Hähnle and Martin Hentschel and Vladimir Klebanov and Wojciech Mostowski and Christoph Scheben and Peter Schmitt and Mattias Ulbrich. *The KeY Platform for Verification and Analysis of Java Programs* 6th Working Conference on Verified Software: Theories, Tools and Experiments 2014, Vienna, Austria, 2014.

4 Christoph Gladisch and Shmuel Tyszberowicz. *Specifying a Linked Data Structure in JML for Formal Verification and Runtime Checking*. Brazilian Symposium on Formal Methods (SBMF), Brasilia, Brasil, 2013
5 Christoph Gladisch. *Model Generation for Quantified Formulas with Application to Test Data Generation*. International Journal on Software Tools for Technology Transfer (STTT), Volume 14, Number 4, 2012.
6 Christoph Gladisch. *Verification-based Software-fault Detection*. KIT Scientific Publishing, Karlsruhe, 2011.

3.10 Static Analysis of Energy Consumption

Manuel Hermenegildo (IMDEA Software – Madrid, ES)

Joint work of López, Pedro; Haemmerlé, Remy; Hermenegildo, Manuel V.; Klemen, Maximiliano; Liqat, Umer; Serrano, Alejandro; Eder, Kerstin; Georgiou, Kiryakos; Kerrison, Steve
Main reference U. Liqat, S. Kerrison, A. Serrano, K. Georgiou, P. López-Garcia, N. Grech, M. V. Hermenegildo, K. Eder, "Energy Consumption Analysis of Programs based on XMOS ISA-Level Models," in Proc. of the 23rd Int'l Symp. on Logic-Based Program Synthesis and Transformation (LOPSTR'13), to appear; pre-print available from author's webpage.
URL http://clip.dia.fi.upm.es/papers/isa-energy-lopstr13-final.pdf

Energy consumption is a major concern in data centers and high-performance computing, and there is also an increased demand for energy savings in devices which operate on batteries and other limited power sources, such as implantable/portable medical devices, sensors, or mobile phones. Beyond the advances in hardware power efficiency, significant additional energy savings can be achieved by improving the software. Static inference of the energy consumed by programs during execution is instrumental in this task, having important applications in the optimization and verification of such consumption by programs, and in general in energy-aware software development. At the same time it is an area that presents a number of interesting challenges.

We present an approach to the inference and verification of upper- and lower-bounds on the energy consumption of programs, as well as some current results from our tools. The bounds we infer and check are functions of the sizes of the input data to the program. Our tools are based on translating the program to a block-based intermediate representation, expressed as horn clauses, deriving cost equations, and finding upper- and lower-bound cost solutions. We also present some recent improvements to resource bounds inference, including casting the cost analysis more fully within abstract interpretation frameworks and using sized shapes as data abstractions. The energy analysis makes use of ISA- and LLVM-level models of the cost of instructions or sequences of instructions. The inferred bounds compare well to measurements on the hardware and open up new avenues for future research and application.

See the Dagstuhl Seminar slides (http://www.dagstuhl.de/mat/Files/14/14352/14352.HermenegildoManuel.Slides.pdf) for a full classified bibliography.

References

1 U. Liqat, S. Kerrison, A. Serrano, K. Georgiou, P. López-Garcia, N. Grech, M.V. Hermenegildo, K. Eder. Energy Consumption Analysis of Programs based on XMOS ISA-Level Models. In *Pre-proceedings of the 23rd International Symposium on Logic-Based Program Synthesis and Transformation (LOPSTR 13)*, September 2013.

2 A. Serrano, P. López-Garcia, M. Hermenegildo. Resource Usage Analysis of Logic Programs via Abstract Interpretation Using Sized Types. In *Theory and Practice of Logic Programming, 30th Int'l Conference on Logic Programming (ICLP'14) Special Issue*, Vol. 14, Num. 4–5, pages 739–754, Cambridge U. Press, 2014.

3.11 Insides and Insights of Commercial Program Analysis

Ralf Huuck (NICTA – Sydney, AU)

Joint work of Huuck, Ralf; Cassez, Franck; Fehnker, Ansgar

In this work we give an overview of the technologies underpinning our commercial C/C++ program analyzer Goanna and we share some of the experiences in applying these technologies to large industrial code bases. In particular, we highlight the core technologies of model checking, abstract interpretation and SMT-based automatic trace refinement as well as their interplay within the Goanna tool. We present some commercial experiences ranging from runtime metrics to tool comparison, and we highlight some of the challenges that we have been facing and as well as opportunities ahead.

3.12 Steps towards usable verification

Francesco Logozzo (Microsoft Research – Redmond, US)

Main reference M. Fähndrich, F. Logozzo, "Static contract checking with Abstract Interpretation," in Proc. of the 2010 Int'l Conf. on Formal Verification of Object-oriented Software (FoVeOOS'10), LNCS, Vol. 6528, pp. 10–30, Springer, 2010.
URL http://dx.doi.org/10.1007/978-3-642-18070-5_2

We describe our experience with the CodeContracts static checker (cccheck), probably the most successful verification tool out there: The analyzer has been downloaded over 150K times, and it is used in Microsoft product groups.

The cccheck is based on abstract interpretation, and it does not use any out-of- the-box SMT solver. In the talk I explain the rationale for this decision. Briefly, abstract interpretation allows us to have a very fine grain control on the precision/cost ration and it provides a level of automation unmatched by other approaches (e. g., for the inference of loop invariants, the suggestion of code fixes, or the generation of *sound* contracts). Furthermore, we avoid all kind of problems that come from using external generic tools, as, e. g., timeouts, non-monotonicity of the analysis, randomizations, non-determinism, etc.

3.13 Viper – Verification Infrastructure for Permission-based Reasoning

Peter Mueller (ETH Zürich, CH)

Joint work of Juhasz, Uri; Kassios, Ioannis; Müller, Peter; Novacek, Milos; Schwerhoff, Malte; Summers, Alexander

Main reference U. Juhasz, I. T. Kassios, P. Müller, M. Novacek, M. Schwerhoff, A. J. Summers, "Viper: A Verification Infrastructure for Permission-Based Reasoning," Technical Report, ETH Zürich, 2014.

URL http://pm.inf.ethz.ch/publications/getpdf.php?bibname=Own&id=JKMNSS14.pdf

The automation of verification techniques based on first-order logic specifications has benefited greatly from verification infrastructures such as Boogie and Why. These offer an intermediate language that can express diverse language features and verification techniques, as well as back-end tools such as verification condition generators.

However, these infrastructures are not well suited for verification techniques based on separation logic and other permission logics, because they do not provide direct support for permissions and because existing tools for these logics often prefer symbolic execution over verification condition generation. Consequently, tool support for these logics is typically developed independently for each technique, dramatically increasing the burden of developing automatic tools for permission-based verification.

In this talk, we present a verification infrastructure whose intermediate language supports an expressive permission model natively. We provide tool support, including two back-end verifiers, one based on symbolic execution, and one on verification condition generation; this facilitates experimenting with the two prevailing techniques in automated verification. Various existing verification techniques can be implemented via this infrastructure, alleviating much of the burden of building permission-based verifiers, and allowing the developers of higher-level techniques to focus their efforts at the appropriate level of abstraction.

3.14 Static Analysis Modulo Theory

Andreas Podelski (Universität Freiburg, DE)

A recent approach to static analysis can be described in analogy with SMT solving. Satisfiability here corresponds to the existence of an error path in the program, or: unsatisfiability corresponds to the emptiness of an automaton. Each time the tool finds an error path, i. e., a word accepted by the automaton, it analyzes the word in the theory of the data domain. If the word is infeasible, it learns a new automaton which rejects the word (and many others). It then adds the new automaton to the intersection of the already existing automata. We can extend the approach from sequential to recursive, parallel or unboundedly parallel programs. The emptiness check always amounts to a static analysis over a "theory-free" abstract domain.

3.15 Static Analysis Blind Spots in Automotive Systems Development

Hendrik Post (Robert Bosch GmbH – Stuttgart, DE)

Sound Static Analysis has become a mature and powerful technique for source code analysis. It is now time to stop and to contemplate whether scalability is still the most important problem or whether industrial applications lack other non-trivial contributions beyond improving performance. In this talk, we give an overview about interests and blockers for static analysis from an industrial perspective. Based on these findings, we give input for the workshop discussions.

3.16 Construction of modular abstract domains for heterogeneous properties

Xavier Rival (ENS – Paris, FR)

In this talk, we study the construction of shape-numeric static analysers. We set up an abstract interpretation framework that allows to reason about simultaneous shape-numeric properties by combining shape and numeric abstractions into a modular, expressive abstract domain. Such a modular structure is highly desirable to make its formalisation, proof and implementation easier to perform and to get correct. Furthermore, we extend this modular abstract domains so as to combine different memory abstractions, for better scalability and greater expressiveness. This framework is implemented in the MemCAD static analyser.

3.17 Efficiently Intertwining Widening with Narrowing

Helmut Seidl (TU München, DE)

Joint work of Seidl, Helmut; Apinis, Kalmer; Vojdani, Vesal
Main reference K. Apinis, H. Seidl, V. Vojdani, "How to combine widening and narrowing for non-monotonic systems of equations," in Proc. of the 34th ACM SIGPLAN Conf. on Programming Language Design and Implementation (PLDI'13), pp. 377–386, ACM 2013.
URL http://dx.doi.org/10.1145/2491956.2462190

Non-trivial analysis problems require posets with infinite ascending and descending chains. In order to compute reasonably precise post-fixpoints of the resulting systems of equations, Cousot and Cousot have suggested accelerated fixpoint iteration by means of widening and narrowing.

The strict separation into phases, however, may unnecessarily give up precision that cannot be recovered later. While widening is also applicable if equations are non-monotonic, this is no longer the case for narrowing. A narrowing iteration to improve a given post-fixpoint, additionally, must assume that all right-hand sides are monotonic. The latter assumption, though, is not met in presence of widening. It is also not met by equation

systems corresponding to context-sensitive interprocedural analysis, possibly combining context-sensitive analysis of local information with flow-insensitive analysis of globals.

As a remedy, we present a novel operator that combines a given widening operator with a given narrowing operator. We present adapted versions of round-robin as well as of worklist iteration, local and side-effecting solving algorithms for the combined operator and prove that the resulting solvers always return sound results and are guaranteed to terminate for monotonic systems whenever only finitely many unknowns are encountered.

3.18 Automating Software Analysis at Large Scale

Michael Tautschnig (Queen Mary University of London, GB)

Main reference Michael Tautschnig, "Automating Software Analysis at Large Scale," in Proc. of the 2014 Doctoral Workshop on Mathematical and Engineering Methods in Computer Science (MEMICS'14), to appear.

Software model checking tools promise to deliver genuine traces to errors, and sometimes even proofs of their absence. As static analysers, they do not require concrete execution of programs, which may be even more beneficial when targeting new platforms. Academic research focusses on improving scalability, yet largely disregards practical technical challenges to make tools cope with real-world code. The Debian/GNU Linux distribution proved to provide a perfect basis for experimenting with those tools. Initial experiments lead to a number of improvements in tools, but also more than 500 bug reports.

3.19 Automatic Inference of Ranking Functions by Abstract Interpretation

Caterina Urban (ENS – Paris, FR)

We present a family of parameterized abstract domains for proving termination of imperative programs by abstract interpretation. The domains automatically synthesize piecewise-defined lexicographic ranking functions and infer sufficient preconditions for program termination. The abstract domains are parameterized by a numerical abstract domain for state partitioning and a numerical abstract domain for ranking functions. This parameterization allows to easily tune the trade-off between precision and cost of the analysis. We describe instantiations of these domains with intervals, octagons, polyhedra and affine functions. We have implemented a prototype static analyzer for proving conditional termination of programs written in (a subset of) C and, using experimental evidence, we show that it is competitive with the state of the art and performs well on a wide variety of benchmarks.

3.20 Byte-Precise Verification of Low-Level List Manipulation

Tomas Vojnar (Technical University of Brno, CZ)

Main reference K. Dudka, P. Peringer, T. Vojnar, "Byte-Precise Verification of Low-Level List Manipulation," Brno University of Technology, Technical Report, No. FIT-TR-2012-04, 2013.
URL http://www.fit.vutbr.cz/~vojnar/Publications/FIT-TR-2012-04.pdf

We propose a new approach to shape analysis of programs with linked lists that use low-level memory operations. Such operations include pointer arithmetic, safe usage of invalid pointers, block operations with memory, reinterpretation of the memory contents, address alignment, etc. Our approach is based on a new representation of sets of heaps, which is to some degree inspired by works on separation logic with higher-order list predicates, but it is graph-based and uses a more fine-grained (byte-precise) memory model in order to support the various low-level memory operations. The approach was implemented in the Predator tool and successfully validated on multiple non-trivial case studies that are beyond the capabilities of other current fully automated shape analysis tools.

The result is a joint work with Kamil Dudka and Petr Peringer. The work was originally published at SAS'13. The Predator tool is available here: http://www.fit.vutbr.cz/research/groups/verifit/tools/predator/

3.21 System LAV and Automated Evaluation of Students' Programs

Milena Vujosevic-Janicic (University of Belgrade, RS)

Joint work of Vujosevic-Janicic, Milena; Kuncak, Viktor;
Main reference M. Vujošević-Janičić, V. Kuncak, "Development and Evaluation of LAV: An SMT-Based Error Finding Platform," in Proc. of the 4th Int'l Conf. on Verified Software: Theories, Tools, Experiments (VSTTE'12), LNCS, Vol. 7152, pp. 98–113, Springer, 2012.
URL http://dx.doi.org/10.1007/978-3-642-27705-4_9

In this talk, we give a short overview of a software verification tool LAV [1], we present challenges and experiences in applying verification techniques in automated evaluation of students' programs [2] and discuss our ongoing work on regression verification of students' programs.

LAV is an open-source tool for statically verifying program assertions and locating bugs such as buffer overflows, pointer errors and division by zero. It integrates into the popular LLVM infrastructure for compilation and analysis. Combining symbolic execution and SAT encoding of program's behaviour, LAV generates polynomial-size verification conditions for loop-free code, while for modelling loops it can use both under- or over- approximation techniques. Generated verification conditions are passed to one of the several SMT solvers: Boolector, MathSAT, Yices, and Z3.

Software verification tools are not commonly applied in automated evaluation of students' programs, although precise and reliable automated grading techniques are of big importance for both classical and on-line programming courses. We ran LAV on a corpus of students' programs, observed advantages and challenges of using verification in this context and we showed that verification techniques can significantly improve the automated grading process. LAV outperformed (concerning time, bugs found and false alarms) black-box fuzzing

techniques that are commonly used for bug finding in students' programs and successfully met all the specific requirements posed by evaluation process.

Our ongoing research focuses on functional correctness of small-sized programs written by students at introductory courses. We explore automatic assessment of functional correctness by regression verification (where the specification of a student's program is given as a teacher's program). Although this problem is undecidable in general, regression verification can give useful results and enhance automated grading process in some cases.

References

1 M. Vujošević-Janičić and V. Kuncak. Development and Evaluation of LAV: An SMT-Based Error Finding Platform. In *VSTTE*, volume 7152 of *LNCS*, pages 98–113. Springer, 2012.
2 M. Vujošević-Janičić, M. Nikolić, D. Tošić, and V. Kuncak. Software verification and graph similarity for automated evaluation of students' assignments. *Information and Software Technology*, 55(6):1004–1016, 2013. *Elsevier*.

4 Tool Demonstrations

4.1 aiT Worst-Case Execution Time Analysis

Christian Ferdinand (AbsInt – Saarbrücken, DE)

AiT WCET Analyzers statically compute tight bounds for the worst-case execution time (WCET) of tasks in real-time systems. They directly analyze binary executables and take the intrinsic cache and pipeline behavior into account.

4.2 The Goanna Static Analyzer

Ralf Huuck (NICTA – Sydney, AU)

We present some of the features and capabilities of our source code analyser Goanna. In particular, we show the IDE integration and usage in Visual Studio. This includes analysis features such as interprocedural tracing, selection for various coding standards and the bug management dashboard. Furthermore, we explain the Linux command line interface and demonstrate some exemplary bug finding capabilities. Finally, we present a number of benchmark detection results and open questions for future work.

4.3 Cccheck/Clousot

Francesco Logozzo (Microsoft Research – Redmond, US)

I demo cccheck, the popular abstract interpretation-based verifier for .NET. The demo includes:

1. finding bugs in C# programs;
2. provide *automatic* code fixes for such bugs;
3. show how Clousot infers contracts and it is also able to prove that a method computes the max of an array.

4.4 LLBMC: The Low-Level Bounded Model Checker

Carsten Sinz (KIT – Karlsruher Institut für Technologie, DE)

URL http://llbmc.org

We present LLBMC, the low-level bounded model checker. LLBMC implements a bounded model checking algorithm complemented by a rewriting approach to simplify verification conditions and prove simple properties. LLBMC is fully automatic and requires minimal preparation efforts and user interaction. It supports all C constructs, including not so common features such as bitfields. LLBMC models memory accesses (heap, stack, global variables) with high precision and is thus able to find hard-to-detect memory access errors like heap or stack buffer overflows. LLBMC can also uncover errors due to uninitialized variables or other sources of non-deterministic behavior. Due to its precise analysis, LLBMC produces almost no false alarms (false positives).

We demonstrate the features of LLBMC on three examples: the first is on checking equivalence of two programs containing many bit-wise logical operations; the second explains LLBMC's precise modeling of memory on a program where writing to memory has an unexpected effect on the control flow; the third example presents LLBMC's ability to derive lambda-expressions for loops with array updates.

For further information on LLBMC see http://llbmc.org. An evaluation / academic version can also be downloaded from this URL.

4.5 FuncTion

Caterina Urban (ENS – Paris, FR)

We present FuncTion, a research prototype static analyzer able to infer piecewise-defined ranking functions for programs written in (a subset of) C language. In particular, we present FuncTion's web interface and we demonstrate the features and capabilities of the analyzer by means of a few exemplary programs.

4.6 Predator: A Shape Analyzer Based on Symbolic Memory Graphs

Tomas Vojnar (Technical University of Brno, CZ)

URL http://www.fit.vutbr.cz/research/groups/verifit/tools/predator

Predator is a shape analyzer that uses the abstract domain of symbolic memory graphs (SMGs) in order to support various forms of low-level memory manipulation commonly used in optimized C code. Predator is implemented as a GCC (GNU Compiler Collection) plug-in. Predator is freely available at http://www.fit.vutbr.cz/research/groups/verifit/tools/predator.

5 Discussion Session

5.1 Discussion on "The current limitations of static analysis tools"

On the last day of the seminar a discussion session was held involving people from both academia and industry. The discussion focussed on the current limitations of static analysis tools and how to overcome them.

More specifically, we asked participants to think about:

- How can we improve the usability of static analysis tools and bring them to more users?
- What is the best way to combine algorithms and tools? Do we need a standardized exchange format?
- What is the "best" language for specifying properties and the environment, in which a program is run?
- What has to be done to bring static analysis tools to new fields such as security and privacy?

During the discussion it was observed that the most notable problems of current tools are:

- Annotations and specifications, which are essential to obtain precise analysis results and fewer false positives, are not standardized and can often not be exchanged between tools.
- There is a lack in detailed comparisons between static analysis tools, which makes it more difficult for a possible user to decide which tool to apply.
- Static analysis tools and compilers are not sufficiently integrated.

It was also argued that standardized specification and annotation languages exist, but are insufficient and thus not frequently used. It was also brought forward that even though there are benchmarks and competitions for analysis tools, it is still hard for a non-expert to decide which tool is most appropriate for a specific purpose. More work is needed to comparatively describe features and strengths of individual tools.

Participants

- Roberto Bagnara
BUGSENG & University of Parma, IT
- Dirk Beyer
Universität Passau, DE
- Mehdi Bouaziz
ENS, Paris, FR
- Patrick Cousot
ENS, Paris, FR
- Tomasz Dudziak
Universität des Saarlandes, DE
- David Faragó
KIT – Karlsruher Institut für Technologie, DE
- Christian Ferdinand
AbsInt, Saarbrücken, DE
- Jérôme Feret
ENS, Paris, FR
- Marcelo Frias
University of Buenos Aires, AR
- Vijay Ganesh
University of Waterloo, CA
- Roberto Giacobazzi
University of Verona, IT
- Christoph Gladisch
KIT – Karlsruher Institut für Technologie, DE
- Udo Gleich
Daimler Research, Ulm, DE
- Manuel Hermenegildo
IMDEA Software, Madrid, ES
- Ralf Huuck
NICTA, Sydney, AU
- Daniel Kroening
University of Oxford, GB
- K. Rustan M. Leino
Microsoft Res., Redmond, US
- Francesco Logozzo
Microsoft Res., Redmond, US
- Peter Müller
ETH Zürich, CH
- Filip Niksic
MPI-SWS, Kaiserslautern, DE
- Andreas Podelski
Universität Freiburg, DE
- Hendrik Post
Robert Bosch GmbH – Stuttgart, DE
- Francesco Ranzato
University of Padova, IT
- Xavier Rival
ENS, Paris, FR
- Helmut Seidl
TU München, DE
- Carsten Sinz
KIT – Karlsruher Institut für Technologie, DE
- Michael Tautschnig
Queen Mary University of London, GB
- Shmuel Tyszberowicz
Academic College of Tel Aviv Yaffo, IL
- Caterina Urban
ENS, Paris, FR
- Tomas Vojnar
Technical University of Brno, CZ
- Milena Vujosevic-Janicic
University of Belgrade, RS
- Reinhard Wilhelm
Universität des Saarlandes, DE

Report from Dagstuhl Seminar 14361

Computational Aspects of Fabrication

Edited by

Marc Alexa[1], Bernd Bickel[2], Sara McMains[3], and Holly E. Rushmeier[4]

1 **TU Berlin, DE,** marc.alexa@tu-berlin.de
2 **Disney Research – Zürich, CH,** bernd@disneyresearch.com
3 **University of California – Berkeley, US,** mcmains@me.berkeley.edu
4 **Yale University, US,** holly@acm.org

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 14361 "Computational Aspects of Fabrication".

Seminar August 31 to September 5, 2014 – http://www.dagstuhl.de/14361
1998 ACM Subject Classification I.3.8 [Computer Graphics] Applications, J.2 [Computer Applications] Physical Sciences and Engineering
Keywords and phrases Fabrication, Manufacturing Process Characterization (Additive Machining), 3D printing, Layered Manufacturing, 3D Rasterization, Geometric Modeling, Shape Perception, Material Properties (Mechanical/Optical)
Digital Object Identifier 10.4230/DagRep.4.8.126

1 Executive Summary

Marc Alexa
Bernd Bickel
Sara McMains
Holly E. Rushmeier

As manufacturing goes digital, the current understanding of industrial production will change fundamentally[1]. The digital age in manufacturing is coupled with new output devices that allow rapid customization and rapid manufacturing, revolutionizing the way we design, develop, distribute, fabricate, and consume products. We need to find computational models that support this new way of production thinking and lead its technological understanding. This opens challenges for many areas of science research, such as material science, chemistry, and engineering, but also and perhaps foremost computer sciences. The currently available digital content creation pipelines, algorithms, and tools cannot fully explore new manufacturing capabilities. To meet these demands, we need a deep understanding of computer graphics fundamentals: Shape, appearance of shape and materials, and physically-based simulation and animation. When designing an object, there is an inherent interplay among all these fundamental aspects.

[1] Special report: manufacturing and innovation. *The Economist* 403(8781):46, 2012.

Computational Aspects of Fabrication, *Dagstuhl Reports*, Vol. 4, Issue 8, pp. 126–150
Editors: Marc Alexa, Bernd Bickel, Sara McMains, and Holly E. Rushmeier
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

The purpose of this seminar is to bring together leading experts from academia and industry in the area of computer graphics, geometry processing, and digital fabrication. The goal is to address fundamental questions and issues related to computational aspects of fabrication and jump-start collaborations that will pioneer new approaches in this area.

2 Table of Contents

3 Overview of Talks

3.1 From Digital to Physical: My Biased View

Moritz Baecher (Disney Research – Zürich, CH)

Joint work of Bächer, Moritz; Whiting, Emily; Bickel, Bernd; Sorkine-Hornung, Olga; James, Doug L.; Pfister, Hanspeter; Otaduy, Miguel A.; Lee, Hyunho R.; Matusik, Wojciech; Gross, Markus;
Main reference M. Bächer, E. Whiting, B. Bickel, O. Sorkine-Hornung, "Spin-It: Optimizing Moment of Inertia for Spinnable Objects," ACM Trans. on Graphics, 33(4):96, 2014.
URL http://dx.doi.org/10.1145/2601097.2601157

Additive Manufacturing (AM) technologies have advanced enough to enable 3D printing at high resolution, in full-color, and with mixtures of soft and hard materials. As opposed to subtractive manufacturing (SM) such as milling or drilling, they can fabricate highly complex assemblies without the need for a manual assembly of individual components. Yet, one of the major issues holding back widespread use of AM is the lack of efficient algorithms for the automated fabrication of digital CG, the reproduction of physical, and the computational design of content. I will talk about one instance of each: (1) a method to reproduce elastic deformation properties of real world objects, (2) the automated fabrication of articulated characters from skinned meshes, and (3) the computational design of spinnable objects by optimizing their moment of inertia.

3.2 Blue Sky – Computational Tissue Fabrication

Bernd Bickel (Disney Research – Zürich, CH)

In my talk I will report on very recent developments in the area of tissue engineering and 3D printing of tissue. I will give an overview of grand challenges current research groups are focusing on, such as creating a functional kidney and the technical challenges along this way, both from a hardware/materials and software perspective. Finally, I will highlight several computational challenges in this area and give my biased view on how the Computer Graphics community could contribute towards a BioCAD system, an essential component for designing and fabricating functional organs.

3.3 Adobe Research: 3D Printing for the Masses

Nathan Carr (Adobe Inc. – San José, US)

Joint work of Tim Reiner, Nathan Carr, Radomír Měch, Ondrej Stava, Carsten Dachsbacher, Gavin Miller
Main reference T. Reiner, N. Carr, R. Měch, O. Stava, C. Dachsbacher, G. Miller, "Dual-Color Mixing for Fused Deposition Modeling Printers," Computer Graphics Forum, 33(2):479–486, 2014; pre-print available from author's webpage.
URL http://dx.doi.org/10.1111/cgf.12319
URL https://cg.ivd.kit.edu/publications/2014/DCM/DualColorMixing.pdf

3D Printing is starting to become available for the masses. In this talk I will cover Adobe's efforts in this space and detail the 3D printing capabilities inside of its flagship product Photoshop. The capabilities inside of Photoshop were built upon technology developed

inside of Adobe Research. I will cover a number of these technologies including Dual-Color Mixing for fused depositing Modeling Printers which provides a cool way to get continuous tone prints from two headed FDM print devices such as the Makerbot Replicator 2x. I will conclude with frontiers and challenges that I see this industry facing and what might be done to address some of these problems.

3.4 Design of Functional Models

Duygu Ceylan (EPFL – Lausanne, CH)

Joint work of Ceylan, Duygu; Li, Wilmot; Mitra, Niloy J.; Agrawala, Maneesh; Pauly, Mark
Main reference D. Ceylan, W. Li, N. J. Mitra, M. Agrawala, M. Pauly, "Designing and Fabricating Mechanical Automata from Mocap Sequences," ACM Trans. on Graphics, 32(6):186, 2013.
URL http://dx.doi.org/10.1145/2508363.2508400

Mechanical assemblies are collections of interconnected parts that move together to achieve specific functional goals. Such assemblies arise in various forms in our daily lives such as convertible furniture, kitchen supplies, mechanical toys etc. Enabling casual users to design such functional models requires to explore effective and intuitive ways of specifying the desired functionality. In this presentation, I will talk about our journey in experimenting with different interaction metaphors for specifying the desired functionality. This journey has led us to the development of an automatic system that can generate mechanical automata capable of realizing input motion sequences such as dancing or walking.

3.5 Research and Thoughts on 3D Printing

Yong Chen (University of Southern California, US)

The advent of 3D printing (additive manufacturing) and its use in rapid prototyping has drastically changed the design and manufacturing practice by enabling companies to prototype products faster and cheaper. With the price of 3D printers dramatically dropping in recent years, their accessibility is on the increase. However, significant challenges remain to be addressed in order for 3D printing to be used in direct digital manufacturing. My talk will introduce some of our research results on 3D printing processes including: (1) a complex internal structure design system for additive manufacturing; (2) a smooth surface fabrication process that can significantly improve the surface finish of curved surfaces; (3) a non-layer based 3D printing process named CNC accumulation for better part properties and building around inserts; (4) a deformation control strategy for 3D printing processes based on closed loop control and deformation simulation; (5) a support generation system for one of 3D printing processes (SLA); and (6) a digital material design method for multi- material 3D printing processes. I will also share some of my thoughts on some challenges and opportunities for computer graphics researchers.

3.6 Fabricating in constrained settings (with a human help)

Paolo Cignoni (CNR – Pisa, IT)

In the talk I will discuss three main topics: fabricating illustrative shapes with planar slices (Mesh Joinery), decomposing objects into interlocking pieces that can be manufactured by CNC and the use of anisotropic Voronoi diagrams to design a class of architectural structures called grid shells. Beside the details of these topics I will try to highlight "human in the end" issues that arise in digital fabrication techniques discussing how computational methods can improve the assembly problem.

3.7 From animated characters to legged robots

Stelian Coros (Carnegie Mellon University – Pittsburgh, US)

Computer graphics techniques allow artists to realize their imaginative visions, leading to immersive virtual worlds that capture the imagination of audiences world-wide. And now, thanks to advancements in rapid manufacturing devices, tangible links between these vivid virtual worlds and our own can be created. In order to unleash the full potential of this technology, however, a key challenge lies in determining the fundamental principles and design paradigms that allow digital content to be processed into forms that are suitable for fabrication. A particularly challenging task is that of creating physical representations of animated virtual characters in the form of complex robotic systems. In this talk, I will present evidence that control algorithms developed for physics-based character animation can also be applied to legged robots, allowing them to move with skill and purpose.

3.8 Design and Fabrication Using Wire Meshes

Bailin Deng (EPFL – Lausanne, CH)

Joint work of Akash, Garg; Sageman-Furnas, Andrew O.; Deng, Bailin; Yue, Yonghao; Grinspun, Eitan; Pauly, Mark; Wardetzky, Max
Main reference A. Garg, A. O. Sageman-Furnas, B. Deng, Y. Yue, E. Grinspun, M. Pauly, M. Wardetzky, "Wire mesh design," ACM Trans. on Graphics, 33(4):66, 2014.
URL http://dx.doi.org/10.1145/2601097.2601106

Wire meshes consist of interwoven metal wires arranged in a regular grid. Despite their widespread use in art, architecture, and engineering, it is challenging to design and fabricate freeform shapes using wire mesh material. One major difficulty is the global nature of wire mesh shapes: small local changes might have drastic global effects. In this talk, I will show how wire meshes can be modeled as discrete Chebyshev nets, which helps us to gain insights into their shape space and develop a computational design system for wire meshes. Moreover, I will present a method to physically realize freeform wire mesh shapes with the help of digital fabrication. Finally, I will discuss some open problems in this domain.

3.9 Perceptually-driven fabrication

Piotr Didyk (MIT – Cambridge, US)

In recent years there has been a tremendous development of new manufacturing technologies such as 3D printing. Despite the high quality of produced objects and the possibility of using multiple materials, reproducing 3D hardcopies of real objects is still a challenging task. Also the current level of understanding with regards to how these objects influence user experience is insufficient to fully utilize this kind of technology. In this talk, I will discuss importance of better understanding and modeling of human visual and haptic perception. I will present a few examples of how such knowledge combined with carefully designed computational techniques may lead to improved quality of manufactured objects.

3.10 Computational Fabrication Education at MIT

Piotr Didyk (MIT – Cambridge, US)

Joint work of Matusik, Wojciech; Levin, David; Didyk, Piotr

Hardly a 21st century day goes by without a reference to 3D printing and the revolution it is projected to cause in engineering and manufacturing. This enthusiasm is shared among hobbyists and within a vibrant maker community. During the last two years professor Wojciech Matusik and Dr. David Levin taught a novel course on computational fabrication at the Electrical Engineering and Computer Science Department at MIT. This talk provides an overview of the course material, programming assignments, and labs. Formally designated computational fabrication, 6.S079 provides a broad overview of both hardware and software for additive manufacturing. In particular, students are introduced to methods for parametric modeling of solid objects that take into account fabrication constraints. In the lab, they master computing models of physical objects using real-time 3D scanning. They also study and implement advanced physically- based simulation methods. The students explore the kinematics of mechanisms, such as four-bar-linkages, and finite element methods in the context of deformable solids. These techniques are fundamental in the engineering community and are crucial for designing highly predictive tools for 3D print preview. Using these tools, many variations of a virtual solid object can be interactively analyzed without committing to fabrication. The course also covers optimization methods that are applied to automate the design process. After learning this basic toolset, students analyze many instances of recent computational fabrication systems that seamlessly blend interactive design, simulation, and optimization, for example, the interactive designing of printable automata. In the second part of the semester, students work in groups on large open-ended projects. The primary goal of the class is to give students both a practical and a theoretical knowledge of every stage in the computational fabrication pipeline – raising awareness about this rapidly expanding field.

3.11 On the Challenges of Manufacturing

Gershon Elber (Technion – Haifa, IL)

In this two-parts talk I will first consider the major deficiencies I see in modern geometric modeling tools, and give some insights into what the next generation geometric modeling abilities should be, as I see them.

Then, I will exemplify the difficulties of manufacturing via a sequence of artifacts that went through the design-to-manufacturing process.

3.12 Fabricating Optics

Wolfgang Heidrich (KAUST – Thuwal, SA)

Modern fabrication methods show great promise for prototyping components for optical systems in computational imaging and display. In this talk I will report on some recent work in designing freeform lenses for goal-driven caustics, and then describe several experiments for fabricating these shapes on polyjet 3D printers as well as inexpensive 3-axis mills. I will report on issues with finishing the resulting shapes to optical grade. Finally I will descibe some initial approaches for fabricating custom diffractive optical elements in KAUST's Nanofabrication Lab.

3.13 Zometool Shape Approximation

Leif Kobbelt (RWTH Aachen, DE)

Joint work of Zimmer, Henrik; Florent, Lafarge; Alliez, Pierre; Kobbelt, Leif
Main reference H. Zimmer, F. Lafarge, P. Alliez, L. Kobbelt, "Zometool Shape Approximation," Graphical Models, 76(5):390–401, 2014.
URL http://dx.doi.org/10.1016/j.gmod.2014.03.009

We present an algorithm that approximates 2-manifold surfaces with Zometool models while preserving their topology. Zometool is a popular hands-on mathematical modeling system used in teaching, research and for recreational model assemblies at home. This construction system relies on a single node type with a small, fixed set of directions and only 9 different edge types in its basic form. While being naturally well suited for modeling symmetries, various polytopes or visualizing molecular structures, the inherent discreteness of the system poses difficult constraints on any algorithmic approach to support the modeling of freeform shapes. We contribute a set of local, topology preserving Zome mesh modification operators enabling the efficient exploration of the space of 2-manifold Zome models around a given input shape. Starting from a rough initial approximation, the operators are iteratively selected within a stochastic framework guided by an energy functional measuring the quality of the approximation. We demonstrate our approach on a number of designs and also describe parameters which are used to explore different complexities and enable coarse approximations.

3.14 Natural User Interfaces for Digital Fabrication

Manfred Lau (Lancaster University, GB)

Joint work of Lau, Manfred; Saul, Greg; Mitani, Jun; Igarashi, Takeo; Weichel, Christian; Kim, David; Villar, Nicolas; Gellersen, Hans

I discuss three natural user interface tools for digital fabrication. The motivation for such interfaces is that easy-to-use tools for 3D modeling and fabrication is still lacking, despite years of research in developing modeling tools for novice users. The recent trend of rapid prototyping technologies such as 3D printers will lead to an increased demand for these interfaces for fabrication purposes. The first system (SketchChair) is a sketch-based interface where the end-user can participate in the whole process of designing, modeling, and fabricating chairs with a laser cutter. The second system (Situated Modeling) has a tangible interface for modeling 3D shapes in an augmented reality environment. The user can immersively create and edit 3D shapes with a small number of physical primitive shapes, and with the guidance of the real-world environment and existing objects. The third system (MixFab) is a hand gesture based interface that takes advantage of a mixed reality environment to model small everyday objects that can be 3D printed. I end by speculating potential future interfaces for fabrication with the concept of Embodied Modeling and Fabrication.

3.15 Slicing for additive manufacturing: A computer graphics point of view

Sylvain Lefebvre (LORIA & INRIA – Nancy, FR)

Joint work of Lefebvre, Sylvain; Hergel, Jean; Dumas, Jérémie
URL http://www.aracknea-core.com/sylefeb/research/

In this talk I will introduce our work on slicing for additive manufacturing. We build upon recent GPU rendering techniques to directly slice objects specified using a Constructive Solid Geometry language. Our technique directly produces the code to drive the printer, without having to produce an intermediate mesh.

This led us to several software improvements to make printing on low cost printing more reliably, in particular for multi-material prints and for the generation of stable support structures.

References

1 Bridging the Gap: Automated Steady Scaffoldings for 3D Printing, J. Dumas, J. Hergel, S. Lefebvre, ACM SIGGRAPH, 2014
2 Clean color: Improving multi-filament 3D prints, J. Hergel, S. Lefebvre, EUROGRAPHICS 2014
3 IceSL: A GPU Accelerated modeler and slicer, S. Lefebvre, 18th European Forum on Additive Manufacturing, 2013

3.16 Rendering, Animating, and Fabricating Volumetric Materials

Steve Marschner (Cornell University – Ithaca, US)

Joint work of Bala, Kavita; Bickel, Bernd; Jakob, Wenzel; James, Doug; Jarosz, Wojciech; Kaldor, Jon; Matusik, Wojciech; Papas, Marios; Yuksel, Cem; Zhao, Shuang

For rendering and animation of textiles, detailed models of the material structure are useful in producing highly realistic results.

For coarse knit fabrics, we have simulated both the structure and deformation of intricate lace patterns and complex garments in terms of the geometry and motion of individual yarns, producing texture and motion that closely resemble the real materials. For woven fabrics, we have used micro CT scans to model the geometric arrangement of fibers at the microscopic scale, leading to highly realistic images that exhibit the distinctive texture and sheen of these materials. In both cases, the needs of realism have driven us to work directly in terms of the descriptions used to fabricate the materials: patterns for hand knitting and the binary images used by industrial Jacquard looms.

This is a general trend in realistic rendering and animation, and translucent materials are another case where the descriptions used for rendering are quite direct specifications for the material itself. We have developed a system that calibrates a particular pigmentation system, then can measure a translucent material and uses rendering techniques inversely to compute a recipe for making a material that matches it.

3.17 Geometric Analysis for Manufacturing: Conventional vs. Additive Manufacturing

Sara McMains (University of California – Berkeley, US)

Joint work of McMains, Sara; Yasui, Yukuse; Li, Wei

I present computational geometry algorithms that we developed to support manufacturing process planning for waterjet cleaning (used to remove manufacturing byproducts in conventional manufacturing processes). Parallels with process planning for additive manufacturing are discussed.

3.18 Creating Works-Like Prototypes Of Mechanical Objects

Niloy Mitra (University College London, GB)

Joint work of Koo, Bongjin; Li, Wilmot; Yao, JiaXia; Agrawala, Maneesh; Mitra, Niloy
Main reference B. Koo, W. Li, J. Yao, M. Agrawala, N. J. Mitra, "Creating Works-Like Prototypes of Mechanical Objects," ACM Trans. on Graphics, 33(6):217, 2014; pre-print available from author's webpage.
URL http://dx.doi.org/10.1145/2661229.2661289
URL http://geometry.cs.ucl.ac.uk/projects/2014/works-like/

Designers often create physical works-like prototypes early in the product development cycle to explore possible mechanical architectures for a design. Yet, creating functional prototypes

requires time and expertise, which discourages rapid design iterations. Designers must carefully specify part and joint parameters to ensure that the parts move and fit and together in the intended manner. We present an interactive system that streamlines the process by allowing users to annotate rough 3D models with high-level functional relationships (e. g., part A fits inside part B). Based on these relationships, our system optimizes the model geometry to produce a working design. We demonstrate the versatility of our system by using it to design a variety of works-like prototypes.

References

1 Bongjin Koo, Wilmot Li, JiaXian Yao, Maneesh Agrawala, Niloy J. Mitra. *Creating Works-Like Prototypes of Mechanical Objects*. ACM Transactions on Graphics (Special issue of SIGGRAPH Asia 2014), 2014.

3.19 Interacting with Personal Fabrication Devices – Current challenges from an HCI perspective.

Stefanie Müller (Hasso-Plattner-Institut – Potsdam, DE)

Joint work of Müller, Stefanie; Baudisch, Patrick; Guimbretière, François
Main reference S. Müller, S. Im, S. Gurevich, A. Teibrich, L. Pfisterer, F. Guimbretiere, P. Baudisch, "WirePrint: Fast 3D Printed Previews," in Proc. of the 27th ACM User Interface Software and Technology Symposium (UIST'14), pp. 273–280, ACM, 2014; pre-print available from author's webpage.
URL http://dx.doi.org/10.1145/2642918.2647359
URL https://www.hpi.uni-potsdam.de/baudisch/home.html

In this talk, I discuss three challenges when interacting with personal fabrication devices: (1) Personal fabrication machines, such as 3D printers, are so slow that many objects require printing overnight. This limits designers to a single design iteration per day even though the actual design work between each iteration was only a couple of minutes. With our projects, faBrickation and WirePrint, we address this problem by allowing designers to fabricate intermediate versions of a prototype as fast, low-fidelity previews, and to only create the final version as a full 3D print. (2) Currently users use a digital editor to design physical objects. There are good reasons for doing this since the digital world allows for precise interaction and editing steps can be easily undone. However, having the input and output space separated is not intuitive for novice users. With our projects constructable and LaserOrigami, we show how to merge input and output space by letting users work directly on the physical workpiece and by creating physical output after every step. (3) Personal fabrication tools allow us to create more and more things. With our project Scotty, we question that more is always better as having more affects how we value a single object. With Scotty, we show how to relocate physical objects across distances by ensuring that there is never more than one copy at a time, i. e. the object disappears on the sender side and reappears on the receiver side, thereby preserving its value.

3.20 Computational Caustics

Mark Pauly (EPFL – Lausanne, CH)

Joint work of Pauly, Mark; Yuliy Schwartzburg, Romain Testuz, Andrea Tagliasacchi
Main reference Y. Schwartzburg, R. Testuz, A. Tagliasacchi, M. Pauly, "High-contrast Computational Caustic Design," ACM Trans. on Graphics, 33(4):74, 2014.
URL http://dx.doi.org/10.1145/2601097.2601200
URL http://lgg.epfl.ch/publications/2014/HighContrastCaustics.pdf

We present a new algorithm for computational caustic design. Our algorithm solves for the shape of a transparent object such that the re- fracted light paints a desired caustic image on a receiver screen. We introduce an optimal transport formulation to establish a correspon- dence between the input geometry and the unknown target shape. A subsequent 3D optimization based on an adaptive discretization scheme then finds the target surface from the correspondence map. Our approach supports piecewise smooth surfaces and non-bijective mappings, which eliminates a number of shortcomings of previous methods. This leads to a significantly richer space of caustic images, including smooth transitions, singularities of infinite light density, and completely black areas. We demonstrate the effectiveness of our approach with several simulated and fabricated examples.

3.21 Interactive modeling with developable NURBS surfaces

Helmut Pottmann (KAUST – Thuwal, SA)

Joint work of Tang, Chengcheng; Pottmann, Helmut

Developable surfaces play an important role in various manufacturing technologies, since they model the natural behavior of materials which do not stretch. Although there is a rich literature on modeling these surfaces, the interactive design of developable NURBS surfaces is still a challenge. Employing a combination of the standard and the dual representation, we propose an efficient numerical constraint solver which overcomes the limitations of previous work. While the user manipulates the B-spline control structure, the surfaces get automatically corrected in real time towards developable NURBS surfaces with high numerical accuracy. We illustrate our framework by various types of developable strip models and present initial results on models with curved folds. This is ongoing unpublished research.

3.22 Computational / Physical Light Routing: 3D Printed Fiber Optics

Szymon Rusinkiewicz (Princeton University, US)

Joint work of Pereira, Thiago; Rusinkiewicz, Szymon; Matusik, Wojciech
Main reference T. Pereira, S. Rusinkiewicz, W. Matusik, "Computational Light Routing: 3D Printed Fiber Optics for Sensing and Display," ACM Trans. on Graphics, 33(3):24, 2014.
URL http://dx.doi.org/10.1145/2602140

Despite recent interest in digital fabrication, there are still few algorithms that provide control over how light propagates inside a solid object. Existing methods either work only

on the surface or restrict themselves to light diffusion in volumes. We use multi-material 3D printing to fabricate objects with embedded optical fibers, exploiting total internal reflection to guide light inside an object. We introduce automatic fiber design algorithms together with new manufacturing techniques to route light between two arbitrary surfaces. Our implicit algorithm optimizes light transmission by minimizing fiber curvature and maximizing fiber separation while respecting constraints such as fiber arrival angle. We also discuss the influence of different printable materials and fiber geometry on light propagation in the volume and the light angular distribution when exiting the fiber. Our methods enable new applications such as surface displays of arbitrary shape, touch-based painting of surfaces and sensing a hemispherical light distribution in a single shot.

3.23 3D Printing Tools in Meshmixer (Support Structures, Strength Analysis, etc)

Ryan Schmidt (AUTODESK Research – Toronto, CA)

3D printing is becoming increasingly practical as a technique for one-off fabrication of customized objects and small-batch manufacturing. In many applications, the 3D printer provides a means to solve problems that would be prohibitively expensive to address in any other way. I will present three example use cases in aerospace, manufacturing, and prosthetics that illustrate how both professional and consumer-level 3D printers are being used today. However, the premise of automatically fabricating an arbitrary 3D design remains more hype than reality. Based on extensive observation of users of Autodesk meshmixer, discussion with domain experts, and personal experience, I have noted many issues which could be addressed with the mathematical tools of computer graphics and geometry processing. I will describe a set of current challenges in the 3D printing pipeline that hinder user creativity and prevent many seemingly-straightforward applications.

3.24 User-Guided Inverse 3D Modeling

Carlo H. Sequin (University of California – Berkeley, US)

Joint work of Andrews, James; Sequin, Carlo H.
Main reference J. Andrews, "User-Guided Inverse 3D Modeling," Ph. D. Thesis, U. C. Berkeley, May 2013.
URL http://www.eecs.berkeley.edu/Pubs/TechRpts/2013/EECS-2013-103.pdf

Few designs start from scratch in a vacuum. Often there is a previous artifact that provides inspiration or may even be close enough so that some high-level redesign might be the most effective approach. Unfortunately there may be no CAD files available or they may be at such a low level (100'000 triangles) that it is not a good starting point for a major redesign. "User-Guided Inverse 3D Modeling" is an approach to re-create a well-structured, high-level, parameterized, procedural description of some geometry very close to the inspirational artifact. Its hierarchical structure and the degree of its parameterization are imposed with some high-level instructions by the designer, so that the resulting description is most appropriate to make the intended design changes.

3.25 Use the winding number to determine "inside" and "outside" for each level in layered manufacturing

Carlo H. Sequin (University of California – Berkeley, US)

Models of trusses consisting of many joining or intersecting beams, often with non-trivial cross sections, may result in a solid model that does not have a proper, oriented, 2-manifold boundary representation. Cleaning up such a 3D model is laborious and difficult, even when the individual components (beams) themselves have perfectly good B-Reps. Sending such "unclean" models to a layered manufacturing machine will often produce unexpected result – or will simply be refused by the machine's software. This problem could easily be remedied by slicing all components individually, paying careful attention to the orientation of the surface normals, and then forming the Boolean union of all the extracted slicing contours on each level while summing up their respective winding numbers.

3.26 Modular Models of Mathematical Knots

Carlo H. Sequin (University of California – Berkeley, US)

The LEGO-Knot system is a collection of tubular parts with a 1" square cross-section that are compatible with the LEGO DUPLO blocks. It was originally inspired by the modular "Borsalino" sculpture by Henk van Putten presented in the art exhibition of the Bridges 2013 conference in Enschede, Netherlands. Currently this set of prototype parts is rich enough so that a wide variety of free-form tubular sculptures can be assembled, and simple mathematical knots like the Trefoil Knot and the Figure-8 Knot can be modeled in a graceful manner with the sweep curve closing smoothly onto itself. Recently I have given myself the challenge to design a single tubular element that would be versatile enough for modeling many of the simple mathematical knots in a similar graceful manner. The resulting tubular module is based on a cross-section in the form of a regular 16-gon and bends through an angle of 30 degrees. Now the remaining challenge is to develop an effective search algorithm that can find elegant solutions for most of the simple knots. Through manual search stretching over a few hours, good solutions have been found for the Trefoil Knot, the Figuee-8 Knot, and the Knots 5-1 and 7-4, as well as for the Borromean Link. The challenge remains to find a good automated search algorithm. This problem was presented at this seminar, because it has similarities with the design problem discussed by Leif Kobelt: How to best approximate an arbitrary 2-manifold such as the surface of the Stanford Bunny with a mesh made solely from Zome Tool parts.

3.27 Computational Design-to-Fabrication

Kristina Shea (ETH Zürich, CH)

In this talk I will introduce Computational Design-to-Fabrication from an engineering and product design viewpoint, give an overview of our research on design automation and optimization of structural and mechanical systems and a glimpse into new research making the link to automated fabrication via additive manufacturing, mainly multi-material.

3.28 Flat Fabrication

Karan Singh (University of Toronto, CA)

Joint work of Singh, Karan; McCrae, James
Main reference J. McCrae, N. Umetani, K. Singh, "FlatFitFab: Interactive Modeling with Planar Sections," in Proc. of the 27th Annual ACM Symp. on User Interface Software and Technology (UIST'14), pp. 13–22, ACM, 2014.
URL http://dx.doi.org/10.1145/2642918.2647388
URL http://flatfab.com/

Assembled planar section structures are common in art and engineering. This talk presents the state of the art on the computational abstraction of 3D shape using planar sections. It also describes a comprehensive drawing interface to author planar section structures from scratch, based on principles of inter-plane orthogonality, procedural regularity and fronto-parallel drawing. Finally, a number of open problems and issues specific to flat fabrication are described, along with speculated solutions and directions for future work.

3.29 Computational Design and Fabrication of Deformable Objects

Melina Skouras (Disney Research – Zürich, CH)

Joint work of Skouras, Melina; Thomaszewski, Bernhard; Bickel, Bernd; Coros, Stelian; Kaufmann, Peter; Bradley, Derek; Beeler, Thabo; Jackson, Phil; Marschner, Steve; Matusik, Wojciech; Gross, Markus

Deformable objects have a plethora of applications: they can be used for entertainment, advertisement, engineering or even medical purposes. However designing custom deformable objects remains a difficult task. The designer must foresee and invert effects of external forces on the behavior of the figure in order to take the proper design decisions. In this talk, I will present novel approaches based on physics-based simulation and inverse optimization techniques which alleviate these difficulties and propose a complete framework to design custom deformable objects by automating some of the most tedious aspects of the design process. This framework is tailored to the design of various objects such as rubber balloons, skin for animatronics figures and custom actuated characters, for which optimization of diverse variables including rest shape, materials and actuation system is alternately considered. Validation of our method is performed by fabricating representative sets of physical prototypes designed with our method and compared to the results predicted by simulation.

References

1 M. Skouras, B. Thomaszewski, B. Bickel, and M. Gross. Computational Design of Rubber Balloons, *Proceedings of Eurographics (Cagliari, Italy, May 13–18, 2012), Computer Graphics Forum*, vol. 31, no. 2, pp. 835–844, 2012.

2 B. Bickel, P. Kaufmann, M. Skouras, B. Thomaszewski, D. Bradley, T. Beeler, P. Jackson, S. Marschner, W. Matusik and M. Gross. Physical Face Cloning, *Proceedings of ACM SIGGRAPH (Los Angeles, USA, August 5-9, 2012), ACM Transactions on Graphics*, vol. 31, no. 4, pp. 118:1–118:10, 2012.

3 M. Skouras, B. Thomaszewski, S. Coros, B. Bickel and M. Gross. Computational Design of Actuated Deformable Characters, *Proceedings of ACM SIGGRAPH (Anaheim, USA, July 21-25, 2013), ACM Transactions on Graphics*, vol. 32, no. 4, pp. 82:1–82:10, 2013.

3.30 Designing Mechanical Characters

Bernhard Thomaszewski (Disney Research – Zürich, CH)

The design of virtual characters has been in the focus of graphics research since its very beginnings. With the advent of 3D printers we now have the machinery to create tangible, physical counterparts of digital characters that can be used as input devices for video games, mechanical toys, or even animatronics. But besides progress in manufacturing technology, we need progress in software tools to facilitate the translation from virtual to real characters. This talk addresses a number of challenges that arise when designing mechanical characters, in particular the question of how to design mechanisms that are able to reproduce a desired motion.

3.31 Interactive Design of Functional Shapes

Nobuyuki Umetani (Disney Research – Zürich, CH)

URL http://www-ui.is.s.u-tokyo.ac.jp/~ume

Physical simulation allows validation of geometric designs without tedious physical prototyping. However, since geometric modeling and physical simulation are typically separated, simulations are mainly used for rejecting bad design, and, unfortunately, not for assisting creative exploration towards better designs. In this talk, I introduce several interactive approaches to integrate physical simulation into geometric modeling to actively support creative design process. More specifically, I demonstrate the importance of (i) presenting the simulation results in real-time during user's interactive shape editing so that the user immediately sees the validity of current design, and to (ii) providing a guide to the user so that he or she can efficiently explore the valid deign space. I present novel algorithms to achieve these requirements.

3.32 Crafting Light by Hacking Pixels

Gordon Wetzstein (Stanford University, US)

Compressive image acquisition and display is an emerging architecture for consumer electronics that explores the co-design of optics, electronics, applied mathematics, and real-time computing. Together, these hardware/software systems exploit compressibility of the recorded or presented data to facilitate new device form factors and relax requirements on electronics and optics. For instance, light field or glasses-free 3D displays usually show different perspectives of the same 3D scene to a range of different viewpoints. All these images are very similar and therefore highly compressible. By combining multilayer hardware architectures and directional backlighting with real-time implementations of light field tensor factorization, limitations of existing displays, for instance in resolution, contrast, depth of field, and field of view, can be overcome. A similar design paradigm also applies to light field and multi-spectral image acquisition, super-resolution and high dynamic range display, glasses-free 3D projection, computational lithography, microscopy, and many other applications. In this talk, we review the fundamentals of compressive camera and display systems and discuss their impact on future consumer electronics, remote sensing, scientific imaging, and human-computer interaction.

3.33 Crafting Light by Hacking Pixels

Gordon Wetzstein (Stanford University, US)

Joint work of Wetzstein, Gordon; Raskar, Ramesh; Lanman, Douglas; Hirsch, Matthew
Main reference M. Hirsch, D. Lanman, G. Wetzstein, R. Raskar, "Tensor Displays," in Proc. of the ACM SIGGRAPH 2012 Int'l Conf. on Computer Graphics and Interactive Techniques (SIGGRAPH'12), 24 pages, ACM, 2012; pre-print available from author's webpage.
URL http://dx.doi.org/10.1145/2343456.2343480
URL http://web.media.mit.edu/~gordonw/TensorDisplays/

With the invention of integral imaging and parallax barriers in the beginning of the 20th century, glasses-free 3D displays have become feasible. Only today – more than a century later – glasses-free 3D displays are finally emerging in the consumer market. The technologies being employed in current-generation devices, however, are fundamentally the same as what was invented 100 years ago. With rapid advances in optical and digital fabrication, digital processing power, and computational models for human perception, a new generation of display technology is emerging: computational displays exploring the co-design of optical elements and computational processing while taking particular characteristics of the human visual system into account. This technology does not only encompass 3D displays, but also next-generation projection systems, high dynamic range displays, perceptually-driven devices, and computational probes.

This talk serves as an introduction to the emerging field of computational display fabrication. We will discuss a wide variety of different applications and hardware setups of computational displays as well as their fabrication, including high dynamic range displays, advanced projection systems as well as glasses-free 3D display. We will only briefly review conventional technology and focus on practical and intuitive demonstrations of how an

interdisciplinary approach to display design and fabrication encompassing optics, perception, computation, and mathematical analysis can overcome the limitations for a variety of applications.

3.34 Appearance Fabrication

Tim Weyrich (University College London, GB)

Joint work of Weyrich, Tim; Hullin, Matthias B.; Ihrke, Ivo; Fuchs, Martin; Papas, Marios; Jarosz, Wojciech; Jakob, Wenzel; Rusinkiewicz, Szymon; Matusik, Wojciech

Appearance fabrication aims at creating custom reflectance properties on real- world surfaces. As the inverse of appearance acquisition, it starts from a digital description of spatio-angular reflectance properties and seeks to alter physical surface to match that description. My talk provides an overview over working principles employed by prior art, and raises a number of questions on future directions of appearance fabrication.

3.35 Structurally-Informed Geometry

Emily Whiting (Dartmouth College – Hanover, US)

Joint work of Whiting, Emily; Prevost, Romain; Lefebvre, Sylvain; Sorkine-Hornung, Olga; Baecher, Mortiz; Bickel, Bernd; Deuss, Mario; Panozzo, Daniele; Liu, Yang; Block, Philippe; Pauly, Mark

In computer graphics, many of the 3D objects we are interested in modeling are physically-inspired or are designed with the intent of being built or manufactured. Yet many of the tools developed for geometric modeling are unaware of structural considerations, largely based on geometric surface characteristics alone. The motivation for my work is to use structural soundness and stability properties to enhance the traditional modeling pipeline. I will discuss recent work investigating structurally-informed design of 3D printed objects. Time permitting, I may also review topics in the design of masonry structures.

3.36 Small-scale Structure and Material Properties

Denis Zorin (New York University, US)

Joint work of Zhou, James; Panetta, Julian; Zorin, Denis

Complex structures are widely used to achieve optimal mechanical performance (e. g., minimize weight for given strength). In traditional manufacturing, the cost is proportional to complexity; additive manufacturing makes complexity almost independent from cost, making much more complex with novel properties possible. In particular, one can create small-scale structures, approximating homogeneous material properities. Optimizing such structures directly is difficult. I describe a multistage process process for optimizing small-scale structures,

usually referred to as "optimization by homogenization": a map from material properties to structure parameters is constructed, then continuous material properties are optimized and mapped to a tiled structure using this map. Experiments show that a number of structure optimization problems can be solved in this way.

4 Open Problems and Panel Discussion

Based on the ideas from the individual presentations, there were group discussions on four topics. First there was a discussion of particular problems that have been posed in fabrication, and the applications that motivate these problems. This list of problem definitions was a major outcome of the meeting. The group discussed particular issues that need to be resolved for these specific problems. Second, on a more general level, the role of computer graphics research in fabrication was discussed, in light of the activities in other disciplines. Third, the fabrication processes that the computer graphics community should consider working on was discussed. Finally, the need for systems to assist in the design of custom fabricated objects emerged as a major theme in the meeting. The group tried to define the different types of user populations, and their needs for design systems.

4.1 Problems/Applications

Many overlapping ideas were identified from the various presentations. These were combined to identify specific computations problems and the applications that motivate them.

Problem: A general problem is the production of a particular shape from fixed set of physical primitives. Problems related to this include how to produce a set of primitives that can express a wide variety of shape and how to produce instructions for assembling shape from primitives.
Applications: Toys (e. g. Lego-type systems), sets of physical objects for illustrating concepts.
Discussion: This is actually a large group of problems, and designing with orthogonal planar pieces is also in this set, as well as design with specific geometric constraints. As well as the toy and illustrative applications, systems based on fixed primitives are useful in architectural design exploration.

Problem: Design structures to achieve spatially varying physical properties (Young's modulus, Poisson ratio). One approach to this is to emulate complex designs previously found only in nature (e. g. bird bones). Solutions need to include not just one structure, but the description of the rest geometry and functionality, and these must stay within limits of manufacture. Achieving particular properties can be achieved with structure and by using mixes of material.
Applications: Helmets, custom footwear.
Discussion: To study material variation, we need collaborators in materials science and mechanical structures. Computer graphics can contribute models of humans that these objects need to fit. Computer graphics can also contribute simplified simulations of very complex structures. Manipulating and representing very complex geometries is an area of computer graphics. We need to create tools to design combinations for hard and soft materials – e. g. cases for phones etc.

Problem: Design assemblies with specific motions. The motions may be controlled in different ways – for example with a single input crank or by being propelled by something. There are different ways that motions can be described. Existing libraries of mechanisms can be used in the solutions. Beyond just rigid linkages, the motion of deformable surfaces can also be designed.
Applications: Toys, animatronics, design of artificial limbs.
Discussion: We need to model deformations in many settings to understand the manufacturing process. We now have the opportunity to put complex mechanisms in ANY object – how can we exploit this to make new useful things?

Problem: Fabrication of new optical components. This can provide new dimensions in viewing (stereo/hi dynamic range, light field). Displays can be made in arbitrary shapes. Are deposition methods good enough, or are they just a proof of concept technology?
Applications: New displays, cameras.
Discussion: Optical systems need post-processing, i. e. polishing and assembly. These operations need to be included in the design of the fabrication process. The human is part of the fabrication process. We need to model the role the human can play to a greater extent. We need to design work flows between stages in the fabrication process. In this an other applications, we need a place to document methods and best practices.

Problem: UI's for: Design of objects by function and manufacture technique (furniture, orthogonal planes). Design of objects in context of real objects (augmented reality). Design of objects with performance analysis built into design system (musical instruments).
Applications: Universal.
Discussion: We need UI's for designing UI's. We need more exploration of data driven techniques in place of simulations in the interface. We need a common language to communicate designs to printers, to communicate with modules for doing clean-up of models. Is the answer STL? Or AMF (Hod Lipson's initiative)? In the UI, when we are designing to fit in the physical world, do we think of bringing the digital into the physical world, or the physical into the digital world?

Problem: Drawbacks of FDM – Design alternative additive manufacture systems. Redesign system for obtaining smooth surfaces. Many manufacturing processes are not supported computationally.
Applications: Optical systems where surface quality is critical.
Discussion: In the near future more patents will expire, and additional types of fabrication will become cheap. However, those techniques may have their own drawbacks as far as materials required (cost, stability). In general we need to document the limitations of machines. This is a moving target with the technology changing. How do we abstract to the correct level? What can we learn from the more mature field of 2D printing on how they have characterized machines?

Problem: Appearance Reproduction. Is FDM too limited? Simulation for other manufacturing processes? What will be the role of perception?
Applications: Prototype appearance, assist traditional manufacture.
Discussion: While appearance reproduction is limited, as long as machines are being used for prototyping and evaluation it is important to do the best job possible reproducing appearance and haptic/tactile properties.

4.2 Computational Fabrication and Computer Graphics

Computer graphics is still finding its place in computational fabrication research. To some extent we are solving new problems, but to some extent we are rediscovering problems other people have solved.

A good way to speed up productive work is through community software initiatives. An open source library that can support low-end fabrication devices would quickly benefit from the efforts of many people who are studying fabrication.

Currently G-code is the current common tool for software. Can we do better? Many people write their own slicers, but access to machine control is making this harder. Should we start a coordinated effort to develop open source software?

CNC machines became easier to control and more open, will the same happen for 3D printing? Standardizing on a poor scheme like G-code could be dangerous and limiting. Other examples of systems that opened up are GPUs, and Epson printers. What motivates companies to make software that controls more open? We need to find a way to demonstrate to them the financial benefit of openness.

We need to provide a place (perhaps a wiki or similar site) to capture documentation of methods and best practices. Useful software could evolve from these methods and practices.

4.3 Targeted Fabrication Processes

A lot of current fabrication work in computer graphics centers around FDM – fused deposition modeling. Other computer controlled machining processes, as well as hybrid FDM and computer controlled making processes could benefit from advanced computational methods.

In the 1980's, computer graphics was closely tied with the CADAM (computer aided design and manufacturing) communities. The current interest in computational fabrication is an opportunity to renew ties with that community. The computer graphics community should avoid focussing on a single manufacturing technology.

4.4 Design Systems

A major issue in computation for fabrication is the development of UI's (user interfaces). UI's need to be designed for different types of users, and for different phases of design.

In fabrication, there are at least three different groups of users – novices, "do it yourself" users with a high level of expertise, and professional designers.

For novices, a very restricted design space is useful. It may be that novices are really just customizing, not designing, objects.

No one stays a novice. Is it the generality of the design vs a very constrained space? We need "scalable" UI's that move from general concepts to specifics.

Design systems can simulate the fabrication process and the performance of the final object. In the user interface we need to take into account what the users wants – just an object sent to them, or an object in which they have participated in the production.

Professional designers move from general concepts to detail specification. Design exploration needs rapid generation of variations. Specification needs precision. We need bridges between systems to do this, to flow from concept to precision. Consider that another group that uses systems are decision makers. We need a seamless tool chain.

4.5 Open Questions and Future Research Directions

During the final session of the workshop we again discussed important open questions and potential future research directions. Ten areas were discussed, as listed below, building on the discussion during Ryan Schmidt's earlier presentation.

Open questions that in many cases participants have done research on in the past, but that didn't lead to much additional discussion during this session, included:

- Better surface quality
- Better support structures
- Light-weighting and hollowing
- Multi-objective orientation optimization

Open questions that inspired more discussion:

- Generation of G-code paths. Optimal paths for 3-D printing differ from optimal paths for machining, where toolpath generation has been widely studied. The patterns are very different for additive versus subtractive G-code. There are also aspects of fluids that are relevant; approximate models could be very useful. We need a good simulator for 3-D printing (perhaps along the lines of the Vericut software that is widely used for CNC). Right now some designers will have the software generate a G-code path from a preliminary design, examine the path to gain insight about whether the design will print well, and redesign based on what they can see visually in the toolpath, but an actual simulator would give even more useful feedback.
- Ensuring printability. One example where there has been some research but more can be done is in the area of selective thickening of geometry. One example where this is very useful is for architectural models, where when an architect scales down a model of a house uniformly in order to print it, the banisters all break because they become too thin for the target printer.
- Tolerance and clearance analysis. Here one of the important differences compared to conventional manufacturing is the need for anisotropy-aware algorithms.
- Nesting and packing. There is lots of literature in two dimensions, but less in 3-D. The commercial system Magics supports 3-D packing, but it is slow.
- Residual stress estimation and correction. There has been some work on estimation, but no one at the workshop was aware of much work on correction. Active positioning with feedback is not enough, because the warping often happens later, as subsequent layers cool/solidify. With cheap printers, people have seen awful warping. Perhaps printing a calibration part for testing, evaluating it with a computer vision system, and setting parameters based on current working conditions surmised from the behavior of the calibration test part is an approach that is worth studying.
- Planning support material removal. This could include planning escape holes, or a series of orientations for emptying the part interior. For example, injection molds need to be designed with cooling channels. Now that people are 3-D printing molds, the geometry of the channels is suddenly getting much more complex because 3-D printing doesn't have tooling accessibility issues. But now the un-sintered powder or dissolved support material needs to be removed. Accessibility for cleaning the part to remove this excess material after printing becomes the issue instead.

Participants

- Marc Alexa
TU Berlin, DE
- Moritz Bächer
Disney Research – Zürich, CH
- Bernd Bickel
Disney Research – Zürich, CH
- Nathan Carr
Adobe Inc. – San José, US
- Duygu Ceylan
EPFL – Lausanne, CH
- Yong Chen
Univ. of Southern California, US
- Paolo Cignoni
CNR – Pisa, IT
- Stelian Coros
Carnegie Mellon University, US
- Bailin Deng
EPFL – Lausanne, CH
- Piotr Didyk
MIT – Cambridge, US
- Gershon Elber
Technion – Haifa, IL
- Wolfgang Heidrich
KAUST – Thuwal, SA
- Kristian Hildebrand
TU Berlin, DE
- Leif Kobbelt
RWTH Aachen, DE
- Manfred Lau
Lancaster University, GB
- Sylvain Lefebvre
LORIA & INRIA – Nancy, FR
- Steve Marschner
Cornell University, US
- Sara McMains
University of California – Berkeley, US
- Eder Miguel
University Rey Juan Carlos – Madrid, ES
- Niloy Mitra
University College London, GB
- Stefanie Müller
Hasso-Plattner-Institut – Potsdam, DE
- Mark Pauly
EPFL – Lausanne, CH
- Nico Pietroni
ISTI-CNR – Pisa, IT
- Helmut Pottmann
KAUST – Thuwal, SA
- Holly E. Rushmeier
Yale University, US
- Szymon Rusinkiewicz
Princeton University, US
- Ryan Schmidt
AUTODESK Res. – Toronto, CA
- Carlo H. Sequin
University of California – Berkeley, US
- Kristina Shea
ETH Zürich, CH
- Karan Singh
University of Toronto, CA
- Melina Skouras
Disney Research – Zürich, CH
- Bernhard Thomaszewski
Disney Research – Zürich, CH
- Nobuyuki Umetani
Disney Research – Zürich, CH
- Gordon Wetzstein
Stanford University, US
- Tim Weyrich
University College London, GB
- Emily Whiting
Dartmouth College – Hanover, US
- Denis Zorin
New York University, US

Report from Dagstuhl Seminar 14362

Augmenting Human Memory – Capture and Recall in the Era of Lifelogging

Edited by

Mark Billinghurst[1], Nigel Davies[2], Marc Langheinrich[3], and Albrecht Schmidt[4]

1 **University of Canterbury – Christchurch, NZ,** mark.billinghurst@hitlabnz.org
2 **Lancaster University, GB,** nigel@comp.lancs.ac.uk
3 **Università della Svizzera italiana (USI), CH,** marc.langheinrich@usi.ch
4 **Universität Stuttgart, DE,** albrecht.schmidt@vis.uni-stuttgart.de

Abstract

Recent developments in capture technology and information retrieval allow for continuous and automated recordings of many aspects of our everyday lives. By combining this with basic research in memory psychology, today's memory augmentation technologies may soon be elevated from a clinical niche application to a mainstream technology, initiating a major change in the way we use technology to remember and to externalize memory. Future capture technologies and corresponding control mechanisms will allow us to automate the acquisition of personal memories and subsequently trigger feedback of such memories through ambient large displays and personal mobile devices in order to aid personal memory acquisition, retention, and attenuation. The emergence of this new breed of memory psychology-inspired capture and recall technology will represent a radical transformation in the way we understand and manage human memory acquisition and recall. This report documents the program and the outcomes of Dagstuhl Seminar 14362 "Augmenting Human Memory – Capture and Recall in the Era of Lifelogging", which brought together 28 researchers from multiple disciplines both within computer science – mobile computing, privacy and security, social computing and ethnography, usability, and systems research – as well as from related disciplines such as psychology, sociology, and economics, in order to discuss how these trends are changing our existing research on capture technologies, privacy and society, and existing theories of memory.

Seminar August 31 to September 5, 2014 – http://www.dagstuhl.de/14362
1998 ACM Subject Classification H.1.2 [Models and Principles] User/Machine Systems, H.3.0 [Information Storage and Retrieval] General, H.5.0 [Information Interfaces and Presentation] General J.4 Social and Behavioral Sciences, K.4.1 [Computers and Society] Public Policy Issues
Keywords and phrases human memory interaction, lifelogging, memory augmentation
Digital Object Identifier 10.4230/DagRep.4.8.151

Augmenting Human Memory – Capture and Recall in the Era of Lifelogging, *Dagstuhl Reports*, Vol. 4, Issue 8, pp. 151–173
Editors: Mark Billinghurst, Nigel Davies, Marc Langheinrich, and Albrecht Schmidt

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

1 Executive Summary

Mark Billinghurst
Nigel Davies
Marc Langheinrich
Albrecht Schmidt

Technology has always had a direct impact on how and what humans remember. This impact is both inevitable and fundamental – technology radically changes the nature and scale of the cues that we can preserve outside our own memory in order to trigger recall. Such change is not new – we have seen the transition from story-telling to written books, from paintings to photographs to digital images and from individual diaries to collective social networks. However, in recent years three separate strands of technology have developed to the extent that collectively they open up entirely new ways of augmenting human memory:

1. near-continuous collection of memory cues has become possible through the use of technologies such as Microsoft's SenseCam, social networks and interaction logs.
2. advances in data storage and processing now enables widespread mining of stored cues for proactive presentation, both in terms of cues collected by an individual and in terms of complex networks of related cues contributed by others.
3. the presence of ubiquitous displays (both in the environment and via personal devices such as Google Glasses) provides many new opportunities for displaying memory cues to trigger recall.

It is self-evident that we do not effectively encode all of the information that we encounter, nor are we able to retrieve at will, all of that content that we do encode. When trying to recall known facts, many of our day-to-day memory failures result from a temporary failure to retrieve memories rather than from their permanent erasure. Our ability to recall target information is particularly vulnerable to transient changes in accessibility that arise through changes in the contents of our short-term memories and the cues in our environment. That memory can be improved with effective cues is beyond doubt: whilst a typical participant might be able to recall only 38 out of a set of 100 words that had been continually studied and sorted over many minutes, this accuracy increases to 96% when the most effective cues are presented at test. One experiences these temporary failures to retrieve memories everyday when we might remark "I cannot recall X (e. g., his name, the company, the town, etc.), but if I saw it, I would recognise it". Tellingly, we are unlikely to experience or say the converse.

One of the most frustrating features of human memory is that we are particularly vulnerable at remembering to do something in the future (the area of memory research known as prospective memory). Prospective memory failures readily occur for remembering time-based future events (hence the value of setting computer alarms reminding us of meetings), and for remembering event-based future events (remember to post a letter on the way to work, remember to pick up a takeaway for the family tonight). Research suggests that whereas there is a general decline in memory with increasing old age, it is prospective memory and retrieval in the absence of cues that are particularly impaired, whereas cued recall and recognition are more preserved.

The Dagstuhl Seminar 14362 "Augmenting Human Memory – Capture and Recall in the Era of Lifelogging" focused on a vision of the world in which augmented memory systems make everyday use of peripheral, ambient multi-media content – delivered via large wall-mounted

displays, smartphone wallpapers, or wearable in-eye projectors – to intelligently integrate, display, and enable the review of life-relevant personal data. Such memory augmentation technologies have the potential to revolutionise the way we use memory in a wide range of application domains.

Behaviour Change: Effecting behaviour change is an important objective in many important areas such as health (e. g. lifestyle changes such as increasing exercise or stopping smoking) and sustainable transport (e. g. encouraging people to make more environmentally-friendly transport choices). Unfortunately, despite good intentions, many people experience difficulty in implementing planned behaviour: for example, it is well known that many people are reluctant to make a trip to the gym despite paying large gym membership fees. Psychological theory stresses that intentional behaviours are more likely to be implemented when individuals are reminded of their own attitude towards such behaviours (e. g., the positive gains that will result), and the attitudes of significant others to the behaviour (what loved ones, family, friends, peers, and society in general think of the behaviour and its outcomes). In addition, realistic scheduling is important: planned behaviour is more likely to be performed if it is timetabled with the transition from immediately preceding activities in mind. Finally, behaviour is more likely if it is perceived to be more achievable and more enjoyable. Memory augmentation can help with the realistic scheduling and reminding of the planned activities, and can remind people at the point at which decision making is necessary (e. g., at the planned time to visit the gym) of the positive benefits from the behaviour, the previous good experience of the behaviour and the progress that is being made.

Learning: Such technologies can be used as part of a learning environment. In particular, through the use of ambient displays it might be possible to cue recall, and hence reinforce learning of a wide range of skills. For example, the acquisition of a new language could be supported by providing appropriate cues to facilitate recall of vocabulary. Similarly, a class teacher could be encouraged to remember the names of their pupils, and a study abroad student could learn culturally-significant facts as they explore a new city.

Supporting Failing Memories: Research has shown that as we age, our ability to perform uncued recall is particularly vulnerable to age-related decline. Memory augmentation technologies could be used to help remedy this memory loss by providing older users with time-relevant and context-appropriate cues. In this way, older individuals could enjoy greater self-confidence and greater independence by being reminded of moment-by-moment situated details of where they were, what they were intending to do, and how they could get home. They may also enjoy better relationships if they could be reminded of the autobiographical details of their loved ones (such as the names and ages of their loved ones' children), or if they could review and then be reminded of the details of a recent conversation or event (e. g., a recent day out or family gathering).

Selective Recall: Through appropriate selection of memory cues that are presented to the user, memory augmentation technologies might also be used to facilitate selective recall. According to the psychological theory of retrieval-induced forgetting, the act of reviewing memories not only enhances the probability of spontaneously retrieving these reviewed memories in the future, but it can also attenuate the spontaneous retrieval of related but unreviewed memories. The study of retrieval-induced forgetting has largely been confined to the laboratory using lists of categorised words. It is of both pure and applied interest (e. g., the desired attenuation of unwanted, outdated, or traumatic memories; and the undesired attenuation of wanted but unreviewed memories) to see if this phenomenon can be observed when reviewing a subset of "real world" memories, and if so, we will be

able to measure the extent to which unreviewed memories could be attenuated through selective reviewing.

Memory Based Advertising: While many of the application domains for memory augmentation technologies are for the public good, the same technologies can also be employed in the commercial context. For example, such technologies could be used to support a new form of advertising in which users have memories triggered explicitly to drive purchasing decisions. For example, when passing a shop selling luggage a cue could be presented that causes a passer-by to remember a specific experience from their past in which their own luggage didn't work satisfactorily. This may then cause the user to enter the shop and purchase some new luggage.

Collectively, the seminar participants explored the scientific foundations for a new technology eco-system that can transform the way humans remember in order to measurably and significantly improve functional capabilities while maintaining individual control. At its heart lies the creation of memory augmentation technology that provides the user with the experience of an extended and enhanced memory. Such technology is based on recent improvements in the collection, mining, and presentation of appropriate information to facilitate cued memory recall. This research is inherently multidisciplinary and combines elements of pervasive computing, information retrieval and data privacy with psychology and sociology.

2 Table of Contents

3 Overview of Talks

The research questions we pursued included the following:

1. Collection and Control. What is the best mix of technologies for capturing relevant human experiences to improve human memory? How can we create a novel class of capture systems that specifically support human memory functions while offering a fine-grained level of control over their recording and fully respecting personal privacy? This is likely to entail novel approaches to the integration of captured data across devices, domains, and owners.
2. Presentation. What are appropriate tools and methods for integrating, correlating, and visualizing captured sensor data and other information sources into a coherent "memory prosthetics" streams? Such streams will be based on theoretical principles of human memory organization, in order to positively influence the acquisition, retention and attenuation of knowledge from personal experiences.
3. Theory. On a theoretical level, we wish to explore validation of human memory theory in these new systems, targeting the feasibility of targeted attenuation of unwanted memories.

Most participants presented short talks centered around one of these questions in order to start off discussions.

3.1 Robust regularities of human memory function

C. Philip Beaman (University of Reading, GB)

A short bullet-point list is presented of "principles" of human memory relevant to lifelogging. The list is personal and not exhaustive but is indicative of the phenomena of human memory revealed from investigation within the experimental psychology lab. No attempt is made to provide theoretical justification for the principles – rather, they summarise robust and generally agreed empirical regularities that may be of use when considering what aspects of memory can or should be augmented, supplemented or enhanced.

3.2 Augmenting memory – a means to what end?

Michel Beaudouin-Lafon (University of Paris South XI, FR)

While I do not work directly on lifelogging, I have worked in several areas that can inform this field.

Twenty years ago, I worked on media spaces [1], which were designed to help people share a workspace at a distance. We learned a number of lessons related to the social aspects of human communication when it is mediated by technology, such as the law of reciprocity, the need to control accessibility or the transitions between different levels of engagement [2]. Similar principles could be applied to lifelogging as it also provides a similar window into people lives and activities.

Ten years later, I worked on the European InterLiving project, where we developed technology to help distributed families to live together. We created a number of technology probes [3], among which the VideoProbe [4], which automatically captured and shared photos based on activity. We discovered the importance of intimate social networks, the difficulty of designing technology that meshes with people's everyday life, and the active co-adaptation of people and technology. Similar co- adaptation is at play with lifelogging systems, and should be taken into account in their design.

More recently, I have worked on large interactive spaces that support distributed and collaborative interaction [5]. Part of this work involves creating gesture-based interfaces with larger vocabularies than the typical touch-based interfaces of today. We have created a number of techniques, based on dynamic guides and the combination of feedforward and feedback, to help users learn and remember gestures [6, 7]. This leads to the notion of co-adaptive instruments that users can learn, but that can also shape users' behaviors. Can lifelogging be used to augment not only memory, but other skills as well?

My challenge to the participants of this workshop is to go beyond the recording of events "just because we can" to a better understanding of how we can create new ways of externalizing experiences so that they can be shared with others and so that we can learn from them, both individually and collectively. Augmenting memory is a noble goal, but to what end? Memory enables us to avoid repeating mistakes and lets us learn and transmit skills. The mere recordings of our lives provided by current lifelogging systems are still insufficient to provide these functions. But shouldn't the goal also be to support new functions? The previous externalization of memory was written language: unlike an audio recording of spoken language, it enabled skimming, summarizing, commenting, etc. Similarly, we need to create appropriate instruments [8] to represent, access, share and manipulate past experiences so that we can imagine new ways to enhance our collective intelligence.

References

1 Wendy Mackay. Media Spaces: Environments for Informal Multimedia Interaction. In Michel Beaudouin-Lafon (ed), *Computer Supported Co-operative Work*. John Wiley & Sons, 1999.

2 Nicolas Roussel. From analog to digital, from the office to the living room: why I happily worked in a media space but don't live in one. In *Media Space: 20+ Years of Mediated Life*. Springer, 2009.

3 Hilary Hutchinson, Wendy Mackay, Bosse Westerlund, Benjamin Bederson, Alison Druin, Catherine Plaisant, Michel Beaudouin-Lafon, Stéphane Conversy, Helen Evans, Heiko Hansen, Nicolas Roussel, Björn Eiderbäck, Sinna Lindquist and Yngve Sundblad. Technology Probes: Inspiring Design for and with Families. In *Proc. of ACM CHI 2003 Conf. on Human Factors in Computing Systems*, pp. 17–24, AMC, 2003.

4 Stéphane Conversy, Wendy Mackay, Michel Beaudouin-Lafon and Nicolas Roussel. VideoProbe: sharing pictures of everyday life. *LRI research report 1409*, Université Paris-Sud, France, 8 pages, April 2005.

5 Michel Beaudouin-Lafon, Olivier Chapuis, James Eagan, Tony Gjerlufsen, Stéphane Huot, Clemens Klokmose, Wendy Mackay, Mathieu Nancel, Emmanuel Pietriga, Clément Pillias, Romain Primet and Julie Wagner. Multi-surface Interaction in the WILD Room. *IEEE Computer*, 45(4):48–56, 2012.

6 Olivier Bau and Wendy Mackay. OctoPocus: A Dynamic Guide for Learning Gesture-Based Command Sets. In *Proc. of ACM Symp. on User Interface Software and Technology* (UIST 2008), pp. 37–46, ACM, 2008.

7 Emilien Ghomi, Stéphane Huot, Olivier Bau, Wendy E. Mackay and Michel Beaudouin-Lafon. Arpège: Learning Multitouch Chord Gestures Vocabularies. In *ACM Int'l Conf. on Interactive Tabletops and Surfaces*, pp. 209–218, ACM, 2013.
8 Michel Beaudouin-Lafon. Designing interaction, not interfaces. In *Proc. of the Conference on Advanced Visual Interfaces* (AVI 2004), pp. 15–22, ACM, 2004.

3.3 Re-Live the moment: Using run visualizations to provide a positive feedback loop for workouts

Agon Bexheti (Università della Svizzera italiana (USI), CH)

Behaviour change is considered an important objective in many areas such as health (e. g. quit smoking, increasing exercise level, etc.) and sustainable transport. Nevertheless, for most of us it is very difficult to implement planned behaviour. We plan to exploit a finding from psychology research which states that planned behaviour is more likely to be achieved when individuals are reminded of their own attitudes towards such a behaviour (e. g. the positive benefits of running regularly). For this reason, we have developed a prototype that records among others images and background music of running sessions and assembles a multimedia slideshow which is played afterwards. The aim of this work is to encourage individuals to run more by remembering the fun they had during their last run and the post-run satisfaction they felt.

3.4 Empathic Computing

Mark Billinghurst (University of Canterbury – Christchurch, NZ)

In recent years there has been considerable research on how to build user interfaces that recognise and respond to emotion. However there has been less research on how to create shared understanding or empathy between users, using technology to allow one person to better understand the emotions of another. In this presentation we talk about Empathic Interfaces that are designed to go beyond understanding user input and to help create shared understanding and emotional experiences. We explore how Augmented Reality (AR) can be used to convey that emotional state and so allow users to capture and share emotional experiences. In this way AR not only overlays virtual imagery on the real world, but also can create deeper understanding of user's experience at particular locations and points in time. The recent emergence of truly wearable systems, such as Google Glass, provide a platform for Empathic Communication using AR. Examples will be shown from research conducted at the HIT Lab NZ and other research organizations, and key areas for future research described.

3.5 On the potential of human visual behaviour for memory augmentation and life logging

Andreas Bulling (MPI für Informatik – Saarbrücken, DE)

Joint work of Bulling, Andreas; Zander, Thorsten
Main reference A. Bulling, T. O. Zander, "Cognition-aware computing," IEEE Pervasive Computing, 13(3):80–83, July 2014.
URL http://dx.doi.org/10.1109/MPRV.2014.42

In his talk, Dr Bulling discussed the potential of eye tracking and eye movement analysis for life logging and human memory augmentation. He pointed out that what makes our eyes so interesting is the fact that they are with us wherever we go and whatever we do. Their movements are closely linked to our activities, goals, intentions, they indicate attention and what we are interested in, and they are linked to a large number of cognitive processes. This link to human cognition makes the eyes a particularly interesting sensing modality and a promising means to implement the vision of cognition-aware computing. Cognition-aware computing systems sense and adapt to the so-called cognitive context of the person that is comprised of all aspects related to mental information processing, such as reasoning, memory and learning. He then summarised recent work by his group on eye-based activity and context recognition, long-term visual behaviour analysis and automatic inference of visual memory recall. He concluded by briefly summarising the state-of-the-art in mobile eye tracking technology and challenges in using eye tracking in daily life settings.

References

1 A. Bulling and D. Roggen. Recognition of visual memory recall processes using eye movement analysis. In *Proc. of the 13th Int'l Conf. on Ubiquitous Computing (UbiComp 2011)*, pp. 455–464, ACM Press, 2011.
2 A. Bulling, D. Roggen, and G. Tröster. Wearable EOG goggles: Seamless sensing and context-awareness in everyday environments. *Journal of Ambient Intelligence and Smart Environments*, 1(2):157–171, May 2009.
3 A. Bulling, J. A. Ward, H. Gellersen, and G. Tröster. Eye Movement Analysis for Activity Recognition Using Electrooculography. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(4):741–753, April 2011.
4 A. Bulling, C. Weichel, and H. Gellersen. Eyecontext: Recognition of high-level contextual cues from human visual behaviour. In *Proc. of the 31st SIGCHI Int'l Conf. on Human Factors in Computing Systems*, pp. 305–308, 2013.
5 A. Bulling and T. O. Zander. Cognition-aware computing. *IEEE Pervasive Computing*, 13(3):80–83, July 2014.
6 P. Majaranta and A. Bulling. Eye tracking and eye-based human–computer interaction. In *Advances in Physiological Computing*, pp. 39–65. Springer, 2014.
7 M. Vidal, A. Bulling, and H. Gellersen. Pursuits: Spontaneous interaction with displays based on smooth pursuit eye movement and moving targets. In *Proc. of the 2013 ACM Int'l Joint Conf. on Pervasive and Ubiquitous Computing (UbiComp 2013)*, pp. 439–448, 2013.
8 M. Vidal, J. Turner, A. Bulling, and H. Gellersen. Wearable eye tracking for mental health monitoring. *Computer Communications*, 35(11):1306–1311, 2012.
9 Y. Zhang, A. Bulling, and H. Gellersen. Sideways: A gaze interface for spontaneous interaction with situated displays. In *Proc. of the 31st SIGCHI Int'l Conf. on Human Factors in Computing Systems*, pp. 851–860, 2013.

3.6 The Big Picture: Lessons Learned from Collecting Shared Experiences through Lifelogging

Nigel Davies (Lancaster University, UK)

The emergence of widespread pervasive sensing, personal recording technologies and systems for quantified self are creating an environment in which it is possible to capture fine-grained traces of many aspects of human activity. Such traces have uses in a wide variety of application domains including human memory augmentation, behaviour change and healthcare. However, obtaining these traces for research purposes is non-trivial, especially when they contain photographs of everyday activities. In order to source traces for our own work we created an experimental setup in which we collected de- tailed traces of a group of researchers for a period of 2.75 days. We share our experiences of this process and present a series of lessons learned that can be used by other members of the research community proposing to conduct similar experiments in order to obtain appropriately detailed traces that include photographic images.

3.7 The quantified self: Understanding and augmenting human memory

Simon Dennis (University of Newcastle, AU)

An overview of dynamical systems, behavioural, and neurscientific approaches to understanding human memory using lifelong data as well as a brief description of an initial attempt to build a Google for your life context retrieval system.

3.8 Designing Knowledge Acquisition Points: Speeding up Reading Tasks on Electronic Devices

Tilman Dingler (Universität Stuttgart, DE)

Joint work of Dingler, Tilman; Alireza Sahami

Reading is an ancient activity traditionally taken up for information gain and pleasure. With the advent of the information age and the rising popularity of electronic reading devices our reading behavior has been changing and we are facing an abundance of text on a daily basis. However, our reading strategy has mainly remained the same since our formal reading education stopped in a young age. Naturally people develop their innate reading, skimming and skipping strategies. Rapid Serial Visual Presentation (RSVP) has been proposed as a reading technique to push a reader through a text by displaying single or groups of words sequentially in one focal point. Other techniques include the use of a kinetic stimulus (such as a moving pen or finger) to guide a reader consistently across lines of text. We implemented

a number of different kinetic stimuli to guide the reader's eye across text passages on a computer screen. In a user study with 36 participants we assessed the effects of different stimuli (including RSVP) on comprehension level, perceived mental load and eye movements. We envision electronic devices to be able to detect the skill level of readers, take into account the text type and apply different reading strategy options to facilitate reading tasks.

3.9 Sensing People

Christos Efstratiou (University of Kent, GB)

Joint work of Efstratiou, Christos; Leontiadis, Ilias; Picone, Marco; Mascolo, Cecilia; Rachuri, Kiran; Crowcroft; Jon
Main reference C. Efstratiou, I. Leontiadis, M. Picone, K. Rachuri, C. Mascolo, J. Crowcroft, "Sense and Sensibility in a Pervasive World," in Proc. of the 10th Int'l Conf. on Pervasive Computing (PERVASIVE'12), LNCS, Vol. 7319, pp. 406–424, Springer, 2012.
URL http://dx.doi.org/10.1007/978-3-642-31205-2_25

The popularity of online social networking services has increasingly transformed them into lifelogging services where personal activities are manually logged and shared with others. The availability of a wide range of sensing technologies in our everyday environment presents an opportunity to further enrich social networking systems with fine-grained real-world sensing. The introduction of real-world sensing into a social networking applications can allow the unbiased logging of life activities, enabling the recording of a more accurate picture of our daily lives. At the same time, passive sensing disrupts the traditional, user-initiated input to social services, raising both privacy and acceptability concerns. In this work we present an empirical study of the introduction of a sensor-driven social sharing application within the working environment of a research institution. Our study is based on a real deployment of a system that involves location tracking based on smartphone indoor localisation, conversation monitoring using microphones in the environment, and interaction with physical objects augmented with sensors, such as desks, coffee machines, ect. The system allowed the detection of social activities, such as co-location of colleagues and participation in conversation, feeding them into a private web-based social networking platform. Following a 2 week deployment of the system involving 21 participants, we report on findings regarding privacy and user experience issues, and significant factors that can affect acceptability of such services by the users. Our results suggest that such systems deliver significant value in the form of self reflection and comparison with others, while privacy concerns are raised primarily by the limited control over the way individuals are projected to their peers.

References

1 Efstratiou, Christos and Leontiadis, Ilias and Picone, Marco and Rachuri, Kiran and Mascolo, Cecilia and Crowcroft, Jon (2012) *Sense and Sensibility in a Pervasive World.* In: Proc. of 10th Int'l Conf. on Pervasive Computing (PERVASIVE 2012), June 2012, Newcastle, UK.

2 Rachuri, Kiran and Efstratiou, Christos and Leontiadis, Ilias and Mascolo, Cecilia and Rentfrow, Peter J. (2013) *METIS: Exploring mobile phone sensing offloading for efficiently supporting social sensing applications.* In: Proc. of the IEEE Int'l Conf. on Pervasive Computing and Communications (PERCOM 2013).

3.10 Augmenting Human Memory For UX Evaluation

Evangelos Niforatos (Università della Svizzera italiana (USI), CH)

We present a set of work for supporting episodic memory recall with cues for obtaining feedback on User Experience (UX). EmoSnaps is a mobile app that captures unobtrusively pictures of one's facial expressions throughout the day and uses them for later recall of her momentary emotions. We found that that people are better able to infer their past emotions from a self-face picture the longer the time has elapsed since capture. Then, a follow up study demonstrates the eMotion mobile app, which collects memory cues during commute for measuring drivers' anger and frustration levels retrospectively. Next, we presented a planned study for contrasting the traditional, limited film camera capturing with the new digital cameras and automatic life logging tools. We created a mobile app (MyGoodOldKodak) to investigate the effect of capture limitation on the user picture capturing behavior. Last, we demonstrated Atmos, a tool for crowdsourcing user estimations about current and future weather conditions.

3.11 Lifelogging at Dublin City University

Cathal Gurrin (Dublin City University, IE)

Main reference C. Gurrin, A. F. Smeaton, A. R. Doherty, "LifeLogging: personal big data," Foundations and Trends in Information Retrieval, 8(1):1–125, 2014.
URL http://dx.doi.org/10.1561/1500000033

We have recently observed a convergence of technologies to foster the emergence of lifelogging as a mainstream activity. In this talk I provided an introduction to the vision of lifelogging that we hold in DCU, which is a form of pervasive computing which utilises software and sensors to generate a permanent, private and unified multimedia record of the totality of an individual's life experience and makes it available in a secure and pervasive manner.

The four core components of our view are sensing, segmenting, indexing and interacting [1] and they were combined in a number of demonstrator descriptions covering digital visual dairies, QS visualisations and object deception and search. Finally, the issue of privacy of lifelogs was considered and one proposal presented that could help to maintain privacy of subjects and bystanders in a world of ubiquitous lifelogging.

References

1 Gurrin, Cathal and Smeaton, Alan F. and Doherty, Aiden R. *LifeLogging: personal big data*. Foundations and Trends in Information Retrieval, 8(1):1–125.

3.12 Lifelogging and Wellbeing for Care Home Residents

Vicki Hanson (Rochester Institute of Technology, US)

My interest in lifelogging stems from two research projects I have underway: Portrait and BESiDE, both of which deal with the wellbeing of care home residents. Portrait seeks to provide care staff (those charged with residents' day-to-day care such as bathing and dressing) with information about the personal and social life experiences of these residents. The goal is to provide care staff with a simple means of getting information to support person-centered care. BESiDE is an interdisciplinary effort that seeks to inform building design about features that facilitate physical and social activities of residents. This talk considers lifelogging methods that can inform these efforts with respect to positive interactions for care home residents.

3.13 Augmenting Food with Information

Niels Henze (Universität Stuttgart, DE)

Eating is not only one of the most fundamental human needs but also among the most regular human activities. Consequently, preparing meals is deeply rooted in all human cultures. Food, however, can not only serve to satisfy hunger but also to ubiquitously communicate information. Through food augmentation, a dinner can communicate its characteristics such as the ingredients. Food can provide instructions, for example, the recipe of a meal or communicate arbitrary information such the eater's schedule in a way that can hardly be avoided.

3.14 Activity-Enriched Computing: Retrieving and Restoring Context

James D. Hollan (University of California – San Diego, US)

The intertwining of computers with virtually every aspect of life brings many benefits, but also a growing stream of interruptions. Even though the fragmenting of activity is an increasingly accepted part of modern life, a critical research challenge remains: how to smooth and mitigate its impact and assist in resuming interrupted activities.

3.15 Photobooks for Memories

Christoph Korinke (OFFIS – Oldenburg, DE)

Many photos are taken at events to share experiences later on with others or for personal remembrance. Photobooks are used to preserve events and focus on the most pressures

moments of an event. It is common to take several shots of the same scene. Looking at thousands of them is tedious. The selection, arrangement and annotation of photos is time consuming and can be a major obstacles preventing the creation of a photobook. Algorithms can support the user by clustering and ranking images. These algorithms can even take aesthetic qualities into account by enabling (semi-) automatic creation of books with respect to e. g. golden ratio. Enrichment can be supported with scene understanding algorithms. A major challenge in case of lifelogging is to distinguish between "picture vs. moment", i. e. a photo can be an emotional moment for one person, but just be a picture for another.

3.16 Augmenting the Human Mind

Kai Kunze (Osaka Prefecture University, JP)

The talk gives an overview about the emerging field of smart glasses and how they can be used to augment our mind (e. g. how to improve our brain with technology). The talk will focus mostly on how to quantify cognitive tasks in real world environments. I also present a first application scenarios on how to use smart eyewear (e. g. Google Glass or J!NS MEME) for short term memory augmentation and cognitive activity recognition.

Considering the last centuries, major scientific breakthroughs aimed at overcoming our physical limitations (faster transportation, higher buildings, longer, more comfortable lifes). Yet, I believe the coming big scientific breakthroughs will focus on overcoming our cognitive limitations. Smart glasses can play a vital role in

1. understanding our cognitive actions and limitations by quantifying them
2. helping us design interventions to improve our mind.

The talk will focus mostly on the first point, what kind of cognitve tasks can we track already with the smart glasses that are available in the market and what will happen in the near future. I will discuss application examples for Google Glass and J!NS MEME. J!NS MEME is the first consumer level device measuring eye movements using electrodes also called Electrooculography (EOG). The MEME glasses not a general computing platform. They can only stream sensor data to a computer (e. g. smart phone, laptop, desktop) using Bluetooth LE. Sensor data includes vertical and horizontal EOG channels and accelerometer + gyroscope data. The runtime of the device is 8 hours enabling long term recording and, more important, long term real-time streaming of eye and head movement. They are unobtrusive and look mostly like normal glasses. For Google Glass I present an open sensor-logging platform (including the infrared sensor to count eye blinks) and a fast interface to do lifelogging. We will discuss which eye movements correlate with brain functions and how this fact can be used to estimate the cognitive task a user is performing, from fatigue detection, over reading segmentation to cognitive workload and the advances to track attention and concentration. Challenges discussed in the talk include how to get ground truth and how to evaluate performance in general.

3.17 On the Design of Digital Mementos

Daniela Petrelli (Sheffield Hallam University, UK)

Do we have precious memories captured in digital form? How can we give those "files" a presence in our life? How would digital mementos look like and how would people interact with them to reminisce about their life? In a number of studies conducted with families I have explored these questions trying to tease out what are the principles behind the keeping of mementos and how we could reproduce the same affective level of engagement with relevant digital belongings. Findings relevant to Lifelogging technology include: as heterogeneous objects are clustered together in memory boxes, so data collected through multiple lifelog sources be fused; the physical presence of mementos in our environment makes them salient and memorable over an extended period of life; there is a strong emotional power in forgetting and rediscovery autobiographical memories. I have discussed two examples of digital mementos, bespoke devices that enable users to easily access their digital belonging in an engaging a playful way to support meaning making over time.

3.18 Digital Sign Analytics in the Context of Memory Augmentation

Mateusz Mikusz (Lancaster University, GB)

Joint work of Mikusz, Mateusz; Clinch, Sarah; Davies, Nigel

Digital signs are already present in many public places and are integrated in urban life. Researchers believe that these screens are not being used to their full potential, supporting human memory is one example of how these displays could be used in future application areas. To enable these applications we can combine displays with devices such as mobile phones and other wearables to support the aim of showing personalized and personal content. A study at Lancaster University showed that displaying personal images on pervasive displays for the individual passing by helps to trigger memory recall, and can also improve attention for digital signs. In order to allow developers and researches to analyze their display applications, we believe that digital sign analytics will be of essential relevance in future. Analytics are important to provide detailed information about movement patterns and behavior in front of the sign and across devices, for example how the walking path or destination changed after seeing an advert or other content on a screen. For memory applications, analytics are required to measure the success of new applications, the result of memory recall, and understanding whether the shown content on a screen leads to the expected behavior, measured across devices.

3.19 Reinstating the Context of Interrupted Activities

Adam C. Rule (University of California – San Diego, US)

Joint work of Rule, Adam; Tabard, Aurélien; Hollan, Jim

How can we help people re-establish the context of activities that have been interrupted? One approach is Activity-Based Computing, which groups computing resources such as documents and applications by activity so they can be closed and restored en masse. However, activities are more than just collections of documents but include fragile mental states. Tracking and visualizing computer activity could help cue these mental states. We developed a simple tool to track computer activity using key logging and screen recording. This process raised three questions: 1) What are the tradeoffs between implicit and explicit tracking? 2) How should we handle tracking and visualizing similar activities? 3) Where should tracking stop?

3.20 Emotional Memories: Cueing, Forgetting and Digital Disposal

Corina Sas (Lancaster University, GB)

Joint work of Sas, Corina; Whittaker, Steve
Main reference C. Sas, S. Whittaker, "Design for forgetting: disposing of digital possessions after a breakup," in Proc. of the 2013 ACM SIGCHI Conf. on Human Factors in Computing Systems (CHI'13), pp. 1823–1832, ACM, 2013.
URL http://dx.doi.org/10.1145/2470654.2466241

Episodic memories lie at the core of our sense of identity, and emotions play important role in organizing and cuing them. This talk introduces AffectCam, a wearable system integrating SenseCam and BodyMedia SenseWear capturing galvanic skin response as a measure of bodily arousal, with the aim to explore the value of arousal arising naturally in daily life, as a mechanism for cuing episodic recall. The system was tested with 14 participants who were asked at the end of a day of wearing the sensor to recall events cued by top high and low arousal-stamped pictures. Findings suggest the value of arousal as a filtering mechanism, with 50% richer recall cued by high arousal photos.

Most theorists argue for retaining all these possessions to enhance 'total recall' of our everyday lives, but there has been little exploration of the negative role of digital possessions. The second part of the talk focused on digital disposal and intentional forgetting following digital breakup. We interviewed 24 participants and found that digital possessions were often evocative and upsetting in this context, leading to distinct disposal strategies with different outcomes. We advance theory by finding strong evidence for the value of intentional forgetting and provide new data about complex practices associated with the disposal of digital possessions. Our findings led to a number of design implications to help people better manage this process, including automatic harvesting of digital possessions, tools for self-control, artifact crafting as sense-making, and digital spaces for shared possessions.

References

1 Corina Sas; Thomasz Fratczak; Matthew Rees; Hans Gellersen; Vaiva Kalnikaite; Alina Coman; Kristina Höök. *AffectCam: arousal-augmented Sensecam for richer recall of episodic memories.* . In CHI'13 Ext. Abstracts, pp. 1041–1046, ACM, 2013.

2 Elise van den Hoven; Corina Sas; Steve Whittaker. *Introduction to this Special Issue on Designing for Personal Memories: Past, Present and Future*. Human-Computer Interaction, 27(1–2):1–12, 2012.
3 Nigel Davies, Adrian Friday, Sarah Clinch, Corina Sas, Marc Langheinrich, Geoffrey Ward, and Albrecht Schmidt (2015). Security and Privacy Implications of Pervasive Memory Augmentation. *IEEE Pervasive Computing*. 14(1):44–53, 2015.
4 Corina Sas; Alan Dix. *Designing for Collective Remembering* In Extended Abstracts of CHI, pp. 1727–1730, ACM, 2006.

3.21 Déjà Vu – Technologies that make new Situations look Familiar

Albrecht Schmidt (University of Stuttgart, DE)

Joint work of Davies, Nigel; Langheinrich, Marc; Ward, Geoffrey
Main reference A. Schmidt, N. Davies, M. Langheinrich, G. Ward, "Déjà Vu – Technologies That Make New Situations Look Familiar: Position Paper," in Adjunct Proc. of the 2014 Int'l Joint Conf. on Pervasive and Ubiquitous Computing (UbiComp'14) – Workshop on Ubiquitous Technologies for Augmenting the Human Mind (WAHM'14), pp. 1389–1396, ACM, 2014.
URL http://dx.doi.org/10.1145/2638728.2641720

We envision a technology concept for making new situations and encounters more familiar and less threatening. Going to new places, interacting with new people and carrying out new tasks is part of everyday life. New situations create a sense of excitement but in many cases also anxiety based on a fear of the unknown. Our concept uses the metaphor of a pin board as peripheral display to automatically provide advance information about potential future experiences. By providing references to and information about future events and situations we aim at creating a "feeling of having already experienced the present situation" (term Déjà Vu as defined in the Oxford Dictionary) once people are in a new situation. This draws on the positive definition of the concept of déjà vu. We outline the idea and use scenarios to illustrate its potential. We assess different ways the concept can be realized and chart potential technology for content creation and for presentation. We also present a discussion of the impact on human memory and how this changes experiences.

References

1 Nigel Davies, Adrian Friday, Sarah Clinch, Corina Sas, Marc Langheinrich, Geoffrey Ward, and Albrecht Schmidt (2015). Security and Privacy Implications of Pervasive Memory Augmentation. *IEEE Pervasive Computing*. 14(1):44–53, 2015.

3.22 Visualisation for activity based computing

Aurelien Tabard (University Claude Bernard – Lyon, FR)

We present our ongoing efforts to visualize computer based activity. We aim at leveraging visual perception and visual memory to explore past activities. This approach avoids the problems of activity recognition (trying to infer automatically activities) and activity specification (letting users specify what their ongoing activity is).

One premise of our work is that such visualizations need only be intelligible to the person whose activity was capture as s/he will be able to make sense of it in a much richer way based on the past experiences. Another important point is our strategies to compress time in order to enable more efficient browsing of highly visual temporal data (i. e. screenshots of desktops). To this end, we explore animations, key point within screenshots, and activity breakpoints as important memory cues.

3.23 Long-term activity recognition to provide a memory to others

Kristof Van Laerhoven (TU Darmstadt, DE)

Main reference E. Berlin, K. Van Laerhoven, "Detecting Leisure Activities with Dense Motif Discovery," in Proc. of the 2012 ACM Conference on Ubiquitous Computing (UbiComp'12), pp. 250–259, ACM, 2012.
URL http://dx.doi.org/10.1145/2370216.2370257

The concept of activity recognition, in which machines are able to capture what a person is doing, is usually described in terms of where the system is situated (worn by the user, in the user's environment) or what system's purpose is (to improve ones' health, to serve as a memory aid for later, etc.). In my presentation, I presented the numerous challenges that lie in focusing on long-term activity recognition in which the system is recording the day-to-day activities of someone else.

Many studies that involve tracking behaviour related to lifestyle, mental state, or mood are currently based on questionnaires that rely on study participants remembering what they were doing in the past days or past weeks. This reliance on human memory poses several dangers: recall can be notoriously bad and is often biased. Having a machine detect activities automatically, with a minimum of intervention required by the study participant, would be an attractive alternative. Sleep researchers could monitor the activities that their patients would perform during the day and the effect they have on subsequent nights, or psychiatrists could analyse the interplay between certain episodes such as depressions or mania, and regularly- performed activities.

As a case study of such research, I presented work from a wrist-worn activity recognition system [1], in which long-term 24/7 inertial data is logged and analysed for specific leisure activities such as practicing yoga, or playing the guitar. For creating such a memory system, I stressed on the current challenge of making such a system powerful enough so that it can parse the huge amount of data that long-term studies of weeks to months of inertial data (taken at 100Hz) produce. A second challenge that remains to be solved, I argued, is the proper visualisation of said data so that it can be used by others.

References

1 Berlin, Eugen and Van Laerhoven, Kristof *Detecting Leisure Activities with Dense Motif Discovery*. UbiComp 2012, ACM Press.

3.24 The act of recall will affect future recall: Early insights into intelligent reviews

Geoff Ward (University of Essex, GB)

We entertain the possibility that one might wish to review a daily digest of one's day, and we consider what effect such a review might that have on one's spontaneous ability to later access one's memory for past events. We present three experiments that examine whether the phenomenon of retrieval-induced forgetting (RIF) could be observed in the real world, as well as in the psychology laboratory. According to this literature, actively retrieving a subset of memories will lead to enhanced later accessibility to these practiced items, but decreased accessibility to related but unreviewed memories. We replicate and extend RIF in the laboratory using words and pictures as stimuli, but our initial experiments using more real-world stimuli (i. e., fictitious holiday events) suggest that the accessibility to real life events increases by being reviewed but reviewing does not decrease the later accessibility of related but unpracticed memory events.

3.25 Lifelogging image presentation and navigation

Katrin Wolf (Universität Stuttgart, DE)

In this talk, characteristics of lifelogging images are discussed and initial design ideas for lifelogging image navigation are presented.

Image characteristics changed a lot over the history of image production. Capture devices for photography and movie production were in the beginning expensive and required expert knowledge. As also the image production cost money, the images and movies produced in the old times were often showing important events, such as weddings. Moreover, portraits were professionally captured in special and rare moments in rather wealthy peoples' life. During the last hundred year image capture technology became affordable for many people, the procedure of picture taking does not require special skills anymore, the pictures do not need to be developed but are immediately accessible, and the process of recording pictures do not cost money anymore. Thus, image capturing is done by everybody and massive amounts of images are produced, need to be stored, achieved, and novel methods of image organization are needed to be able to browse through archives and find specific pictures or movies. Lifelogging cameras allow for passively take pictures. Thus, large data is produced that is difficult to be archived and even harder to search in for certain information and memory recall.

A method of navigating through large image data is to use meta information, which allows to apply computational search algorithms for information navigation and recall. In autobiographical memory, time, place, persons, events, and emotions are used as memory cues, and we propose to use the same cues as meta information to navigate through autobiographical image data. Finally initial mockups were presented that build on the idea of allowing for navigating through by using the autobiographical memories mentioned before.

4 Working Groups

The seminar included a series of break-out sessions to foster more in-depth discussion around 3 central topics: Visualization, Applications of Lifelogging, and Social Impact.

4.1 Visualization

This group discussed the topic of *visualizations* and basically took on 3 key questions:

- How do we support in-the-moment annotating?
- How do we visualize multi-dimensional information?
- How do we search for, select, and filter the recordings

Their output included a number of key thoughts, outlined as follows: the way a visualization is designed depends on the **intended application**. Systems designed for personal reflection will differ significantly from those designed for social communication. We can use **abstractions** or parts of an image to convey information about what was happening. For example, tracking a single pixel of a doorway can show occupancy. **Space** can be used to show time. For example, using the periphery of a screen to show what happened in the past, or what will happen next. Another example is using small multiples. **Multiple interactive views** can be used for search, selection, and filtering Visualizations should support **multi-scale viewing and navigation**. For example, viewing locations at the country, state, and city level Techniques from **cinematography** can inspire life-log visualizations. For example, storyboards and panoramas. Certain applications will want to preserve a storytelling element.

4.2 Applications of Lifelogging

This group engaged with the topic of identifying and sketching out the application space for lifelogging. Key questions included:

- Why do people engage in lifelogging?
- What lifelogging activities are there, what goals and benefits people envisage from them?
- How can lifelogging practice be motivated?
- Which are the challenges of adopting lifelogging and how can we address them?

The group discussed main activities of lifelogging, such as information retrieval, event reminiscing and the reflection on successive events. Thereby the granularity of the activity's details may differ: in some cases retrieval of specific aspects or factual data about an event may be required, such as the name of a speaker or person met. In other cases reminiscing over an entire event may be the goal, such as remembering the last holiday or Christmas celebration. One abstraction higher there is the reflection on patterns of data extracted from successive events, i. e. behavior change: a person may have run longer and quicker over the last 6 months. Hence, the main benefits of lifelogging activities seem to be 1) Retrieval (supporting daily functioning for both work and leisure), 2) reminiscing (supporting mood regulation and group cohesion) and 3) reflection (supporting self-awareness for self-regulation and behaviour change). The group defined a **3x3 taxonomy** of lifelloging technologies, which combines these 3 main lifelogging activities with 3 societal contexts, namely: individuals, small group of individuals knowing each others and large group of individuals with loose ties (Table 1).

Table 1 A 3x3 taxonomy for lifelogging activities in their societal context.

	Retrieval	**Reminiscing**	**Reflection**
Individual	todo lists, photos, videos, business cards	photos, text messages	quantified self technologies
Small Group	private wikis	photos, artifacts	group forums
Large Group	Wikipedia, dictionaries	war memorabilia	smart cities reflecting on environmental changes and impact

4.3 Social Impact

This group discussed the social impact of lifelogging and memory augmentation on 3 levels: for 1) individuals, 2) groups or closed communities and 3) societies or countries. Key focus of the discussion were negative and positive aspects of massive data collection. In the following some considerations shall be outlined:

4.3.1 Risks

What has been described in several literary pieces is the question of mass surveillance and the resulting lack of privacy. Scenarios include authoritarian governments demanding all people's data. Hence, society would not be able to deny its actions. High instance control would be given the power to change history by modifying or forging the data and hence influence people (digital propaganda). On an individual's level, people could get psychotic if they are taken the right and possibility to forget. Others may develop tendencies to live in the past and get drawn into a behavior of neurotically checking their own history/memories. On a different notion, if lifelogging produces simply too much data and if we as society and individuals don't know what to do with it, it just collects dust.

4.3.2 Opportunities

If lifelogs and recordings were ubiquitous and designed to be preserved long-term, one obvious advantage is the facilitation of historical research. Researchers could go back in history and for example experience wars or economic highs in order to learn about conflict resolution or better understand (their own or foreign) cultures. Further, a comprehensive long-term dataset may create greater awareness of sustainability issues, CO_2 levels, environmental impact, social science or lifestyle and wellbeing. These datasets would create a level or transparency that could help us understand how other communities live, democracy improve and thus teach us important lessons about citizen science. Mining such dataset could reveal patterns which could be used to in many ways (e. g. increase societal security by predicting negative behavior and preventing it).

Participants

- C. Philip Beaman
University of Reading, GB
- Michel Beaudouin-Lafon
University of Paris South XI, FR
- Agon Bexheti
University of Lugano, CH
- Mark Billinghurst
University of Canterbury – Christchurch, NZ
- Andreas Bulling
MPI für Informatik – Saarbrücken, DE
- Ozan Cakmakci
Google Inc. – Mountain View, US
- Nigel Davies
Lancaster University, GB
- Simon Dennis
University of Newcastle, AU
- Tilman Dingler
Universität Stuttgart, DE
- Christos Efstratiou
University of Kent, GB
- Niforatos Evangelos
University of Lugano, CH
- Scott Greenwald
MIT – Cambridge, US
- Cathal Gurrin
Dublin City University, IE
- Vicki Hanson
Rochester Institute of Technology, US
- Niels Henze
Universität Stuttgart, DE
- James D. Hollan
University of California – San Diego, US
- Christoph Korinke
OFFIS – Oldenburg, DE
- Kai Kunze
Osaka Prefecture University, JP
- Wendy E. Mackay
University of Paris South XI, FR
- Mateusz Mikusz
Lancaster University, GB
- Daniela Petrelli
Sheffield Hallam University, GB
- Adam C. Rule
University of California – San Diego, US
- Corina Sas
Lancaster University, GB
- Albrecht Schmidt
Universität Stuttgart, DE
- Aurélien Tabard
University Claude Bernard – Lyon, FR
- Kristof Van Laerhoven
TU Darmstadt, DE
- Geoff Ward
University of Essex, GB
- Katrin Wolf
Universität Stuttgart, DE