DAGSTUHL
REPORTS

Volume 1, Issue 1, January 2011

ISSN 2192-5283

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

Publication date
June, 2011

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.1.1.i

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

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

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

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

www.dagstuhl.de/dagrep

Report from Dagstuhl Seminar 11011

Multi-Core Memory Models and Concurrency Theory

Edited by

Hans J. Boehm[1], Ursula Goltz[2], Holger Hermanns[3], and Peter Sewell[4]

1 **Hewlett Packard Labs - Palo Alto, US,** Hans.Boehm@hp.com
2 **TU Braunschweig, DE,** goltz@ips.cs.tu-bs.de
3 **Universität des Saarlandes, DE,** hermanns@cs.uni-sb.de
4 **University of Cambridge, GB,** peter.sewell@cl.cam.ac.uk

Abstract

This report documents the programme and the outcomes of Dagstuhl Seminar 11011 "Multi-Core Memory Models and Concurrency Theory".

Seminar 03.–01. January, 2011 – www.dagstuhl.de/11011
1998 ACM Subject Classification C.1.2 [Multiple Data Stream Architectures (Multiprocessors)]: Parallel processors; D.1.3 [Concurrent Programming]: Parallel programming; F.3.1 [Specifying and Verifying and Reasoning about Programs]
Keywords and phrases Relaxed Memory Models, Concurrency Theory, Multi-Core, Semantics, Parallel Programming, Cache Coherence
Digital Object Identifier 10.4230/DagRep.1.1.1
Edited in cooperation with Christian Eisentraut and Malte Lochau

1 Executive Summary

Hans J. Boehm (Hewlett Packard Labs - Palo Alto, US)
Ursula Goltz (TU Braunschweig, DE)
Holger Hermanns (Universität des Saarlandes, DE)
Peter Sewell (University of Cambridge, GB)
Christian Eisentraut (Universität des Saarlandes, DE)
Malte Lochau (TU Braunschweig, DE)

The current and future trend to multi-core and many-core computing systems suggests that within the next decade, concurrent multi-threaded programming will continue to replace sequential programming as the standard programming paradigm. However, concurrency and modern computer architecture do not go together easily:

- Current programming language memory models are still incomplete. Mainstream languages such as Java increasingly promise sequential consistency for data-race-free programs. However, how data races can be handled in a way that supports reasonable performance, security, and debugability, is currently completely unknown.
- Hardware specifications are so informal that it is very hard to know whether we have a correct implementation of the language specs (if we knew how to specify those fully). It is not clear that existing ISAs, which have a long history, are a good match for the language

semantics we are developing. There is an argument that this misalignment continues to encourage languages to support complex low-level constructs.
- The concurrent algorithms that are now being developed, and which are key to exploiting multiprocessors (via high-performance operating systems, hypervisor kernels, and concurrency libraries), are very subtle, so informal reasoning cannot give high confidence in their correctness.
- While there is a rich literature on concurrency theory, and extensive prior work on software verification for concurrency software (including process calculi, model-checking, temporal logics, rely-guarantee reasoning, and separation logic), almost all of it assumes a global notion of time, unrealistic though that is for these systems and algorithms.

Concurrency theory has a long tradition in investigating the foundational principles of concurrent computation, devoted to parallelism and causality of operations. It has created a wealth of models and notions readily applicable in this semantic challenge. Recent developments in the research communities of programming languages and concurrency theory, respectively, indeed show a growing trend towards cross-fertilization.

This seminar has fostered cross-fertilization of different expertises that will help to develop novel practical and, at the same time, theoretically sound paradigms for multi-core programming. It brought together concurrency theorists, memory model experts, computer systems architects, compiler experts, and formal verification researchers. The aim of the seminar was to adress in particular:

1. Survey of problem domain: state of the practice in multi-core-programming and state of the art in memory consistency models.
2. Application of concurrency theory approaches to reveal underlying concepts of parallelism, reordering, causality, and consistency.
3. Cross-fertilization across formal approaches to memory consistency models and semantics of multithreaded programming.
4. Attack points for formal analysis and computer aided programmer support.

Many of the questions that stood at the outset of this seminar have not been conclusively answered, thus yielding many potentials for further investigation. However, what makes this seminar uniquely successful is that it initiated a vivid exchange between a multitude of different scientific and industrial communities. During the seminar, it became clear that the current and future challenges of multi-core programming and memory models design in software and hardware can only be solved if the communities continue to exchange ideas and will learn from each other.

Schedule

Monday

Evening Session

19:00 Welcome
19:30-20:30 5 min intro talks

Tuesday

Morning session

9:00 Hans Boehm – *Memory Models for Threads in Mainstream Programming Languages*
10:40 Coffee Break
11:00 more 5 min intro talks
11:30 Scott Owens – *x86-TSO*

12:15 Lunch

14:00 rest of 5 min intro talks
15:00 Milo Martin – *InvisiFence: Performance-Transparent Memory Ordering in Conventional Multiprocessors*
15:40 Coffee Break
16:30 Doug Lea – *Some Weak Idioms*
17:15 Mark Hill – *Calvin: Deterministic or Not? Free Will to Choose*
17:30 Vijay Saraswat – *RAO Memory Model*

Wednesday

Morning session (chair: Rob van Glabbeek)

9:00 Gerard Boudol – *Why is True Concurrency Theory (Possibly) Relevant to the Study of Relaxed Memory Models*
9:45 Sarita Adve – *Rethinking Parallel Languages and Hardware for Disciplined Parallelism*
10.30 Coffee Break
11:15 Heike Wehrheim – *Verifying Linearisability of a Lazy Concurrent Set*
11:35 Maged Michael – *Memory Ordering Tradeoffs*
11:55 Alexey Gotsman – *Modular Verification of Preemptive OS Kernels*

12:15 Lunch

14:00 Hike

Afternoon session (chair: Hans Boehm)

15:30 Coffee Break
16:00 Susmit Sarkar – *Understanding POWER Multiprocessors*
16:45 Madan Musuvathi – *A Case for a SC-Preserving Compiler*
17:05 Luis Ceze – *A Case for Concurrency Exceptions*

18:00 Dinner

Evening session

19:30 Arvind – *Commit-Reconcile and Fences (CRF): A Memory Model for Compiler Writers and Architects*
20:15 Brandon Lucia – *Detecting, Avoiding and Understanding Errors in Concurrent Programs*

Thursday

Morning session (chair: Doug Lea)

9:00 Mark Batty – *Mathematizing C++ Concurrency*
9:45 James Riely – *Generative Operational Semantics for Relaxed Memory Models*
10:10 Andreas Lochbihler – *A Unified Machine-Checked Model for Multithreaded Java*
10:30 Coffee Break
11:15 Stephan Diestelhorst – *Extending the AMD64 Memory Model with an Advanced Synchronization Facility*
11:55 Erik Hagersten – *Selling Your Memory Model Ideas to HW Architects*

12:15 Lunch

Afternoon session

14:00 Jaroslav Sevcik – *Validity of Program Transformations (in the Java Memory Model)*
14:45 Michael Kuperstein – *Partial Coherence Abstractions*
15:05 Laura Effinger-Dean – *Simpler Reasoning About Variables in Multithreaded Programs*
15:30 Coffee Break
16:15 Paul McKenney – *Multi-Core Memory Models and Concurrency Theory: A View from the Linux Community*
16:40 Jade Alglave – *Fences in Weak Memory Models*
17:00 Sela Mador-Haim – *Generating Litmus Tests for Contrasting Memory Consistency Models*
17:20 Sebastian Burckhardt – *Concurrent Revisions: A Strong Alternative to SC*

18:00 Dinner

Evening session

19:30 Samuel Midkiff – *Compilers and Memory Models – Selected Topics*

Friday

Morning session

9:00 Lisa Higham – *Specifying Memory Consistency Models for Write-Buffer Multiprocessors and Proving Them Correct*
9:20 Gustavo Petri – *Speculation for Relaxed Memory: Using "True Concurrency"*
9:40 Sibylle Froeschle – *True-Concurrency: Foundational Concepts*
10:00 Richard Bornat and Mike Dodds – *Program Logics: We have come a Long Way, Baby*
10:20 Coffee Break
11:00 Open Problem session

12:15 Lunch

2 Table of Contents

3 Overview of Talks

3.1 Rethinking Parallel Languages and Hardware for Disciplined Parallelism

Sarita Adve (University of Illinois - Urbana, US)

Main reference DeNovo: Rethinking Hardware for Disciplined Parallelism, Byn Choi, Rakesh Komuravelli, Hyojin Sung, Robert Bocchino, Sarita V. Adve, and Vikram V. Adve, Second USENIX Workshop on Hot Topics in Parallelism (HotPar), 2010.
URL http://denovo.cs.illinois.edu/Pubs/10-hotpar-denovo.pdf

For parallelism to become tractable for mass programmers, shared memory programming languages and environments must evolve to enforce disciplined practices that ban "wild shared-memory behaviors;" e.g., unstructured parallelism, arbitrary data races, and ubiquitous non-determinism. This software evolution is also a rare opportunity for hardware designers to rethink hardware from the ground up to exploit opportunities exposed by such disciplined software models. Such a co-designed effort is more likely to achieve many-core scalability and large gains in power/performance than a software-oblivious hardware evolution. In this talk, we first briefly overview the Deterministic Parallel Java (DPJ) language that provides a disciplined shared-memory programming model. We then discuss DeNovo, a hardware architecture motivated by languages such as DPJ. We show how a disciplined parallel programming model can greatly simplify the cache coherence protocol and memory consistency model, while enabling a more efficient communication and cache architecture.

Our programming model ensures data-race-freedom, so DeNovo does not need to deal with protocol races and transient states, making it simpler and more extensible than software-oblivious protocols. Our programming model makes shared-memory side effects explicit, so each core keeps its cache coherent, eliminating the need for sharer-lists in a directory and associated invalidation traffic. To evaluate simplicity, we verify a base version of DeNovo with model checking and compare it with an otherwise equivalent state-of-the-art MESI protocol. The DeNovo version entails 25X fewer reachable states and takes 30X less time to verify. To evaluate extensibility, we add two sophisticated performance-enhancing optimizations to DeNovo: flexible bulk transfers and cache-to-cache direct data transfers. These add significant races and new protocol states with MESI, but not with DeNovo. With a simple extension, DeNovo's storage overhead breaks even with efficient MESI versions at about 30 core systems, but unlike MESI, it maintains scalability beyond that point. The net result is a system that seamlessly integrates message passing-like interactions within a shared memory programming model with improved design complexity, up to 67% reduction in memory stall time, and up to 70% reduction in network traffic.

This talk covers joint work with Rob Bocchino, Nicholas P. Carter, Byn Choi, Ching-Tsun Chou, Nima Honarmand, Rakesh Komuravelli, Robert Smolinski, Hyojin Sung, Vikram S. Adve, Tatiana Shpeisman, Marc Snir, and Adam Welc.

3.2 Fences in Weak Memory Models

Jade Alglave (University of Oxford, GB)

Joint work of Luc Maranget, Susmit Sarkar and Peter Sewell

We present a class of relaxed memory models, defined in Coq, parameterised by the chosen permitted local reorderings of reads and writes, and the visibility of inter- and intra-processor communications through memory (e.g. store atomicity relaxation). We prove results on the required behaviour and placement of memory fences to restore a given model (such as Sequential Consistency) from a weaker one. Based on this class of models we develop a tool, diy, that systematically and automatically generates and runs litmus tests to determine properties of processor implementations. We detail the results of our experiments on Power and the model we base on them. This work identified a rare implementation error in Power 5 memory barriers (for which IBM is providing a workaround); our results also suggest that Power 6 does not suffer from this problem.

3.3 Commit-Reconcile and Fences (CRF): A Memory Model for Compiler Writers and Architects

Arvind (MIT, Cambridge, MA, USA)

The rise of multicores is forcing us to readdress one of the fundamental problems of multi-threaded programming, namely memory models. Sequential Consistency (SC), in spite of its simple definition, has been found wanting because of exorbitant performance overheads in hardware implementations as well as opacity in reasoning at a high-level. Several weaker memory models have been proposed, all of which are motivated primarily by implementation concerns – in fact, the semantics of these memory models are given by the implementation rather than by any abstract rigorous method and understood only by experts. In view of this difficulty, many people think that we should just stick to Sequential Consistency as the interface between hardware and software. The use of various synchronization primitives and programming libraries has already made SC irrelevant for application programmers. There is no reason for compilers and architectures to adhere to SC if the higher-layers do not require it. We think that the Commit-Reconcile and Fences (CRF) memory model which was proposed earlier has the desirable properties for our purpose: its definition is given in terms of algebraic rules and it can directly model all the behaviors permitted by other relaxed memory models; CRF can be implemented in a variety of ways by the underlying machine; whether it be an in-order machine preserving sequential consistency or an aggressively out-of-order machine. So the high level programs can be ported across machines with different memory models. The talk will describe CRF and show how it could be used to give a meaningful semantics to a memory model for Java.

3.4 Mathematizing C++ Concurrency

Mark Batty (University of Cambridge, UK)

The next versions of C++ (C++0x) and of C (C1X) will have new concurrency features, aiming to support high-performance code with well-defined semantics. Unfortunately, as we near the end of the long standardization process, not all has been well. Unsurprisingly, the prose specification style of the draft standards is poorly suited to describe the complex design of the relaxed memory model, and in fact there have been significant problems.

I will discuss work on formalization of the memory model, what was broken, and some resulting improvements to the C++0x draft standard. In addition I will present a tool, Cppmem, for graphically exploring the semantics of small concurrent C++0x programs, and describe a proof of the correctness of a compilation strategy targeting x86-TSO.

3.5 Memory Models for Threads in Mainstream Programming Languages

Hans J. Boehm (HP Labs - Palo Alto, US)

Although multithreaded programming languages are common, there has been a surprising amount of confusion surrounding the basic meaning of shared variables. There is finally a growing consensus, both that programming languages should by default guarantee an interleaving-based semantics for data-race-free programs, and on what that should mean. We discuss the motivation for, and both implementation and user consequences of, this guarantee.

Unfortunately, it is also increasingly clear that such a guarantee, though useful, is insufficient for languages intended to support sand-boxed execution of untrusted code. The existing solution in Java is only partially satisfactory. The solution to this problem is far less clear.

3.6 Program Logics: We have come a Long Way, Baby

Richard Bornat (Middlesex University, GB) and Mike Dodds (Cambridge University, GB)

My brief, given at the last minute, was to persuade the audience that program proof has come a long way since Hoare logic was promoted in the 1970s. We were (and are) woefully unprepared to give a fair survey of world work in the area: we are simple program provers. The talk ended for an apology for that deficiency, and this abstract must start with one.

Hoare logic was not a solution to the 'Software Engineering Problem'. Loops were (necessarily) hard; there was no effective treatment of procedures; arrays were difficult to

deal with, and there was no treatment of pointers. Nevertheless it was the basis of successful program verification work in the 80s and 90s, notably in B, Spark, JML and others.

Reynolds and O'Hearn, in 2000, built on work of Burstall's (separated lists can be dealt with separately) to make separation logic, an extension of Hoare logic which deals with access and assignment to the heap, with *local reasoning*. An effective treatment of pointers led, quite quickly, to a treatment of concurrency (O'Hearn's CSL).

Separation logic has had an impact on program analysis and may (because of local reasoning) be able to make an impact on the problem, much discussed in the workshop, of detecting data races in Java programs. At present the Abducter tool (DiStefano) can analyse the whole Linux kernel in 20 minutes, finding space leaks in list-manipulating procedures; Windows device drivers can be handled in seconds, finding safety bugs (Yang, Cook).

In concurrency there are separation logic hand proofs of several intricate algorithms, and work on automatic proof of linearisability.

In summary, there is much useful work which can be exploited. We also mentioned an important negative point: ownership types don't work. Ownership, as various proofs of simple algorithms demonstrates, is dynamic not static.

3.7 Why is True Concurrency Theory (Possibly) Relevant to the Study of Relaxed Memory Models

Gerard Boudol (INRIA Sophia Antipolis, FR)

I briefly introduce "true concurrency" theory, and in particular the formalism of labelled transition systems with independence which, I think, are appropriate to provide abstract formalization of relaxed memory models.

3.8 Concurrent Revisions: A Strong Alternative to SC

Sebastian Burckhardt (Microsoft Research - Redmond, US)

Main reference Sebastian Burckhardt, Alexandro Baldassin, Daan Leijen, *"Concurrent programming with revisions and isolation types"*, pp. 691–707, OOPSLA'10, ACM, 2010.
URL http://dx.doi.org/10.1145/1869459.1869515

Although considered a strong memory model, sequential consistency (SC) is in fact still weaker than desirable in the sense that the semantics is based on nondeterministic interleavings. In contrast, Concurrent Revisions, a programming model for shared-memory concurrency/parallelism, is deterministic: the semantics of shared memory accesses are based on deterministic replication & conflict resolution rather than nondeterministic arbitration as in SC.

3.9 A Case for Concurrency Exceptions

Luis Ceze (University of Washington, US)

We argue in this talk that concurrency errors should be treated as exceptions, i.e., have fail-stop behavior and precise semantics. We propose an exception model based on conflict of synchronization-free regions, which precisely detects a broad class of data-races. We show that our exceptions provide enough guarantees to simplify high-level programming language semantics and debugging, but are significantly cheaper to enforce than traditional data-race detection. We also propose a new exception that enforces disciplined shared-memory communication in a code-centric manner. To make the performance cost of enforcement negligible, we propose architecture support for accurately detecting and precisely delivering concurrency exceptions.

3.10 Extending the AMD64 Memory Model with an Advanced Synchronization Facility

Stephan Diestelhorst (AMD - Dornbach, DE)

Joint work of Diestelhorst, Stephan; Christie, David; Hohmuth, Michael; Pohlack, Martin;
Main reference Jae-Woong Chung, Luke Yen, Stephan Diestelhorst, Martin Pohlack, Michael Hohmuth, David Christie, Dan Grossman: ASF: AMD64 Extension for Lock-Free Data Structures and Transactional Memory. MICRO 2010: 39-50
URL http://dx.doi.org/10.1109/MICRO.2010.40

Current general-purpose, high-performance microprocessor cores aggressively reorder instructions to increase throughput. With the wide availability of shared memory multi-core systems, architects have to pay attention to the architecturally visible memory semantics provided by these systems.

In our talk, we will look at how current AMD microprocessors provide strong memory semantics while still allowing aggressive reordering on the microarchitectural level. We will also highlight the functionality of the HyperTransport (tm) interconnect that connects processors and provides cache coherence and contributes significantly to the overall memory semantics.

On top of this foundation, we introduce AMD's Advanced Synchronization Facility (ASF), an experimental AMD64 ISA extension. ASF provides applications with speculative regions, akin to transactions in transactional memory, which allow atomic read-modify-write constructs on multiple independent memory locations. In the course of our presentation of ASF, we highlight various design decisions and illustrate why and how extensions to existing complex microarchitectures need to be simple.

References

1 Jaewoong Chung, Luke Yen, Stephan Diestelhorst, Martin Pohlack, Michael Hohmuth, David Christie, Dan Grossman. *ASF: AMD64 Extension for Lock- free Data Structures and Transactional Memory*. Proceedings of the 43rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO-43), 2010

2 Felber, P.; Riviere, E.; Moreira, W.M.; Harmanci, D.; Marlier, P.; Diestelhorst, S.; Hohmuth, M.; Pohlack, M.; Cristal, A.; Hur, I.; Unsal, O.S.; Stenström, P.; Dragojevic, A.;

Guerraoui, R.; Kapalka, M.; Gramoli, V.; Drepper, U.; Tomic;, S.; Afek, Y.; Korland, G.; Shavit, N.; Fetzer, C.; Nowack, M.; Riegel, T. *The VELOX Transactional Memory Stack*. IEEE Micro Special Issue on European Multicore Processing Projects, 2010

3 Stephan Diestelhorst, Michael Hohmuth, Martin Pohlack. *Sane Semantics of Best-effort Hardware Transactional Memory*. 2nd Workshop on the Theory of Transactional Memory (WTTM'10), 2010

4 Dave Christie, Jae-Woong Chung, Stephan Diestelhorst, Michael Hohmuth, Martin Pohlack, Christof Fetzer, Martin Nowack, Torvald Riegel, Pascal Felber, Patrick Marlier, Etienne Riviere *Evaluation of AMD's Advanced Synchronization Facility within a Complete Transactional Memory Stack*. EuroSys, 2010

5 Jaewoong Chung, David Christie, Martin Pohlack, Stephan Diestelhorst, Michael Hohmuth, Luke Yen. *Compilation of Thoughts about AMD Advanced Synchronization Facility and First-Generation Hardware Transactional Memory Support*. 5th ACM SIGPLAN Workshop on Transactional Computing (TRANSACT'10), 2010

6 Stephan Diestelhorst, Martin Pohlack, Michael Hohmuth, Dave Christie, Jae-Woong Chung, Luke Yen. *Implementing AMD's Advanced Synchronization Facility in an Out-of-order x86 Core*. 5th ACM SIGPLAN Workshop on Transactional Computing (TRANSACT'10), 2010

7 Conway, P.; Hughes, B. *The AMD Opteron Northbridge Architecture*. Micro, IEEE , vol.27, no.2, pp.10-21, March-April 2007

8 Conway, P.; Kalyanasundharam, N.; Donley, G.; Lepak, K.; Hughes, B. *Cache Hierarchy and Memory Subsystem of the AMD Opteron Processor*. Micro, IEEE , vol.30, no.2, pp.16-29, March-April 2010

9 *AMD "Advanced Synchronization Facility" Proposal*. http://developer.amd.com/cpu/asf/Pages/default.aspx

3.11 Simpler Reasoning About Variables in Multithreaded Programs

Laura Effinger-Dean (University of Washington - Seattle, US)

Joint work of Effinger-Dean, Laura; Boehm, Hans-J.; Chakrabarti, Dhruva; Joisha, Pramod

We present a simple characterization of code regions for which we can prove that a given shared variable cannot be modified by other threads in the system.

Following the new C++0x standard, we assume data-race-freedom. Our "interference-free regions" expand on simple synchronization-free regions to include limited patterns of synchronization operations. In particular, our observations demonstrate that it is legal to eliminate a redundant variable access across a lock or unlock operation. By coalescing the interference-free regions for multiple accesses to the same variable, we can extend this characterization to complex control flow.

3.12 True-Concurrency: Foundational Concepts

Sibylle Froeschle (Universität Oldenburg, DE)

In this talk I will give a short introduction to models and foundational concepts in concurrency theory with an emphasis on true-concurrency.

In particular, I will mention the computational dichotomy of true-concurrency: truly-concurrent versions of concepts such as behavioural equivalences can be computationally hard but for classes with good structural properties they are often efficiently decidable.

I will speculate on some connections to the challenges of multi-core memory models.

3.13 Modular Verification of Preemptive OS Kernels

Alexey Gotsman (IMDEA Software - Madrid, ES)

Most major OS kernels today run on multiprocessor systems and are preemptive: it is possible for a process running in the kernel mode to get descheduled. Existing modular techniques for verifying concurrent code are not directly applicable in this setting: they rely on scheduling being implemented correctly, and in a preemptive kernel, the correctness of the scheduler is interdependent with the correctness of the code it schedules. This interdependency is even stronger in mainstream kernels, such as Linux, FreeBSD or XNU, where the scheduler and processes interact in complex ways.

We propose the first logic that is able to decompose the verification of preemptive multiprocessor kernel code into verifying the scheduler and the rest of the kernel separately, even in the presence of complex interdependencies between the two components. This is achieved by establishing a novel form of refinement between an operational semantics of the real machine and an axiomatic semantics of OS processes, where the latter assumes an abstract machine with each process executing on a separate virtual CPU. The refinement is local in the sense that the logic focuses only on the relevant state of the kernel while verifying the scheduler. Our logic soundly inherits proof rules from concurrent separation logic to verify OS processes thread-modularly.We illustrate its power by verifying an example scheduler, modelled on the one from Linux 2.6.11.

3.14 Selling Your Memory Model Ideas to HW Architects

Erik Hagersten (University of Uppsala, SE)

First, I will discuss how architect tells good ideas from bad ideas by weighing tradeoffs between complexity (implying schedule slip), performance and power consumptions. Of the 1000s of the published architecture papers promising 10% performance improvement each year, only a fraction can possible have impact since the performance improvement per year

due to architectural improvement is roughly 20% and not 10 000%. This is often due to the overwhelming complexity of the proposals.

Second, I will ask the formal method's community for favors. 1) Do not define how to implement the memory models: I will show examples of correct implementations for coherence and sequential consistency that violate such definitions. 2) When verifying a protocol, consider all transactions, not just the LD/ST/Atomics: They tend to only represent 10-20% of all transactions in a protocol. 3) Do consider all interacting protocols in the system, not just one at a time. 4) Finally, I urge protocol verification to not only verify coherence and memory models. Just as important (and often harder) properties to verify include livelock and deadlock properties.

3.15 Specifying Memory Consistency Models for Write-Buffer Multiprocessors and Proving Them Correct

Lisa Higham (University of Calgary, CA)

Joint work of Higham, Lisa; Kawash, Jalal; Jackson, LillAnne
Main reference Lisa Higham, LillAnne Jackson, Jalal Kawash: Specifying memory consistency of write buffer multiprocessors. pp. 2–42, ACM Trans. Comput. Syst. 25(1), 2007.
URL http://doi.acm.org/10.1145/1189736.1189737

Multiprocessor architectures employ various hardware components such as multiple busses, write-buffers, and replicated memory to enhance the efficiency.

As a consequence, the computations that can arise satisfy guarantees substantially weaker than sequential consistency. Given a specific multi-processor architecture, our goals are: 1) to develop a simple, precise predicate that specifies exactly what computations can arise from a multiprocessing program that executes on that architecture; and 2) prove that the predicate is correct.

This talk illustrates our techniques for achieving this goal through a case study of shared memory architectures that use write buffers.

3.16 Calvin: Deterministic or Not? Free Will to Choose

Mark D. Hill (University of Wisconsin - Madison, US)

Calvin implements optional determinism with Total Store Order (TSO) compatibility for about a 20% performance loss relative to a conventional system. This talk is a short adaptation for non-architects of an upcoming HPCA 2011 paper

3.17 Partial Coherence Abstractions

Michael Kuperstein (Technion - Haifa, IL)

Joint work of Kuperstein, Michael; Vechev, Martin; Yahav, Eran;

We present an approach for automatic verification of concurrent programs running under relaxed memory models. Verification under relaxed memory models is a hard problem. Given a finite state program and a safety specification, verifying that the program satisfies the specification under a sufficiently relaxed memory model is undecidable. For somewhat stronger memory models, the problem is decidable but has non-primitive recursive complexity. In this paper, we focus on models that have store-buffer based semantics, e.g. SPARC TSO and PSO.

We use abstract interpretation to provide an effective verification procedure for programs running under this type of models. Our main contribution is a family of novel partial-coherence abstractions, specialized for relaxed memory models, which partially preserve information required for memory coherence and consistency.

3.18 Some Weak Idioms

Doug Lea (SUNY - Oswego, US)

The design and implementation of core libraries and runtime systems entail several idiomatic non-sequentially-consistent constructions. This talk illustrates some such techniques encountered in publishing and transferring objects and messages.

3.19 A Unified Machine-Checked Model for Multithreaded Java

Andreas Lochbihler (KIT - Karlsruhe Institute of Technology, DE)

I present a machine-checked formalisation of the Java memory model (JMM) and connect it to an operational semantics for source and bytecode. It extends previous formalisations by dynamic memory allocations, thread spawns and joins, infinite executions, the wait-notify mechanism and thread interruption. I proved that the model provides the Java data race freedom (DRF) guarantee.

I instantiated the JMM with JinjaThreads, a large subset of Java (bytecode), thereby providing the missing link between operational semantics on statement and instruction level and the JMM. To discharge the assumptions of the DRF proof, I constructed sequentially consistent executions of source and bytecode.

3.20 Detecting, Avoiding and Understanding Errors in Concurrent Programs

Brandon M. Lucia (University of Washington - Seattle, US)

Joint work of Lucia, Brandon M.; Ceze, Luis;
Main reference Brandon, L.; Ceze, L.; "Finding concurrency bugs with context-aware communication graphs." Proc. of 42nd Int'l Symp. on Microarchitecture, pp. 553–563. Dec. 2009.
URL http://dx.doi.org/10.1145/1669112.1669181

Programmers are the weak link in the chain of concurrent software development.

People make mistakes that lead to subtle concurrency errors that are difficult to find and fix. These errors degrade the reliability of software, and can lead to costly system failures. The movement of accessible concurrent programming to the mainstream is dependent on a solution to the problems posed by concurrency errors.

In this talk I will describe two different approaches to this problem.

First, I will discuss Bugaboo, a technique for automatically isolating patterns of inter-thread communication that lead to buggy program behavior. The key contribution of this work is the introduction of context-aware communication graphs, a new graphical data-flow abstraction for program execution. These graphs encode communication between instructions, and an abstract view of the order of communication events. By collecting context-aware communication graphs from many different execution, we can use statistical reasoning to identify communication events most likely related to a failure. Our reasoning is simple: communication that occurs often in buggy executions, and rarely or never in correct executions is likely related to the failure. We develop supervised and unsupervised approaches to error identification. We also describe hardware extensions that enable graph collection with negligible performance overhead.

Second, I will discuss ColorSafe and Atom-Aid, two similar techniques for avoiding atomicity violation bugs. Atomicity errors are the result of incorrect program synchronization: a programmer should have prevented a particular buggy interleaving of operations from different threads, but failed to do so. Prior work has shown that the manifestation of atomicity violations is characterized by an unserializable interleaving of memory accesses. We leverage serializability analysis to find atomicity bugs. Each thread maintains a history of locally performed memory accesses and a history of memory accesses performed by other threads. Periodically, at the end of fixed length execution epochs, threads analyze the serializability of recent memory access interleavings to find likely atomicity bugs. Upon finding a likely violation, our system uses dynamic atomic regions (viz. transactions) to prevent interleaving, thereby avoiding the buggy program behavior. In addition we describe a technique for generalizing serializability analysis to simultaneously consider multiple memory locations instead of single locations only. We develop simple hardware support to make these techniques efficient.

3.21 Generating Litmus Tests for Contrasting Memory Consistency Models

Sela Mador-Haim (University of Pennsylvania, US)

Well-defined memory consistency models are necessary for writing correct parallel software. Developing and understanding formal specifications of hardware memory models is a challenge due to the subtle differences in allowed reorderings and different specification styles. To facilitate exploration of memory model specifications, we have developed a technique for systematically comparing hardware memory models specified using both operational and axiomatic styles. Given two specifications, our approach generates all possible multi-threaded programs up to a specified bound, and for each such program, checks if one of the models can lead to an observable behavior not possible in the other model. When the models differs, the tool finds a minimal "litmus test" program that demonstrates the difference. A number of optimizations reduce the number of programs that need to be examined. Our prototype implementation has successfully compared both axiomatic and operational specifications of six different hardware memory models. We describe two case studies: (1) development of a non-store atomic variant of an existing memory model, which illustrates the use of the tool while developing a new memory model, and (2) identification of a subtle specification mistake in a recently published axiomatic specification of TSO.

3.22 InvisiFence: Performance-Transparent Memory Ordering in Conventional Multiprocessors

Milo M. K. Martin (University of Pennsylvania, US)

Joint work of Blundell, Colin; Martin, Milo M.K.; Wenisch, Thomas F.;
Main reference Blundell, C.; Martin, M.M.K.; Wenisch, T.F.; InvisiFence: performance-transparent memory ordering in conventional multiprocessors, Proc. of 36th Int'l Symp. on Comp. Architecture (ISCA'09).
URL http://dx.doi.org/10.1145/1555754.1555785

A multiprocessor's memory consistency model imposes ordering constraints among loads, stores, atomic operations, and memory fences. Even for consistency models that relax ordering among loads and stores, ordering constraints still induce significant performance penalties due to atomic operations and memory ordering fences. Several prior proposals reduce the performance penalty of strongly ordered models using post-retirement speculation, but these designs either (1) maintain speculative state at a per-store granularity, causing storage requirements to grow proportionally to speculation depth, or (2) employ distributed global commit arbitration using unconventional chunk-based invalidation mechanisms.

In this paper we propose INVISIFENCE, an approach for implementing memory ordering based on post-retirement speculation that avoids these concerns. INVISIFENCE leverages minimalistic mechanisms for post-retirement speculation proposed in other contexts to (1) track speculative state efficiently at block-granularity with dedicated storage requirements independent of speculation depth, (2) provide fast commit by avoiding explicit commit arbitration, and (3) operate under a conventional invalidation-based cache coherence protocol. INVISIFENCE supports both modes of operation found in prior work: speculating only when

necessary to minimize the risk of rollback-inducing violations or speculating continuously to decouple consistency enforcement from the proces- sor core. Overall, INVISIFENCE requires approximately one kilobyte of additional state to transform a conventional multiprocessor into one that provides performance-transparent memory ordering, fences, and atomic operations.

3.23 Multi-Core Memory Models and Concurrency Theory: A View from the Linux Community

Paul McKenney (IBM - Beaverton, US)

Main reference McKenney, Paul E., "Is Parallel Programming Hard, And, If So, What Can You Do About It?", kernel.org, Corvallis, OR, USA 2011.
URL http://kernel.org/pub/linux/kernel/people/paulmck/perfbook/perfbook.html

This talk presents multi-core memory models and concurrency theory from the viewpoint of a practitioner with 20 years experience producing production-quality shared-memory parallel code and with 10 years with the Linux kernel community. This experience has led to the following conclusions: (1) although abstraction is valuable, intimate knowledge of underlying software and hardware layers is equally valuable, (2) although general techniques are valuable, specialized techniques resulting in excellent performance, scalability, real-time response, and energy efficiency are equally valuable, (3) although concurrent-software validation techniques are valuable, techniques drawn from hardware validation may well prove more valuable, and (4) different levels of abstraction required different programming paradigms.

That said, within the Linux kernel community, organizational mechanisms are at least as important as specific techniques and tools. These organizational mechanisms include the maintainership hierarchy with its focus on quality assurance, an informal apprenticeship/-mentoring model, a strong tradition of design and code review, and aggressive pursuit of modularity and simplicity.

These mechanisms have the important side effect of focusing attention on techniques that are known to work well in a given situation, which allows ordinary practitioners to successfully produce production-quality parallel designs and code.

3.24 Memory Ordering Tradeoffs

Maged Michael (IBM TJ Watson Research Center - Yorktown Heights, US)

Joint work of Attiya, Hagit; Guerraoui, Rachid; Hendler, Danny; Kuznetsov, Petr; Michael, Maged; Vechev, Martin;
Main reference Hagit Attiya, Rachid Guerraoui, Danny Hendler, Petr Kuznetsov, Maged M. Michael, Martin T. Vechev: Laws of order: expensive synchronization in concurrent algorithms cannot be eliminated. Proceedings of the 38th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2011, 487-498
URL http://doi.acm.org/10.1145/1926385.1926442

It is often the case that in designing concurrent algorithms that it appears that avoiding both atomic operations and store-load ordering is difficult. It is shown that it is impossible such patterns in algorithms for methods that are strongly non-commutative. The SNC condition can be avoided by specification relaxations, such as limiting concurrency , limiting the API,

relaxing determinism, or relaxing the requirement of linearizability. This result opens the door for trade-offs between synchronization overhead and specification.

3.25 Compilers and Memory Models – Selected Topics

Samuel P. Midkiff (Purdue University, US)

This talk discusses three important topics related to memory models. First, we argue that static compiler analyses are insufficient, and will always be insufficient, to certify programs as race-free. Thus, memory models for general purpose languages that depend on programs being race-free for correctness will almost certainly require some runtime checking. Second, we present some results showing the effectiveness of a tool that helps programmers to identify transaction regions, and thus help the programmer in the creation of race free programs. We also argue that analyses such as this one, lock assignment and other transformations handling programs with pointers benefit from the use of semantic information about standard library code. Semantic information we used included information about destructive data structure operations, commutativity of operations on the data structure, and whether the data structure is thread-safe. These properties follow directly from the specification of the code, and thus are trivially easy to specify for routines that implement the specification. Finally, we discuss the benefit of hardware that provides isolation for code regions in threads. We show that allowing the compiler to set and know regions that are excited in isolations, as is done in the BulkSC architecture, allows a sequentially consistent memory model to be implemented with an increase in performance.

3.26 A Case for a SC-Preserving Compiler

Madan Musuvathi (Microsoft Corp. - Redmond, US)

In this talk, I will argue against relaxing memory models in the compiler. I will demonstrate that a SC-preserving compiler, one which guarantees that every SC behavior of the binary is a SC behavior of the source program, is feasible with acceptable performance overhead. Time permitting, I will also demonstrate how static and dynamic analyses can further reduce this overhead.

3.27 x86-TSO

Scott Owens (University of Cambridge, GB)

Exploiting the multiprocessors that have recently become ubiquitous requires high-performance and reliable concurrent systems code, for concurrent data structures, operating

system kernels, synchronisation libraries, compilers, and so on. However, concurrent programming, which is always challenging, is made much more so by two problems. First, real multiprocessors typically do not provide the sequentially consistent memory that is assumed by most work on semantics and verification. Instead, they have relaxed memory models, varying in subtle ways between processor families, in which different hardware threads may have only loosely consistent views of a shared memory. Second, the public vendor architectures, supposedly specifying what programmers can rely on, are often in ambiguous informal prose (a particularly poor medium for loose specifications), leading to widespread confusion.

In this talk we focus on x86 processors. We review several recent Intel and AMD specifications, showing that all contain serious ambiguities, some are arguably too weak to program above, and some are simply unsound with respect to actual hardware. We present a new x86-TSO programmer's model that, to the best of our knowledge, suffers from none of these problems. It is mathematically precise (rigorously defined in HOL4) but can be presented as an intuitive abstract machine which should be widely accessible to working programmers. We illustrate how this can be used to reason about the correctness of a Linux spinlock implementation and describe a general theory of data-race-freedom for x86-TSO. This should put x86 multiprocessor system building on a more solid foundation; it should also provide a basis for future work on verification of such systems.

3.28 Speculation for Relaxed Memory: Using "True Concurrency"

Gustavo Petri (INRIA Sophia Antipolis, FR)

We present a semantic framework to describe speculative computations of parallel programs. An important ingredient of our framework is the definition of *valid* speculation, for which we use standard true concurrency techniques. Interestingly, the effects of speculations are similar to those observed in relaxed memory models. We briefly show how to instantiate the TSO, PSO and RMO memory models of Sparc using this framework.

3.29 Generative Operational Semantics for Relaxed Memory Models

James Riely (DePaul University - Chicago, US)

Joint work of Jagadeesan, Radha; Pitcher, Corin; Riely, James;
Main reference Jagadeesan, R.; Pitcher, C.; Riely, J.; "Generative Operational Semantics for Relaxed Memory Models", ESOP 2010, pp. 307–326.
URL http://dx.doi.org/10.1007/978-3-642-11957-6_17

The specification of the Java Memory Model (JMM) is phrased in terms of acceptors of execution sequences rather than the standard generative view of operational semantics. This creates a mismatch with language-based techniques, such as simulation arguments and proofs of type safety.

We describe a semantics for the JMM using standard programming language techniques that captures its full expressivity. For data-race-free programs, our model coincides with the

JMM. For lockless programs, our model is more expressive than the JMM. The stratification properties required to avoid causality cycles are derived, rather than mandated in the style of the JMM.

The JMM is arguably non-canonical in its treatment of the interaction of data races and locks as it fails to validate roach-motel reorderings and various peephole optimizations. Our model differs from the JMM in these cases. We develop a theory of simulation and use it to validate the legality of the above optimizations in any program context.

3.30 RAO Memory Model

Vijay A. Saraswat (IBM TJ Watson Research Center - Hawthorne, US)

Talk describes essence of the RAO Model presented in PPoPP 2007.

3.31 Understanding POWER Multiprocessors

Susmit Sarkar (University of Cambridge, GB)

IBM POWER multiprocessors have a very relaxed memory model (ARM is similar), including instances in which programmers can observe non-atomicity of stores, register shadowing, and speculative executions past branches. I will describe joint (ongoing) work with Peter Sewell, Jade Alglave, Luc Maranget, and Derek Williams on our extensive testing of Power G5, 5, 6, and 7, producing some perhaps-surprising results, and on an abstract-machine semantics that explains these results.

3.32 Validity of Program Transformations (in the Java Memory Model)

Jaroslav Sevcik (University of Cambridge, GB)

In this talk, I will give an overview of validity of compiler transformations in the Java Memory Model, Sequential Consistency and the DRF guarantee. I will also explain several interesting examples of transformations that are problematic for the Java Memory Model. Finally, I will show that there are program transformations that are valid under SC but not under TSO or DRF.

3.33 Verifying Linearisability of a Lazy Concurrent Set

Heike Wehrheim (Universität Paderborn, DE)

Joint work of Derrick, John; Schellhorn, Gerhard; Wehrheim, Heike;

Linearisability is the key correctness criterion for concurrent implementations of data structures shared by multiple processes.

In the seminar we presented a proof of linearisability of the *lazy* implementation of a set due to Heller et al. The lazy set presents one of the most challenging issues in verifying linearisability: a linearisation point of an operation (here, *contains*) set by a process other than the one executing it. For this we have developed a proof strategy based on refinement which uses thread local simulation conditions (the proof obligations talk about at most two processes at a time) and the technique of potential linearisation points. The former allows us to locally prove linearisability for arbitrary numbers of processes, the latter permits disposing with backward reasoning. The operation *contains* may get several potential linearisation points (including some set by e.g. the *remove* operation of another process), only the last one is a valid one. As *contains* is not modifying the abstract data structure (the set), this type of reasoning is sound for proving linearisability.

All proofs (including the one showing soundness of our proof obligations) have been mechanically carried out using the interactive prover KIV.

References

1 J. Derrick, G. Schellhorn, and H. Wehrheim. Mechanically verified proof obligations for linearizability. ACM Transactions on Programming Languages and Systems (TOPLAS), vol. 33, no. 1 (2011).

2 J. Derrick, G. Schellhorn, and H. Wehrheim. Mechanizing a correctness proof for a lock-free concurrent stack. In *FMOODS 2008*, volume 5051 of *LNCS*, pages 78–95. Springer, 2008.

4 State of the Art and Open Problems

The field of relaxed memory models, unusually, cuts across many areas of Computer Science: Computer Architecture, Programming Languages, Compilation, Concurrent Algorithms, Concurrency Theory and Semantics, Hardware Verification, and Software and Algorithm Verification. Real progress in the development and understanding of memory models in a multi-core/shared-memory world can only be achieved by taking a holistic view. In this concluding note we summarise some discussion at the workshop of the state of the art and open problems in the area. This is at best an outline and idiosyncratic summary, far from complete.

Models for Mainstream Hardware

A basic question is that of establishing usable and rigorous models for today's mainstream multiprocessors, which include ARM, Itanium, Power, Sparc, x86, and IBM zSeries. The state of the art varies for each of these: Sparc TSO has long had a precise model; for x86, Power and ARM we heard in the workshop about recent work that covers common-case programming, though further work is needed to cover the full architecture in each case;

and for Itanium the vendor specification is reasonably precise, and there has been various academic work, though we are not aware of extensive empirical testing.

One might also consider more 'exotic' hardware: Tilera, the Intel SCC, Cray machines (influencing OpenMP), and various GPUs.

Some other hardware architectures are of less current interest: Alpha is no longer made, though a few machines remain and it has influenced the Linux kernel memory barriers; Sparc RMO and PSO are not used for Sparc systems; PA-RISC is no longer made (but said to be SC).

Simple Hardware Deltas

A line of discussion throughout the workshop concerned 'simple' proposals for hardware developments, i.e., those which a hardware vendor might reasonably incorporate into near-term revisions of their architecture. For example, there was discussion of: adding explicit synchronisation variables to an ISA; support for race detection, using the above (perhaps not so 'simple'); faster store fences; and lightweight (limited) transaction memory support.

Radical Hardware Deltas

More long-term questions include: determining how much one gains from a relaxed hardware model (difficult to establish as a vast amount of hardware and software has been co-designed in the current ecosystems); considering different APIs for hardware coherence mechanisms, exposing more of the hardware (between coherent shared memory and message passing); and considering how scalable coherent memory can be made.

Hardware Verification

Multiprocessor relaxed memory model behaviour is an emergent property of an entire multiprocessor design, involving both core behaviour (especially speculation) and memory communication fabrics. This leads us to ask how one can *prove* that a realistic microarchitecture does implement an architectural memory model, whether better memory model testing can be introduced into the hardware design process, and, more generally, whether it would be useful to work routinely with the formal 'abstract microarchitecture' models that one would use for such verification.

Models for Mainstream Programming Languages

Despite very extensive work on programming language memory models, it is arguable that none are yet fully satisfactory.

The semantically simplest option is SC, but providing SC requires one to limit compiler optimisation and to make use of hardware synchronisation primitives (barriers etc.). Folklore suggests that, at least on high-performance compilers, the cost of those is substantial, but there were interesting preliminary results to the contrary presented at the workshop.

A more efficiently implementable alternative is "DRF and catch-fire": guaranteeing SC behaviour for programs that are race-free (in some sense) in an SC semantics, but providing no guarantees at all for other programs. This permits a wide range of optimisations (the verification of those was discussed at the meeting), and arguably matches programmer intuition for most code except low-level concurrent algorithms, but raises development problems: pragmatically, given that almost all software is buggy, how can race-freedom be ensured?

That brings us to "DRF and abort" models, which are similar except that one has a guarantee that any such race (or, in some cases, any race that matters) will force termination or raise an exception; there was discussion of how this can be implemented.

Statically, one might instead have type systems that exclude races altogether, or even that establish determinacy, but that are flexible enough for a wide range of code, and in general one might establish better concurrent programming models, perhaps with explicit ownership transfer and explicit treatment of data and computation locality.

DRF and catch-fire is at the heart of the C++0x/C1x memory model, a formal development of which was presented at the meeting; an open technical question here is how one should allow close-to-the-hardware non-SC synchronization, and yet still forbid causal cycles and thin-air reads.

The Java Memory Model was also discussed. Java aims to provide SC behaviour for race-free code, but additionally (and in contrast to C++/C) to guarantee some basic security properties for arbitrary code. The current JMM is flawed in the sense that it prohibits certain compiler optimisations that implementations do in fact perform; how best to admit those optimisations while still ensuring the security properties is an open question.

Semantics, Verification, and Reasoning

Basic Semantic Questions

For any well-motivated relaxed memory model, we should ask some basic questions: how to characterise observational congruences; whether we can make effective use of noninterleaving semantics; what the expressiveness of various sublanguages is; and the decidability and complexity of various problems.

Concurrent Algorithm Verification

For algorithm verification, several directions are being pursued: interactive proof (directly above a semantics), program logics, automatic methods, and linearisability proofs. There is a great need for effective compositional methods in this domain, especially to show that libraries using weakly ordered atomic operations can encapsulate that weakness, giving SC behaviour as far as any client can observe.

Compiler Verification

We heard in the meeting of first steps in verifying complete compilers, e.g., from a C-like language with TSO semantics to x86 with its TSO-based semantics, and of verifying compilation of executions from C++0x to x86. There is much more to be done, on verifying optimisations w.r.t. particular models, targetting the weaker and more subtle hardware models, and extending these results to full-scale languages.

Dynamic Verification

There are many properties that are difficult to prove but useful to check, especially if this can be done with reasonable performance: better dynamic race detection, atomicity, and more general concurrency error detection.

Radical models

Finally, when one looks beyond the largest conventional coherent shared-memory machine that one can reasonably build, one can simultaneously ask about the (co)design of hardware, programming model and language, compilation techniques, algorithms, and semantics.

Participants

- Sarita Adve
University of Illinois - Urbana, US
- Jade Alglave
University of Oxford, GB
- Arvind
MIT - Cambridge, US
- Mark Batty
University of Cambridge, GB
- Hans J. Boehm
HP Labs - Palo Alto, US
- Peter Boehm
University of Oxford, GB
- Richard Bornat
Middlesex University, GB
- Ahmed Bouajjani
LIAFA - Université Paris VII, FR
- Gérard Boudol
INRIA Sophia Antipolis, FR
- Sebastian Burckhardt
Microsoft Research - Redmond, US
- Luis Ceze
University of Washington, US
- Stephan Diestelhorst
AMD - Dornbach, DE
- Mike Dodds
University of Cambridge, GB
- Laura Effinger-Dean
University of Washington - Seattle, US
- Christian Eisentraut
Universität des Saarlandes, DE
- Sibylle Fröschle
Universität Oldenburg, DE
- Ursula Goltz
TU Braunschweig, DE
- Alexey Gotsman
IMDEA Software - Madrid, ES
- Richard Grisenthwaite
ARM Ltd. - Cambridge, GB
- Erik Hagersten
University of Uppsala, SE
- Holger Hermanns
Universität des Saarlandes, DE
- Lisa Higham
University of Calgary, CA
- Mark D. Hill
University of Wisconsin - Madison, US
- Marieke Huisman
University of Twente, NL
- Bengt Jonsson
University of Uppsala, SE
- Joost-Pieter Katoen
RWTH Aachen, DE
- MIchael Kuperstein
Technion - Haifa, IL
- Doug Lea
SUNY - Oswego, US
- Malte Lochau
TU Braunschweig, DE
- Andreas Lochbihler
KIT - Karlsruhe Institute of Technology, DE
- Michele Loreti
University of Firenze, IT
- Brandon M. Lucia
University of Washington - Seattle, US
- Gerald Lüttgen
Universität Bamberg, DE
- Sela Mador-Haim
University of Pennsylvania, US
- Luc Maranget
INRIA - Le Chesnay, FR
- Milo M. K. Martin
University of Pennsylvania, US
- Paul McKenney
IBM - Beaverton, US
- Michael Mendler
Universität Bamberg, DE
- Maged Michael
IBM TJ Watson Research Center - Yorktown Heights, US
- Samuel P. Midkiff
Purdue University, US
- Madan Musuvathi
Microsoft Corp. - Redmond, US
- Uwe Nestmann
TU Berlin, DE
- Scott Owens
University of Cambridge, GB
- Gustavo Petri
INRIA Sophia Antipolis, FR
- James Riely
DePaul University - Chicago, US
- Vijay A. Saraswat
IBM TJ Watson Research Center - Hawthorne, US
- Susmit Sarkar
University of Cambridge, GB
- Jens-Wolfhard Schicke
TU Braunschweig, DE
- Jaroslav Sevcik
University of Cambridge, GB
- Peter Sewell
University of Cambridge, GB
- Rob van Glabbeek
NICTA - Kensington, AU
- Heike Wehrheim
Universität Paderborn, DE
- Derek Williams
IBM Research Lab. - Austin, US
- Sascha Zickenrott
Universität des Saarlandes, DE

Report from Dagstuhl Seminar 11021

Feature-Oriented Software Development (FOSD)

Edited by

Sven Apel[1], William Cook[2], Krzysztof Czarnecki[3], and Oscar Nierstrasz[4]

1 **Universität Passau, DE,** apel@uni-passau.de
2 **University of Texas - Austin, US,** wcook@cs.utexas.edu
3 **University of Waterloo, CA,** czarnecki@acm.org
4 **Universität Bern, CH,** oscar@iam.unibe.ch

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 11021 "Feature-Oriented Software Development (FOSD)". FOSD is an emerging paradigm that aims at increasing the level of automation, reuse, and variation in software development. The main goal of the Dagstuhl Seminar on FOSD was to gather researchers and practitioners who are active in different communities to discuss the roots, state of the art, and future directions of FOSD research and practice. Additional goals were to strengthen the identity of the feature orientation community and to relate FOSD to other software development paradigms. The report contains an executive summary, abstracts of the talks held during the seminar, and summaries of special sessions.

Seminar 09.–14. January, 2011 – www.dagstuhl.de/11021
1998 ACM Subject Classification D.2.13 Reusable Software
Keywords and phrases FOSD, automation, software family
Digital Object Identifier 10.4230/DagRep.1.1.27
Edited in cooperation with Kacper Bąk

1 Executive Summary

Sven Apel
William R. Cook
Krzysztof Czarnecki
Oscar M. Nierstrasz

Seminar Motivation

Feature orientation is an emerging paradigm of software development. It supports partially or completely automated generation of software systems in a domain based on features—units of functionality covering the domain. The key idea of feature orientation is to emphasize the similarities of a family of software systems for a given application domain (e.g., database systems, banking software, and text processing systems) by treating features as first-class entities throughout the entire development lifecycle and across all the software artifacts, with the goal of reusing software artifacts among the family members. For example, features of a database system could be transaction management, query optimization, and multi-user operation, and those of a text processing system could be printing, spell checking, and document format conversions.

Feature-Oriented Software Development (FOSD), *Dagstuhl Reports*, Vol. 1, Issue 1, pp. 27–41
Editors: Sven Apel, William Cook, Krzysztof Czarnecki, and Oscar Nierstrasz
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

A key software engineering challenge is that a feature does not necessarily map cleanly to an isolated module of code. Rather, it may affect ("cut across") many components of a modular software system. For example, the feature transaction management would affect many parts of a database system, e.g., query processing, logical and physical optimization, and buffer and storage management.

The concept of feature orientation is still in its infancy. However, a growing community of researchers have been working on it for years, and there are related, well-known concepts of software engineering with well-populated research communities, e.g., software product lines, aspect-oriented software development, service-oriented architecture, and model-driven engineering. The main goal of the Dagstuhl seminar on FOSD was to gather researchers and practitioners who are active in these different communities to discuss the roots, state of the art, and future directions of FOSD research and practice and to strengthen the identity of the feature orientation community. We think that this seminar met this goal. A overview of the seminar organization and a summary of results are given below.

Seminar Organization

As a warm-up for the seminar we conducted a survey on FOSD. The idea was to ask the emerging community what they think FOSD was about. We asked the following seven questions:

1. What do you think are the distinguishing concepts and ideas of FOSD?
2. What do you think are the major challenges in FOSD?
3. Which success stories of FOSD do you know?
4. What is missing in FOSD to adopt it in industry?
5. Is FOSD sufficiently visible in the software engineering community?
6. What do you expect to get out of the week?
7. What format and what kind of activities are you interested in (tutorials, demos, talks, breakout groups, brainstorming, social events, etc.)?

Based on the responses of 27 participants (available at the seminar's website), we prepared an introductory presentation on FOSD that aimed at "synchronizing" the participants, which is especially important in a field that is still in its infancy. After the self-introductions of all of the 49 participants and the introductory presentation, we allocated slots for the "hot" topics in the field of FOSD. On Monday, we had a discussion session of feature modularity. Tuesday was dedicated entirely to feature interactions. On Thursday, we had a mix of discussions sessions on industry adoption, the relation of FOSD to other development paradigms, as well as automatic product generation based on FOSD. On Tuesday and Thursday, we had demo sessions in the evening; on Wednesday, we had breakout sessions and a social event. Finally, on Friday, we had two wrap-up sessions, one concluding the individual discussions of the breakout groups and one summarizing the seminar and discussing results and further action items.

Seminar Results

From the organizers' perspective, the seminar was successful, although the large number of participants pushed the Dagstuhl concept to its limits. The topic attracted a lot of

interest (the seminar was fully booked), and during the seminar there were many very lively and sometimes quite controversial discussions. Many participants contributed actively by organizing mini-tutorials, discussion sessions, and breakout groups. The results of the discussion sessions and the breakout groups are available at the seminar's website.

The participants used the seminar as an opportunity to learn about each others work and to establish collaborations, which will bear fruit in the years to come. As a first tangible outcome, we would like to point out the list of resources that the seminar's participants developed in a team effort:

- Key FOSD papers
- Annotated bibliographies in the portal researchr.org
- A suite of benchmark problems
- Teaching material on FOSD

The details of this list are described on the seminar's website. Further discussion points were how to promote FOSD in the future, how to further strengthen the community, and how to collaborate in an efficient manner.

In summary, we conclude that the seminar was constructive and largely met its goals. Dagstuhl provided a productive and interactive atmosphere. It was certainly a key event in the maturation of the FOSD community.

Acknowledgement

The editors would like to thank Kacper Bąk for his help in collecting the summary material from the participants and compiling this report.

2 Table of Contents

3 Overview of Talks

3.1 Feature-Aware Verification

Sven Apel (University of Passau, DE)

Joint work of Apel, Sven; Speidel, Hendrik; Boxleitner, Stefan; von Rhein, Alex; Wendler, Philipp; Beyer, Dirk

The verification of a software system is challenging. Verifying a software product line adds two aspects to the challenge. First, a software product line represents a set of software systems, and thus we have to verify that all products (distinguished in terms of their features) fulfill their specifications. Second, we have to make sure that the features of a product work properly together. Feature interactions —situations in which the combination of features leads to emergent and possibly critical behavior— are a major source of failures in software product lines. We propose feature-aware verification, a verification approach that is aware of features and their combinations. It supports the specification of feature properties along with the features in separate and composable units. We encode the verification problem (i.e., checking for critical feature interactions) such that we can apply off-the-shelf model-checking technology. Furthermore, we integrate the technique of variability encoding, which allows us to verify a product line, without generating and checking a possibly exponential number of feature combinations. To demonstrate the feasibility of our approach, we implemented feature-aware verification based on standard model-checking and applied it to a case study that incorporates the domain knowledge of AT&T on e-mail systems.

3.2 Feature Interactions: the Good, the Bad, and the Ugly

Joanne Atlee (University of Waterloo, CA)

This talk will be an overview of the research on feature interactions from the perspective of the FI community. It will include the "feature interaction problem" (which is NOT that features sometimes interact), general approaches to detecting and resolving interactions, and the trend toward architectural solutions.

3.3 Clafer Demo

Kacper Bąk (University of Waterloo, CA)

Joint work of Bąk, Kacper; Czarnecki, Krzysztof; Wąsowski, Andrzej
Main reference Kacper Bąk, Krzysztof Czarnecki und Andrzej Wąsowski, "Feature and Meta-Models in Clafer: Mixed, Specialized, and Coupled," pp. 102–122, Software Language Engineering, LNCS 6563, Springer-Verlag, 2011.
URL http://dx.doi.org/10.1007/978-3-642-19440-5_7

In the demo we present Clafer, a meta-modeling language with first-class support for feature modeling. We designed Clafer as a concise notation for meta-models, feature models, mixtures

of meta- and feature models (such as components with options), and models that couple feature models and meta-models via constraints (such as mapping feature configurations to component configurations or model templates). Clafer also allows arranging models into multiple specialization and extension layers via constraints and inheritance. We show how to model a car telematics product line in Clafer, and how to generate an instance of it in the Alloy analyzer. Clafer models are translated to Alloy via clafer2alloy.

3.4 FOSD—A Science of Software Design: A Personal and Historical Perspective

Don Batory (University of Texas at Austin, US)

I present a summary of my involvement with FOSD for the last 30 years, from a tutorial perspective.

3.5 Model checking feature-based specifications with SNIP

Andreas Classen (University of Namur, BE)

In software product line engineering, systems are developed in families and differences between systems of a product line are expressed in terms of features. The model checking problem for product lines is more difficult than for single systems because a product line with n features yields up to 2^n individual systems to verify.

We present SNIP, a tool for model checking product lines against temporal properties. Contrary to existing approaches, SNIP relies on an efficient mathematical structure for product line behaviour, that exploits similarities and represents the behaviour of all systems in a compact manner. This structure is used to model check all systems of the product line in a single step.

The tool comes together with an intuitive specification language based on Promela. Variability in the behaviour is specified by guarding statements with features. A guarded statement is only executed in products that contain the features of its guard. When checking properties, SNIP takes these guards and the feature model into account. For this, it uses the TVL feature modelling language.

SNIP puts the theoretical results of our ICSE 2010 paper "Model Checking Lots of Systems" into practice and makes them available to engineers through an intuitive specification language for behaviour and for feature models.

3.6 Context-oriented Programming: First-class layers at runtime

Pascal Costanza (Vrije Universiteit Brussel, BE)

Main reference Robert Hirschfeld, Pascal Costanza, Oscar Nierstrasz: "Context-oriented Programming", in Journal of Object Technology, vol. 7, no. 3, March-April 2008, pp. 125–151.
URL http://dx.doi.org/10.5381/jot.2008.7.3.a4

Context-oriented Programming is strongly related to Feature-oriented Programming in that both approaches use layers as abstractions for expressing feature increments and behavioral variations of a software system. However, in Context-oriented Programming, layers are available as first-class entities at runtime that can be activated and deactivated with well-defined scopes, in order to enable a software system to change its behavior depending on the context of use.

I have presented a short overview of the notion of Context-oriented Programming, its essential elements, and a concrete realization in the form of the programming language ContextL.

3.7 Choice Calculus Mini Tutorial

Martin Erwig (Oregon State University, US)

Joint work of Erwig, Martin; Walkinshaw, Eric
Main reference The Choice Calculus: A Representation for Software Variation, ACM Transactions on Software Engineering and Methodology, 2011, to appear
URL http://web.engr.oregonstate.edu/~erwig/papers/abstracts.html#TOSEM11

We describe the basic elements of the *choice calculus*, a formal representation for software variation that can serve as a common, underlying representation for variation research, playing a similar role that lambda calculus plays in programming language research. We will sketch the syntax and semantics of the choice calculus and present several applications.

At the core of the choice calculus are *choices*, which represent different alternatives that can be selected. Choices are annotated by names, which group choices into *dimensions*. Dimensions provide a structuring and scoping mechanism for choices. Moreover, each dimension introduces the number of alternatives each choice in it must have and tags for selecting those alternatives. The semantics of the choice calculus is defined via repeated elimination of dimensions and their associated choices through the selection of a tag defined by that dimension.

The choice calculus obeys a rich set of laws that give rise to a number of normal forms and allow the flexible restructuring of variation representations to adjust to the needs of different applications.

Among the potential applications of the choice calculus are feature modeling, change pattern detection, property preservation, and the development of change IDEs.

3.8 Re-Thinking Product Line Verification as a Constraints Problem

Kathi Fisler (Worcester Polytechnic Institute, US)

Main reference Colin Blundell, Kathi Fisler, Shriram Krishnamurthi, Pascal Van Hentenryck, "Parameterized Interfaces for Open System Verification of Product Lines," pp. 258–267, Proceedings of the 19th IEEE International Conference on Automated Software Engineering, 2004.
URL http://dx.doi.org/10.1109/ASE.2004.53

Software product-lines view systems as compositions of features. Each component corresponds to an individual feature, and a composition of features yields a product. Feature-oriented verification must be able to analyze individual features and to compose the results into results on products. Since features interact through shared data, verifying individual features entails open system verification concerns. To verify temporal properties, features must be open to both propositional and temporal information from the remainder of the composed product. This talk argues that we can handle both forms of openness by viewing product-line verification as a two-phase process of constraint generation and discharge, rather than as a conventional verification problem.

3.9 Type Checking entire Product Lines

Christian Kästner (Universität Marburg, DE)

Joint work of Kästner, Christian; Apel, Sven

I propose to give a very brief (6-10 minute) overview of concept of checking the product line instead of checking the derived products and how this can be applied to type systems. The idea is to highlight a common research trend and stipulate discussions about adopting all kinds of checks to check the entire product line.

In a nutshell, the idea of a variability-aware type system to check the product-line implementation containing all variability (before product derivation) instead of checking each (of potentially billions of) derived product in isolation. Although the number of potential products may be exponential, variability-aware type systems exploit knowledge about variability implementation and can separate intrinsic complexity from accidental complexity. We are currently trying to scale such type system to type check the Linux kernel with over 8000 features. We believe that moving analysis from derived products to the product-line implementation is a common theme that can and should also be applied to many other approaches.

3.10 Modularity vs. Virtual Separation of Concerns in Feature-Oriented Implementations

Christian Kästner (Universität Marburg, DE)

When implementing features, the question of modularity arises. Interactions between feature (intended and accidental) at domain level and implementation level pose serious questions of how to divide a system into modules and what benefits a strict modularization brings. At the same time, we argue that tool support can provide many benefits of modularity also for other forms of implementation, such as the notorious conditional compilation used frequently by practitioners.

In this talk, I like to give a brief overview of different implementation strategies, but mostly focus on open questions regarding modularity to initiate a discussion. As one potential solution I discuss some results of our experience with virtual separation of concerns and related approaches.

3.11 A Survey of Product Family Algebra

Bernhard Möller (Universität Augsburg, DE)

Main reference Peter Höfner, Ridha Khédri, Bernhard Möller, "Supplementing Product Families with Behaviour." Accepted for publication in the International Journal of Software and Informatics.

Using the well-known concept of a semiring we develop an algebra of product families with + standing for union and . for composition, like e.g. in regular expressions. The distributive semiring laws allow factoring out commonality and, conversely, algebraically calculating all possible products in a family. The algebra is enriched by a relation expressing that presence of one feature necessitates or excludes that of another one. Finally the "syntactic" view working only with feature names is supplemented by a semantic one, in which features are programs with a soundly defined semantics, in such a way that the axioms of product family algebra are still satisfied.

3.12 Testing Software Product Lines

Sebastian Oster (TU Darmstadt, DE)

Main reference Sebastian Oster, Florian Markert, Philipp Ritter, "Automated incremental pairwise testing of software product lines," Proceedings of the 14th International Conference on Software product lines: going beyond, pp. 196–210, Jeju Island, South Korea, Springer-Verlag, Berlin, Heidelberg, 2010.
URL http://dx.doi.org/10.1007/978-3-642-15579-6_14

Testing Software Product Lines is very challenging due to a high degree of variability leading to an enormous number of possible products. The vast majority of today's testing approaches for Software Product Lines validate products individually using different kinds of reuse

techniques for testing. Due to the enormous number of possible products, individual product testing becomes more and more unfeasible.

Combinatorial testing offers one possibility to test a subset of all possible products. In this contribution we provide a detailed description of a methodology to apply combinatorial testing to a feature model of a Software Product Line. We combine graph transformation, combinatorial testing, and forward checking for that purpose. Additionally, our approach considers predefined sets of products.

3.13 Event-B and Rodin Overview

Michael Poppleton (University of Southampton, UK)

We present the Event-B language for formal modelling, development and verification, and its supporting Rodin toolkit. This is now a leading state-based Formal Method, supported by a series of EU Framework STREP/IP grants. As introduction, the value added to a V-model is outlined. An example development is outlined to demonstrate the modelling, refinement, and V&V techniques, using the tools. Various Rodin plugin tools are described, those for decomposition/ composition being of particular interest (to a FOSD audience)

3.14 FOSD Adoption in Industry (Plenary Session Summary)

Rick Rabiser (Universität Linz, AT)

In the discussion session *FOSD Adoption in Industry* three impulse talks were given about experiences of using FOSD and product line engineering in industry. More specifically, Christa Schwanninger reported experiences of feature-oriented development from Siemens AG; Charles Krueger presented experiences of introducing feature-oriented tools by BigLever Inc. in different large-scale organizations; and Paul Grünbacher reported project experiences of maturing a PLE research tool for industrial application within Siemens VAI.

The presenters discussed different industrial challenges motivating the use of FOSD. For example, the high number of configuration options in systems leading to long release durations; the lack of agreement on the scope of systems; or the competition forcing companies to perform explicit variability management. Main goals when applying FOSD in practice are to harness complexity, to improve the efficiency of production, and to reduce the time-to-market.

The session then focused on two particular areas for discussion, i.e., *What's already done in industrial practice?* and *What are the industrial challenges FOSD researchers should work on?*. The presentations and the discussion showed that FOSD is performed in different degrees in practice.

Variability management is done for diverse artifacts and using different tools ranging from simple spreadsheets to fully fledged PLE tool suites. Companies have their own and often very specific ways to perform FOSD. However, participants reported many challenges are remaining for FOSD researchers:

Handling of non-functional properties; feature-driven build management; standardization (e.g., of variability modeling, cf. CVL); simplification (esp. of user interfaces); integration with COTS apps and existing environments; support for industrial workflows; sharing of models and project management data; security; deployment support.

3.15 Feature-Oriented Requirements Engineering—An Integrated Modeling, Aspect and Decision Based Approach

Reinhard Stoiber (Universität Zürich, CH)

Main reference Reinhard Stoiber, Martin Glinz, "Supporting Stepwise, Incremental Product Derivation in Product Line Requirements Engineering," pp. 77–84, Proceedings of VaMoS'10: Fourth International Workshop on Variability Modelling of Software-intensive Systems, Linz, Austria, 2010.
URL http://www.vamos-workshop.net/proceedings/VaMoS_2010_Proceedings.pdf

Considering variability modeling in requirements engineering is essential for developing, maintaining and evolving a software product line.

In this talk we introduce current techniques how practitioners and researchers typically deal with variability and point out open challenges in requirements engineering. Based on a real-world example, we demonstrate how our work on integrated requirements modeling, aspect-oriented modeling and Boolean decision modeling can help to overcome these problems. Our approach avoids information scattering of features over multiple diagrams, automates clerical as well as intellectual tasks for extracting and composing the variability, and automatically resolves and propagates variability constraints. This allows an easier creation of software product line models, improves the understanding of variability and its impact, and leads to an efficient and intuitive derivation of products that are correct by construction. Ongoing work addresses tool optimizations and a comprehensive validation.

4 Working Groups

4.1 Variability Representation (Breakout Session Summary)

Kacper Bąk (University of Waterloo, CA)

The variability representation session focused on languages for modeling variability in source code and models. It started with an introduction to the choice calculus, a fundamental language for representing software variations. The choice calculus aspires to play a similar role for variation research as lambda calculus does in programming language research. The presentation covered basic ideas behind the choice calculus (choices and dimensions) and explained how choices get eliminated to derive a single product. The designers of the choice calculus were challenged to show how to translate feature models to the choice calculus. Even though the resulting models may be verbose, the translation is indeed possible.

Later the session turned to the Common Variability Language, which is the upcoming OMG standard for introducing variability into existing models. CVL is composed of several layers, of which two were particularly relevant to the session: variability realization and

variability abstraction. The former includes variation points that have concrete impact on the base model, while the latter shows variation at a higher level.

The main conclusion was that the choice calculus might be a suitable theory for variation points.

The last part of the session concerned static versus dynamic variability. Static variability is resolved at compile time, while dynamic features are loaded at the run-time. One question discussed was this: What should happen when user loads a dynamic feature which is in conflict with an existing feature? Participants discussed soft and hard constraints in the context of mobile applications. Soft constraints can be violated, while hard constraints must always hold. Mobile applications are user-centric, thus user preferences should also be taken into account when two features are in conflict. A proposed solution was to use probabilistic feature models to attach user preferences to features.

4.2 Non-functional versus Function Properties in Feature-oriented Product Line Generation (Breakout Session Summary)

Charles Krueger (BigLever, Austin, US)

With feature-oriented product line engineering, stakeholders exploit feature models to understand and manage functional properties during the generation of a system. Although functional behavior of a system is the primary concern, the associated non-functional properties – such as resource usage, performance, cost and perceived product value – are also important concerns in industry settings.

The common scenarios include (1) adjusting functionality in order to bring one or more non-functional properties with specified bounds, (2) expanding functionality to increase system value by taking advantage of underutilized capacity in a non-functional property. The relationship between features and non-functional properties may be hard or easy to measure and control, additive, binary, non-monotonic or chaotic.

How can feature-oriented approaches simultaneously manage the desired functional and non-functional properties of a system under generation? Measurements and statistical modeling of non-functional properties relative to feature selections can be incorporated into generators to provide guidance.

Adjacent fields that can contribute include operations research, AI and satisfaction solvers.

4.3 Feature Location (Breakout Session Summary)

Christian Kästner (Universität Marburg, DE)

In a breakout session on feature location, we assembled a group of people who pursue different location strategies. After introducing the respective concepts or tools, we soon realized that the goals differ significantly. We identified three main higher level goals: locating features for maintenance tasks, for design decisions, and for reuse. Some participants pursue mainly one of these goals, for example CodeX for maintenance, Portolio and CIDE for reuse, and

Feature Unweaving for Design, whereas several tools also pursued multiple goals at the same time, e.g., Linux Analysis and the Requirements to Test mappings.

We concluded that many approaches are complementary and we discussed possible integrations. A main outcome was the agreement to share case studies and the formulation of a vision of an entire process that spans feature-location tasks from design to implementation.

Participants

- Sven Apel
Universität Passau, DE
- Joanne Atlee
University of Waterloo, CA
- Kacper Bak
University of Waterloo, CA
- Don Batory
Univ. of Texas at Austin, US
- Thorsten Berger
Universität Leipzig, DE
- Götz Botterweck
University of Limerick, IE
- Andreas Classen
University of Namur, BE
- William R. Cook
University of Texas - Austin, US
- Pascal Costanza
Vrije Universiteit Brussel, BE
- Krzysztof Czarnecki
University of Waterloo, CA
- Ben Delaware
University of Texas - Austin, US
- Zinovy Diskin
University of Waterloo, CA
- Christoph Elsner
Universität Erlangen-Nürnberg, DE
- Martin Erwig
Oregon State University, US
- Kathi Fisler
Worcester Polytechnic Institute, US
- Martin Glinz
Universität Zürich, CH
- Gerhard Goos
KIT - Karlsruhe Institute of Technology, DE
- Mark Grechanik
Accenture Labs - Chicago, US
- Paul Grünbacher
Universität Linz, AT
- Florian Heidenreich
TU Dresden, DE
- Robert Hirschfeld
Hasso-Plattner-Institut - Potsdam, DE
- Christian Kästner
Universität Marburg, DE
- Kyo-Chul Kang
POSTECH - Pohang, KR
- Hans Körber
FH Kaiserslautern-Zweibrücken, DE
- Shriram Krishnamurthi
Brown University - Providence, US
- Charles Krueger
BigLever- Austin, US
- Ingolf Krüger
University of California - San Diego, US
- Ralf Lämmel
Universität Koblenz-Landau, DE
- Jaejoon Lee
Lancaster University, GB
- Thomas Leich
METOP GmbH - Magdeburg, DE
- Christian Lengauer
Universität Passau, DE
- Bernhard Möller
Universität Augsburg, DE
- Oscar M. Nierstrasz
Universität Bern, CH
- Sebastian Oster
TU Darmstadt, DE
- Holger Papajewski
pure-systems GmbH - Magdeburg, DE
- Joe Politz
Brown University - Providence, US
- Michael Poppleton
University of Southampton, GB
- Christian Prehofer
LMU München, DE
- Rick Rabiser
Universität Linz, AT
- Jorge Ressia
Universität Bern, CH
- Bernhard Rumpe
RWTH Aachen, DE
- Klaus Schmid
Universität Hildesheim, DE
- Christa Schwanninger
Siemens AG - Erlangen, DE
- Steven She
University of Waterloo, CA
- Olaf Spinczyk
TU Dortmund, DE
- Reinhard Stoiber
Universität Zürich, CH
- Salvador Trujillo
Ikerlan Research Centre - Arrasate-Mondragón, ES
- Eric Walkingshaw
Oregon State University, US
- Andrzej Wasowski
IT University of Copenhagen, DK

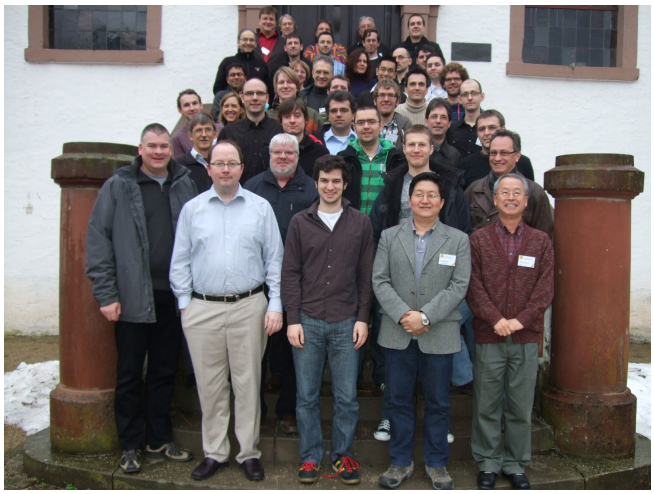

Report from Dagstuhl Seminar 11031

Bidirectional Transformations "bx"

Edited by

Zhenjiang Hu[1], Andy Schürr[2], Perdita Stevens[3], and James Terwilliger[4]

1 **National Institute of Informatics - Tokyo, JP,** hu@nii.ac.jp
2 **TU Darmstadt, DE,** Andy.Schürr@es.tu-darmstadt.de
3 **University of Edinburgh, GB,** perdita@inf.ed.ac.uk
4 **Microsoft Corporation - Redmond, US,** james.terwilliger@microsoft.com

Abstract

Bidirectional transformations (*bx*) are a mechanism for maintaining the consistency of two (or more) related sources of information. Researchers from many different areas of computer science including databases (DB), graph transformations (GT), software engineering (SE), and programming languages (PL) are actively investigating the use of *bx* to solve a diverse set of problems. Although researchers have been actively working on bidirectional transformations in the above mentioned communities for many years already, there has been very little cross-discipline interaction and cooperation so far. The purpose of a first International Meeting on Bidirectional Transformations (GRACE-BX), held in December 2008 near Tokyo, was therefore to bring together international elites, promising young researchers, and leading practitioners to share problems, discuss solutions, and open a dialogue towards understanding the common underpinnings of *bx* in all these areas. While the GRACE-BX meeting provided a starting point for exchanging ideas in different communities and confirmed our believe that there is a considerable overlap of studied problems and developed solutions in the identified communities, the Dagstuhl Seminar 11031 on "Bidirectional Transformations" also aimed at providing a place for working together to define a common vocabulary of terms and desirable properties of bidirectional transformations, develop a suite of benchmarks, solve some challenging problems, and launch joint efforts to form a living *bx* community of cooperating experts across the identified subdisciplines. This report documents the program and the outcomes of Dagstuhl Seminar 11031 with abstracts of tutorials, working groups, and presentations on specific research topics.

Seminar 17.–21. January, 2011 – www.dagstuhl.de/11031
1998 ACM Subject Classification D.2 Software Engineering, D.3 Programming Languages, H.1 Models and Principles
Keywords and phrases Bidirectional Languages, Transformation, Model/Data Synchronisation
Digital Object Identifier 10.4230/DagRep.1.1.42
Edited in cooperation with Anthony Anjorin

Bidirectional Transformations "bx", *Dagstuhl Reports*, Vol. 1, Issue 1, pp. 42–67
Editors: Zhenjiang Hu, Andy Schürr, Perdita Stevens, and James Terwilliger

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

1 Executive Summary

Zhenjiang Hu
Andy Schürr
Perdita Stevens
James Terwilliger

This Dagstuhl Seminar was the second meeting bringing together 39 researchers from 13 countries across disciplines that study bidirectional transformations. The first one was the GRACE International Meeting on Bidirectional Transformations held in December 2008 near Tokyo, Japan [1]. The GRACE meeting consisted of short introductions from each of the participants on their background and work, followed by some longer presentations and demonstrations on some representative technologies from each field, concluding in some open discussion time. A primary takeaway from the GRACE meeting was an opportunity for each discipline to get some initial exposure to each other.

bx Dagstuhl Seminar 11031 Programme, Variant 2

	Monday	Tuesday	Wednesday	Thursday	Friday
09:00-10:30	Seminar Opening Round of Introductions	Basic DB Tutorial, P. 2 WG Proposals	Working Groups	Working Groups	4 Position Statements WG Report
10:30-11:00	Coffee Break				
11:00-12:00	General bx Intro Basic PL Tutorial	Basic GT Tutorial, P. 1	6 Position Statements	6 Position Statements	Closing Session
12:15-14:00	Lunch				
14:00-15:30	Basic PL Tutorial, Part 2	Basic GT Tutorial, P. 2 Basic SE Tutorial, P. 1	Excursion to Trier Dinner	WG Reports 3 Position Statements	
15:30-16:00	Coffee Break			Coffee Break	
16:00-17:30	Basic DB Tutorial Part 1	Basic SE Tutorial, P. 2		Working Groups	
18:00-.....	Dinner			Dinner	

Seminar Opening	Information about background, organizational matters, ...
Round of Introductions	about 1 minute per person without any slides
Basic Tutorials	Introduction to one subdomain; often combination of several tutorial proposals
WG Proposals	Presentation of ideas for working groups (up to 10 minutes)
Working Groups	Parallel working groups on different topics
WG Reports	Summaries of Working Group Activities
Position Statements	Short presentations of up to 10 minutes length

Figure 1 Time Table of Seminar - Updated Final Version

The Dagstuhl seminar intended to go a step further and begin to identify commonalities between the disciplines and start to set a cross-disciplinary research agenda. The first part of the seminar (cf. Fig 1 with a time table) consisted of tutorials from each of the four represented disciplines on the various bidirectional transformation solutions that field has to offer. The second part consisted of cross-disciplinary working groups dedicated to investigating specific examples of commonality between solutions or identifying requirements, terminology, or scenarios that may reach across fields. There were also a number of sessions reserved for participants to give position statements regarding their individual fields of research.

Participation at both the Dagstuhl and GRACE seminars came from four disciplines: (1) Programming Languages (PL), (2) Graph Transformations (GT), (3) Software Engineering

(SE), and (4) Databases (DB). Each of the first three disciplines made up about 2/7 of the participants, while databases took the remaining 1/7 out of about 39 participants. Representation from the database field was, nevertheless, an improvement over the turnout from the GRACE meeting.

Tutorials

The first part of the workshop was allocated for representatives from each of the four disciplines giving deep tutorials of about two hours length on the various solutions to bidirectional transformation problems offered by that discipline. In all cases these tutorials were presented by groups of scientists, who were responsible to offer the attendants of the other disciplines a basic introduction into their field as well as to highlight different aspects of ongoing research activities. The PL-related tutorial consisted of five parts presented by Nate Foster, Robert Glück, Zhenjiang Hu, Benjamin Pierce, and Janis Voigtländer. It contained a general introduction to the whole research area of bx. Furthermore, it addressed a large variety of different aspects ranging from specific types of bx programming languages on one hand to model/graph transformation approaches in the software-engineering world and lense-based view definition approaches in the database system world on the other hand. As a consequence this set of (sub-)tutorials also was responsible for highlighting already existing links between the existing bx subcommunities. The DB-tutorial offered by Jean-Luc Hainaut and James Terwilliger nicely complemented the preceding tutorial. A survey of various opportunities for using bx transformations in database engineering processes as well as for specifying views by means of bx technologies was presented. The GT-related tutorial afterwards given by Andy Schürr and Frank Hermann focused on one special category of bidirectional graph transformations. Its two parts dealt with practical as well as theoretical aspects of bx graph transformations. Finally, the SE-related tutorial of Krzysztof Czarnecki and Stephan Hildebrandt concluded the first 2.5 days long introduction to the field of bx with a discussion of model synchronization and incremental change propagation techniques from the point of view of the software engineering community. For a more detailed description of the four (sets of) tutorials the reader is referred to their abstracts collected in this Proceedings.

Working Groups

After the ground had been prepared by the above mentioned list of tutorials the participants had the opportunity to present specific research results and challenges in the form of 27 short presentations. Furthermore, six working groups were organized related to cross-disciplinary topics that were identified during the first days of the Seminar. The summaries of three of these working groups only finally made into the Proceedings, but all of them were very helpful to deepen our understanding of common challenges and triggered new research cooperations across the four disciplines. The list of addressed topics included aspects like the design of bx benchmarks, the identification of typical bx application scenarios and their requirements as well as options for the joint development of a new generation of bx languages that would incorporate and combine elements from solutions developed in different disciplines. The individual position statements covered a very broad spectrum of topics; they were used to intensify already started discussions of the first tutorial days and gave further input for the identification of new workshop topics. Again the reader is referred to the abstracts collected here for more details.

Summary and Future Work

We went into the seminar knowing that longer-term ideas like a common research agenda or a benchmark would take longer than a single week. The participants decided on several follow-up actions to keep work progressing:

- A follow-up meeting in the same style as the GRACE and Dagstuhl seminars to continue collaborating on a cross-disciplinary research agenda
- Workshops at conferences associated with each of the disciplines to work toward specific, targeted goals (a first one has already been scheduled associated with GTTSE 2011, and will focus on developing a benchmark[1]; a second follow-up event has just been accepted as a satellite workshop for ETAPS.)
- Tutorials and other assorted smaller, education-minded events at conferences to continue bringing awareness of bidirectional solutions from other disciplines, as well as awareness of the general BX effort
- Smaller-scale research cooperations that combine techniques from different fields like merging concepts from bidirectional programming languages and triple graph grammars as envisaged in one of the seminar's working groups.

In particular, one goal of the upcoming seminars and workshops is to increase participation from the database community. The bidirectional transformation problem has origins deep in the database community, but now has grown so that solutions are being driven from many different directions in different fields across computer science. The plan is to hold some of the tutorials or workshops at database venues to help solicit more ideas and opportunities for collaboration; details will be made available once they are scheduled.

References

1 K. Czarnecki, J. N. Foster, Z. Hu, R. Lämmel, Andy Schürr, and J. F. Terwilliger. Bidirectional Transformations: A Cross-Discipline Perspective.*ICMT 2009*, 260–283.

[1] http://www.di.univaq.it/CSXW2011/

2 Table of Contents

3 Overview of Tutorials

3.1 Languages for Bidirectional Transformations

Nate Foster (Cornell University, US)

In recent years, a wide variety of approaches for describing bidirectional transformations have been proposed in the programming languages community. This multi-part tutorial surveyed recent work in this area and presented some of most promising approaches in detail. The first talk, by Nate Foster, introduced the key semantic issues and identified some connections to prior work on the view update problem in databases. The second talk, by Robert Glück, described *reversible languages*, which are languages in which every program denotes an injective function and can be effectively inverted. The third talk, by Janis Voigtländer, described a technique called *semantic bidirectionalization* that constructs a backward function from a polymorphic forward function, using parametricity and free theorems to prove well-behavedness. The fourth talk, by Benjamin Pierce, described *lens combinators* and focused on the use of type systems to establish well-behavedness as well as the issues surrounding the handling of ordered data. The final talk, by Zhenjiang Hu, presented a language for describing *bidirectional graph transformations* that uses trace information to guide the reverse transformation.

3.2 Complement-Based Bidirectionalization

Janis Voigtländer (Universität Bonn, DE)

This was part of the programming languages tutorial. The slides used are here:
http://www.dagstuhl.de/mat/Files/11/11031/11031.VoigtlaenderJanis1.Slides.pdf.

3.3 Trace-based Bidirectionalization

Zhenjiang Hu (NII - Tokyo, JP)

Traces play an important role in bidirectional transformation, widely used for data exchanges between two databases, such as relational database and relational database, relational database and tree database, tree database and tree database, and even graph database and graph database. Traced-based bidirectionalization effectively handles updates on the view whose destination has origins (can be traced back) to updatable input data, by first checking whether the updates can be traced back to its origin and then propagating the updates to the source in a safe way. In this mini-tutorial, I explain the basic idea of trace-based bidirectionalization, and show how to compute traces effectively and how to propagate the view updates to the original database safely.

3.4 Principles of Reversible Programming Languages

Robert Glück (University of Copenhagen, DK)

Reversible computing is the study of computing models that exhibit both forward and backward determinism. Understanding the fundamental properties of such models is not only naturally relevant for reversible programming, but has also been found important in bidirectional model transformation and program transformations such as inversion. The unique design features of these languages include explicit postcondition assertions, direct access to an inverse semantics (e.g. uncall), and the possibility of clean (garbage-free)computation of injective functions.

References

1 Yokoyama T., Axelsen H. B., Glück R., Principles of a reversible programming language. In: Conference on Computing Frontiers. 43-54, ACM 2008. http://doi.acm.org/10.1145/1366230.1366239

2 Axelsen H. B., Glück R., Yokoyama T., Reversible machine code and its abstract processor architecture. Lecture Notes in Computer Science, LNCS 4649, 2007. http://dx.doi.org/10.1007/978-3-540-74510-5_9

3.5 Bidirectional Transformations in Database Research and Practice

James Terwilliger (Microsoft Corporation - Redmond, US)

The primary tool for transformations in databases is the query.

Substantial research has been done on attempting to turn this basic tool into a bidirectional transformation. These attempts have developed a substantial amount of theoretical progress but hardly any uptake in practical applications.

This tutorial covered four research topics introduced in the past five years that provide new solutions with practical applications to the bidirectional transformation problem in databases. PRISM is a tool - with a formal predicate calculus backing - designed to handle the schema versioning problem in applications. Object-Relational Mappings are a way to bridge the traditional impedance mismatch between application and database, where an application must reliably be able to store and update data.

Guava is a framework that allows an application the freedom to query and update against a virtual schema, as well as evolve that schema, with a flexible mapping to its physical storage. Finally, the field of data exchange has invested research in being able to invert mappings expressed using predicate calculus.

3.6 The transformational approach to data-intensive system engineering and evolution

Jean-Luc Hainaut (University of Namur, BE)

Joint work of Hainaut, Jean-Luc; Anthony, Cleve
Main reference Hainaut, J-L, The Transformational Approach to Database Engineering, in Lämmel, R., Saraiva, J., Visser, V., (Eds), Generative and Transformational Techniques in Software Engineering, pp. 95-143, LNCS 4143, Springer, 2006
URL http://dx.doi.org/10.1007/11877028_4

Transformation-based software engineering has long been considered a major scientific approach to build reliable and efficient programs. According to this approach, abstract specifications (or model) can be converted into correct, compilable and efficient programs by applying selected, correctness-preserving operators called transformations.

In the database engineering realm, an increasing number of bodies (e.g., OMG through the MDE proposal) and of authors recognize the merits of transformational approaches, that can produce in a systematic way machine-processable database structures from abstract models. Transformations that are proved to preserve the correctness of the source specifications have been proposed in virtually all the activities related to schema engineering: schema normalization, schema quality evaluation, logical design, schema integration, views derivation, schema equivalence, database migration, data conversion, reverse engineering, schema optimization, ETL, wrapper generation and others. The proposed tutorial addresses both basic and practical aspects of database transformation techniques. The concept of transformation is developed, together with its properties of semantics-preservation. Major database engineering activities are redefined in terms of transformation techniques, and the impact on CASE technology is discussed. We also show in this tutorial that schema transformations can be used as a formal basis for deriving data transformations as well as program transformation.

3.7 Triple Graph Grammars in a Nutshell

Andy Schürr (TU Darmstadt, DE)

Main reference F. Klar, M. Lauder, A. Königs, A. Schürr: "Extended Triple Graph Grammars with Efficient and Compatible Graph Translators", in: A. Schürr, C. Lewerentz, G. Engels, W. Schäfer, B. Westfechtel (eds.): Graph Transformations and Model Driven Enginering - Essays Dedicated to Manfred Nagl on the Occasion
URL http://dx.doi.org/10.1007/978-3-642-17322-6_8

The first part of the Triple Graph Grammar (TGG) tutorial is a gentle introduction to the basic ideas of and motivations for the development of a rule-based and declarative specification formalism that allows for the high-level description of functional and non-functional relationships between pairs of graphs (models). A family of graph (model) transformations is derived from such a TGG specification that supports batch transformation and incremental change propagation scenarios in both directions as well as checking the consistency of given pairs of models or graphs. Traceability relationships between elements of related pairs of graphs are created and updated as a side effect. The presentation prepares the ground for the 2nd part of the tutorial which sketches the formal background of TGGs

based on category theory and presents techniques for the verification of important properties of TGGs and derived graph (model) translators.

3.8 Analysis of Model Transformations based on TGGs

Frank Hermann (TU Berlin, DE)

Main reference Frank Hermann, Hartmut Ehrig, Ulrike Golas, and Fernando Orejas: Efficient analysis and execution of correct and complete model transformations based on triple graph grammars. In: Proc. MDI '10, Jean Bezivin, Richard Mark Soley, and Antonio Vallecillo (Eds.), ACM 2010.
URL http://dx.doi.org/10.1145/1866272.1866277

Triple graph grammars (TGGs), introduced by Schürr in 1994, are a formal and intuitive concept for the specification of bidirectional model transformations. The key advantages of TGGs are the automatic derivation of operational rules, which simplifies specification, and powerful analysis capabilities based on the underlying formal foundation. This tutorial presents several automated analysis techniques concerning important properties of model transformations based on TGGs. In the first part, we present execution formalisms, which closely correspond to available implementations and ensure syntactical correctness and completeness of model transformations. Moreover, termination of the execution is guaranteed by static checks of non-restrictive conditions.

In the second part, we present how functional behaviour and information preservation of model transformations are checked by automated techniques. These techniques are additionally used for detection and resolution of conflicts between model transformation rules as well as for improving efficiency of the execution.

3.9 Incremental Model Synchronization

Holger Giese (Hasso-Plattner-Institut GmbH, DE)

In system development the size of models and their proper synchronization can become a major problem. In this tutorial we will discuss what is theoretically wanted for model synchronization and what additional needs and limitations come into play in practice. The problem and existing solutions that extend triple graph grammars are outlined using an example from an industry project on model synchronization.

3.10 Model Synchronization: Theory and Practice

Krzysztof Czarnecki (University of Waterloo, CA)

Model-driven engineering often requires many overlapping models of a system, each supporting a particular kind of stakeholder or task. The consistency among these models needs to be

managed during system development. Consistency management occurs in the space of multiple model replicas, versions over time, and different modeling languages, which make the process complex and challenging. In this tutorial, I will introduce the key concepts and mechanisms of consistency management, including mapping definitions to specify overlap, model alignment to establish differences, reconciliation to resolve conflicts, and change propagation to re-establish consistency. Most importantly, I show how bidirectional transformation using asymmetric and symmetric lenses used to check consistency and to propagate changes. I will illustrate these concepts and mechanisms with practical examples of model synchronization, including synchronizing architectural models and code, business process specifications and implementations, and structural and behavioral models in UML.

3.11 Update propagation via tiles

Zinovy Diskin (University of Waterloo, CA)

Main reference Diskin, Z., Model synchronization: Mappings, Tiles, and Categories. In Generative and Transformational Techniques in Software Engineering III, Springer LNCS'6491/2011, 92-165
URL http://dx.doi.org/10.1007/978-3-642-18023-1_3

The tutorial presents a novel algebraic framework for specifying architectures of model synchronization tools. The basic premise is that synchronization procedures, and hence algebraic operations modeling them, are *diagrammatic*: they take a configuration (diagram) of models and mappings as their input and produce a diagram as output. Many important synchronization scenarios are based on diagram operations of square shape. Composition of such operations amounts to their *tiling*, and complex synchronizers can thus be assembled by tiling together simple synchronization blocks. This gives rise to a visually suggestive yet precise notation for specifying synchronization procedures and reasoning about them. And the last but not least, specification and design with tiles are enjoyable.

Details can be found in my paper *Model synchronization: Mappings, Tiles, and Categories. In GTTSE'2009, Springer LNCS'6491/2011, pp. 92-165.*

4 Overview of Position Statements

4.1 Inconsistency detection and resolution in heterogenous model-based specifications using Maude

Artur Boronat (University of Leicester, GB)

Joint work of Boronat, Artur; Meseguer, José
Main reference A. Boronat, J. Meseguer: Automated Model Synchronization: a Case Study on UML with Maude. Tenth International Workshop on Graph Transformation and Visual Modeling Techniques. Collocated with ETAPS 2011. Saarbrücken (Germany). April 2-3, 2011. Pre-Proceedings.

Software-intensive systems comprise a wide range of heterogeneous software artefacts, which are usually developed by distributed teams. In model-based development, these software artefacts are given as models, abstracting relevant features of a system from implementation details, enhancing development, maintenance and evolution processes. In heterogeneous

software specifications, these models describe parts of the same system from different points of view and from different levels of abstraction, hence overlapping in various ways. Parallel updates in different models may result in unanticipated inconsistencies in the overlapped part. An approach based on rewriting logic for detecting inconsistencies in heterogeneous model-based specifications and to synchronise their constituent models will be presented. In this approach, strategies are used to guide the synchronisation process in an optimal way, filtering out many non-desirable solutions.

4.2 On the role of Triple Graph Grammars Concerning Requirements for Enterprise Modeling

Christoph Brandt (University of Luxembourg, LU)

Joint work of Brandt, Christoph; Hermann, Frank
Main reference Christoph Brandt and Frank Hermann, How Far Can Enterprise Modeling for Banking Be Supported by Graph Transformation?, ICGT, Springer, LNCS, 6372, 2010, 3-26
URL http://dx.doi.org/10.1007/978-3-642-15928-2_2

Reconstructed requirements for enterprise modeling are presented that came up during an ongoing evaluation of enterprise modeling practices at Credit Suisse. They will help to reorganize, reframe and reconstruct today's modeling approaches leading to better results built on top of sound formal techniques. First investigations show that triple graph grammars are suitable to provide solutions for several of these requirements.

4.3 Co-evolving schemas and programs using coupled transformations

Anthony Cleve (University of Namur, BE)

Joint work of Cleve, Anthony; Hainaut, Jean-Luc
Main reference Anthony Cleve and Jean-Luc Hainaut,Co-transformations in database applications evolution. In Ralf Lämmel, João Saraiva, and Joost Visser, editors, Generative and Transformational Techniques in Software Engineering, volume 4143 of Lecture Notes in Computer Science, pages 409421. Springer, 2006.
URL http://dx.doi.org/10.1007/11877028_17

In this short position statement we will elaborate on the use of coupled transformations to support the co-evolution of inter-dependent artefacts in data-intensive systems. We will particularly focus on the co-evolution of database schema and programs in such scenarios as schema refactoring, database migration and database design.

4.4 Type-safe Evolution of Spreadsheets

Jacome Cunha (Universidade de Minho - Braga, PT)

Main reference Jãcome Cunha, Joost Visser, Tiago Alves, and João Saraiva. Type-safe evolution of spreadsheets. In FASE '11: Proc. of the 13th Int'l Conf. on Fundamental Approaches to Software Engineering: Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2011.
URL http://alfa.di.uminho.pt/ jacome/down/fase11.pdf

Spreadsheets are notoriously error-prone. To help avoid the introduction of errors when changing spreadsheets, models that capture the structure and interdependencies of spreadsheets at a conceptual level have been proposed. Thus, spreadsheet evolution can be made safe within the confines of a model. As in any other model/instance setting, evolution may not only require changes at the instance level but also at the model level. When model changes are required, the safety of instance evolution can not be guarded by the model alone.

We have designed an appropriate representation of spreadsheet models, including the fundamental notions of formula and references. For these models and their instances, we have designed coupled transformation rules that cover specific spreadsheet evolution steps, such as the insertion of columns in all occurrences of a repeated block of cells. Each model-level transformation rule is coupled with instance level migration rules from the source to the target model and vice versa. These coupled rules can be composed to create compound transformations at the model level inducing compound transformations at the instance level. This approach guarantees safe evolution of spreadsheets even when models change.

4.5 From State- to Delta-Based Bidirectional Model Transformations: Unweaving Alignment and Update Propagation

Zinovy Diskin (University of Waterloo, CA)

Joint work of Diskin, Zinovy; Xiong, Yingfei; Czarnecki, Krzysztof
Main reference Diskin, Zinovy; Xiong, Yingfei; Czarnecki, Krzysztof. From State- to Delta-Based Bidirectional Model Transformations. Theory and Practice of Model Transformations, Third International Conference, ICMT 2010. Proceedings
URL http://dx.doi.org/10.1007/978-3-642-13688-7

Existing bidirectional model transformation languages are mainly state- based: model alignment is hidden inside update propagating procedures, and model deltas are implicit. Weaving alignment with update propagation complicates the latter and makes it less predictable and less manageable.

We propose to separate concerns and consider two distinct operations: delta discovery and delta propagation. This architecture has several technological advantages, but requires a corresponding theoretical support.

We present an algebraic framework of delta lenses, and discuss an emerging landscape of delta-based model synchronization. An essentially updated version of our ICMT paper is submitted to JOT.

4.6 Propagation of Constraints along Model Transformations

Hartmut Ehrig (TU Berlin, DE)

Main reference Hartmut Ehrig, Frank Hermann, Hanna Schölzel and Christoph Brandt: Propagation of Constraints along Model Transformations Based on Triple Graph Grammars. In: Proc. Graph Transformation and Visual Modeling Techniques (GT-VMT 2011), EC-EASST (To appear 2011).

Within the approach of model transformations based on TGGs, it is an interesting problem to transform not only source language to target language models, but to transform also properties of models given by graph constraints. This technique is called propagation of constraints and should satisfy the following property: If the source model satisfies the source constraint, then also the target model should satisfy the propagated target constraint. First investigations show that it is useful to construct first a propagated integrated constraint, which is satisfied by the integrated model obtained as an intermediate step of the model transformation.

4.7 HOT Topics in Bidirectional Programming

Jeremy Gibbons (University of Oxford, GB)

URL http://www.dagstuhl.de/mat/Files/11/11031/11031.GibbonsJeremy1.Slides.pdf

There is a lot of interesting structure in bidirectional programming problems - structure that should be exploited in solving those problems. The kinds of structure I have in mind are higher-order and typed techniques from PL theory - what Bob Harper called "HOT" topics; Janis Voigtländer's work on exploiting parametricity is a good example. Other techniques that look promising are higher-order representations, indexed- and dependently-typed programming, and relational algebra. Perhaps it is evidence of my own ignorance, but it seems to me that this rich structure is not currently being exploited to the full. I will explain why I think HOT techniques show promise, and why I hope that they will lead to a more semantic rather than syntactic approach to bidirectional programming.

4.8 Lenses are Coalgebras for the Costate Comonad

Jeremy Gibbons (University of Oxford, GB)

URL http://www.dagstuhl.de/mat/Files/11/11031/11031.GibbonsJeremy2.ExtAbstract.pdf

The three laws of a "very well-behaved lens" are exactly the condition that the lens (that is, the pairing of its get and put functions) is a coalgebra for the costate comonad. This is my explanation of an observation due to Russell O'Connor.

4.9 Lenses, Coalgebraically

Jeremy Gibbons (University of Oxford, GB)

Main reference Submitted for publication
URL http://web.comlab.ox.ac.uk/publications/publication4599-abstract.html

Lenses are a heavily studied form of bidirectional transformation with diverse applications including database view updating, software development and memory management. Recent work has explored lenses category theoretically, and established that the category of lenses for a fixed "view" V is, up to isomorphism, the category of algebras for a particular monad on Set/V. In this paper we show that in addition lenses are the coalgebras for the comonad generated by the cartesian closure adjunction on Set. We present a fully constructive proof of the coalgebra correspondence, we note that the algebra correspondence extends to arbitrary categories with products and that the coalgebra correspondence extends to arbitrary cartesian closed categories, and we show that both correspondences extend to isomorphisms of categories. The resulting isomorphism between a category of algebras and a category of coalgebras is unexpected, and we analyze it isolating its underlying generality, and also the particularity that restricts its applicability. We end with remarks about the utility of the two different treatments of lenses,especially for obtaining further, more realistic, generalizations of the notion of lens.

4.10 Direction Neutral Language Transformation with Metamodels

Martin Gogolla (Universität Bremen, DE)

The aim of this work is to sketch a general metamodel-based frame for describing potentially bidirectional transformations between software languages. We propose to describe a single language with a metamodel consisting of a UML class diagram with classes, attributes and associations and accompanying OCL constraints. A language description is separated into a syntax and a semantics part. The allowed object diagrams of the syntax part correspond to the syntactically allowed words of the language. The semantics can associate with every word of the language certain semantical objects. However, only finite fractions of the semantics can be handled in metamodel tools. Having two languages described by their metamodels, a transformation between them is established by another metamodel, a transformation model. The transformation model can associate a syntax object from one language with a syntax object from the other language in a direction neutral way and opens the possibility for bidirectionality. Analogously, semantical objects can be connected. Transformation properties like ‘equivalence’ or ‘embedding’ can and must be expressed with constraints. Thus, the approach describes syntax and semantics of the languages, their transformation and their properties in a uniform way by means of metamodels.

4.11 Unified (Bidirectional) Transformation Language

Joel Greenyer (Universität Paderborn, DE)

Main reference J. Greenyer and E. Kindler: Comparing relational model transformation technologies: implementing Query/View/Transformation with Triple Graph Grammars. Software and Systems Modeling (SoSyM) 2010, Vol. 9, 21–46, 1
URL http://dx.doi.org/10.1007/s10270-009-0121-8

In the recent past, the declarative, relational transformation languages QVT-Relations (part of the OMG's Query/View/Transformation standard) and TGGs (Triple Graph Grammars) have gained in popularity in the domain of model-based software engineering. That is partly because many transformation problems can be described very conveniently with these languages, and because, in principle, the same set of transformation rules can be interpreted both in a forward and backward transformation direction. Even though the principles of QVT-Relations and TGGs are different (specifying relations between model patterns vs. rules for producing valid corresponding graphs/models), the languages are very similar [GK10]. However, neither TGGs nor QVT-Relations have yet found a broad application in industry. Over the past years, we have developed a transformation tool for TGGs, called the TGG-Interpreter (http://www.cs.uni-paderborn.de/index.php?id=tgg-interpreter&L=1), and we have gained experience in a wide variety of transformation examples. From this experience, we identified some requirements that are crucial for a successful application of TGGs (and also QVT-Relations) in practice. These requirements ranges from techniques for analyzing for example the confluence and bi-directionality of a TGG to support for testing, debugging, to an extensible and maintainable transformation engine architecture. Especially, we feel that OCL, as it is used in QVT-Relations as well as TGGs in the TGG-Interpreter today, is insufficient for specifying attribute constraints in bi-directional transformations. Therefore, it should be investigated how to integrate TGGs (and QVT-Relations) with other bidirectional model transformation approaches, for example reversible programming, in the future.

4.12 Model Integration and Synchronization

Frank Hermann (TU Berlin, DE)

Main reference Hartmut Ehrig, Karsten Ehrig, Frank Hermann: From Model Transformation to Model Integration based on the Algebraic Approach to Triple Graph Grammars. ECEASST Vol. 10, EASST 2008.
URL http://journal.ub.tu-berlin.de/index.php/eceasst/article/view/154

Within the approach of model transformations based on TGGs, model integration and synchronization are important problems. Let us call a source and a target model consistent, if there exists an integrated model generated by a given TGG, such that the source and target components of the integrated model are equal to the given models. The integration problem is the following: Find out whether, for a given source and target model, both models are consistent and construct an integrated model in case of consistency. If both models are not consistent, the synchronization problem is to modify the source and/or the target model, such that both become consistent. If the modification of both models is allowed, then the problem is to find a synchronization, such that the consistent modified models can be constructed from the given ones with a "minimal number of changes". In contrast to several

approaches in the literature, we are not so much interested in efficient heuristic, but more in correct formal solutions.

4.13 Bidirectional Graph Transformations based on Structural Recursion

Soichiro Hidaka (NII - Tokyo, JP)

Joint work of Hu, Zhenjiang; Inaba, Kazuhiro; Kato, Hiroyuki; Matsuda, Kazutaka; Nakano, Keisuke;
Main reference Soichiro Hidaka, Zhenjiang Hu, Kazuhiro Inaba, Hiroyuki Kato, Kazutaka Matsuda, and Keisuke Nakano. 2010. Bidirectionalizing graph transformations. In Proceedings of the 15th ACM SIGPLAN international conference on Functional programming (ICFP '10). ACM, New York, NY, USA, 205-216.
URL http://doi.acm.org/10.1145/1863543.1863573

We have been challenging bidirectional transformations problems within the context of graphs, by proposing a formal definition of a well-behaved bidirectional semantics for UnCAL, i.e., a graph algebra for the known UnQL graph query language. We utilize both the recursive and bulk semantics of structural recursion on graphs.

Existing forward evaluation of structural recursion has been refined so that it can produce sufficient trace information for later backward evaluation. In this position statement I will introduce our framework of bidirectional graph transformations and how they are implemented based on structural recursion with trace information, followed by optimization opportunities and remaining challenges.

4.14 Incremental Bidirectional Model Synchronization

Stephan Hildebrandt (Hasso Plattner Institut - Potsdam, DE)

Main reference Giese, Holger and Neumann, Stefan and Hildebrandt, Stephan; Model Synchronization at Work: Keeping SysML and AUTOSAR Models Consistent; Graph Transformations and Model Driven Enginering - Essays Dedicated to Manfred Nagl on the Occasion of his 65th Birthday; LNCS Volume 5765; Pages 555-579
URL http://dx.doi.org/10.1007/978-3-642-17322-6_24

Pure model transformations are not sufficient to cope with concurrent changes, different information details in different modeling languages, and increasing model sizes in practical software engineering. Therefore, model synchronization technologies are required that only propagate changes between models. This avoids overwriting additional information in target models and can be executed faster than a complete model transformation. This talk presents MoTE, an incremental bidirectional model synchronization system based on triple graph grammars. Two case studies are presented, where the system has been successfully put to practice. Some open issues are discussed that were encountered in the case studies.

4.15 Symmetric lenses

Martin Hofmann (LMU München, DE)

Main reference Martin Hofmann, Benjamin C. Pierce, Daniel Wagner: Symmetric lenses. Proc. ACM Symp. POPL 2011: 371-384, ACM Press, 2011.

Symmetric lenses have been proposed by Pierce, Wagner, and myself as a completely symmetric generalisation of lenses which in turn are a category-theoretic abstraction of a pair of a view extraction and update functions. This position statement gives a taste of definition, basic category-theoretic properties, datatypes and iterators, as well as a representation theorem. Details in our paper at POPL2011.

4.16 Triple Graph Grammars: Concepts, Extensions, Implementations, and Application Scenarios

Ekkart Kindler (Technical University of Denmark, DK)

Main reference Technical Report tr-ri-07-284, Software Engineering Group, Department of Computer Science, University of Paderborn, June 2007.
URL http://www2.cs.uni-paderborn.de/cs/ag-schaefer/Veroeffentlichungen/Quellen/Papers/2007/tr-ri-07-284.pdf

Triple Graph Grammars (TGGs) are a technique for defining the correspondence between two different types of models in a declarative way. The power of TGGs comes from the fact that the relation between the two models can not only be defined, but the definition can be made operational so that one model can be transformed into the other in either direction; even more, TGGs can be used to synchronize and to maintain the correspondence of the two models, even if both of them are changed independently of each other; i. e., TGGs work incrementally.

TGGs have been introduced more than 10 years ago by Andy Schürr.

Since that time, there have been many different applications of TGGs for transforming models and for maintaining the correspondence between these models. To date, there have been several modifications, generalizations, extensions, and variations. Moreover, there are different approaches for implementing the actual transformations and synchronizations of models.

In this paper, we present the essential concepts of TGGs, their spirit, their purpose, and their fields of application. We also discuss some of the extensions along with some of the inherent design decisions, as well as their benefits and caveats. All these are based on several year's of experience of using TGGs in different projects in different application areas.

4.17 Some challenges of integrating bidirectional transformation technologies

Ekkart Kindler (Technical University of Denmark, DK)

There are different kinds of uni- and bi-directional transformation mechanisms, which are based on different philosophies, paradigms, and principles. In concrete situations, one often would like to combine different technologies in order to achieve certain goals; but this is not possible for fundamental principal differences, some conceptual differences, or simply for technical reasons. This position statement heads at identifying some issues and triggering the work on a clear conceptual framework and interfaces that would allow comparing and integrating different transformation technologies.

4.18 Towards Systematic Development of Bidirectional Transformations

Jochen M. Kuester (IBM Research Zürich - Rüschlikon, CH)

In order to adopt bidirectional model transformations in industry, an approach for their systematic development is required. Such an approach includes methods and tools for requirements specification, design and implementation of bidirectional model transformations.

In this presentation, we first report on lessons learnt from the development of a solution based on unidirectional model transformations and then establish several requirements for systematic development of bidirectional model transformations.

4.19 Right Inverses in Bidirectionalization

Kazutaka Matsuda (Tohoku University, JP)

Bidirectionalization is a program transformation that derives a bidirectional transformation from a unidirectional transformation automatically or manually.

We emphasize on the importance of right-inverse construction in syntactic bidirectionalization. Right inverses are useful not only for construction of so-called "create" functions but also for construction of bidirectional transformations whose behavior can be controlled by users.

4.20 Bidirectional transformations and inter-modelling

Richard F. Paige (University of York, GB)

Main reference Esther Guerra, Juan de Lara, Dimitrios S. Kolovos, Richard F. Paige: Inter-modelling: From Theory to Practice. MoDELS (1) 2010: 376-391, LNCS, Springer-Verlag.
URL http://dx.doi.org/10.1007/978-3-642-16145-2_26

We have been developing, implementing and deploying model management tools for a number of years. Model management focuses on manipulating models once they have been constructed. Typical model management operations include model transformation (including model-to-model, model-to-text, text-to-model and update-in-place transformations), model merging, model comparison, model validation (e.g., inter-model consistency checking), and model migration. We have provided tool support for these operations in the Epsilon model management framework , which provides a suite of interoperable task-specific languages. We have also used this framework and these languages in a large number of industrial applications of MDE; some of these applications involved applying several of the languages in concert. As a result, we have improved our framework, increased our understanding of what makes MDE challenging and tractable, and what are some of the obstacles impeding its uptake.

A number of our applications of MDE have involved model transformation (of various flavours), and some of these applications have either included, or generated, requirements for some kind of support for bidirectionality. Bidirectionality, when it appeared in requirements, was usually associated with specific use cases; moreover, when investigating the requirement more carefully, it was often the case that a bidirectional transformation was only one of a number of possible ways in which the requirement could be supported.

As a result of our work on practical applications of MDE, we have determined that we have more questions than answers about bidirectional transformation. Some of these questions are as follows.

- What is the problem engineers are trying to solve with a bidirectional transformation (bx)?
- When is a bx a good solution to a problem, and when is it a bad solution?
- More specifically, what is the sweet spot in choosing whether to use a bx or to use more traditional MDE approaches (unidirectional transformation, inter-model consistency, etc)?
- What are scenarios in which bx applies?
- What patterns appear in a bx?
- Are there sufficient recurring patterns in a bx that suggests that we need dedicated languages for programming bidirectional transformations?

We have been thinking, initially, about the first question: what problem are engineers trying to solve? Fundamental in MDE is the notion of inter-modelling: you build different models (some of which are of languages or domains, i.e., metamodels) that are inter-related in various ways. The inter-relationships between models are interesting and diverse and come in many different flavours, but they are typically implemented as trace-links (or traceability links). Some example inter-model relationships include: simple dependencies, code generation, consistency, bx, synchronisation, etc. It is possible to define very strange inter-model relationships, including security and privacy relationships, for example: users with a particular role cannot see both X and Y. It is also possible to define very complicated inter-model relationships, e.g., between a PDF document and an EJB implementation where parts of the PDF document define concepts that are designed and implemented in the EJBs. This latter relationship is not really something that is a transformation problem, nor is it

really a consistency problem, though conceivably parts of the problem could be captured in suitable transformation and consistency tools. Overall, we think that the richness of inter-modelling relationships (and types of trace-links that can implement these relationships) illustrates the challenges that we have in understanding where bx is really useful and usable.

It may be helpful to have a unified conceptual model of MDE like that of inter-modelling, in order to better understand where bx are useful, where (for example) approaches like inter-model consistency checking are suitable, and what benefits can be obtained through bidirectional transformations directly. At the moment, it seems that there is much confusion in the bx-MDE field, possibly because we are not yet clear what are the different kinds of relationships there are that we want to construct between our models, and how they should best be constructed.

4.21 Bidirectional and change propagating transformations in MDE

Alfonso Pierantonio (Univ. degli Studi di L'Aquila, IT)

Main reference A. Cicchetti, D. Di Ruscio, R. Eramo and A. Pierantonio, JTL: a bidirectional and change propagating transformation language, 3rd International Conference on Software Language Engineering (SLE 2010), Eindhoven (The Netherlands)
URL http://dx.doi.org/10.1007/978-3-642-19440-5_11

In Model Driven Engineering bidirectional transformations are considered a core ingredient for managing both the consistency and synchronization of two or more related models. However, while non-bijectivity in bidirectional transformations is considered relevant, current languages still lack of a common understanding of its semantic implications hampering their applicability in practice. In this work, the Janus Transformation Language (JTL) is presented, a bidirectional model transformation language specifically designed to support non bijective transformations and change propagation. In particular, the language propagates changes occurring in a model to one or more related models according to the specified transformation regardless of the transformation direction. Additionally, whenever manual modifications let a model be non reachable anymore by a transformation, the closest model which approximate the ideal source one is inferred. The language semantics is also presented and its expressivity and applicability are validated against a reference benchmark. JTL is embedded in a framework available on the Eclipse platform which aims to facilitate the use of the approach, especially in the definition of model transformations.

4.22 Reversible Higher Order Pi Calculus

Alan Schmitt (INRIA Rhône-Alpes, FR)

Reversible HOPi is a higher-order process calculus where every action can be undone. We are developing this calculus as a foundation for dependable distributed application. We will give an intuition as how this approach is different from sequential reversible languages, and how it may be extended with controlled rollback to better control when undoing occurs.

4.23 Efficiency of Bidirectional Transformations

Janis Voigtländer (Universität Bonn, DE)

I raise some questions related to the efficiency of programs coming out of linguistic approaches to bidirectional transformation. Do we even care (yet) about the efficiency of these programs, or are we still at a point where we have to struggle too much even to get their semantic behaviour right? If we do care about efficiency, what are the criteria/setups to measure? Can we hope to rely on standard program transformations (from the functional languages community) to improve efficiency, or do we have to invent new techniques tailored to bidirectionalization or specific DSLs?

The slides used are here:
http://www.dagstuhl.de/mat/Files/11/11031/11031.VoigtlaenderJanis.Slides.pdf.

4.24 Change-based incremental updates

Meng Wang (University of Oxford, GB)

Joint work of Wang, Meng; Gibbons, Jeremy

To handle large data that are subject to relatively small modifications, incrementality of updates is the key. Incrementality has been explored in the settings of relational databases and graph transformations, where the data is typically not typed (in the sense not restricted to a particular shape). This flexibility in structure makes it relatively easy to divide the data into separate parts that can be transformed and updated independently. The same is not true if the data is to be encoded with more general purpose algebraic datatypes, with transformations defined as functions: dividing data into well-typed separate parts is tricky, and recursions typically create interdependencies. In this work, we look at the identification of transformations that support incremental updates, and device a constructive process to achieve it.

4.25 Bx 'killer applications' in health care

Jens-Holger Weber-Jahnke (University of Victoria, CA)

Nations around the world are creating infrastructures for maintaining electronic health records (EHRs) at massive scale. These infrastructures need to translate between a large number of heterogenous data bases. The richness and sensitivity of medical data provide new challenges and opportunities for the bx transformation research community. Current transformation technologies applied in the eHealth sector are immature and lack fundamentally important properties. This results not only in cost overruns of EHR infrastructure projects, but also jeopardizes patient safety. My presentation has introduced typical variations of EHR

infrastructures and motivated fundamental required properties of bx approaches suitable for this domain.

4.26 Fix Generation

Yingfei Xiong (University of Waterloo, CA)

Bidirectional transformation maintains consistency between two pieces of data automatically. However, consistency cannot always be maintained automatically, and existing tools present users with a list of choices of modifying data, called "fixes". Fixes are widely used in many areas, but there is little support for fix generation in general. I will show that fix generation can be formalized in a similar way to BX, as two functions and laws governing their behavior, and ideas in BX could be used as a starting point to develop general support for fix generation.

4.27 A reversible programming language

Tetsuo Yokoyama (Nanzan University, JP)

Reversible programming languages, which I focus on, are imposed restrictions in a way that no program loses information and thus the reconstruction of the previous states is always possible. In common with other programming paradigms, reversible programming has its own programming methodology. We introduced Janus, a high-level reversible language, and discussed the features of the languages such as cleanliness, structured programming and r-Turing completeness, and a guideline for the systematic development of efficient reversible programming.

5 Overview of Working Groups

5.1 Model Consistency Management in Software Engineering

Krzysztof Czarnecki (University of Waterloo, CA)

Joint work of Czarnecki, Krzysztof; Diskin, Zinovy; Antkiewicz, Michal; Xiong, Yingfei
URL http://gsd.uwaterloo.ca/publications

Model-driven engineering often requires many overlapping models of a system, each supporting a particular kind of stakeholder or task. The consistency among these models needs to be managed during system development. Consistency management occurs in the space of multiple model replicas, versions over time, and different modeling languages, which make the process complex and challenging. In this tutorial, I will give an overview of the main consistency management scenarios, such as involving single model vs. N models, models

in the same language or in different languages, and precise automated alignment or semi-automated one, introduce the key concepts, and establish some basic terminology of model synchronization.

5.2 Relationships between BX and View Updates

Soichiro Hidaka (NII - Tokyo, JP)

In bidirectionalization, which derives backward transformation from forward transformation, different forward transformations that can produce identical result may lead to different acceptance of the modifications on the view. For example, in the author's framework of bidirectional graph transformation, variable references to the entire source can accept any modification while copying transformation by replacing edge by edge with pattern matching, accepts limited editing operation in case of edge renaming.

In this working group we discussed the impact of the description of forward transformation on the accepted view updates using concrete example like a duplication function. Elements that are likely to make backward transformation harder are enumerated.

In summary, tight coupling of trace information and syntactic definition of the transformation is considered to be the source of the problem. Reflecting research subjects of the participants, concepts that play a role similar to trace like complement and postcondition had also been discussed. Decoupling, abstract postconditions or related concepts like contracts, could be investigated as a future direction.

Participants: Robert Glück, Soichiro Hidaka, Michael Johnson, Kazutaka Matsuda, Meng Wang, Tetsuo Yokoyama

5.3 Toward a Bidirectional Transformation Benchmark and Taxonomy

James Terwilliger (Microsoft Corporation - Redmond, US)

Joint work of Terwilliger, James; Kindler, Ekkart

The benchmark and taxonomy working group gathered for the purpose of documenting commonality across disciplines. The group heard from representatives across all four disciplines present at the seminar. The need for such a benchmark became abundantly clear over the course of the discussion as the participants realized that, despite working towards similar and sometimes identical research goals, the words used to describe concepts in those disciplines were vastly different. For instance, the word "model" in one discipline corresponds to "meta-model" in another discipline, and to "instance" in yet another discipline. In the short time that the group was gathered, the group primarily focused on documenting these differences in language rather than resolving them. In addition, the group also documented the various settings in which bidirectional transformations occur, the problems that Bx solutions are intending to solve, and the set of solutions currently available.

Participants

- Anthony Anjorin
TU Darmstadt, DE
- Artur Boronat
University of Leicester, GB
- Christoph Brandt
University of Luxembourg, LU
- Anthony Cleve
University of Namur, BE
- Jácome Cunha
Univ. de Minho - Braga, PT
- Krzysztof Czarnecki
University of Waterloo, CA
- Zinovy Diskin
University of Waterloo, CA
- Hartmut Ehrig
TU Berlin, DE
- Nate Foster
Cornell University, US
- Jeremy Gibbons
University of Oxford, GB
- Robert Glück
University of Copenhagen, DK
- Martin Gogolla
Universität Bremen, DE
- Joel Greenyer
Universität Paderborn, DE
- Jean-Luc Hainaut
University of Namur, BE
- Frank Hermann
TU Berlin, DE
- Soichiro Hidaka
NII - Tokyo, JP
- Stephan Hildebrandt
Hasso Plattner Institut - Potsdam, DE
- Martin Hofmann
LMU München, DE
- Zhenjiang Hu
NII - Tokyo, JP
- Michael Johnson
Macquarie Univ. - Sydney, AU
- Ekkart Kindler
Technical Univ. of Denmark, DK
- Jochen M. Küster
IBM Research - Zürich, CH
- Ralf Lämmel
Universität Koblenz-Landau, DE
- Marius Lauder
TU Darmstadt, DE
- Kazutaka Matsuda
Tohoku University, JP
- Richard F. Paige
University of York, GB
- Alfonso Pierantonio
Univ. degli Studi di L'Aquila, IT
- Benjamin C. Pierce
University of Pennsylvania, US
- Jan Rieke
Universität Paderborn, DE
- Michael Schlereth
Wilhermsdorf, DE
- Alan Schmitt
INRIA Rhône-Alpes, FR
- Andy Schürr
TU Darmstadt, DE
- Perdita Stevens
University of Edinburgh, GB
- James Terwilliger
Microsoft Res. - Redmond, US
- Janis Voigtländer
Universität Bonn, DE
- Meng Wang
University of Oxford, GB
- Andrzej Wasowski
IT Univ. of Copenhagen, DK
- Jens-Holger Weber-Jahnke
University of Victoria, CA
- Yingfei Xiong
University of Waterloo, CA
- Tetsuo Yokoyama
Nanzan University, JP

Report from Dagstuhl Seminar 11041

Multimodal Music Processing

Edited by

Meinard Müller[1], Masataka Goto[2], and Simon Dixon[3]

1 **Saarland University and MPI Informatik - Saarbrücken, DE,** meinard@mpi-inf.mpg.de
2 **AIST – Ibaraki, JP,** m.goto@aist.go.jp
3 **Queen Mary University of London, GB,** simon.dixon@eecs.qmul.ac.uk

Abstract

From January 23 to January 28, 2011, the Dagstuhl Seminar 11041 "Multimodal Music Processing" was held at Schloss Dagstuhl – Leibniz Center for Informatics. During the seminar, we discussed various aspects of the automated processing of music-related documents. These documents may describe a musical work in different ways comprising visual representations (e. g., sheet music), symbolic representations (e. g., MIDI, tablatures, chords), acoustic representations (CD recordings), audio-visual representations (videos), or text-based metadata. In this report, we give an overview of the main contributions and results of the seminar. We start with an executive summary, which describes the main topics, goals, and group activities. Then one finds a list of abstracts giving a more detailed overview of the participants' contributions as well as of the ideas and results discussed in the group meetings and panels of our seminar.

Seminar 23.–28. January, 2011 – http://www.dagstuhl.de/11041
1998 ACM Subject Classification H.5.5 Sound and Music Computing, I.5.4 Applications—Waveform analysis, H.5.1 Multimedia Information Systems, J.5 Arts and Humanities
Keywords and phrases Music information retrieval, music processing, multimodality, audio, sheet music, content-based analysis, signal processing, user interaction
Digital Object Identifier 10.4230/DagRep.1.1.68
Edited in cooperation with Sebastian Ewert

1 Executive Summary

Meinard Müller
Masataka Goto
Simon Dixon

Music can be described, represented, and experienced in various ways and forms. For example, music can be described in textual form not only supplying information on composers, musicians, specific performances, or song lyrics, but also offering detailed descriptions of structural, harmonic, melodic, and rhythmic aspects. Furthermore, music notation can be encoded in text-based formats such as MusicXML, or symbolic formats such as MIDI. Music annotations, metadata, and social tags are also widely available. Beside textual data, increasingly more types of music-related multimedia data such as audio, image or video data are widely available. For example, there are numerous MP3 and CD audio recordings, digitized images of scanned sheet music, and an increasing number of video clips of music performances. In this seminar we discussed and studied various aspects of the

automated processing of music-related documents that vary in their modalities and formats, i.e., text, symbolic data, audio, image and video. Among others, the topics dealt with aspects of data analysis, retrieval, navigation, classification, and generation of music documents while considering music-specific characteristics, exploiting multiple information sources, and accounting for user-specific needs.

In this executive summary, we give a brief overview of the main topics addressed in this seminar. We start by briefly describing the background of the seminar participants and the overall organization. We then give an overview of the presentations and the results from working groups and panels. Finally, we reflect on the most important aspects of this seminar and conclude with future implications.

Participants, Interaction, Activities

In our seminar, we had 35 participants, who came from various countries around the world including North America (7 participants), Japan (4 participants), New Zealand, Singapore, and Europe (Austria, France, Germany, Netherlands, Portugal, Spain, United Kingdom). Most of the participants came to Dagstuhl for the first time and expressed enthusiasm about the open and retreat-like atmosphere. Besides its international character, the seminar was also highly interdisciplinary. While most of the participating researchers are working in computer science and its neighboring fields, a large number of participants also have a strong background in music and musicology. This made the seminar very special in having not only interactive and provoking scientific discussions, but also numerous social activities including common music making. One particular highlight of such social activities was a three-hour spontaneous concert on Thursday evening, where various participants performed in changing ensembles a wide variety of music including popular music, jazz, and classical music.

Overall Organization and Schedule

Dagstuhl seminars are known for having a high degree of flexibility and interactivity, which allow participants to discuss ideas and to raise questions rather than to present research results. Following this tradition, we fixed the schedule during the seminar asking for spontaneous contributions with future-oriented content, thus avoiding a conference-like atmosphere, where the focus is on past research achievements. The first day was used to let people introduce themselves and express their expectations and wishes for the seminar. We then had a brainstorming session on central topics covering the participants' interests while discussing the overall schedule and format of our seminar. In particular, we identified a total of six topics for discussion. For four of these topics, we divided into four groups, each group discussing one of the topics in greater depth in parallel sessions on Tuesday. The results and conclusions of these group meetings were then presented to the plenum on Wednesday. For the remaining two topics, we decided on having panel-like discussions within the plenum with introductory stimulus talks (Thursday). Finally, group and panel discussions were interleaved with regular sessions that allowed participants to present their personal research to the plenum. This mixture of presentation elements gave all participants the opportunity for presenting their ideas to the plenum while avoiding a monotonous conference-like presentation format.

Main Topics and Results

We discussed various topics that address the challenges of organizing, understanding, and searching music-related information in a robust, efficient, and intelligent manner. Here, a particular focus was put on music-specific aspects, the fusion of multiple information sources, as well as the consideration of user-specific needs. After a joint brainstorming session, we agreed on discussing six central topics which fitted in the overall theme of the seminar and reflected the participants' interests. We now give a brief summary of these topics, which were discussed within four parallel group meetings and two panels. Then, we give an overview of additional contributions made by the participants in the regular sessions of the seminar.

1. The *"Model"* group discussed the issue of how signal models can be developed that exploit multimodal information. Here, one main goal was to review strategies for combining different sources of information to support music analysis. In particular, various early and late fusion approaches were identified and advantages and weaknesses of the respective approaches were discussed. Particular attention was paid to the aspect of data uncertainty and its propagation in the fusion processes.
2. The *"User"* group addressed the topic of user-aware music information retrieval. Here, a central question was how contextual information can be integrated into the retrieval process in order to account for short-term user interests and long-term user behavior. Additionally, it was discussed how search engines may yield satisfying results in terms of novelty, popularity, and serendipity of the retrieved items.
3. The *"Symbol"* group discussed the question of how to bridge the gap between visual, symbolic, and acoustic representations of music. Here, particular attention was given to the problem referred to as *Optical Music Recognition* (OMR) with the goal of converting an image-based sheet music representation into a symbolic music representation where note events are encoded explicitly. In this context, user interfaces were reviewed that allow for a synchronized presentation of visual and acoustic music content.
4. The *"Meaning"* group addressed the subject of how musical meaning can be derived from musical data and, in particular, from musical sound. Here, the path from the given low-level (acoustic) raw data to high-level musical models was traced from a human-based as well as from a machine-based perspective.
5. In the *"Ground Truth"* panel, fundamental issues related to the interpretation, usage, and generation of ground truth annotations were discussed. Some of the questions raised during the panel were: What is ground truth in music? How can one handle inter- and intra-annotator variations? How can the quality of ground truth be evaluated? Are there alternatives to manually generated ground truth annotation?
6. Finally, in the *"Grand Challenges"* panel we discussed in which way music information research may and should impact our daily lives and our society in the future. Here fundamental questions were how to provide the best music for each person, how to predict specific effects of music on our society, and how to enrich human relationships by music.

Beside the extensive discussion of these six topics, we had a number of additional contributions where participants presented more specific research results. These contributions covered a number of different topics such as audio parameterization, music alignment and synchronization, singing voice processing, crowd music listening, music tagging, music indexing, interfaces for music exercises, personalization issues in music search, analysis of ethnic music, and many more.

These topics were complemented by some more interdisciplinary contributions relating the field of music processing to neighboring fields such as speech processing, musicology, music

perception, and information retrieval. For example, we discussed the ways in which the field of music processing has benefitted from older fields such as speech processing and how music processing might give something back to these fields. Furthermore, a musicologist reported on the difficulties and resistance experienced when introducing novel computer-based methods into traditional humanistic sciences such as musicology. Another highlight of our seminar was a keynote presentation given by Hannah Bast on her CompleteSearch Engine that allows for very fast processing of complex queries on large text collections.

Conclusions

In our seminar, we addressed central and groundbreaking issues on how to process music material given in various forms corresponding to different musical aspects and modalities. In view of the richness and complexity of music, there will be no single strategy that can cope with all facets equally well. Therefore unifying frameworks and fusion approaches are needed which allow for combining, integrating, and fusing the various types of information sources to support music analysis and retrieval applications. Also, to further enhance our field, one needs to understand better the complex relationships within music as well as the complex effects of music on the human mind, thus requiring interdisciplinary research efforts. The Dagstuhl seminar gave us the opportunity for discussing such issues in an inspiring and retreat-like atmosphere. The generation of novel, technically oriented scientific contributions was not the focus of the seminar. Naturally, many of the contributions and discussions were on a rather abstract level, laying the groundwork for future projects and collaborations. Thus the main impact of the seminar is likely to take place in the medium to long term. Some more immediate results, such as plans to share research data and software, also arose from the discussions. As measurable outputs from the seminar, we expect to see several joint papers and applications for funding (e.g. to the European Union) proceeding from the discussions held at Dagstuhl.

Beside the scientific aspect, the social aspect of our seminar was just as important. We had an interdisciplinary, international, and very interactive group of researchers, consisting of leaders and future leaders in our field. Most of our participants visited Dagstuhl for the first time and enthusiastically praised the open and inspiring atmosphere. The group dynamics were excellent with many personal exchanges and common activities. Younger scientists mentioned their appreciation of the opportunity for prolonged discussions with senior researchers—something which is often impossible during conference-like events.

In conclusion, our expectations of the seminar were not only met but exceeded, in particular with respect to networking and community building. Last but not least, we heartily thank the Dagstuhl board for allowing us to organize this seminar, the Dagstuhl office for their great support in the organization process, and the entire Dagstuhl staff for their excellent services during the seminar.

2 Table of Contents

3 Overview of Talks

3.1 Multimodal Processing in a Digital Music Library Context

David Bainbridge (University of Waikato, NZ)

In the digital library research group at the University of Waikato we have been (or are in the process of) working with several aspects of multimodal music processing to assist in the formation of digital libraries, and the augmentation of the services they offer. In my presentation, I have addressed the following issues.

- A digital music stand that is fully immersed in the digital library, and whose functionality is enhanced through automated content processing.
- A spatial-hypermedia approach to organizing and managing "musical moments" that occur as part of the process of composition.
- Using digital library collections that combine linked-data with music analysis to capture (and ultimately publish) the outcome of MIR experiments.

3.2 The CompleteSearch Engine

Hannah Bast (Universität Freiburg, DE)

Joint work of Bast, Hannah; Weber, Ingmar
Main reference Hannah Bast and Ingmar Weber: "The CompleteSearch Engine: Interactive, Efficient, and Towards IR& DB Integration," CIDR 2007, pp. 88–95.

CompleteSearch is a new search engine technology for very fast processing of complex queries on large text collections. Supported query types are: prefix search, faceted search, error-tolerant search, synonym search, database-like search, and semantic search. CompleteSearch is highly interactive: hits are displayed instantly after each keystroke, along with suggestions for possible continuations / refinements of the query. In my talk, I showed exciting demos of the various features of CompleteSearch and also said a few words about their realization.

3.3 What can fMRI tell us about MIR?

Michael Casey (Dartmouth College - Hanover, US)

Music representation requires making decisions about what information to extract, or declare. We outline details of a set of studies that are designed to test the predictive power of different music representations via-a-vis human neural responses in fMRI experiments. Our methods combine audio music feature extraction, latent variable analysis, linguistic and symbolic representations, and multi-variate pattern analysis on neural population codes. To this end, we designed and implemented a Python software framework, named Bregman, built upon OMRAS2 AudioDB C++ framework, that facilitates the design of experiments combining

MIR and fMRI methods. The software is currently being used by a group of 19 graduate and undergraduate students studying music, information, and neuroscience.

3.4 Ground Truth: Where Aggregates Fail

Elaine Chew (USC - Los Angeles, US)

Differences in composer vs. listener annotations of music structure in improvised performances with the Mimi system was shown, and presented evidence for individual emotion perception variability. The Mimi structure analysis project is joint work with Alexandre François, Isaac Schankler, and Jordan Smith; the latter project is a Master's thesis by Merrick Mosst; both are work done at the MuCoaCo lab at USC.

3.5 Towards Automated Processing of Multimodal Representations of Music

Michael Clausen (Universität Bonn, DE)

Joint work of Clausen, Michael; Fremerey, Christian; Müller, Meinard
Main reference David Damm, Christian Fremerey, Frank Kurth, Meinard Müller, and Michael Clausen: "Multimodal presentation and browsing of music," Proc. 10th International Conference on Multimodal Interfaces (ICMI), pp. 205–208, 2008.

There are many different types of digital music representations that capture aspects of music on various different levels of abstraction. For example, a piece of music can be described visually by means of sheet music that encodes abstract high-level parameters such as notes, keys, measures, or repeats in a visual form. Because of its explicitness and compactness, most musicologists discuss and analyze the meaning of music on the basis of sheet music. On the other, most people enjoy music by listening to audio recordings, which represent music in an acoustic form. In particular, the nuances and subtleties of musical performances, which are generally not written down in the score, make the music come alive. In this contribution, we discussed various strategies towards automated processing of music data across different representations. In particular, we showed how one can bridge the gap between the sheet music domain and the audio domain discussing aspects on music representations, music synchronization, and optical music recognition, while indicating various strategies and open research problems.

3.6 Multimodality in Human Computer Music Performance

Roger B. Dannenberg (CMU - Pittsburgh, US)

Human Computer Music Performance (HCMP) is the integration of computer performers into popular live music. At present, HCMP exists only in very limited forms, due to a lack of understanding of how computer performers might operate in the context of live music and a lack of supporting research and technology. I believe we need to envision a future performance practice that involves computers as musicians. An immediate realization is that non-audio communication is of the utmost importance in real-world music performance. By thinking about HCMP, we discover many interesting multimodal problems. These include: music notation and its use as a two-way communication channel in live performance, musical gestures for conducting and cuing, processing live multi-channel audio, coordination with video and projections in performance, multiples sensors and modalities for beat and measure tracking, and semantic processing concerning intention, emotion, and creation in multiple media.

3.7 Multimodal Music Exercises

Christian Dittmar (Fraunhofer Institute for Digital Media Technology - Ilmenau, DE)

URL http://www.songs2see.eu

We presented the project Songs2See that combines multiple music related modalities in order to assist users while learning a musical instrument. The developed software helps budding musicians to do their exercises on the instrument they are acquiring. The concept is similar to music video games; the main difference is the usage of real music instruments instead of game controllers. The most important modalities are the music score and the corresponding music recording, but additional, interactive visualizations supplement the application. We demonstrated the creation of new exercise content by means of semi-automatic music transcription, where backing tracks, which can be used to play along with, are extracted from real-world music recordings. This is achieved using symbolic note transcriptions to initialize the time-varying parameters of a source separation algorithm. At any stage of the content creation process, the user can intervene and correct possibly erroneous transcription results.

3.8 Association and Linking of Related Music Information

Simon Dixon (Queen Mary University of London, GB)

Many types of music information are available on the Web and in libraries and music databases, including structured and unstructured text (e.g. biographical metadata); complete or partial scores (e.g. chords, tabs, lead sheets, lyrics); and recordings of performances (e.g.

mixed, multi-track, video). For any given work, many such instances might exist, and an ideal music information system would allow navigation between the various representations at both the document and fragment levels, based on the relationships between and within documents that it has discovered. Such a system would enhance both the human experience (e.g. browsing, search) and automatic analysis (e.g. informed transcription, multimodal similarity computation). To realise this vision, the open research questions involve: finding related documents; synchronising them; representing the relationships in ways that are meaningful to users and allow reuse by other systems; and licensing/IP issues.

3.9 Score-Informed Audio Parameterization

Sebastian Ewert (Universität Bonn, DE)

Joint work of Ewert, Sebastian; Müller, Meinard

In recent years, the processing of audio recordings by exploiting musical knowledge as specified by a musical score has turned out to be a promising research direction. Here, one assumes that, additionally to the audio recording to be analyzed, one is given a MIDI file (representing the score) of the same underlying piece of music. The note event information specified by the MIDI file can then be used to support audio analysis tasks such as source separation or instrument equalization. In our contribution, we consider the problem of score-informed audio parameterization with the objective to successively adapt and enrich the note event information provided by the MIDI file to explain the given audio recording. More precisely, our goal is to parameterize the spectrogram of the audio recording by exploiting the score information. Our parameterization approach works iteratively proceeding in several steps. In the first step, we compute a temporal alignment between the MIDI file and the audio recording. Since the alignment accuracy is of major importance, we employ a refined synchronization method that exploits onset information. In the next steps, we successively adapt model parameters referring to dynamics, timbre, and instrumentation such that the spectrogram described by the model reflects the audio spectrogram as accurately as possible.

3.10 Time and Perception in Music and Computation

Alexandre R.J. François (Harvey Mudd College - Claremont, US)

Main reference Alexandre R.J. François: "Time and Perception in Music and Computation," in New Computational Paradigms for Computer Music, G. Assayag and A. Gerzso (Eds.), Editions Delatour France, IRCAM, pp. 125–146, 2009.

Music has long been a fascinating domain of application for computer scientists. It is a particularly challenging and inspiring one, as it only exists at the confluence of creation, representation and performance. Two "un-computational" properties characterize musical tasks. First, they are geared towards human perception and cognition. Computation aims at super-natural exactness, boring consistency, and absolute reproducibility; on the other hand, perception and cognition operate in the realm of natural variability, exciting unmet expectations, and approximate live reproductions. Secondly, musical tasks epitomize the

constant struggle between the desire to stop time and the necessity to live (and experience) in the present; traditional computing paradigms abstract and collapse these two sides of time onto arbitrarily antagonistic sets of compromises. The resulting abstract manifestation of time in computing is enforced as a strong invariant, universally and implicitly relied upon. Hermes/dl, a design language created specifically for the specification of complex dynamic systems, addresses many of the challenges posed by the computational modeling of musical tasks. As a design language, Hermes/dl consists of a collection of primitives, a set of organizing principles, and collections of qualifying situations. A graphical notation shields users from traditional algebraic notation, in an effort to make Hermes/dl more accessible and appealing to potential users in creative and scientific fields.

3.11 Groundtruthing for SALAMI and Billboard Projects

Ichiro Fujinaga (McGill University - Montreal, CA)

I presented two recent annotation projects undertaken at McGill University. One of them is a structural analysis of about 1000 pieces of variety of music as part of the Structural Analysis of Large Amount of Musical Information (SALAMI) project. The other one is a harmonic analysis of popular music listed on the Billboard Hot 100 list from 1950 to 1990. We have annotated about 600 songs so far. At the end of my talk, I posed the following questions to the participants: What is ground truth in music? What are the uses of ground truth in music research? Are there different kinds of ground truth? Are there different qualities of ground truth? How does one evaluate the quality of ground truths? Do we know when we have enough amount of ground truth? Should all ground truth be created by human beings? By any human beings? What are the effects of inter- and intra-annotator variability? Can ground truth be generated by machines in certain applications?

3.12 Crowd Music Listening: Internet-Based Music Listening with Shared Semantic Information

Masataka Goto (AIST - Ibaraki, JP)

URL http://en.wikipedia.org/wiki/Nico_Nico_Douga

I introduced a rapidly-growing video sharing web service in Japan, called "NICO NICO DOUGA." Although this service was not invented by ourselves, I informed the Dagstuhl participants of some advanced features that can enhance music listening experiences, shared interesting phenomena we observed on this service, and discussed whether essential advanced ideas on this service can be spread to other countries/cultures in the future. The NICO NICO DOUGA is a very popular web service managed by a Japanese company, Niwango. This service started in December 2006 and now has 19,610,000 registered users, which is more than 15 % of Japanese citizens. It won the Japanese Good Design Award in 2007, and an Honorary Mention of the Digital Communities category at Prix Ars Electronica 2008. On this service, users can upload, share, and view video clips like YouTube, but it

supports networked text communication where comments by anonymous users are overlaid on video clips and are synchronized to a specific playback time. In this way, comments can be related to events in the video. Since it can create a sense of shared watching experience called *"Pseudo-Synchronized Communication"* (coined by Satoshi Hamano), users can feel as if they enjoy together. We observed a lot of interesting ways of using comments and social tags, such as barrage ("DANMAKU"), ASCII art, original/parody lyrics, impressions, interpretations, and feelings. Since those time-synchronous comments can be considered semantic information about video content (music, for example), such internet-based music listening with shared semantic information can enhance our music listening experiences. In my keynote talk at AdMIRe 2009 of IEEE ISM 2009, I coined a new term *"Crowd Music Listening"* for this new way of music listening. You can enjoy music together with the crowd. You are not alone anymore while listening to music.

3.13 A Method for Obtaining Semantic Facets of Music Tags

Fabien Gouyon (INESC Porto, PT)

Main reference Mohamed Sordo, Fabien Gouyon, and Luís Sarmento: "A Method for Obtaining Semantic Facets of Music Tags," Workshop on Music Recommendation and Discovery, ACM Conference on Recommender Systems, Barcelona, 2010.

URL http://sunsite.informatik.rwth-aachen.de/Publications/CEUR-WS/Vol-633/wom2010_paper4.pdf

Music folksonomies have an inherent loose and open semantics, which hampers their use in structured browsing and recommendation. In my Dagstuhl talk, I presented a method (detailed in more length in a WOMRAD 2010 paper) for automatically obtaining a set of semantic facets underlying a folksonomy of music tags. The semantic facets are anchored upon the structure of the dynamic repository of universal knowledge Wikipedia. We also illustrated the relevance of the obtained facets for the automatic classification of Last.fm tags.

3.14 Why and How should MIR Research Open up to all Music Modalities?

Fabien Gouyon (INESC Porto, PT)

There are diverse modalities to music, namely, auditory (through hearing), natural language (through hearing, sight), gesture (through sight), physiology (through touch, temperature, and other inner senses), and so on. That is, humans can make sense of, or associate a meaning to a musical phenomenon transmitted via any of the above modalities (e.g. recognize a musical instrument when hearing it, understanding a song's topic by reading its lyrics). Furthermore, there are interactions between modalities (e.g. shivering when hearing a given song). Understanding how humans achieve these associations is certainly an interesting scientific goal, and one of the ways to try and understand this is the computational way: building machines that should do the same. In my Dagstuhl talk, I intended to bring forward the notion that in most "classic" MIR research topics (e.g. genre classification), which often

focused on the auditory modality, we may gain valuable insights by considering a multimodal approach. This may be a worthy approach to the understanding of multimodal associations we make when listening to music. But embracing more modalities in our research is certainly not an easy task. There are several fundamental issues including data issues, methodologies issues, and researcher training issues. Regarding data issues, I put forward a few questions:

- How can researchers obtain data from a particular modality? Are there inherent difficulties?
- Which are inherent issues with data cleaning and ground truth annotations for each particular modality?
- How can general data from a given modality help leverage music-specific research in that same modality? For instance, how can the availability of general text (e.g. on Internet) leverage improvements in analyses of music-related textual data (e.g. last.fm tags)?
- How can we relate data from diverse modalities? For instance, is it the best course of research to try mapping low-level audio features (auditory) to written labels of arbitrary high levels of semantics (natural language) via machine learning? Does it even really make sense to try to do so?

In this Dagstuhl talk, I did not intend to answer all the above questions of course, but rather tried to foster discussions on these. I also showed some of the work we do at the Sound and Music Computing group in Porto (http://smc.inescporto.pt), where some of the above questions are—albeit not directly addressed—latent scientific concerns in our research.

3.15 Integrated Content-based Audio Retrieval Framework

Peter Grosche (MPI für Informatik - Saarbrücken, DE)

Joint work of Grosche, Peter; Müller, Meinard

Even though there is a rapidly growing corpus of audio material, there still is a lack of efficient systems for content-based audio retrieval, which allow users to explore and browse through large music collections without relying on manually generated annotations. To account for the various user requirements, we plan to develop a content-based retrieval framework that supplies and integrates various functionalities for flexibly adjusting and tuning the underlying retrieval strategy. Based on the query-by-example paradigm, the user will be able to mark an arbitrary passage within a music representation. This passage is then used as query to retrieve all documents from the music collection containing parts or aspects similar to this passage. The new approach is to develop a framework that facilitates flexible and intuitive control mechanisms for adjusting various aspects in the search process. Firstly, the user may specify the musical properties to be considered in the similarity search. This allows to search for rhythmic, melodic, or harmonic patterns. Secondly, the framework will integrate various retrieval strategies ranging from high-specificity audio identification, over mid-specificity audio matching to low-specificity cover song and genre identification. Here, the goal is to supply the user with a control mechanism that allows to seamlessly adjust the specificity level in the search process. Thirdly, the retrieval framework will provide functionalities that account for subdocument as well as document-level retrieval. In combination with suitable visualization, navigation, and feedback mechanisms, the user is then able to successively refine and adjust the query formulation as well as the retrieval strategy.

3.16 Integrating Different Knowledge Sources for the Computational Modeling of Flamenco Music

Emilia Gómez (Universitat Pompeu Fabra - Barcelona, ES)

Main reference Joaquín Mora, Francisco Gómez, Emilia Gómez, Francisco Escobar-Borrego, and José Miguel Díaz-Báñez: "Characterization and Melodic Similarity of A Cappella Flamenco Cantes," Proc. 11th Int'l Society for Music Information Retrieval Conference (ISMIR), pp. 351–356, 2010.
Catherine Guastavino, Francisco Gómez, Godfried Toussaint, Fabrice Marandola, and Emilia Gómez: "Measuring Similarity between Flamenco Rhythmic Patterns," Journal of New Music Research 38(2), pp. 129–138, 2009.

URL http://mtg.upf.edu/research/projects/cofla

There is a wealth of literature on music research that focuses on the understanding of music similarity from different viewpoints and on the computation of similarity distances to cluster different pieces according to composer, performer, genre or mood. This similarity measure is often based on comparing musical excerpts in audio format and measuring the distance of a set of content descriptors representative of different musical facets (e.g. the used instruments, rhythmic pattern or harmonic progression). Alternative approaches are based on comparing context information from the contrasted pieces (e.g. influences, temporal and geographical coincidences), which is usually extracted from the web or manually labelled. A combination of these two approaches (content and context) seems to be the most adequate solution, but there is still a limitation on current approaches. This might be due to the fact that there are still other information sources to consider, such as the listening conditions. State-of-the-art research has mainly focused on the analysis of music from the so-called "Western tradition," given that most music retrieval systems are targeted toward this kind of music. Nevertheless, some studies are now considering if the available descriptors and similarity distances are suitable when dealing with music from other traditions. In this situation, the notion of similarity is also affected by the listener's cultural background and his previous exposure to the considered musical structures.

We focus here in the study on flamenco music. Flamenco is a music tradition mostly originally from Andalusia, in southern Spain. The origin and evolution of flamenco styles and variants have been studied by different disciplines, mainly ethnomusicology, anthropology or literature. Prior studies have mainly focused on artists' biographies, lyrics and social context, and there are few works on music analysis. There are some difficulties and motivations for developing computational models of similarity in flamenco music: being an oral tradition, there are no written scores; there exist few quality historical recordings and music collections are spread and not consistently documented; flamenco is not as present on the web as other musical styles; cultural institutions and music platforms are concerned with the preservation and spreading of flamenco music, given its commercial and cultural interest.

The goal of our project is to develop computational models for computer-assisted description, similarity computation, comparative analysis and processing of flamenco music (http://mtg.upf.edu/research/projects/cofla). We want to integrate different knowledge sources: content description (computational analysis of recordings and music similarity algorithms), context information (expert analyses) and user modeling (human judgements). As a case study, we deal with flamenco a capella singing. We focus on melodic similarity, and we evaluate state-of-the-art algorithms for automatic transcription and similarity computation. We work with a music collection of the most representative performances from 4 different a capella singing styles, mainly Debla, Martinete and Toná. We focus on analyzing the melodic exposition, and we approach both inter-style classification and intra-style similarity. Some

of the challenges of this project are strongly connected to three of the topics discussed at Dagstuhl Seminar on Multimodal Music Processing: Multimodality, Evaluation and Ground Truth, and User Modeling.

- **Multimodality.** Flamenco research is highly multimodal, as we can combine different inputs when analyzing a certain performance: audio, video, context information (style and singer) and lyrics. These modalities are sometimes complementary and sometimes contradictory. In addition, we do not always have access to all of them simultaneously, so the main research to be done here is related to their integration in current algorithms. As mentioned before, this project intends to combine some of these modalities and their related knowledge sources by means of combined measures of similarity.
- **Evaluation and Ground Truth.** There are some difficulties in this project when evaluating algorithms for two different tasks: melodic transcription and similarity measurement. Regarding melodic transcription, we have adopted a two-stage evaluation procedure (annotations vs. corrections): first, we have collected manual melodic contour transcriptions from flamenco experts, where time information is not relevant and ornaments are removed; then, we have asked them to perform manual corrections and refinements of detailed transcriptions provided by the computational model. This allows us to gather and contrast both annotations of overall melodic contour and ornaments (melisma information). In the same way, we have gathered different sources of ground truth information for music similarity: list of relevant features, similarity ratings and validation of clusters/trees generated by computational models. We have observed that different ground-truth sources are complementary, and there is always a degree of subjectivity in each of them. Computational models should then integrate them and adopt procedures for interactive-validation and user-adapted annotation.
- **User Modeling.** Our project deals with two different user profiles: musicians with little knowledge of flamenco music and experts with high knowledge of flamenco music. We have observed a low correlation among their similarity ratings and the features they use to compare flamenco performances. Here, one particular challenge is to implement user-adapted similarity measures.

In conclusion, the integration of distinct knowledge sources as well as user adaptation are two key aspects in developing a computer-assisted model of flamenco music retrieval.

3.17 Audio Signal Representations for Temporal Structure Segmentation

Florian Kaiser (TU Berlin, DE)

Music structural segmentation is the core of many MIR applications and remains the focus of many research activities. Most solutions proposed for this task were designed to either detect repetitive patterns or homogeneous acoustical segments in a self-similarity matrix of the audio signal. Depending on the parametrization, it is thus assumed that feature frames extracted over musical sections present some sort of statistical invariance, or that the sequence of features will be exactly repeated while the section is repeated. This assumption is however rarely fulfilled, and similarity matrices are often unable to yield a proper visualization of the

actual structure. Modeling features over a local short-time horizon, some work has shown that introducing contextual information in the measure of similarity between two time instants of the audio signal could enhance the structure visualization. We follow this idea and propose to concatenate local chroma sequences in "Multi-Probe Histograms." Transitions between major chroma bins of adjacent frames are mapped to the bins of a histogram. This yields a fixed-size representation of the chroma sequence that condenses local temporal relations between feature frames and summarizes the whole sequence. Preliminary results show that embedding this modeling in a similarity matrix, structure visualization can be strongly enhanced. We would like to discuss the musical interpretation of such representations of chroma sequences, especially with regard to their relation to tonality and harmony.

3.18 Between two Domains: Synergies of Music and Speech Information Retrieval

Frank Kurth (Fraunhofer FKIE, Wachtberg - Bonn, DE)

From its early days on, research in Music Information Retrieval (MIR) has adapted various established technologies from speech processing. Some popular examples are the use of MFCC (Mel Frequency Cepstral Coefficients) features, dynamic time warping or hidden Markov models. Subsequently, those technologies were suitably adopted to the requirements of the music domain, and then considerably extended and complemented by a wide variety of novel methods. In the last years, we were interested in transferring technology in the opposite direction: How can research in Speech Retrieval profit from the—now more mature—methods developed in MIR? In particular, we investigated the domains of feature design, matching techniques, and exploitation of multimodal representations. As a concrete example, we derived novel speech features (HFCC-ENS) from chroma-based music features to develop a method for detecting short sequences of spoken words within a speech recording ("keyphrase spotting"), where the underlying search technology is based on an audio matching approach from MIR. In another example, we adapted a strategy from multimodal MIR to time-align a speech recording with a corresponding textual transcript. In contrast to methods from classical speech processing, the use of MIR-based strategies allowed us to develop methods which work in an unsupervised manner.

3.19 Cross-Disciplinary Perspectives on Music Information Retrieval Challenges

Cynthia Liem (TU Delft, NL)

Many tasks in Music Information Retrieval research take place at a conceptually high level. As the real-world "non-digital" versions of such tasks involve human understanding and interpretation of music, when computational methods are developed for performing these same tasks, methodological challenges occur. A clear-cut solution to overcoming these

challenges has not yet been found, which especially poses problems in the stages of result evaluation and interpretation.

Music Information Retrieval is not the first research field in which such challenges occur. Longer-established fields such as Artificial Intelligence, Computational Linguistics, Content-Based Image Retrieval, Multimedia and Interactive Information Retrieval have faced similar problems. The initial plan was to discuss whether we could identify specific directions from such external neighbouring fields that the Music-IR community could learn and benefit from in its own development. For example, are there any success stories or epic failures from other fields than we can follow or avoid? Are there common pitfalls, or open questions that have not been sufficiently resolved yet in any of those fields? Such considerations would stimulate a cross-disciplinary dialogue and exchange of ideas, and indicate opportunities where Music-IR can take initiative and start to play an exemplary role for neighbouring fields. While in the end, this topic did not become an official main discussion point at the seminar, throughout the whole week it became clear that cross-disciplinary knowledge transfers both to and from the MIR domain will indeed be valuable. Concrete neighbouring fields that were mentioned as being of interest and inspiration included Text Information Retrieval and Speech Recognition, but also Linguistics and Musicology, indicating a broad basis for further discussions.

3.20 Analysis of Guitar Tabs and Chord Sequences

Robert Macrae (Queen Mary University of London, GB)

With over four and a half million tablatures and chord sequences (tabs), the world wide web holds vast quantities of hand annotated scores wrapped up in non-standardised ASCII text files. Theses scores are typically so incomplete, with errors, noise and duplicates that simply filtering out the correct tabs is challenging. Despite this, tabs are by far the most popular means of sharing musical instructions on the Internet. We are interested in developing tools that use text analysis and alignment for the automatic retrieval, interpretation and analysis of such files in order to filter out or recreate the true original music score from the multitude of tabs available. Ultimately it is hoped this work could allow for tab recommendation, tab synthesis, tab synchronisation/score following and integration with other music information for music processing.

3.21 Lyrics-to-Audio Alignment: Methods of Integrating Textual Chord Labels and an Application

Matthias Mauch (AIST - Tsukuba, JP)

Joint work of Mauch, Matthias; Fujihara, Hiromasa; Goto, Masataka
Main reference Matthias Mauch, Hiromasa Fujihara, and Masataka Goto: "Lyrics-to-Audio Alignment and Phrase-Level Segmentation Using Incomplete Internet-Style Chord Annotations," Proc. 7th Sound and Music Computing Conference (SMC), pp. 9–16, 2010.
URL http://smcnetwork.org/node/1328

Aligning lyrics to audio is by definition a multimodal task. It has a wide range of applications such as the automatic generation of karaoke scores, song-browsing by lyrics, and the generation of audio thumbnails. Existing methods are restricted to using only lyrics and match them to phoneme features extracted from the audio (usually mel-frequency cepstral coefficients). Our novel idea is to integrate the textual chord information provided in the paired chords-lyrics format known from song books and Internet sites into the inference procedure. By doing so the text–audio bimodality is complemented by a second form of multimodality: timbre–harmony.

We proposed two novel methods that implement our idea. Firstly, assuming that all chords of a song are known, we extended a hidden Markov model (HMM) framework by including chord changes in the Markov chain and an additional audio feature (chroma) in the emission vector. Secondly, for the more realistic case in which some chord information is missing, we presented a method that recovers the missing chord information by exploiting repetition in the song. We conducted experiments with five changing parameters and showed that with accuracies of 87.5% and 76.0%, respectively, both methods perform better, with statistical significance, than the baseline. Furthermore, we demonstrated the "Song Prompter" software system, which acts as a performance assistant by showing horizontally scrolling lyrics and chords in a graphical user interface together with an audio accompaniment consisting of bass and MIDI drums. This application allowed us to show that the automatic alignment is accurate enough to be used in a musical performance.

3.22 Synchronization-based Music Audio/Video Annotation

Meinard Müller (MPI für Informatik - Saarbrücken, DE)

For general music, there still is a significant gap between the demand of descriptive high-level features and the capability of existing feature extractors to automatically generating them. In particular, the automated extraction of high-level metadata from audio representations such as score parameters, timbre, melodies, instrumentation, or lyrics constitutes an extremely difficult problem with many yet unsolved problems. To bridge this gap, we suggest a knowledge-based approach for metadata extraction by exploiting the fact that one and the same piece of music often exists in several versions on different descriptive levels. For example, one version may encode explicit high-level information (score, lyrics, tablature, MIDI) and another version low-level information (audio, CD recording, video clip). Then the strategy is to use the information given by a high-level version in order to support localization and extraction of corresponding events in the low-level version. As a special case, this approach

has been successfully applied to align musical onset times of notes (given by some score representation) with corresponding physical onset times in some audio representation—a process, which can be regarded as an automatic knowledge-based annotation of the audio data. For the future, we will systematically refine and extend such a knowledge-based approach to automatically generate various kinds of metadata annotations and linking structures between music documents including new extraction techniques for higher level semantic descriptors such as harmonic progressions, tablature, lyrics, rhythm patterns, or motivic patterns.

3.23 A Multi-Perspective Analysis of Chord Labeling Procedures

Meinard Müller (MPI für Informatik – Saarbrücken, DE)

Joint work of Verena Konz; Müller, Meinard
Main reference Verena Konz, Meinard Müller, and Sebastian Ewert: "A Multi-Perspective Evaluation Framework for Chord Recognition," Proc. 11th International Conference on Music Information Retrieval (ISMIR), pp. 9–14, 2010.

The automated extraction of chord labels from audio recordings constitutes a major task in music information retrieval. To evaluate computer-based chord labeling procedures, one requires ground truth annotations for the underlying audio material. However, the manual generation of such annotations on the basis of audio recordings is tedious and time-consuming. On the other hand, trained musicians can easily derive chord labels from symbolic score data. In our contribution, we bridge this gap by describing a procedure that allows for transferring annotations and chord labels from the score domain to the audio domain and vice versa. Using music synchronization techniques, the general idea is to locally warp the annotations of all given data streams onto a common time axis, which then allows for a cross-domain evaluation of the various types of chord labels. As a further contribution, we extend this principle by introducing a multi-perspective evaluation framework for simultaneously comparing chord recognition results over multiple performances of the same piece of music. The revealed inconsistencies in the results do not only indicate limitations of the employed chord labeling strategies but also deepen the understanding of the underlying music material. Our multi-perspective visualization based on a musically meaningful time axis has turned out to be a valuable analysis tool for musicological tasks. In collaboration with musicologists, we are now investigating how recurrent tonal centers of a certain key can be determined automatically within large musical works.

3.24 VocaListener: Synthesis of Human-Like Singing by Using User's Singing and Its Lyrics

Tomoyasu Nakano (AIST – Tsukuba, JP)

Joint work of Nakano, Tomoyasu; Goto, Masataka
Main reference Tomoyasu Nakano and Masataka Goto: "VocaListener: A Singing-to-Singing Synthesis System Based on Iterative Parameter Estimation," Proc. 6th Sound and Music Computing Conference (SMC), pp. 343–348, 2009.
URL http://smc2009.smcnetwork.org/programme/pdfs/175.pdf

Since 2007, many Japanese users have started to use commercial singing synthesis systems to produce music, and the number of listeners who enjoy synthesized singing is increasing. Over 100,000 copies of singing synthesis software based on Yamaha's Vocaloid have been sold, and various compact discs that include synthesized vocal tracks have appeared on commercial music charts in Japan. We presented the singing synthesis system "VocaListener" that iteratively estimates parameters for singing synthesis software from a user's singing voice with the help of song lyrics. Since a natural voice is provided by the user, the synthesized singing voice mimicking it can be human-like and natural without time-consuming manual adjustments. The iterative estimation provides robustness with respect to different singing synthesis systems and singer databases. Moreover, VocaListener has a highly accurate lyrics-to-singing synchronization function, and its interface lets a user easily correct synchronization errors by simply pointing them out. In addition, VocaListener has a function to improve synthesized singing as if the user's singing skills were improved. Demonstration videos including examples of synthesized singing are available at http://staff.aist.go.jp/t.nakano/VocaListener/. In this seminar, we discussed mechanisms used in the VocaListener, presented some synthesized results, and indicated some future directions. Here, one main goal is to be able to synthesize singing voices that can not be distinguished any longer from human singing voices.

3.25 Modeling Novelty in Music Retrieval

Nicola Orio (University of Padova, IT)

One interesting aspect of available search engines is that sometimes, or maybe quite often, they fail and retrieve documents that are only slightly related to the user information needs. So they unwillingly introduce novelty in the retrieved list of documents by presenting unexpected results. These results may promote serendipity, help the user to discover new things, and may enlarge user interest—possibly even modifying the initial information needs. This contribution describes a draft idea, far to be tested, on how game theory results may be applied to model and somehow to control novelty in a music retrieval task. In particular, Prisoner's Dilemma seems to provide an interesting approach that is worth being explored. The idea is that parameters, algorithms, and different modalities behave as players who, at each time step, may choose whether or not to collaborate. The goal of each player is to maximize his or her payoff, while the goal of the envisaged system is to encourage players to change their strategy over time, to promote unexpected results to be presented to the user, and to hopefully improve the user's appreciation. Multimodal music retrieval is naturally

based on the concept of cooperation between subsystems. Yet, cooperation is a complex phenomenon that encompasses also the possibility that individual strategies are based on betraying, at least from time to time, from which the overall system may be able to retrieve original items.

3.26 Towards a Definition of the Description of Annotated M.I.R. Corpora

Geoffroy Peeters (Ircam - Paris, FR)

This presentation concerned a proposal for the definition of the description of annotated M.I.R. corpora. Given that the availability of annotated data within a given M.I.R. task usually corresponds to the start of a growing number of research activities (this has been the case for music genre classification, chord recognition, or music structure analysis), accessibility to annotated data must be considered as a major issue. Today, annotation data is often provided by various research labs (or companies) each one using its own annotation methodology and own concept definitions. This is not a problem by itself. However, the lack of descriptions of these methodologies or concepts (What is actually described? How is the data annotated?) is a problem in view of sustainability, usability, and sharing of such corpora. Therefore, it is essential to exactly define what and how annotations of M.I.R. corpa should be supplied and described. It should be noted that usual "standards" such as MPEG-7, RDF or Music-XML only concern the storage of the annotation, but not the description of the annotations themselves. The following issues become crucial for the description of corpora.

- What concepts are actually annotated? For example, in the case of structure annotations, which kind of structure is considered? In the case of chord annotations, what instruments or voices are considered? Only accompaniment, the leading voice, or both? For most M.I.R. corpora, the definitions of the concepts that are actually annotated are missing.
- How is the data actually annotated? Is the data derived from meta-data catalogues, from parameters of a synthesis process, or from human annotations? In the latter case, how many people were involved in the annotation process? What is the reliability of the data? Has cross-validation been applied?
- What are the actual annotation rules? For example, what is the temporal precision used for segment annotations? Which type of dictionary has been used for the labels? Are there equivalences between labels?
- How does the corpus look like? For example, what are the properties of the audio being described? Does the corpus consist of real or synthesized audio recordings? For which purposes have specific music-tracks been selected? Are these artificial recordings made specifically for the purpose of an annotation task?
- What are the storage formats for the raw data and annotation data? For example, in the case of audio material, are the files encoded as PCM or in some compressed form? Was the audio data synthesized from MIDI files? Are the annotations given in MPEG-7, RDF, or Music-XML?

3.27 "Copy and Scale" Method for Doing Time-Localized M.I.R. Estimation: Application to Beat-tracking

Geoffroy Peeters (Ircam - Paris, FR)

Main reference Geoffroy Peeters: " 'Copy and Scale' Method for Doing Time-Localized M.I.R. Estimation: Application to Beat-tracking." Journal of New Music Research, Special issue on Music and Machine Learning, to appear.

This presentation concerned a new algorithm, named "copy and scale" algorithm, with the goal to first estimate the local parameters of an unknown item (such as beat positions) simply by locating the closest item in a database and then to copy and scale the annotations of the located item to serve as the estimations for the unknown item. A nearest neighbour algorithm consists in assigning the information (such as music genre or mood) of the closest item of a pre-annotated database to an unknown target. It can be viewed as a "copy and paste" method. It is usually used for estimating global variables (such as genre or mood). The "copy and scale" method we propose allows for estimating local variables (i.e. variables that have a specific time location) and consists in "scaling" the closest item to adapt it to the properties of the unknown target. For this, we first represent the content of an audio signal using a sampled and tempo-normalized complex DFT of an onset-energy-function. This representation is used as the code over which the nearest neighbour (NN) search is performed. Along each code of the NN space, we store the corresponding annotated beat-marker positions in a normalized form. Once the closest code is found in the database, its tempo is assigned to the unknown item and the normalized beat-markers are scaled to this tempo in order to provide the estimation of the unknown item beat-markers. We perform a preliminary evaluation of this method and show that, with such a simple method, we can achieve results comparable to the ones obtained with sophisticated approaches.

3.28 Toward Reverted Indexing for Multimodal Music Retrieval

Jeremy Pickens (Catalyst Repository Systems, US)

Joint work of Pickens, Jeremy; Cooper, Matthew; Golovchinsky, Gene
Main reference Jeremy Pickens, Matthew Cooper, and Gene Golovchinsky: "Reverted Indexing for Feedback and Expansion," Proc. 19th ACM International Conference on Information and Knowledge Management (CIKM), pp. 26–30, 2010.
URL http://www.fxpal.com/publications/FXPAL-PR-10-581.pdf

Traditional text interactive information retrieval systems operate by creating inverted lists, or term indexes. For every term in the vocabulary, a list is created that contains the documents in which that term occurs and its relative frequency within each document. Retrieval algorithms then use these term frequencies alongside other collection statistics to identify the matching documents for a query. Recent advances turn the process around: instead of indexing documents, we index query result sets. First, a set of basis queries, representative of the collection as a whole, are chosen. Each query is then run through a retrieval system and the resulting document IDs are treated as terms while the score or rank of the document is used as the frequency statistic. Thus, an index of documents retrieved by basis queries is created. We call this index a reverted index. With reverted indexes, standard retrieval algorithms can retrieve the best basis queries (as results) for a set of documents (used as

queries). These recovered queries can then be used to identify additional related documents or to aid the user in query formulation, selection, and feedback.

In this presentation, we proposed applying the idea of reverted indexing to multimodal music retrieval by extending the set of basis queries to multiple modalities: pitch, rhythm, tempo, harmony, timbre, lyric, and structure to name but a few. As it is the results of each of these query modalities that are being indexed rather than the features and structures that go into ranking the results of each modality, basis queries from different modalities can be combined into a single index. At runtime, standard retrieval algorithms can again be used to retrieve the best basis queries (as results) for a set of documents (used as queries). These multimodal basis queries can then be recombined in an ad hoc manner to identify related pieces of music.

3.29 Future of Music Search and Consumptions—Results of the CHORUS+ Think Tank at MIDEM2011

Andreas Rauber (TU Wien, AT)

CHORUS+ (EU FP7 Coordinating Action) organized a highly focused Think Tank on "Music Search and Consumption" with a small group of key industry experts to identify needs and directions for research in Music IR at MIDEM (Marché international de l'édition musicale). The presentation reported on the key suggestions and findings from this Think Tank, specifically focusing on the gaps between the vision for music search and consumption, what research is working on, and what the market is (planning) to deliver. The Think Tank was supported by an on-line questionnaire, which has also been answered by numerous participants in this seminar. The first part of it is rather standard statistical information while the questions towards the end try to identify the biggest gaps and visions for required tools, technologies and services. Results of both the questionnaire as well as the Think Tank highlight the importance, particularly of personalization and recommender technology, with some disagreement concerning the suitability of existing methods. Further highlights include an anticipated shift from ownership-based access to music to access-based, with business models likely shifting from an object-based towards an attention-based model.

3.30 Multimodal/Multi-level Integration Strategies for Music Signal Processing

Gaël Richard (TELECOM-Paristech – Paris, FR)

Main reference Olivier Gillet, Slim Essid, and Gaël Richard: "On the Correlation of Audio and Visual Segmentations of Music Videos," IEEE Trans. on Circuits and Systems for Video Technology, 17 (2), pp 347-355, 2007.

The purpose of this talk was to highlight the interest and challenges of multimodality or multi-level integration by discussing a selection of our prior studies in multimodal music signal processing. The first challenge addressed consisted in a discussion on which ways audio

source separation could be improved when some additional information about the sources such as the musical score, beats, or the sung melody is known. A potential application of active listening involving audio-informed source separation and watermarking was also discussed.

As second topic, we discussed possible strategies for combining features that are calculated at different frame rates or that represent different semantics levels. In particular, it was shown that, besides the well known early and late fusion strategies, it may be appropriate to consider other approaches that iteratively (or jointly) exploit information given by the various modalities to improve estimates for the other modalities. In the same spirit, an example of cross-modality search was discussed. More specifically, it was shown how cross-correlation between semantic functions automatically extracted from each stream can lead to efficient inter-modality search for information retrieval or content authoring. Here, specific examples on music videos were given including a short demonstration on an audio-based video search prototype.

As a third topic, the strategy of considering the time evolution of audio features simultaneously at different temporal levels was addressed. This so-called multi-level feature integration was discussed with a specific focus on the appropriateness to combine both early and late fusion strategies. Finally, a set of questions was selected intended as initial input for further group discussions.

3.31 Personalization in Multimodal Music Retrieval

Markus Schedl (Universität Linz, AT)

This talk provided an overview of current research endeavors and existing solutions in multimodal music retrieval, where the term "multimodal" relates to two aspects. Accounting for the context of a piece of music or an artist constitutes the first aspect, while the second one relates to incorporating the user context. Adding "traditional" signal-based audio analysis, I proposed a model that incorporates three broad categories of influence factors for music retrieval and music similarity computation: music content, music context, and user context. The music context is introduced as all the information that is important to the music, albeit not directly extractable from the audio signal. Such information includes, for example, editorial or collaboratively assembled meta-data, lyrics in textual form, cultural background of an artist, or images of album covers. The user context, in contrast, is defined by various external factors that influence how a listener perceives music. It is therefore strongly related to user modeling and personalization, both facets of music information research that have not gained large attention by the MIR community so far. In my estimation, adding personalization aspects to existing music retrieval systems (for example, playlist generators, recommender systems, or visual browsers) constitutes one key issue in future MIR research.

3.32 Next Gen Music Analysis: Some Inspirations from Speech

Björn Schuller (TU München, DE)

Numerous ideas have been transferred in the past between the fields of Music Information Retrieval (MIR) and Automatic Speech Processing (ASP) including the usage of synthesized speech/chords for model training, feature brute-forcing, language modelling, bag-of-words, and Mel-Frequency Cepstral Coefficients. In this light, opportunities for further cross-discipline methodology transfer were discussed starting from a unified perspective on speech and music analysis. Such may include next generation MIR tasks such as recognition of age, height, weight, ethnicity, voice quality, likability, personality, as well as verification of performers and sung language identification. Concerning databases, rich transcription including several MIR tasks for single databases and added meta-information allow for combined analysis. Non-prototypical instance selection and partitioning including development sets are further to become more standard in MIR. Next, many classification and regression approaches recently evolved for ASP including tailored Graphical Models, Long-Short-Term-Memory enhanced architectures, and Gaussian Mixture Model supervectors for Support Vector Machines. Multitask Learning, Grid Search or similar optimization techniques and Self-Learning are further candidates for interesting advances in the field of MIR. Considering Multimodal Music Processing (MMP), fusion by synthesis and asynchronous stream processing provide elaborate techniques used in Audio-visual Speech Recognition. In addition, more fully automated MIR studies could be seen, where all data has to be retrieved from the web or derived from mid-level features based on automatic recognition. Ground truth interpretation by kappa and alpha measures and result interpretation by significances or effect power is also yet to become more common in MIR. Overall, several options for transferring recent developments in speech analysis to the MIR field exists. In the future, transfer from the fields of image and video analysis looking in particular at MMP as well as combined music and speech (and sound) analysis to maximally exploit mutual dependencies may become increasingly interesting.

3.33 PROBADO: A Multimodal Music Library Systems

Verena Thomas (Universität Bonn, DE)

Joint work of Thomas, Verena; Damm, David; Fremerey, Christian; Clausen, Michael
URL http://www.probado.de

As part of the PROBADO (Prototypischer Betrieb allgemeiner Dokumente) project, we aim at developing a prototypical digital music library system which is tested and implemented at the Bavarian State Library in Munich. Besides an appealing user interface following the WYSIWYH ("What You See Is What You Hear") paradigm, a widely automated processing workflow including digitization, indexing, annotation, and presentation has been developed. To allow for content based search (audio, lyrics, and symbolic representations) as well as convenient document presentation anb browsing, several state-of-the-art MIR techniques have been integrated into the system.

3.34 A User Interface for Motivic Analysis

Verena Thomas (Universität Bonn, DE)

Joint work of Thomas, Verena; Clausen, Michael

Besides developing methods for identifying patterns in symbolic music documents and determining repetitions of these patterns within the document, it is important to provide a user interface for visualizing the analysis results. We presented a first prototype of such a system offering a piano roll like visualization of the analyzed piece of music. The system identifies repeating patterns within the piece of music, their occurrences throughout the piece as well as the type of occurrence such repetitions with equal musical intervals, retrogrades, inversions and retrograde inversions of musical intervals, as well as repetitions with respect to the note durations. In addition, the hierarchical structure of the identified patterns is analyzed and visualized by the system.

3.35 Mobile Multimodal Music Processing for Edutainment and eHealth

Ye Wang (National University of Singapore, SG)

Multimodal smartphones are powerful enough for real-time music processing and can be easily connected to the Internet wirelessly. In this context, I presented our research endeavors at National University of Singapore, leveraging multimodal music processing on smartphones for edutainment and eHealth applications. We have also built a private cloud computing cluster dedicated for the backend computing and Internet scale MIR.

Edutainment is an important research topic, as it seeks to find new ways for applying technology to the education environment in an entertaining format. Our efforts in this area have produced the MOGCLASS prototype which has been successfully deployed in three primary schools and the Muscular Dystrophy Association (Singapore) with positive feedback. MOGCLASS has received enthusiastic comments from members in the Curriculum Planning and Development Department of the Ministry of Education (Singapore) and has started commercial licensing to schools. e-Health seeks to apply web-based and mobile technologies to create new health solutions. We have begun core research in building systems for elderly-care applications that mix mobile computing and domain specific music retrieval. We want to mention two specific current projects.

- In the MusicalWalk project, we are developing a web based system that allows music therapists to find suitable pieces of music to facilitate rhythmic auditory stimulation (RAS)-based gait training for a patient with Parkinson's disease.
- In the EZSleep project, we are investigating technologies for selecting the best sequence of music for improved sleep quality of an insomniac by learning his or her biophysical indicators and mapping them to content-based musical descriptors.

3.36 What is it Like to be a Musicologist?

Frans Wiering (Utrecht University, NL)

In this presentation, I investigated the relationship between present-day musicology and computational music processing. These two fields seem to be far apart: computing has at best a role as tool supplier in mainstream musicological research and musicology figures mainly as a source of domain knowledge and ground truth in computational music processing. Setting aside the matter of different methodologies, there seem to exist completely different views as to what music is. One important reason why this is the case relates to the crisis that took place in musicology in the late 1980s.

Before that time, musicology used to be first of all about Western classical music. Its central concept was the musical "work," the autonomous artistic creation exemplified by the score that captured the composer's intention. This "positivist" view of music was not so much superseded as stripped of its credibility in the following decades of "critical" musicology. The work concept itself was shown to be a cultural construction, loaded with prejudice against other possible views of music. Critical musicology has today become mainstream musicology, examining musical meaning, subjectivity, culture, context, gender and identity rather than beauty, structure and indeed "the music itself."

Music processing is first of all about music as data, whether notation, audio or otherwise. Its benefits are the well-known ones of computation: size, speed, accuracy and formal modelling, but the mismatch with musicology is obvious. However, a common ground may be found in what David Huron has called a "new empiricist" approach. Studies of context, subjectivity and meaning generation may benefit from data-rich approaches complementing the anecdotic evidence that they generally are based on. For example, in meaning generation high-level musical patterns shared between different pieces play an important role, but to understand this role massive amounts of music need to be analysed.

3.37 What does it Mean to "Process" "Music"?

Geraint A. Wiggins (Goldsmiths, University of London, GB)

There is a significant gap (sometimes called a "semantic" gap, but that is only partly correct in one particularly narrow usage of the word "semantic") between the audio signal, which corresponds with "Sound" in the air, and "Music" which is the experience that arises in an appropriately enculturated human mind as a result of the application of an appropriate audio stimulus. My research is about understanding "Music" from this viewpoint, and identifying representations that properly model the human experience of music, and studying their properties. This entails a certain amount of perceptual and cognitive psychology, since, I claim, the word "Music" has no meaning except where a mind is involved. In this my presentation, I tried to bridge the (non-semantic) gap between audio signals and the human experience of music, via what are sometimes called symbolic, mid-level, or discrete representations. While none of these names are incorrect, none of them captures the point: to be meaningful, the representations need to be perceptually and/or cognitively valid.

3.38 Music Commentator: A System for Generating Comments Synchronized with Music Audio Signals

Kazuyoshi Yoshii (AIST – Ibaraki, JP)

Joint work of Yoshii, Kazuyoshi; Goto, Masataka
Main reference Kazuyoshi Yoshii and Masataka Goto: "MusicCommentator: Generating Comments Synchronized with Musical Audio Signals by a Joint Probabilistic Model of Acoustic and Textual Features," Lecture Notes in Computer Science, Vol. 5709, pp. 85–97, 2009.
URL http://dx.doi.org/10.1007/978-3-642-04052-8_8

In online video sharing services, many users often provide free-form text comments for temporal events in music clips, not for entire clips. Although this is one of the most interesting functions of humans, its mechanism has not fully been investigated so far. How can computers emulate the commenting behavior of humans? We have developed a system called "Music Commentator" that suggests possible natural-language comments on appropriate temporal positions in a musical audio clip. To achieve this, we propose a joint probabilistic model of audio signals and text comments. The model is trained by using existing clips and users' comments given to those clips. Given a new clip and some of its comments, the model is used to estimate what temporal positions could be commented on and what comments could be added to those positions. It then concatenates possible words by taking language constraints into account. Our experimental results showed that using existing comments in a new clip resulted in improved accuracy for generating suitable comments to it. To improve the system, we have to take into account high-level semantic musical features related to melody, rhythm, and harmony. In addition, visual features of music video clips should be dealt with.

3.39 A Statistical Approach to the Ground-Truth Problem

Kazuyoshi Yoshii (AIST – Ibaraki, JP)

So far, humans have put a lot of effort to annotate musical pieces. Such annotations are used as ground-truth data and have been greatly contributing to the progress of music information research. Note that some kinds of annotations are forced to follow humans' arbitrary decisions. For example, how finely should we categorize genres, chords, or moods? There is no way to measure the appropriateness of such decisions in a principled manner. To solve this problem, we could use a statistical "data-driven" approach. Bayesian nonparametrics is considered to be especially promising for opening the door to a meta-level research area, i.e., a new generation of computational musicology.

4 Working Groups

4.1 Group Session: Signal Models for and Fusion of Multimodal Information

Sebastian Ewert, Masataka Goto, Peter Grosche, Florian Kaiser, Kazuyoshi Yoshii, Frank Kurth, Matthias Mauch, Meinard Müller, Geoffroy Peeters, Gaël Richard, Björn Schuller

The *"Model"* group discussed the issue of how signal models can be developed that exploit multimodal information. In most music processing methods a piece of music is analyzed by extracting high-level semantics from a single source of information, for example from an audio recording or a symbolic score representation (MIDI, MusicXML). However, there might be additional data sources available such as corresponding videos, chord transcriptions, lyrics annotations, tags and text comments, or other textual meta-data. Given the complexity of most analysis tasks, it is a very promising idea to combine all available sources of information to support music analysis methods or to validate their results. During the group meeting we discussed several aspects related to this idea.

In particular, we exemplarily selected several music analysis approaches that incorporate multiple information sources, and discussed them in detail with the goal to identify general fusion strategies. On the one hand, we categorized strategies into either early or late fusion approaches and discussed their advantages and weaknesses. Here, the terms *early* and *late* refer to the stage in the processing pipeline where the sources of information are finally combined. On the other hand, we discussed aspects related to the modeling of uncertainty involved in the combination process. Here, uncertainty may arise from missing or contradictory information in the given data.

Overall, we identified several general fusion strategies. Firstly, we discussed the *weighting-of-streams* approach where streams of information are processed individually and the results are merged using a weighting scheme. Secondly, we identified the *iterative-reinforcement approach* where unreliable or coarse information is refined in an iterative manner. Furthermore, we discussed how uncertainty is treated in classical machine learning methods. Here, we considered *confusion aware approaches*, where information about the reliability of individual classifiers is gathered from training data in a first step and is used later to build better classifiers. As one major issue, we found that a measure of uncertainty involved in the combination process is a prerequisite for choosing a suitable fusion strategy.

4.2 Group Session: User-aware Music Information Retrieval

Nicola Orio, Emilia Gómez, Fabien Gouyon, Cynthia Liem, Tomoyasu Nakano, Jeremy Pickens, Andreas Rauber, Markus Schedl, Ye Wang

The *"User"* group addressed the topic of user-aware music information retrieval. Actually, the perception of music is highly subjective and depends on the personal background and taste. Furthermore, the perception also depends on the specific user context, which is determined

by various external factors that influence how a listener perceives music in a certain situation. In this session, we discussed the central question on how contextual information can be integrated into the retrieval process in order to account for long-term user behavior and short-term user interests. In particular, the consideration of personalization aspects in the development of music retrieval systems has been identified as one key issue in future MIR research. Additionally, we also discussed how search engines may yield satisfying results in terms of novelty, popularity, and serendipity of the retrieved items. Here, we conjectured that retrieval results that are only slightly related to the user information needs may actually become interesting for promoting serendipity helping the user to discover new and surprising music.

4.3 Group Session: Symbolic Music Representations and OMR

Christopher Raphael, Ichiro Fujinaga, Simon Dixon, Robert Macrae, David Bainbridge, Michael Clausen, Verena Thomas

In the session of the *"Symbol"* group, we discussed the question of how to bridge the gap between visual, symbolic, and acoustic representations of music. Our discussion of symbolic music representations largely centered around optical music recognition (OMR) as this is the most obvious possibility for getting large quantities of symbolic music data for classical music. We feel the need for OMR especially now, since the International Music Score Library Project offers a very large open library of classical music scores. Several of the group members, Fujinaga and Bainbrige have considerable experience working in the problem, though current technology is not adequate for many commonly encountered situations. We discussed particular ways to pose the problem that allow for useful results with only a partial mastery of the grand challenge.

Much of our discussion focused on the SyncPlayer work of Clausen and Thomas (and others) which registers orchestral audio with scores and allows for simultaneously viewing and listening. All were very much hoping to see this project succeed and become widely available (beyond the University of Bonn and the Bavarian State Library). There still remains quite a bit of work before this can be achieved, including solving both technical issues and intellectual property concerns. We also discussed other symbolic representations such as guitar tabs, which are widely used by guitarists and others interested in a chord/lyric description of a song. There remain many interested challenges here, including synchronization with audio and video material.

4.4 Group Session: Meaning of Music

Geraint A. Wiggins, Frans Wiering, Michael Casey, Elaine Chew, Roger B. Dannenberg, Alexandre R.J. François, Matthias Mauch

The *"Meaning"* group addressed the subject of how musical meaning can be derived from musical data and, in particular, from musical sound. Here, we addressed the gap between the audio signal, which corresponds with "Sound" in the air, and "Music" which is the experience that arises in an appropriately enculturated human mind as a result of the application of an appropriate audio stimulus. To better understand this gap, we sketched a path from the given low-level (acoustic) raw data to high-level musical models from a human-based as well as from a machine-based perspective.

5 Panel Discussions

5.1 Panel Session: Ground Truth

Ichiro Fujinaga, Geoffroy Peeters, Kazuyoshi Yoshii, Meinard Müller, Elaine Chew

In the panel session on *"Ground Truth"* fundamental issues related to the interpretation, usage, and generation of ground truth annotations were discussed. What is ground truth in music? How can one handle inter- and intra-annotator variabilities? How can the quality of ground truth be evaluated? Are there alternatives to manually generate ground truth annotations? These are just some of the questions raised during the panel. There were five presentation interspersed with lively discussions. The presentations were made by Ichiro Fujinaga ("Groundtruthing for SALAMI and Billboard Projects," see Section 3.11), Geoffroy Peeters ("Towards a Definition of the Description of Annotated M.I.R. Corpora," see Section 3.26), Kazuyoshi Yoshii ("A statistical approach to the ground-truth problem," see Section 3.39), Meinard Müller ("A Multi-Perspective Analysis of Chord Labeling Procedures," see Section 3.23), and Elaine Chew ("Ground Truth: Where Aggregates Fail," see Section 3.4).

5.2 Panel Session: Grand Challenges

Masataka Goto, Roger B. Dannenberg

Main reference Masataka Goto: "Music Informatics," IPSJ Magazine, Vol. 51, No. 6, pp. 661–668, 2010.
URL http://staff.aist.go.jp/m.goto/PAPER/IPSJ201006goto.pdf

In the panel session *"Grand Challenges"* we discussed in which way music information research may and should impact our daily lives and our society in the future. Among others, the following challenges were discussed:

- How to provide the best music for each person?

- How to predict music trends?
- How to enrich the relation between humans and music?
- How to push new music evolution forward?
- How to contribute to solving environmental and energy problems?
- How to involve computers in live performances?
- How to predict specific effects of music?

There were two presentations by Masataka Goto and Roger Dannenberg ("Multimodality in Human Computer Music Performance," see Section 3.6), which were accompanied by some passionate discussions between and within the audience and the panelists.

Participants

- David Bainbridge
University of Waikato, NZ
- Hannah Bast
Universität Freiburg, DE
- Michael Casey
Dartmouth College - Hanover, US
- Elaine Chew
USC - Los Angeles, US
- Michael Clausen
Universität Bonn, DE
- Roger B. Dannenberg
CMU - Pittsburgh, US
- Christian Dittmar
Fraunhofer Institut IDMT - Illmenau, DE
- Simon Dixon
Queen Mary University of London, GB
- Sebastian Ewert
Universität Bonn, DE
- Alexandre R. J. Francois
Harvey Mudd College - Claremont, US
- Ichiro Fujinaga
McGill University - Montreal, CA
- Emilia Gómez
Universitat Pompeu Fabra - Barcelona, ES
- Masataka Goto
AIST - Ibaraki, JP
- Fabien Gouyon
INESC Porto, PT
- Peter Grosche
MPI für Informatik - Saarbrücken, DE
- Florian Kaiser
TU Berlin, DE
- Frank Kurth
Fraunhofer FKIE, Wachtberg - Bonn, DE
- Cynthia Liem
TU Delft, NL
- Robert Macrae
Queen Mary University of London, GB
- Matthias Mauch
AIST - Tsukuba, JP
- Meinard Müller
Saarland University and MPI für Informatik - Saarbrücken, DE
- Tomoyasu Nakano
AIST - Tsukuba, JP
- Nicola Orio
University of Padova, IT
- Geoffroy Peeters
Ircam - Paris, FR
- Jeremy Pickens
Catalyst Repository Systems, US
- Christopher Raphael
Indiana University - Bloomington, US
- Andreas Rauber
TU Wien, AT
- Gael Richard
TELECOM-Paristech - Paris, FR
- Markus Schedl
Universität Linz, AT
- Björn Schuller
TU München, DE
- Verena Thomas
Universität Bonn, DE
- Ye Wang
National University of Singapore, SG
- Frans Wiering
Utrecht University, NL
- Geraint A. Wiggins
Goldsmiths, University of London, GB
- Kazuyoshi Yoshii
AIST - Ibaraki, JP

Report from Dagstuhl Seminar 11042

Learning from the Past: Implications for the Future Internet and its Management?

Edited by

Gabi Dreo Rodosek[1], Aiko Pras[2], Henning Schulzrinne[3], and Burkhard Stiller[4]

1 **Universität der Bundeswehr München, DE,** gabi.dreo@unibw.de
2 **University of Twente, NL,** pras@cs.utwente.nl
3 **Columbia University, US,** schulzrinne@cs.columbia.edu
4 **Universität Zürich, CH,** stiller@ifi.uzh.ch

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 11042 "Learning from the Past: Implications for the Future Internet and its Management?". The discussion centered around the question if by analyzing the past - especially why certain technologies did or did not succeed - it is possible to reason about the Future Internet, the challenges and especially the management aspect. Valuable observations have been identified during the discussions. A scientific publication that summarizes the key findings is under preparation.

Seminar 25.–28. January, 2011 – www.dagstuhl.de/11042
1998 ACM Subject Classification C.2 Computer-Communication Networks, C.2.1 Network Architecture and Design, C.2.3 Network Management
Keywords and phrases Future Internet, Management, Lessons learnt
Digital Object Identifier 10.4230/DagRep.1.1.102

1 Executive Summary

Gabi Dreo Rodosek
Aiko Pras
Henning Schulzrinne
Burkhard Stiller

The attendees of the seminar discussed a hot topic of various research initiatives, namely the Future Internet, and its management. Can we make assumptions for the Future Internet if we question ourselves how the present Internet needs to be re-engineered by analyzing its current major limitations? Although discussions like the clean-slate vs. evolutionary approach are taking place, and are certainly vision-based, an alternative approach to investigate why certain decisions failed or succeeded seems to be promising as well.

Looking back, we recognize that the Internet architecture and protocols are, in theory, defined by about 5000 RFCs ("standards"), but in practice depend on a much smaller number of core assumptions, protocols and technologies. The Internet is largely defined by a modest set of protocols with numerous extensions and modifications: IPv4 and IPv6 at the network layer, UDP and TCP at the transport layer, and a handful of standardized and a large number of proprietary application protocols, for applications from email and file transfer to

Learning from the Past: Implications for the Future Internet and its Management?, *Dagstuhl Reports*, Vol. 1, Issue 1, pp. 102–107
Editors: Gabi Dreo Rodosek, Aiko Pras, Henning Schulzrinne, and Burkhard Stiller
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

gaming. Along with these data-delivering protocols, control plane protocols such as BGP, configuration protocols such as DHCP and management protocols such as SNMP are needed to keep the Internet running.

Many of these protocols are now at least a decade, in some cases three decades, old. Some have aged gracefully; others clearly have reached the end of their useful lifetime. Also, numerous other protocols and architectures have been proposed in the literature and many were even standardized over the past decades, but most have had very limited or no practical impact. Unfortunately, the design of new protocols is still largely based on folklore and ideas passed on informally during discussions at conferences, IETF standardization meetings and PhD seminars. It is relatively easy to stuff all ideas accumulated into a new protocol, but it seems much harder to destil the lessons learned into future designs, so that we do not loose some of the core properties that have made the Internet successful. Not only in the design of protocols, the same is true also for the development of management approaches. Although it seems that we have recognized that the management of the Future Internet needs to be addressed in the design phase, it is largely ignored so far, maybe because we do not know how to approach it.

Therefore, learning from the past would give us valuable insights for the design of the Future Internet, and its management. The goal of the seminar was to discuss these questions and identify the key findings.

2 Table of Contents

3 Overview of Talks

The seminar investigated core network architectural issues and core protocols as well as management approaches, to define precisely what made design decisions succeed against their competitors at the time of design, and which choices have turned out to be prescient and where we are reaching the end of the extensibility. The discussion was motivated by the following talks:

- Georg Carle, TU München
  Network Management for Measuring, Processing, and Reacting
- Gabi Dreo Rodosek, Universität der Bundeswehr München
  Learning from the Past: KISS principle for the Future Internet, and Self-Management
- Olivier Festor, INRIA - Nancy Grand Est
  A Biased Focus on 20 Years Research in Network Management
- Sape J. Mullender, Alcatel-Lucent Bell Labs - Antwerpen
  New Protocols for Accessing Services, Files, Content and more
- Aiko Pras, University of Twente
  Learning from the past - why OSI failed
- Danny Raz, Technion - Haifa
  Networking and Computation Resources: Can They Mix?
- Ramin Sadre, University of Twente
  Is The Future Unpredictable?
- Peter Schoo, Fraunhofer Institut SIT, München
  Management in FI needs increasingly more and better means for anomaly detection
- Henning Schulzrinne, Columbia University
  What did we learn? Random thoughts on protocol design
- Burkhard Stiller, Universität Zürich
  Communications + Management/Operations + Security != Communications, Management/Operations, Security**
- Hermann de Meer, Universität Passau
  Problems in cross layer optimizations contrasted to E2E communications

4 Lessons Learnt resp. the Key Findings Identified

In the following the key findings of the seminar during the discussions are summarized:

1. We need a better understanding of design trade-offs.
2. New applications can change traffic characteristics in a few months. In the past decade several applications dramatically changed the way how the Internet is used. Nobody has actually foreseen the success of P2P networks, and especially Youtube and Facebook. Thus, the question is whether it is possible to design a Future Internet without having any ideas what the "next big things" could be. If thus the traffic changes are unpredictable, then we need to establish a fast and stable infrastructure without any assumptions on the traffic.
3. The waterfall model does not work in practice in communications, for sure, software is not a "one-time instance", changes will occur for some time. Thus, versions are needed, and for protocols we may arrive at the same iterative refinement approach.

4. The use of formal specifications in case of OSI was rather harmful than it helped. OSI failed because nobody needed it, and options have been considered to be harmful.
5. Implementations from the beginning are necessary. However, with first to the market means to loose quality.
6. Full decoupling of planes (management, user, control) is good in an "old-style telco world", however, it will not work in the Future Internet.
7. There is no future for a centralized management (in most cases). It is necessary to move the research effort towards self-management approaches.
8. Future protocols should have built-in management possibilities (Management-by-Design).
9. The focus of management must concentrate on algorithms to automate it, the supporting data and data feeding will follow afterwards.
10. Effective management is the management that disappears or is invisible, respectively.
11. The need for self-configuration in access networks, programmable nodes (measurement is an important case on layer 3).
12. Fundamentally new protocols and approaches are needed that are additionally suitable for addressing replicated objects, which are mutable, as well as trust.
13. Assumptions for DiffServ/IntServ changed, there is no need anymore for end-to-end Quality-of-Service (QoS), except for Label Switched Paths.
14. There is no need for Network Address Translation anymore due to IPv6.
15. IP addresses are unsuitable for the mobility aspect. Mobile IP has failed.

5 Conclusions

All discussions have shown that a lot of various concepts that have been developed are nowadays, or will be in the near future, more or less obsolete due to the enormous amount of bandwidth being available in network backbones or migration to IPv6. These insights of those discussions have shown very interesting combinations of effects and consequences as well as mechanisms and their implementation alternatives. Based on these observations a scientific paper is under preparation.

Participants

- Carsten Bormann
Universität Bremen, DE
- Georg Carle
TU München, DE
- Prosper Chemouil
France Telecom, FR
- Hermann de Meer
Universität Passau, DE
- Gabi Dreo Rodosek
Unibw - München, DE
- Olivier Festor
INRIA - Nancy Grand Est, FR
- Brendan Jennings
Waterford Institute of Technology, IE
- Sape J. Mullender
Alcatel-Lucent Bell Labs - Antwerpen, BE
- Aiko Pras
University of Twente, NL
- Danny Raz
Technion – Haifa, IL
- Ramin Sadre
University of Twente, NL
- Jürgen Schönwälder
Jacobs University - Bremen, DE
- Peter Schoo
Fraunhofer Institut SIT – München, DE
- Henning Schulzrinne
Columbia University, US
- Burkhard Stiller
Universität Zürich, CH

Report from Dagstuhl Seminar 11051

Sparse Representations and Efficient Sensing of Data

Edited by

Stephan Dahlke[1], Michael Elad[2], Yonina Eldar[3], Gitta Kutyniok[4], and Gerd Teschke[5]

1 **Universität Marburg, DE,** dahlke@Mathematik.Uni-Marburg.de
2 **Technion – Haifa, IL,** elad@cs.technion.ac.il
3 **Technion – Haifa, IL,** yonina@ee.technion.ac.il
4 **Universität Osnabrück, DE,** kutyniok@math.uni-osnabrueck.de
5 **Hochschule Neubrandenburg, DE,** teschke@hs-nb.de

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 11051 "Sparse Representations and Efficient Sensing of Data". The scope of the seminar was twofold. First, we wanted to elaborate the state of the art in the field of sparse data representation and corresponding efficient data sensing methods. Second, we planned to explore and analyze the impact of methods in computational science disciplines that serve these fields, and the possible resources allocated for industrial applications.

Seminar 30. January–04. February, 2011 – www.dagstuhl.de/11051

1998 ACM Subject Classification F.2 Analysis of Algorithms and Problem Complexity, G.1 Numerical Analysis, I.4 Image Processing and Computer Vision, J.2 Physical Sciences and Engineering

Keywords and phrases Efficient signal sensing schemes, sparse signal representations, efficient signal reconstruction algorithms, impact of the methods in neighboring research fields and applications.

Digital Object Identifier 10.4230/DagRep.1.1.108

1 Executive Summary

Stephan Dahlke
Michael Elad
Yonina Eldar
Gitta Kutyniok
Gerd Teschke

Modeling of data is the crucial point in enabling various processing of it. This modeling can take many forms and shapes: it can be done in a low-level way that ties the data samples directly or in higher levels that search for structures and constellations. The task of modeling data is so fundamental that it is underlying most of the major achievements in the fields of signal and image processing. This is true also for processing of more general data sources. Indeed, the field of machine learning that addresses this general problem also recognizes the importance of such modeling. In this realm of models, there is one that stands out as quite simple yet very important this is a model based on sparse description of the data. The core idea is to consider the data as a sparse linear combination of core elements, referred to as

atoms. This model has attracted huge interest in the past decade, with many mathematicians, computer scientists, engineers, and scientists from various disciplines working on its different facets, and building a set of tools that lead all the way from pure mathematical concepts to practical tools to be used in other computational sciences as well as applications. Using this model, researchers have shown in recent years a wide battery of computational research disciplines and applications that directly benefit from it, leading to state-of-the- art results. Various reconstruction problems, data compression, sampling and sensing, separation of signals, cleaning and purifying data, adaptive numerical schemes, and more, all require the utilization of sparse representations to succeed in their tasks.

The goals of the seminar can be summarized as follows:
- Establish communication between different focusses of research
- Open new areas of applications
- Manifest the future direction of the field
- Introduce young scientists

To reach these seminar goals, the organizers identified in advance the most relevant fields of research:
- Sampling and Compressed Sensing
- Frames, Adaptivity and Stability
- Algorithms and Applications

The seminar was mainly centered around these topics, and the talks and discussion groups were clustered accordingly. During the seminar, it has turned out that in particular 'generalized sensing', 'data modeling', and corresponding 'algorithms' are currently the most important topics. Indeed, most of the proposed talks were concerned with these three issues. This finding was also manifested by the discussion groups. For a detailed description of the outcome of the discussion, we refer to Section 4.

The course of the seminar gave the impression that sparsity with all its facets is definitely one of the most important techniques in applied mathematics and computer sciences. Also of great importance are associated sampling issues. We have seen many different view points ranging from classical linear and nonlinear to compressive sensing. In particular, new results on generalized sampling show how to design effective sampling strategies for recovering sparse signals. The impact of these techniques became clear as they allow an extension of the classical finite dimensional theory of compressive sensing to infinite dimensional data models. Moreover, it was fascinating to see how sampling and sparsity concepts are by now influencing many different fields of applications ranging from image processing / compression / resolution to adaptive numerical schemes and the treatment of operator equations/inverse problems. It seems that the duality between sparse sampling and sparse recovery is a common fundamental structure behind many different applications. However, the mathematical technicalities remain quite challenging. As algorithmic issues were also discussed quite intensively, we could figure out that we are now essentially at some point where ℓ_1-optimization is competitive speed-wise with classical linear methods such as conjugate gradient.

Summarizing our findings during the seminar, we believe that the research agenda can be more focused on the actual bottlenecks, being in problem/signal modeling, design of sampling and recovery methods adapted to specific problems, and algorithmic improvements including performance bounds and guarantees.

2 Table of Contents

3 Overview of Talks

3.1 Multivariate Periodic Function Spaces

Ronny Bergmann (Universität zu Lübeck, DE)

Joint work of Bergmann, Ronny; Prestin, Jürgen

In this talk we present patterns $\mathcal{P}(\mathbf{M})$ based on invertible matrices $\mathbf{M} \in \mathbb{Z}^{d\times d}$ as multivariate generalizations of the equidistant points on $[0, 1)$. We present different properties of the patterns, e.g. the classification of subpatterns and a dual group, that is used to define a discrete Fourier transform with respect to $\mathbf{M}$.

Using the pattern to generate translates of a square integrable function f defined on the d-dimensional 2π-periodic torus, we introduce a translation invariant space V^f. This space can be characterized by the Fourier series of f and the aforementioned Fourier transform on $\mathcal{P}(\mathbf{M})$. The same is possible for subspaces that are generated by different translates of f or even translates of other functions in V^f. Finally this yields properties for a decomposition of V^f into two orthogonal subspaces that are translation invariant with respect to a certain subpattern of $\mathcal{P}(\mathbf{M})$.

3.2 Random tight frames and applications

Martin Ehler (National Institutes of Health – Bethesda, US)

Joint work of Ehler, Martin; Galanis, Jennifer
Main reference M. Ehler, J. Galanis, "Frame theory in directional statistics," Stat. Probabil. Lett., in press.
URL http://dx.doi.org/10.1016/j.spl.2011.02.027

It is known that independent, uniformly distributed points on the sphere approximately form a finite unit norm tight frame. In this talk, we introduce probabilistic frames to more deeply study finite frames whose elements are chosen at random. In fact, points chosen from any probabilistic tight frame approximately form a finite tight frame; they do not have to be uniformly distributed, nor have unit norm. We also observe that classes of random matrices used in compressed sensing are induced by probabilistic tight frames.

Finally, we merge directional statistics with frame theory to elucidate directional statistical testing. Distinguishing between uniform and non-uniform sample distributions is a common problem in directional data analysis; however for many tests, non-uniform distributions exist that fail uniformity rejection. We find that probabilistic tight frames yield non-uniform distributions that minimize directional potentials, leading to failure of uniformity rejection for the Bingham test. We apply our results to model patterns found in granular rod experiments.

Related references are:

References

1 M. Ehler, "Random tight frames," arXiv:1102.4080v1.
2 M. Ehler, K.A. Okoudjou, "Minimization of the probabilistic p-frame potential," arXiv:1101.0140.

3.3 Alternating Direction Optimization for Imaging Inverse Problems with Sparsity-Inducing Regularization

Mario Figueiredo (TU Lisboa, PT)

Joint work of Figueiredo, Mário; Bioucas-Dias, José; Afonso, Manya

Most modern approaches (regularization-based or Bayesian) to imaging inverse problems lead to optimization problems. These (usually convex) problems have features that place them beyond the reach of off-the-shelf optimization methods and have stimulated a significant amount of research. In particular, the presence of regularizers encouraging sparse solutions imply the non-smoothness of the objective function, which together with its typical very high dimensionality constitutes a challenge.

Examples of this include frame-based regularization (either in an analysis or synthesis formulation), where the classical regularizer involves the ℓ_1 norm.

This talk will cover our recent work on the application of a class of techniques known as "alternating direction methods" to several imaging inverse problems with frame-based sparsity-inducing regularization, namely: (a) standard image restoration/reconstruction from linear observations with Gaussian noise; (b) image restoration from Poissonian observations; (c) multiplicative noise removal.

In all these cases, the proposed methods come with theoretic convergence guarantees and achieve state-of-the-art speed, as shown in the reported experiments. To further illustrate the flexibility of this class of methods, we show how it can be used to seamlessly address hybrid analysis/synthesis formulations as well as group-norm regularizers (with or without group overlap).

The work described in this talk was co-authored by José M. Bioucas-Dias and Manya V. Afonso, and reported in the following publications:

References

1 M. Afonso, J. Bioucas-Dias, and M. Figueiredo, "An augmented Lagrangian approach to the constrained optimization formulation of imaging inverse problems", *IEEE Transactions on Image Processing*, vol. 20, pp. 681–695, 2011.

2 M. Afonso, J. Bioucas-Dias, and M. Figueiredo, "Fast image recovery using variable splitting and constrained optimization", *IEEE Transactions on Image Processing*, vol. 19, pp. 2345–2356, 2010.

3 J. Bioucas-Dias, and M. Figueiredo, "Multiplicative noise removal using variable splitting and constrained optimization", *IEEE Transactions on Image Processing*, vol. 19, pp. 1720–1730, 2010.

4 M. Figueiredo and J. Bioucas-Dias "Restoration of Poissonian images using alternating direction optimization", *IEEE Transactions on Image Processing*, vol. 19, pp. 3133–3145, 2010.

3.4 Adaptive wavelet methods for inverse parabolic problems

Ulrich Friedrich (University of Marburg, DE)

We combine adaptive wavelet techniques for well-posed problems and regularization theory for inverse problems. We are concerned with identifying certain parameters in a parabolic reaction-diffusion equation from measured data. The PDE describes the gene concentrations in embryos at an early state of development. The forward problem is formulated as an evolution equation, and the analytical properties of the parameter-to-state operator are analyzed. The results justify the application of an iterated soft-shrinkage algorithm within a Tikhonov regularization approach. The forward problem is treated by means of a new adaptive wavelet algorithm which is based on tensor wavelets. A generalized anisotropic tensor wavelet basis dealing with complicated domains is given. This leads to dimension independent convergence rates. An implementation of this procedure involving the new adaptive wavelet solver is proposed and numerical results are presented.

3.5 A Majorization-minimization algorithm for sparse factorization and some related applications

Onur G. Guleryuz (DoCoMo USA Labs – Palo Alto, US)

Joint work of Guleryuz, Onur G.; Chou, Evan; Zujovic, Jana

Given a data matrix, X, we are interested in its approximations of the form $\hat{X} = TC$ where T and C are two sparse matrices. The optimization setup is

$$\min_{T,C} ||X - TC||_2 \text{ subject to } ||T||_0 + ||C||_0 \leq \kappa.$$

The problem arises in

(a) accelerating matrix multiplication where an a priori known X is to be multiplied with a dense matrix S, with S only available online.

(b) approximation where $X = Y + W$, with $Y = TC$, and the goal is to recover Y (i.e., structured signal under noise, structured signal with missing data, etc.)

(c) compression, with T as the data-adaptive basis, C as the matrix of coefficient vectors, and assuming one is using a nonlinear-approximation-based compression algorithm.

X is specified in a domain spanned by two known orthonormal matrices which we call the presentation basis. Many disciplines approach these applications using the SVD of X, disregarding the last twenty-plus years of research. Our aim is to move toward "generic signal processing" where DSP techniques can be used to match/better existing results with minor or no domain-specific information.

References

1 M. Döhler, S. Kunis, and D. Potts. Nonequispaced hyperbolic cross fast Fourier transform. *SIAM J. Numer. Anal.*, 47:4415 – 4428, 2010.

2 M. Aharon, M. Elad, and A.M. Bruckstein, The K-SVD: An Algorithm for Designing of Overcomplete Dictionaries for Sparse Representation, *IEEE Trans. On Signal Processing* **54(11)** (2006), 4311-4322.

3 J. Zujovic and O. G. Guleryuz, Complexity Regularized Pattern Matching, *Proc. IEEE Int'l Conf. on Image Proc. (ICIP2009), Cairo, Egypt*, Nov. 2009.

4 Rubinstein, R., Zibulevsky, M. and Elad, M., Double Sparsity: Learning Sparse Dictionaries for Sparse Signal Approximation, *IEEE Transactions on Signal Processing* **58(3)** (2010).

5 E. Chou and O. G. Guleryuz, A Majorization-minimization algorithm for sparse factorization and related applications, (in preparation).

6 O. G. Guleryuz, Nonlinear Approximation Based Image Recovery Using Adaptive Sparse Reconstructions and Iterated Denoising: Parts I and II, *IEEE Transactions on Image Processing*, (2006).

3.6 Stable discretizations of the hyperbolic cross fast Fourier transform

Stefan Kunis (Universität Osnabrück, DE)

Joint work of Kämmerer, Lutz; Kunis, Stefan; Potts, Daniel

A straightforward discretization of problems in d spatial dimensions with 2^n, $n \in \mathbb{N}$, grid points in each coordinate leads to an exponential growth 2^{dn} in the number of degrees of freedom.

We restrict the frequency domain to the hyperbolic cross

$$H_n^d = \bigcup_{\mathbf{j}\in\mathbb{N}_0^d, \|\mathbf{j}\|_1=n} (-2^{j_1-1}, 2^{j_1-1}] \times \ldots \times (-2^{j_d-1}, 2^{j_d-1}] \cap \mathbb{Z}^d,$$

and ask for the fast approximate evaluation of the trigonometric polynomial

$$f(\mathbf{x}) = \sum_{\mathbf{k}\in H_n^d} \hat{f}_{\mathbf{k}}\, \mathrm{e}^{2\pi i \mathbf{k}\mathbf{x}}, \tag{1}$$

at nodes $\mathbf{x}_\ell \in \mathbb{T}^d$, $\ell = 1, \ldots, M$.

We note that the reduced problem size is $c_d 2^n n^{d-1} \le |H_n^d| \le C_d 2^n n^{d-1}$ and a classical result states the computation of (1) for all sparse grid nodes takes at most $C_d 2^n n^d$ floating point operations.

This has been generalized for arbitrary spatial sampling nodes and both algorithms are available in the Matlab toolbox `nhcfft`.

▶ **Theorem 1.** *[1] The computation of* (1) *at all nodes* $\mathbf{x}_\ell \in \mathbb{T}^d$, $\ell = 1, \ldots, |H_n^d|$, *takes at most* $C_d 2^n n^{2d-2}(|\log\varepsilon| + \log n)^d$, *where* $\varepsilon > 0$ *denotes the target accuracy.*

More recently, we analyzed the numerical stability of these sampling sets and in sharp contrast to the ordinary FFT which is unitary, we found the following negative result.

▶ **Theorem 2.** *[2] The computation of* (1) *at the sparse grid has condition number*

$$c_d 2^{\frac{n}{2}} n^{\frac{2d-3}{2}} \le \kappa \le C_d 2^{\frac{n}{2}} n^{2d-2}.$$

Although random sampling offers a stable spatial discretization with high probability if $M \ge C|H_n^d| \log |H_n^d|$, the fast algorithm [1] relies on an oversampled sparse grid and thus suffers from the same instability.

Ongoing work [3] considers lattices as spatial discretization for the hyperbolic cross fast Fourier transform. These turn out to have quite large cardinality asymptotically but offer perfect stability and outperform known algorithms by at least one order of magnitude with respect to CPU timings for moderate problem sizes.

References

1 M. Döhler, S. Kunis, and D. Potts. Nonequispaced hyperbolic cross fast Fourier transform. *SIAM J. Numer. Anal.*, 47:4415 – 4428, 2010.
2 L. Kämmerer and S. Kunis. On the stability of the hyperbolic cross discrete Fourier transform. *Numer. Math.*, 117:581 – 600, 2011.
3 L. Kämmerer, S. Kunis, and D. Potts. Interpolation lattices for hyperbolic cross trigonometric polynomials. *Preprint 2010-13, Fakultät für Mathematik, TU Chemnitz*, 2010.

3.7 Optimally Sparse Image Approximations Using Compactly Supported Shearlets

Wang-Q Lim (Universität Osnabrück, DE)

It is now widely acknowledged that analyzing the intrinsic geometrical features of a function/signal is essential in many applications. In order to achieve this, several directional systems have been proposed in the past. The first breakthrough was achieved by Candes and Donoho who introduced curvelets and showed that curvelets provide an optimal approximation property for a special class of 2D piecewise smooth functions, called cartoon-like images. However, only band-limited directional systems providing an optimal approximation property have been constructed so far, except adaptive representation schemes.

In this talk, we will show that optimally sparse approximation of cartoon-like images can be achieved using compactly supported shearlet frames in both 2D and 3D. We then briefly discuss our ongoing work to construct a compactly supported directional system which is not only a tight frame but also provides optimally sparse approximation of cartoon-like images.

3.8 Exact test instances for Basis Pursuit Denoising

Dirk Lorenz (TU Braunschweig, DE)

Main reference Dirk A. Lorenz, "Constructing test instances for Basis Pursuit Denoising," submitted for publication, 2011.
URL http://arxiv.org/abs/1103.2897

The number of available algorithms for the so-called Basis Pursuit Denoising problem (or the related LASSO-problem) is large and keeps growing. Similarly, the number of experiments to evaluate and compare these algorithms on different instances is growing. However, many comparisons lack of test instances for which exact solutions are known. We propose to close this gap by a procedure which calculates a right hand side from a given matrix, regularization parameter and a given solution. It can be shown that this can be accomplished by means of projection onto convex sets (POCS) or quadratic programming. The method has been implemented in MATLAB and is available as part of L1TestPack from http://www.tu-braunschweig.de/iaa/personal/lorenz/l1testpack.

3.9 Cosparse Analysis Modeling – Uniqueness and Algorithms

Sangnam Nam (INRIA – Rennes, FR)

In the past decade there has been a great interest in a synthesis-based model for signals, based on sparse and redundant representations. Such a model assumes that the signal of interest can be composed as a linear combination of few columns from a given matrix (the dictionary). An alternative analysis-based model can be envisioned, where an analysis operator multiplies the signal, leading to a cosparse outcome. In this work, we consider this analysis model, in the context of a generic missing data problem (e.g., compressed sensing, inpainting, source separation, etc.). Our work proposes a uniqueness result for the solution of this problem, based on properties of the analysis operator and the measurement matrix. We also considers two algorithms for solving the missing data problem, an L1-based and a new greedy method. Our simulations demonstrate the appeal of the analysis model, and the success of the pursuit techniques presented.

References

1 M. Elad, P. Milanfar, and R. Rubinstein, "Analysis versus synthesis in signal priors," *Inverse Problems*, vol. 23, no. 3, pp. 947–968, June 2007.
2 I. W. Selesnick and M. A. T. Figueiredo, "Signal restoration with overcomplete wavelet transforms: Comparison of analysis and synthesis priors," *In Proceedings of SPIE*, vol. 7446 (Wavelets XIII), August 2009.
3 E. J. Candès, Y. C. Eldar, D. Needell, and Y. Ma, "Compressed sensing with coherent and redundant dictionaries," Tech. Rep. 1005.2613v1, arXiv, 2010.
4 J. Portilla, "Image restoration through l0 analysis-based sparse optimization in tight frames," *Proceedings of the 16th IEEE International Conference on Image Processing*, pp. 3865–3868, 2009.
5 R. Gribonval and M. Nielsen, "Sparse representations in unions of bases," *IEEE Trans. Inform. Theory*, vol. 49, no. 12, pp. 3320–3325, Dec. 2003.
6 D. L. Donoho and M. Elad, "Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ^1 minimization," *Proc. Nat. Aca. Sci.*, vol. 100, no. 5, pp. 2197–2202, Mar. 2003.
7 Y. M. Lu and M. N. Do, "Sampling signals from a union of subspaces," *IEEE Signal Processing Magazine*, pp. 41–47, mar 2008.
8 T. Blumensath and M. E. Davies, "Sampling theorems for signals from the union of finite-dimensional linear subspaces," *IEEE Transactions on Information Theory*, vol. 55, no. 4, pp. 1872–1882, 2009.
9 M. Grant, S. Boyd, and Y. Ye, "CVX: Matlab Software for Disciplined Convex Programming," http://cvxr.com/cvx, August 2008.

3.10 Space Splittings and Schwarz-Southwell Iterations

Peter Oswald (Jacobs University – Bremen, DE)

Joint work of Griebel, Michael; Oswald, Peter

The talk reviews joint work in progress with M. Griebel [1], which continues our collaboration on iterative solvers for symmetric elliptic variational problems that are based on Hilbert space splittings, so-called additive and multiplicative Schwarz methods. Note that Hilbert space splittings underlying the theory of Schwarz methods have lately reappeared as fusion frames. See [2] for a review of previous work.

While in the standard theory of multiplicative Schwarz methods the order of subproblem traversal is fixed, in the new versions the ordering is chosen in a weak greedy fashion, e.g., according to the size of subproblem residuals, or randomly. For linear systems and Gauss-Seidel methods (a special instance of multiplicative Schwarz methods) the greedy ordering goes back to Gauss and Seidel, and has been popularized by Southwell in the 1940-50ies. The method has been theoretically analyzed in the framework of coordinate descent methods for convex optimization methods, and has lately been revived in the context of sparse approximation.

Given these developments, we decided to first formulate and prove convergence results for Schwarz-Southwell methods for the case of splittings into N subproblems. The main result is an exponential energy error decay estimate of the form

$$\|u - u^{(m+1)}\|_E^2 \leq (1 - \frac{\gamma}{N})\|u - u^{(m)}\|_E^2, \qquad m \geq 0,$$

where γ depends on the spectral bounds characterizing the space splitting, the relaxation parameter ω, and the weakness parameter β of the weak greedy step. The result shows that greedy strategies can slightly improve the performance of multiplicative Schwarz methods. We also state a similar estimate for the expected convergence rate if the subproblem ordering is randomized. Investigations on infinite splittings are still at their beginning, they benefit from the theory of greedy algorithms in infinite-dimensional Hilbert and Banach spaces developed by Temlyakov and others. We hope that a better understanding of this topic will shed new light on adaptive multilevel methods such as the early work by Rüde.

References

1 M. Griebel, P. Oswald, *Schwarz-Southwell Methods*. In preparation.
2 P. Oswald, *Stable space splittings and fusion frames*. In Wavelets XIII (V.K. Goyal, M. Papadakis, D. Van de Ville, eds.), Proc. SPIE Vol. 7446 (SPIE, Bellingham, 2009), 744611.

3.11 Sparse Approximation of Images by the Easy Path Wavelet Transform

Gerlind Plonka-Hoch (Universität Göttingen, DE)

The Easy Path Wavelet Transform (EPWT) has recently been proposed as a tool for sparse representations of bivariate functions from discrete data, in particular from image data.

The EPWT is a locally adaptive wavelet transform. It works along pathways through the array of function values and it exploits the local correlations of the given data in a simple appropriate manner.

Using polyharmonic spline interpolation, we show that the EPWT leads, for a suitable choice of the pathways, to optimal N-term approximations for piecewise Hölder smooth functions with singularities along curves.

3.12 Quadrature errors, discrepancies and variational dithering

Gabriele Steidl (Universität Kaiserslautern, DE)

Main reference M. Gräf, D. Potts and G. Steidl, Quadrature errors, discrepancies and their relations to halftoning on the torus and the sphere, Preprint TU Chemnitz, Fakultät für Mathematik, Preprint 5, 2011.
URL http://www.mathematik.uni-kl.de/ steidl

The stippling technique places black dots such that their density gives the impression of tone.

The original idea for considering minimizers of this functional as 'good' dot positions comes from electrostatic principles. This talk is related to the continuous version of the above attraction-repulsion functional with more general functions $\varphi : [0,\infty) \to \mathbb{R}$:

$$E_\varphi(p) := \frac{\lambda}{2} \sum_{i,j=1}^{M} \varphi(\|p_i - p_j\|_2) - \sum_{i=1}^{M} \int_{[0,1]^2} w(x)\varphi(\|p_i - x\|_2)\, dx, \tag{2}$$

where $w : [0,1]^2 \to [0,1]$ and $\lambda := \frac{1}{M}\int_{[0,1]^2} w(x)\, dx$. The function $\varphi(r) = -r$ was used in as well as $\varphi(r) = -\log(r)$. In another paper, the authors mentioned $\varphi(r) = -r^\tau$, $0 < \tau < 2$ and $\varphi(r) = r^{-\tau}$, $\tau > 0$ for $r \neq 0$. In this talk we relate stippling processes with the classical mathematical question of finding best nodes for quadrature rules. We provide theoretical results on the connection between seemingly different concepts, namely quadrature rules, attraction-repulsion functionals, L_2–discrepancies and least squares functionals. For the later approach we provide numerical minimization algorithms. In the theoretical part, we start with worst case quadrature errors on RKHSs in dependence on the quadrature nodes. While in the literature, this was mainly done for constant weights $w \equiv 1$, see [5], we incorporate a weight function related to the image into the quadrature functional. The corresponding quadrature error $\mathrm{err}_K(p)$ which depends on the reproducing kernel K can be defined for RKHSs on $\mathcal{X} \in \{\mathbb{R}^2, [0,1]^2\}$ as well as for RKHSs on compact manifolds like $\mathcal{X} \in \{T^2, S^2\}$. We aim to minimize this quadrature error in order to obtain optimal quadrature nodes p. It turns out that for special kernels K (on special spaces $\mathcal{X}$) this quadrature error (or at least its minimizers) covers the following approaches:

1. Attraction-Repulsion Functionals
   An interesting case of RKHSs appears for radial kernels $K(x,y) = \varphi(\|x-y\|_2)$ depending only on the distance of the points. We will show that in this case the quadrature error $\mathrm{err}_K(p)$ can be considered as a generalization of (2) which works not only on $[0,1]^2$ but also to compact manifolds. Hence our approach goes far beyond the setting in [1] or [2].
2. L_2–Discrepancies
   We prove that for $\mathcal{X} \in \{[0,1]^2, T^2, S^2\}$ and discrepancy kernels K, the quadrature errors on RKHSs defined by these kernels coincide with L_2–discrepancy functionals. For various applications of L_2–discrepancy functionals, see [5] and the references therein. Note

that a relation between the distance kernels $K(x,y) = \|x-y\|_2$ on T^2 and S^2 and the corresponding discrepancy kernels was shown numerically in [3].

3. Least Squares Functionals
   Finally, we consider RKHSs of bandlimited functions with bandlimited kernels on $\mathcal{X} \in \{T^2, S^2\}$. The reason for addressing these spaces is that we want to approximate functions on $\mathcal{X}$ by bandlimited functions in order to apply fast Fourier techniques. We prove that for these RKHSs the quadrature error can be rewritten as a least squares functional.

In the numerical part we approximate functions and kernels on $\mathcal{X} \in \{T^2, S^2\}$ by their bandlimited versions and minimize the corresponding quadrature error which takes in this case the form of a least squares functional. Due to the page limitation we restrict our attention to the sphere S^2. We are not aware of any results on S^2–stippling in the literature. We propose a nonlinear CG method on manifolds to compute a minimizer of the least squares functional on S^2. This method was also successfully used for the approximation of spherical designs, i.e., for $w \sim 1$ in [4] and is generalized in this paper. In particular, each CG step can be realized in an efficient way by the *nonequispaced fast spherical Fourier transform* (NFSFT). This reduces the asymptotic complexity of the proposed algorithm drastically, e.g., from $\mathcal{O}(MN^2)$ to $\mathcal{O}(N^2 \log^2 N + M \log^2(1/\epsilon))$ arithmetic operations per iteration step, where ϵ is the described accuracy and N corresponds to the bandwidth. In other words, only by the help of the NFSFT the computation becomes possible in a reasonable time.

References

1 T. Teuber, G. Steidl, P. Gwosdek, Ch. Schmaltz and J. Weickert, Dithering by differences of convex functions, *SIAM J. on Imaging Science* **4(1)**, 79 - 108, 2011.

2 C. Schmaltz, P. Gwosdek, A. Bruhn and J. Weickert, Electrostatic halftoning, *Computer Graphics Forum*, **29**, 2313 - 2327, 2010.

3 M. Gräf, D. Potts and G. Steidl, Quadrature errors, discrepancies and their relations to halftoning on the torus and the sphere, *TU Chemnitz, Fakultät für Mathematik, Preprint 5*, 2011.

4 M. Gräf and D. Potts, On the computation of spherical designs by a new optimization approach based on fast spherical Fourier transforms, *Preprint, Technical University Chemnitz*, 2010.

5 E. Novak and H. Woźniakowski, *Tractability of Multivariate Problems Volume II: Standard Information for Functionals*, Eur. Math. Society, EMS Tracts in Mathematics Vol 12, 2010.

3.13 Compressive sensing and inverse problems

Gerd Teschke (Hochschule Neubrandenburg, DE)

Joint work of Teschke, Gerd; Herrholz, Evelyn

Main reference Evelyn Herrholz, Gerd Teschke, "Compressive sensing principles and iterative sparse recovery for inverse and ill-posed problems," Inverse Problems, vol. 26, no. 12, 125012.

URL http://dx.doi.org/10.1088/0266-5611/26/12/125012

We shall be concerned with compressive sampling strategies and sparse recovery principles for linear inverse and ill-posed problems. As the main result, we provide compressed measurement models for ill-posed problems and recovery accuracy estimates for sparse approximations of the solution of the underlying inverse problem. The main ingredients are variational formulations that allow the treatment of ill-posed operator equations in the context of compressively sampled data. In particular, we rely on Tikhonov variational and constrained

optimization formulations. One essential difference to the classical compressed sensing framework is the incorporation of joint sparsity measures allowing the treatment of infinite dimensional reconstruction spaces. The theoretical results are furnished with a number of numerical experiments.

We rely on a signal model, $X_k = \{x \in X,\, x = \sum_{\ell\in\mathcal{I},|\mathcal{I}|=k} \sum_{\lambda\in\Lambda} d_{\ell,\lambda} a_{\ell,\lambda},\ \ d \in (\ell_2(\Lambda))^m\}$, in which we assume that the coefficients $d_{\ell,\lambda}$ share a joint sparsity pattern (only k out of m sequences $\{d_{\ell,\lambda}\}_{\lambda\in\Lambda}$ do not vanish). The space X_k can be seen as a union of (shift invariant) subspaces. One approach to recover d was suggested in [Y. C. Eldar. Compressed Sensing of Analog Signals in Shift-Invariant Spaces. IEEE Trans. on Signal Processing, 57(8), 2009.]. We propose an alternative by solving adequate variational problems. The essential idea to tackle the support set recovery problem is to involve the joint sparsity measure $\Psi_{q,r}(d) = (\sum_{\ell=1}^m (\sum_{\lambda\in\Lambda} |d_{\ell,\lambda}|^r)^{\frac{q}{r}})^{\frac{1}{q}}$. This measure promotes a selection of only those indices $\ell \in \{1,\dots,m\}$ for which $\|\{d_{\ell,\lambda}\}_{\lambda\in\Lambda}\|_{\ell_r(\Lambda)}$ is large enough, i.e. where the size of the coefficients $d_{\ell,\lambda}$ indicates a significant contribution to the representation of x. In order to define an adequate variational formulation, we have to introduce a suitable sensing model. Assume the data y are obtained by sensing Kx through F_s (a compressed version of F_v), i.e. $y = F_sKx = F_sKF_a^*d = AF_{K^*v}F_a^*d$, where K is an ill-posed but bounded linear operator and A a sensing matrix satisfying a $2k$-RIP with isometry constant δ_{2k}. If the ansatz systems Φ_v and Φ_a diagonalize K, we can write a noisy measurement scenario as follows

$$y^\delta = (TD)d + z \quad \text{with} \quad \|z\|_{(\ell_2(\Lambda))^m} \le \delta\ ,$$

where T describes the application of A with respect to each λ and where D describes the diagonal matrix performing the application of K. To derive an approximation d^* to the solution d of the inverse problem, we propose to solve the following constrained optimization problem

$$\min_{d\in B(\Psi_{1,2},R)} \|y^\delta - (TD)d\|^2_{(\ell_2(\Lambda))^p} + \alpha\|d\|^2_{(\ell_2(\Lambda))^m}\ . \tag{3}$$

The minimizing element d^* of (3) is iteratively approximated by

$$d^{n+1} = \mathbb{P}_R\left(D^*T^*(y^\delta - TDd^n)\frac{\gamma^n}{C} + \left(1 - \frac{\alpha\gamma^n}{C}\right)d^n\right)\ . \tag{4}$$

We can provide the following accuracy estimate for d^*.

▶ **Theorem 1.** *Assume R was chosen such that the solution d of problem $y = (TD)d$ does not belong to $B(\Psi_{1,2}, R)$ and that $0 \le \delta_{2k} < \frac{(1+\sqrt{2})\kappa^2_{min} - \kappa^2_{max} + \sqrt{2}\alpha}{(1+\sqrt{2})\kappa^2_{min} + \kappa^2_{max}}$. Then the minimizer d^* of (3) satisfies*

$$\|d^* - d\|_{(\ell_2(\Lambda))^m} \le C_0 k^{-1/2}\Psi_{1,2}(d^k - d) + C_1\|L(d^\dagger - d)\|_{(\ell_2(\Lambda))^m} + C_2\delta + C_3\sqrt{\alpha}R\ ,$$

where the constants C_0, C_1, C_2, and C_3 are given explicitly and where d^k denotes the best k-row approximation.

3.14 Sampling in the Age of Sparsity

Martin Vetterli (EPFL – Lausanne, CH)

Sampling is a central topic not just in signal processing and communications, but in all fields where the world is analog, but computation is digital. This includes sensing, simulating, and

rendering the real world.

The question of sampling is very simple: when is there a one-to-one relationship between a continuous-time function and adequately acquired samples of this function? Sampling has a rich history, dating back to Whittaker, Nyquist, Kotelnikov, Shannon and others, and is an active area of contemporary research with fascinating new results.

The classic result of sampling is the one on bandlimited functions, where taking measurements at the Nyquist rate (or twice the maximum bandwidth) is sufficient for perfect reconstruction. These results were extended to shift-invariant subspaces and multiscale spaces during the development of wavelets, as well as in the context of splines.

All these methods are based on subspace structures, and on linear approximation. Recently, non-linear methods have appeared. Non-linear approximation in wavelet spaces has been shown to be a powerful approximation and compression method. This points to the idea that functions that are sparse in a basis (but not necessarily on a fixed subspace) can be represented efficiently.

The idea is even more general than sparsity in a basis, as pointed out in the framework of signals with finite rate of innovation. Such signals are non-bandlimited continuous-time signals, but with a parametric representation having a finite number of degrees of freedom per unit of time. This leads to sharp results on sampling and reconstruction of such sparse continuous-time signals, namely that 2K measurements are necessary and sufficient to perfectly reconstruct a K-sparse continuous-time signal. In accordance with the principle of parsimony, we call this sampling at Occam's rate. We indicate an order K^3 algorithm for reconstruction, and describe the solution when noise is present, or the model is only approximately true.

Next, we consider the connection to compressed sensing and compressive sampling, a recent approach involving random measurement matrices. This is a discrete time, finite dimensional set up, with strong results on possible recovery by relaxing the ℓ_0 into ℓ_1 optimization, or using greedy algorithms.

These methods have the advantage of unstructured measurement matrices (actually, typically random ones) and therefore a certain universality, at the cost of some redundancy. We compare the two approaches, highlighting differences, similarities, and respective advantages.

Finally, we move to applications of these results, which cover wideband communications, noise removal, distributed sampling, and super-resolution imaging, to name a few. In particular, we describe a recent result on multichannel sampling with unknown shifts, which leads to an efficient super-resolution imaging method.

3.15 Digital Shearlet Transform on Pseudo-Polar Grids

Xiaosheng Zhuang (Universität Osnabrück, DE)

We developed a rationally designed digital shearlet theory which is the digitization of the existing shearlet theory for continuum data. Our implementation of the digital shearlet transform is based on utilization of the pseudo-polar Fourier transform, which provide a natural implementation for digital shearlets on the discrete image domain. The pseudo-polar Fourier transform without weighting is generally not an isometry. Isometry can be achieved by careful weighting of the pseudo-polar grid, yet it is difficult to obtain such a weight function.

We showed how efficient weight functions can be designed and obtained on the pseudo-polar grids so that almost isometry can be achieved. In addition, we discussed the software package ShearLab that implements the digital shearlet transform. The ShearLab provides various quantitative measures allowing one to tune parameters and objectively improve the implementation as well as compare different directional transform implementations.

4 Discussion Groups and Further Challenges

During the seminar we had three discussion groups. The outcome shall be briefly reviewed in the following three subsections.

4.1 Sampling and Compressed Sensing

Martin Vetterli (EPFL – Lausanne, CH)
Hans-Georg Feichtinger (University of Vienna, AT)

A consensus emerged that there is a fundamental difference between discrete, compressed sensing problems, and continuous, sampling problems. Ignoring this difference will lead to low performance. Several talks highlighted this in the workshop (e.g. Hansen, etc.). The issue of 'designer matrices' for compressed sensing split the audience. On the one hand, some people (e.g. Pfander) said they did not want another discussion on what are good sensing matrices, while others, more concerned about applications, pointed out the importance of fixed (deterministic) and structured (fast) sensing matrices. Random but cyclic matrices where mentioned as a possible alternative. The importance of looking at infinite dimensional problems was reiterated, as well as modeling the underlying physics correctly. This comes into play both in acquisition or analysis (inc. astronomy and parameter estimation) and simulation or synthesis. In sampling and compressed sensing, the role of adaptivity was pointed out as an open problem by Teschke and there was an agreement, even if efficient and/or practical methods have yet to be found. The link to information based complexity was made, since this is a general framework for function class approximation from samples. However, this is a non-constructive theory. Learning problems, be it dictionary learning for compressed sensing or smoothing kernels (point spread functions) for sampling, play also an important role for sparse methods to become practical.

In sum, a certain maturity has been reached, and it is much more clear where compressed sensing or sparse sampling can be used and when not. The research agenda can thus be more focused on the actual bottlenecks, being in modeling, design of methods adapted to specific problems, or algorithmic improvements including performance bounds and guarantees.

4.2 Frames, Adaptivity and Stability

Rob Stevenson (University of Amsterdam, NL)
Peter Oswald (Jacobs University - Bremen, DE)

M. Fornasier pointed to adaptive methods for frame expansions and mentioned to consider hybrid frames as redundancy by itself does not improve the rate of convergence. There is an requirement of using additional structure for obtaining faster convergence. A basic idea could be to use principles from compressed sensing, i.e. use union of bases from sufficiently diverse types, e.g smooth and non-smooth, quarkonians. The questions arises how to define optimality. P. Oswald answered that there can't be a proper definition of optimality besides checking for particular examples. The idea came up to add another parameter for redundancy, i.e. just enrich existing frames and use a measure for mutual in-/coherence. M. Fornasier reminds of work of Bruno Torresani (acoustic signal of 'Glockenspiel' including a transient): when a hammer attacks, followed by a harmonic, best N-term approximation needs bases of different types, and alternating optimization with respect to these different bases, i.e. algorithmic issues have to be taken into account. R. Stevenson reported on numerical tests for Laplace problems using an adaptive scheme (effect of Richardson extrapolation, results depend on sparsity structure on coarse scales). P. Oswald asked the question differently: don't ask which frames (union of two bases) should be used, but rather ask: given a frame, which signals can be approximated efficiently. M. Fornasier mentioned that the typical criterion for choosing the next building block is based on checking the residual, which is justified in case of wavelets and other bases due to incoherency. P. Oswald answered that just looking at incoherence doesn't solve the problem, it just reports numerical examples. U. Friedrich said that the choice of 'different' bases is application driven; he also asked whether there is a theoretical founded approach, which also allows to prove optimal approximation rates? P. Maaß answered: for convergence rates for solving inverse problems one requires a source condition of type $A^w \in \partial R_p(u^\dagger)$ has to be satisfied, i.e. the choice of the frame has to be linked to the operator. This would lead to 'natural' frames in the range of A, however they are poor for approximating $u^\dagger$. R. Steveson added to the requirements for choosing frames: the approximation properties have to be better than for each individual bases, computation must be computationally effective, frames must be sufficiently incoherent in order to allow for adaptive schemes (report on numerical tests using Schwarz iteration for overlapping grids (similar to old Randy Bank paper in the late 90's)). T. Raasch said that the 'communication' between the frame and the residuals is of great importance and requires the incoherence of cross-correlation between frame elements. O. Holtz pointed to connections to 'deterministic' compressed sensing approaches (conjecture of Bougain). P. Oswald asked how to design frames which are able of capturing both, transport and diffusion simultaneously in physical application).

4.3 Algorithms and Applications

Michael Unser (EPFL – Lausanne, CH)
Bruno Torresani (Université de Provence – Marseille, FR)

The development of efficient algorithms for large-scale sparse signal recovery is probably one of the greatest practical benefits of the intense research effort around the theme of compressed sensing. In addition to CS, these methods are relevant for a large variety of conventional imaging and inverse problems in general, often delivering state-of-the-art performance. Iterative soft-thresholding algorithms (ISTA) were introduced by image processing pioneers about a decade ago and have been improved systematically since then. An important step was to recognize that a good part of the required mathematical machinery for convex optimization (proximal operators, backward-forward splitting techniques) had been established in the 70s and was directly transposable to the present class of problems. We now have at our disposal a large number of variants of ISTA some of which provide orders of magnitude speed improvements. While the discussion group recognized the pressing need for a general purpose and highly-performant algorithm for ℓ_1-type optimization, they also felt that a point of saturation had been reached and that it would difficult to come up with much faster schemes. What is also required is some careful benchmarking and critical comparison of methods. Yet, the participants also felt that there would still be space for creative engineering to come up with optimized schemes that take advantage of the specificities of certain classes of problems. More important than the algorithm is the problem formulation; in particular, the way of introducing prior information on the class of desirable solutions. It was pointed out that imposing sparsity analysis constraints would generally yield better results than the typical synthesis formulation (sparse generative model) that is central to CS. This is in line with the older concept of regularization as well as the Bayesian formulation of the reconstruction task. The panel did identify the following research challenges:

- the development of efficient large-scale optimization methods for extended classes of non-convex functionals
- the search for better sparsifying transforms and dictionaries (synthesis vs. analysis)
- the design of better regularization functionals
- (non-Gaussian) Bayesian formulations and the derivation/characterization of optimal estimators
- error and sensitivity analysis.

It was also pointed that it is in principle harder to design data-driven dictionaries for analysis purposes rather than synthesis. The concept of sparsity promotes simplicity; it provides a data-processing version of Occam's razor that is most attractive for algorithm design and probably here to stay (once the hype around CS has settled down). There are still many opportunities ahead for applying those techniques to real-world imaging and signal processing problems, beyond the typical tasks of denoising, deblurring and in-painting.

We are now essentially at the point where ℓ_1-optimization is competitive speed-wise with the classical linear methods such as conjugate gradient.

5 Seminar schedule

Monday	Vetterli	Sampling in the Age of Sparsity
	Hansen	Generalized Sampling and Infinite Dimensional Compressed Sensing
	Unser	Stocastic models for sparse and piecewise smooth signals
	Pfander	Sensing, Local Approximation and Quantization of Operators with Bandlimited Kohn-Nirenberg Symbols
	Ehler	Random tight frames in directional statistics
	Pezeshki	Compressed Sensing for High Resolution Image Inversion
Tuesday	Oswald	Space Splittings and Schwarz-Southwell Iterations
	Raasch	Quarkonial frames of wavelet type - Stability and moment conditions
	Starck	Reconstructing and Analyzing Astrophysical Dark Matter Mass Maps using Sparsity
	Lim	Optimally sparse image approximations using compactly supported shearlet frames
	Zhuang	The Digital Shearlet Transform on Pseudo-polar Grids
	Lorenz	Basis pursuit denoising: Exact test instances and exact recovery for ill-posed problems
		Discussion groups
Wednesday	Potts	Parameter estimation for exponential sums by approximate Prony method
	Kunis	Stable discretisations for sparse fast Fourier transforms
	Figueiredo	Alternating Direction Optimization for Imaging Inverse Problems with Sparsity-Inducing Regularization
Thursday	Plonka-Hoch	Sparse Approximation of Images by the Easy Path Wavelet Transform
	Bergmann	Multivariate Periodic Function Spaces
	Steidl	Halftoning, Quadrature Errors and Discrepancies
	Wojtaszczyk	How ℓ_1 -minimisation for RIP matrices reacts to measurement errors?
	Nam	Cosparse Analysis Modeling
	Schiffler	Sparsity for ill-posed and ill-conditioned problems
		Summary of discussion groups
Friday	Guleryuz	A Majorization-minimization algorithm for sparse factorization and some related applications
	Friedrich	Adaptive wavelet methods for inverse parabolic problems

Participants

- Ronny Bergmann
Universität zu Lübeck, DE
- Martin Ehler
National Institutes of Health - Bethesda, US
- Michael Elad
Technion - Haifa, IL
- Hans Georg Feichtinger
Universität Wien, AT
- Mario Figueiredo
TU Lisboa, PT
- Massimo Fornasier
RICAM - Linz, AT
- Ulrich Friedrich
Universität Marburg, DE
- Onur G. Guleryuz
DoCoMo USA Labs - Palo Alto, US
- Anders Hansen
University of Cambridge, GB
- Evelyn Herrholz
Hochschule Neubrandenburg, DE
- Olga Holtz
TU Berlin, DE
- Stefan Kunis
Universität Osnabrück, DE
- Gitta Kutyniok Universität Osnabrück, DE
- Wang-Q Lim
Universität Osnabrück, DE
- Dirk Lorenz
TU Braunschweig, DE
- Peter Maaß
Universität Bremen, DE
- Sangnam Nam
INRIA - Rennes, FR
- Peter Oswald
Jacobs University - Bremen, DE
- Ali Pezeshki
Colorado State University, US
- Götz E. Pfander
Jacobs University - Bremen, DE
- Gerlind Plonka-Hoch
Universität Göttingen, DE
- Daniel Potts
TU Chemnitz, DE
- Thorsten Raasch
Universität Mainz, DE
- Stefan Schiffler
Universität Bremen, DE
- Jean-Luc Starck
CEA - Gif sur Yvette, FR
- Gabriele Steidl
Universität Mannheim, DE
- Rob Stevenson
University of Amsterdam, NL
- Gerd Teschke
Hochschule Neubrandenburg, DE
- Bruno Torresani
Université de Provence, FR
- Michael Unser
EPFL - Lausanne, CH
- Martin Vetterli
EPFL - Lausanne, CH
- Przemyslaw Wojtaszczyk
University of Warsaw, PL
- Xiaosheng Zhuang
Universität Osnabrück, DE