DAGSTUHL
REPORTS

Volume 1, Issue 5, May 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
September, 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.1.5.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 11181

Organic Computing – Design of Self-Organizing Systems

Edited by

Kirstie Bellman[1], Andreas Herkersdorf[2], and Michael G. Hinchey[3]

1 **Aerospace Corp. – Los Angeles, US,** `Kirstie.L.Bellman@aero.org`
2 **TU München, DE,** `herkersdorf@tum.de`
3 **University of Limerick, IE,** `mike.hinchey@lero.ie`

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 11181 "Organic Computing – Design of Self-Organizing Systems".

Seminar 1.–6. May, 2011 – www.dagstuhl.de/11181
1998 ACM Subject Classification F.1.1 Models of computation.
Keywords and phrases Organic Computing, Self-Organization, Adaptivity, Trustworthy System Design
Digital Object Identifier 10.4230/DagRep.1.5.1
Edited in cooperation with Johannes Zeppenfeld

1 Executive Summary

Kirstie Bellman
Andreas Herkersdorf
Mike Hinchey
Christian Müller-Schloer
Hartmut Schmeck
Rolf Würtz

Organic Computing (OC) has become a challenging vision for the design of future information processing systems: As the systems become increasingly more complex, powerful, cheaper and smaller, our environment will be filled with collections of autonomous systems. Autonomous systems are equipped with intelligent sensors and actuators to be aware of their environment, to communicate, and to organize themselves in order to perform the required actions and services. However, these abilities result in even greater system complexity, which we will not be able to explicitly design and manage in every detail, nor are we able to anticipate every possible configuration. Nevertheless, these services have to be as robust, safe, flexible, and trustworthy as possible. In particular, a strong orientation of these systems towards human needs – as opposed to a pure implementation of the technologically possible – is absolutely central.

So far, the OC community, mainly driven by the priority research program of the German Research Foundation (DFG), successfully proposed and – at least partially – established a common nomenclature and terminology for terms like emergence, self-organization, selfadaptation, robustness and flexibility within an interdisciplinary research community.

Quantitative metrics for emergence and self-organization were introduced and applied. Observer controller structures have been established as a common architectural pattern for OC systems within a wide spectrum of applications ranging from traffic control, to Systems on Chip, to collaborative robot systems, to wireless sensor networks. Roles and applicability of different types of supervised and reinforcement-based technical learning techniques were investigated and adapted to OC needs.

Despite the progress in understanding the implications and exploiting the potentials of the OC paradigm, a number of key challenges and research questions still remain. In particular, the planned 2011 OC seminar shall shed light on the various notions of design within the OC context. Design in the classical sense follows a hierarchical top-down constraint propagation starting from a purely functional specification. All eventual environmental influences and disturbances have to be anticipated by the designer at "design time". Due to this anticipatory nature the resulting system is rigid and not able to sufficiently react to run time events.

Complex systems in nature often develop bottom-up due to the self-organizing capabilities of their components. Each component and the system as a whole react to the demands of the environment. In doing so, they are guided by the principles to survive as an individual (selfishness) and the necessity to co-operate (altruism). In technical life-like OC systems we must provide some control by a higher-level entity (finally the user) guiding the bottom-up decisions of the components into a globally desirable direction.

In this way, the former top-down design process dissolves into a balanced run-time negotiation between top-down constraints and bottom-up opportunities. The ultimate consequence of this would mean a total replacement of the design process (at design-time) to controlled self-organization (at runtime).

The 2011 OC seminar was held to answer questions resulting from this shift from design-time to run-time. Is OC a realistic or even desirable vision? How can we replace rigid human designtime control by self-adaptive run-time control without stifling the creativity of the emergent bottom up processes? How can we balance top-down control and bottom-up emergence? Beyond these theoretical questions it is a goal of the seminar to define a number of concrete OC demonstrators – or even a common demonstrator – to be pursued in the sequel.

2 Table of Contents

3 Overview of Talks

3.1 Points of Entry between OC design and traditional and advanced design methodologies

Kirstie Bellman (Aerospace Corp. – Los Angeles, US)

OC has always been aware of the creative tension between the role and capabilities of the human developer and the role and capabilities of the OC processes. We have done an excellent job of building initial OC processes that allow us to research how to substitute for human design, especially, capabilities for a system to respond and self-adapt to an environment at run time. Now is a good time to reconsider broad new roles for OC within the development of complex systems from design through manufacturing.

This talk presents possible new "points of entry" between OC/self- organizational processes and advanced design methodologies. New model- based design-to-manufacturing processes include advancements in design generation, verification, complexity metrics, mathematically formal semantics and other good things. How could OC change/alter the design to manufacturing processes and yet coexist with or even leverage them? How could OC potentially improve the way a human being can interact with/manage a complex system (which includes the use of reflection, language, and collaboration?) The last part of the talk discusses the many challenges for both MBD and OC and suggests that they should be faced with complementary efforts.

3.2 Evolutionary Algorithms to support the design of self-organising manufacturing systems

Jürgen Branke (University of Warwick, GB)

Joint work of Branke, Jürgen; Pickardt Christoph
Main reference Branke, J. and Pickardt, C. Analysing job shop dispatching rules by automatically constructing difficult problem instances. European Journal Of Operational Research (2011)
URL http://www.informs-sim.org/wsc10papers/230.pdf

Abstract: Designing complex, self-organising systems is challenging. It requires to design local, decentralised rules for the agents which result in a good global performance of the overall system. In this talk, two approaches are presented at the example of a self-organising manufacturing system where local dispatching rules are used for decentralised scheduling.

The first approach supports a human designer by revealing the weaknesses of an examined manufacturing system. This is achieved by automatically searching for easy-to-analyse problem instances where the applied dispatching rule performs poorly.

The other approach is to generate the dispatching rules fully automatically by simulation-based Genetic Programming.

3.3 Researching the Artificial Hormone System – Lessons Learned

Uwe Brinkschulte (Universität Frankfurt am Main, DE)

The biological hormone system is a flexible and completely decentralized control mechanism. Therefore it is an excellent blueprint for technical self-organizing systems. We have implemented an artificial hormone system to control task assignment in distributed embedded systems. Hormones are emulated by short messages multicasted and broadcasted in the system. Three types of artificial hormones, eagervalues, suppressors and accelerators are used here. To prove the concept, a hormone simulator has been developed. Furthermore, a hormone based middleware has been implemented. Using self-synchronization mechanisms, the implementation and memory needs of hormone processing could be simplified significantly. Timing guarantees for task allocation could be improved from 2m to m cycles to allocate m tasks. It could be as well shown that the quality of task assignment by the artificial hormone system is better than pure load balancing. Furthermore, tight upper bounds for the communication load and stability conditions to avoid unbound task allocation could be derived. Open research questions are how to find good initial hormone levels in an automated way, how to protect the artificial hormone system against malicious attacks and how to apply the system to other fields of application.

3.4 Designing Open Swarming Systems for Dynamic Runtime Adaptation

Sven A. Brueckner (Jacobs Technology Inc. – Ann Arbor, US)

Traditional software methodologies place the burden of achieving an optimal system response onto the designer, hoping that any scenarios presented at runtime have been accounted for in the chosen optimization solution. In a world where systems are embedded in a dynamically changing environment and where system components have to act autonomously and without complete (or even correct) knowledge of the problem state, complete design-time optimization is no longer feasible. Instead, the role of the designer shifts from developing an optimal solution to developing a self-adaptive system that is capable of dynamically finding the appropriate solution at runtime.

At the seminar, we offer our experience in designing, implementing, and evaluating open self-organizing systems for real-world domains where the required capabilities (including self-adaptation for optimization) emerges from local interactions of many simple agents inspired by the architecture and processes of natural systems. In such open swarming systems, the agents are equipped to respond to changes in their environment to collectively reconfigure their activities for the emergence of optimal system-level patterns and functions. We present example designs from three application domains: dynamic prediction (the swarm induces optimal models of recently observed behavior to extrapolate them into the future), information management (the swarm rearranges models of document content and analyst interest to respond to new information or queries), and IED risk forecasting (the swarm dynamically accounts for new event patterns that may indicate shifts in insurgent operations). In each

example, we discuss the critical design decisions that support dynamic self-organization for optimal operation at runtime.

3.5 Automating Decision Making

Yuriy Brun (University of Washington – Seattle, US)

Self-adaptive systems often evaluate potential adaptations via static model analysis or simulation. I propose an alternative evaluation scheme: trying the adaptation out either on the live, running system or in parallel to the un- adapted system, and observing the effects. This approach likely improves precision of adaptation evaluation, allows for quick adaptation implementation, and reduces costly adaptation undos, albeit, significant technical challenges remain.

3.6 Self-configuration from a Machine-Learning Perspective

Wolfgang Konen (FH Köln, DE)

Main reference W. Konen, "Self-configuration from a Machine-Learning Perspective," Dagstuhl Preprint Archive, arXiv:1105.1951v2
URL http://arxiv.org/abs/1105.1951v2

The goal of machine learning is to provide solutions which are trained by data or by experience coming from the environment. Many training algorithms exist and some brilliant successes were achieved. But even in structured environments for machine learning (e.g. data mining or board games), most applications beyond the level of toy problems need careful hand-tuning or human ingenuity (i.e. detection of interesting patterns) or both. We discuss several aspects how self-configuration can help to alleviate these problems.

One aspect is the self-configuration by tuning of algorithms, where recent advances have been made in the area of SPO (Sequential Parameter Optimization).

Another aspect is the self-configuration by pattern detection or feature construction. Forming multiple features (e.g. random boolean functions) and using algorithms (e.g. random forests) which easily digest many features can largely increase learning speed. However, a full-fledged theory of feature construction is not yet available and forms a current barrier in machine learning.

We discuss several ideas for systematic inclusion of feature construction. This may lead to partly self-configuring machine learning solutions which show robustness, flexibility and fast learning in potentially changing environments.

3.7 Adding Mission Aware Resilience to the Enterprise

Robert Laddaga (Doll Inc. – MA, US)

We envision technology to allow mission requirements to be expressed such that hosts, routers and media configuration can be tailored to those mission requirements in a semi-automated fashion. The mission requirements would accompany the configured system, and later be referred to as monitoring or human intervention indicated a change in requirements, or physical or cyber damage sustained by the system.

Regeneration or reconfiguration of components would then be enabled using the original or modified mission requirements, and the current monitored state of the system. We call the system "a mission-aware adaptive response system" (MARS). A MARS system would significantly improve mission effectiveness and cost, and potentially save lives.

Configuring to mission requirements plus whatever additional capabilities are mandated by policy will ensure that mission needs can be met, without expensive and unnecessary oversupply. More important is the effect of adaptation to mission changes or physical and cyber damage. Repair and reconfiguration could mean the difference between mission failure and success.

3.8 Run-Time System Design

Chris Landauer (Aerospace Corp. – Los Angeles, US)

In this paper, we describe an approach to building autonomous systems for hazardous environments that may be more flexible than the methods currently in use. The system gains this flexibility by doing some of its own system design in response to either environmental activity or internal failures or enhancements, at run-time. The basic idea is that it contains generic models of its own behavioral requirements, which are expected to interact with certain kinds of environmental behavior, and depend on certain capabilities of the hardware behavior. As the hardware degrades, or new software capabilities are provided, or the environmental behavior changes, the specializations used at the beginning of deployment are re- examined, and the decisions revisited. The approach depends on the Wrapping integration infrastructure for software-intensive systems, and on the expectation models for the evolution of the environment. All of the models can be changed remotely (as long as the necessary hardware is available), and the models define the behavior of the system.

3.9 The Art of Organic Programming and Ercatons

Falk Langhammer (Living Pages Research GmbH – München, DE)

Joint work of Langhammer, Falk; Imbusch, Oliver
Main reference Oliver Imbusch and Falk Langhammer and Guido von Walter, Ercatons and organic programming: say good-bye to planned economy, OOPSLA, p.41-52, 2005, Companion to the 20th Annual ACM SIGPLAN Conference on Object-Oriented Programming Systems Languages and Applications OOPSLA
URL http://portal.acm.org/citation.cfm?doid=1094855.1094868

Ercatons are the basic elements for a novel approach to (business) programming. They are simple, document-like building blocks with little a-priori constraints; but with all traditional object-oriented features except algorithmic code, added. Systems grow organically by adding and altering existing ercatons which may be done by other ercatons or humans and which happens at run-time. There is no formal design time. Complexity emerges from an increasing amount of interaction rather than code complication.

The above approach is in industrial production at a few corporations and was invented by us. It is a substitute for conventional and expensive enterprise programming.

The talk highlights the relationship with organic computing and lessons learned. A system grown by organic programming is an organic computing system if programmers and users are considered to be part of the system. Ercatons then enable the emergence of unforeseen features in such systems.

3.10 The Six-Legged Walking Robot OSCAR as a Demonstrator for the Organic Robot Control Architecture ORCA

Erik Maehle (Universität Lübeck, DE)

Joint work of Maehle, Erik; Brockmann, Werner; Grosspietsch, Karl-Erwin; Al-Homsy, Ahmad; El-Sayed-Auf, Adam; Jakimovski, Bojan; Krannich, Stephan; Litza, Marek; Maas, Raphael

Autonomous mobile robots are complex machines being challenging to engineer and to program. In order to master this complexity, the Organic Robot Control Architecture ORCA has been developed making use of organic principles like self-organization, self-reconfiguration and self-healing. The six-legged walking robot OSCAR is introduced as a demonstrator for ORCA. Organic principles are employed on all layers of its hierarchical control system starting at the reflexive layer with gait generation and reflexes over the reactive behavioural layer up to the deliberative planning layer. Experimental evaluations demonstrate that OSCAR thus attempts to continue its mission in the best still possible way by adapting to internal faults as well as to unforeseen environmental situations.

3.11 Organic Computing: From Design-time to Run-time

Christian Mueller-Schloer (Leibniz Universität Hannover, DE)

In 2003, when we discussed OC for the first time, we defined it as a technology, which brings life-like properties into complex technical systems to make them more robust and flexible. OC systems are designed to be self- organized, self-configuring, self-protecting, self-explaining etc.

Now, 8 years down the road, I want to suggest a new definition of OC: OC means to move design-time decisions to run-time. This definition has the advantage that it can be operationalized: We can derive the future OC research program if we follow the steps of classical design processes and ask for each of them what it would mean to move it into run-time.

This talk will sketch out this idea in more detail. First we will briefly review design processes and ex-tract a simplified version of a standard design process.

Next we will transfer this process to run-time arriving at a so-called "seamless redesign process". Now we can identify the main challenges that have to be met.

1. Shift from design-space to configuration space: It is the task of the designer to explore the possibilities of the design space. This is an activity, which requires experience as well as creativity. In an OC system the agent itself has to explore the configuration space. Hence it must be endowed with creativity and experience as well. The configuration space will have to be smaller than the design space in order to bound search-time. Also it should exclude illegal configurations in order to keep the creative agent from proposing configurations, which cannot be realized.
2. Run-time optimization: The OC system can be modeled as a cognitive agent moving through the configuration space in search of the highest fitness.
   This fitness landscape is not only unknown to the agent but also time- dependent and self-referential. i.e. it changes depending on the actions of the agents.
3. Run-time validation: An agent who searches the configuration space has to validate possible solutions in order to sort out detrimental ones and to find the best possible one. Validation could be done by trial-and-error but this means that bad (or illegal) solutions are tried out in reality. This is clearly not acceptable in technical systems. Therefore we have to introduce a sandbox approach where solutions are validated in a simulated environment. Simulation (and verification) is usually time consuming, and it requires an accurate description of the system. Moreover it re-quires stimuli reflecting the future real situation as closely as possible. Finally, we should consider the option of on-line testing.
4. Production vs. run-time reconfiguration: While in the classical design process, at a certain point in time the design is frozen and goes into production, the run-time reconfiguration process is seamless: This means in principle that all, even the higher-level, design decision can be revised. But it also means that only those changes can be actuated, which can be realized in terms of software or hardware reconfigurations.
5. Run-time modeling: All design steps, be it in the classical design process or in the run-time re-configuration process, are model-based. But now we have to consider two different flavors of the-se models: Prescriptive models reflect the classical top-down enforcement. Descriptive models reflect the actual system state. They are not necessarily consistent. We have to find mechanisms, which can minimize the possible contradictions.

This corresponds to a run-time version of what in the classical design process is called Yoyo design.

The talk closes with some remarks on more OC challenges, which have not been sufficiently ad-dressed so far: (1) The relationship between the user and the adaptive system, and (2) the broader view of regarding OC as a part of "Organizational Sciences".

3.12 Self-Organized Self-Improvement: Using Self-Directed Experimentation to Improve Models and Methods

Phyllis R. Nelson (Cal Poly – Pomona, US)

The models and methods that a system uses to choose its behaviors are an important link between the system's current situation and its top-level purposes and goals. If the system lacks models and methods that match its current condition and context, it may no longer be able to appropriately link its resources to its purposes and goals. Biological systems use some of their spare time and energy to explore both their own capabilities and their environment.

We discuss our version of this biological style for self-organizing behaviors that enable a system to improve its models of both its own current capabilities and its environment. We also consider how the system can determine if such self improvement is useful. Finally, we present a problem for this concept (the existence of a hidden variable), and examine the consequences for self organization.

3.13 A proposal how to combine bottom-up emergence and top-down control during runtime

Gabriele Peters (FernUniversität in Hagen, DE)

Joint work of Leopold, Thomas; Kern-Isberner, Gabriele; Peters, Gabriele
Main reference T. Leopold, G. Kern-Isberner, G. Peters, "Belief Revision with Reinforcement Learning for Interactive Object Recognition," Proc. 18th European Conf. on Artificial Intelligence (ECAI 2008), Vol. 178, pp. 65–69, Patras, Greece, July 21-25, 2008.
URL http://dx.doi.org/10.3233/978-1-58603-891-5-65

A system is proposed which combines two levels of learning. Both levels negotiate during runtime and establish a balance between bottom-up emergence and top-down control. They are realized by techniques of reinforcement learning (RL) and belief revision (BR). The RL component is able to react to runtime events and learns behavioral strategies in a flexible way. What is learned by RL is available in a numerical form only. Especially, it is not intuitively understandable by humans. In contrast, the BR component acquires knowledge in the form of rules. These rules control the bottom-up RL process from top-down. They are comprehensible by humans. Thus, the BR component can act as an interface for intervention from outside the system by a system designer or a user in case the system displays undesired behavior. Work in progress is presented where this system design is applied to a computer vision task.

References

1 Thomas Leopold, Gabriele Kern-Isberner, and Gabriele Peters. *Combining Reinforcement Learning and Belief Revision – A Learning System for Active Vision*. Proc. 19th British Machine Vision Conf. (BMVC 2008), Vol. 1, pp. 473–482, Leeds, UK, 2008

3.14 Engineering Proprioception in Computing Systems

Marco Platzner (Universität Paderborn, DE)

In this presentation I will give an introduction to the objectives and working areas of the recently started project "Engineering Proprioception in Computing Systems" (EPiCS). EPiCS is part of the EU FET objective Self-awareness in Autonomic Systems and relies on self-awareness and self-expression as key concepts for enabling complex future computing and communication systems.

Key words: Self-awareness, self-expression, autonomic systems

3.15 Cyber Physical Sytems (CPS) or better Cyber Biosphere(CBS)?

Franz J. Rammig (C-LAB – Paderborn, DE)

As most technical artifacts will be interconnected in some sense ("Internet of Things, Cyber Physical Systems"), IT systems of the future cannot be treated as isolated entities any longer. More or less every technical artifact will be linked to the Internet. Two major tendencies can be observed. The first one takes its inspiration from the technical roots of Embedded Systems. This approach became well known under the name "Cyber Physical Systems (CPS)". The main challenge of this approach is the necessity to bridge two incompatible worlds: this one of highly predictable embedded real-time systems and this one of the stochastically operated Internet. The second approach takes inspirations from the achievements of nature. This approach is being discussed using terms like "Biologically Inspired Systems" or "Organic Computing". We discuss these alternatives building the highly sophisticated Embedded Systems of the future.

The basic challenges to be solved when designing CPS as well as CBS are characterized. Some comparisons of CPS and CBS will be made as well.

3.16 A Decade of experience building adaptive systems (Things that change in the night)

Paul Robertson (Doll Inc. – MA, US)

In today's world computers are constantly connected and participating in the full breadth of human existence. Our world is changing at record breaking speed and yet our software is still largely designed and implemented as it was when programs were static, disconnected, and run on demand by people. My colleagues and I have now been building self-adaptive systems for over a decade and have applied them to some of the most challenging environments including real-time vision, robotics, and more recently to cyber security – where the world

can change in real-time and the need to adapt is most crucial. A lot has been learned over that period about successful approaches, technologies, and research challenges.

3.17 Overview of the DFG priority program Organic Computing

Hartmut Schmeck (KIT – Karlsruhe Institute of Technology, DE)

Joint work of Christian Müller-Schloer; Schmeck, Hartmut; Ungerer, Theo

Originating from a series of workshops on future research topics for Computer Engineering back in 2002, a joint position paper of the Gesellschaft für Informatik and the Informationstechnische Gesellschaft outlined the vision of Organic Computing Systems to cope with the increasing presence of intelligent, interacting devices in various application scenarios by controlled self-organization and an emphasis on realizing robust, adaptive and trustworthy systems showing an "organic" behavior even in unanticipated situations. The talk provides a brief survey of the priority program on Organic Computing which has been funded by the German Research Foundation (DFG) from 2005 to 2011. The results of the program have been compiled into a comprehensive compendium on Organic Computing (see [MSU11]) with chapters on Theoretical Foundations, Methods and Tools, Learning, Architectures, Applications, and an Outlook on recently added projects and potential future directions of research.

References

1 Christian Müller-Schloer, Hartmut Schmeck, Theo Ungerer (eds.) *Organic Computing – A Paradigm Shift for Complex Systems*. Birkhäuser Verlag, Basel, Schweiz, 2011

3.18 Smart Grid, Renewables, Electric Mobility: Opportunities and Challenges for Organic ComputingC

Hartmut Schmeck (KIT – Karlsruhe Institute of Technology, DE)

The strong shift in energy policy towards power generation from renewable sources leads to a number of challenges for an adequate management of the future power grid. In particular, the uncontrollable and only partially predictable fluctuations in energy supply ask for a replacement of the traditional paradigm "supply follows demand" with the new principle "demand follows supply" in addition to a shift from centralized power generation to large scale decentralized supply of power at the low voltage segments of the distribution grid. This relies on detailed information on the current status of the relevant components of the distribution grid and on sufficient knowledge about the available degrees of freedom for demand shifting. Organic Computing has the potential to provide adequate concepts for shaping these new approaches to energy management which should combine electrical and thermal energy. The talk outlines the various individual challenges and describes some experiences in the projects MeRegio and MeRegioMobile, where a smart home has been designed for investigating the intelligent integration of the charging needs of electric vehicles into an environment containing a range of different consumers and suppliers of energy.

3.19 Physics-inspired Self-Organization and Adaptation in Large Dynamic Overlay Networks

Ingo Scholtes (Universität Trier, DE)

Overlay networks are becoming an increasingly important abstraction that facilitates the cost-efficient and scalable provision of novel services in the Internet. However, efficiently constructing, maintaining and managing robust and adaptive overlay topologies in the face of highly dynamic participants is a challenging task.

In this talk, a physics-inspired approach towards the management of self- organizing and self-adaptive overlays will be discussed. It is based on the idea that global-scale network infrastructures like the Internet or our biggest Peer-to-Peer systems are becoming so large that it appears justified to design them along models and abstractions originally developed for the study of many-particle systems in statistical physics. The management schemes that will be presented take advantage of recently uncovered analogies between random graph theory and statistical mechanics. They constitute the basis for what may be called a thermodynamic management of large dynamic overlay networks.

3.20 Quantitative Emergence – An Overview of Recent Measurement Techniques

Bernhard Sick (Universität Kassel, DE)

Joint work of Fisch, Dominik; Jänicke, Martin; Kalkowski, Edgar; Sick, Bernhard

A technical system exhibits emergence when it has certain properties or qualities that can be termed to be irreducible in the sense that they are not traceable down to the constituent parts of the system. The presentation summarizes three techniques for emergence detection and emergence measurement that were proposed by members of the organic computing community. These techniques are based on information- theoretic and probabilistic viewpoints: the discrete entropy difference, the Hellinger distance which is a divergence measure for probability densities, and an iterative approach motivated by divergence measures. Advantages and drawbacks of these measures are demonstrated by means of some simulation experiments using artificial data sets. It is shown that these techniques are able to deal with different kinds of emergent phenomena such as transitions from chaos to order, concept drift, or novelty. That is, with these techniques it is possible to cover a wide range of possible applications. Moreover, it will be possible to build systems that are self- aware, environment-aware, self-reflecting, ...

3.21 Organic Self-organizing Bus-based Communication Systems (OrganicBus)

Jürgen Teich (Universität Erlangen-Nürnberg, DE)

Joint work of Teich, Jürgen; Ziermann, Tobias; Wildermann, Stefan
URL http://www12.informatik.uni-erlangen.de/research/organicbus/

We present an organic computing approach for the analysis, design and optimization of run-time message scheduling for priority-based bus systems such as the industrial CAN (Controller Area Network) standard.

The goal of this new approach is to overcome the major drawbacks of today's pure offline scheduling decisions that are based on worst-case assumptions, are not flexible or overly pessimistic with respect to unknown, uncertain or dynamically changing message request scenarios.

In contrast, our decentralized approach using online self-organization is able to monitor the actual traffic of the shared bus medium and adapt either sending rates, probabilities or offsets to establish fair bandwidth sharing and/or reduced response times.

For messages with high bandwidth demands such as streams, we present a decentral adaptation algorithm called *penalty learning algorithm* (PLA) [1] that is able to achieve a fair bandwidth assignment of a set of sending nodes and reaching this equal utilization provably no matter of the initial priority assignment by applying a game-theoretic analysis.

For periodic messages with (soft) real-time constraints as given by deadlines, we introduce a simple decentralized algorithm for adapting the offsets of message in a bursty time interval such that the average observervable worst case response time is reduced.

As a side effect of observing and manipulating this traffic for a monitor interval in the size of the hyperperiod for soft-real time periodic messages, a bus might be driven with higher utilization rates while still satisfying acceptable response times.

This algorithm called *Dynamic Offset Adaption Algorithm* (DynOAA) [2] for soft real-time tasks may be jointly run with the algorithm PLA for bandwidth type of messages.

In the near future, an implementation an FPGA-based platform coupled to a real CAN-bus shall give experimental evidence of the superiority of these techniques with respect to static scheduling.

References

1 Wildermann, Stefan and Ziermann, Tobias and Teich, Jürgen. *Self-organizing Bandwidth Sharing in Priority-based Medium Access.*
Proceedings of the Third IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO 09), San Franzisco, USA, 2009

2 Ziermann, Tobias and Salcic, Zoran and Teich, Jürgen. *DynOAA – Dynamic Offset Adaptation Algorithm for Improving Response Times of CAN Systems.*
Proceedings of Design, Automation and Test in Europe (DATE), IEEE Computer Society, Grenoble, France, 2011

3.22 Knowledge Representation for Autonomous Systems – The ASCENS Case Study

Emil Vassev (University of Limerick, IE)

Joint work of Vassev, Emil; Mike Hinchey
Main reference E. Vassev, M. Hinchey, "Knowledge Representation and Awareness in Autonomic Service-Component Ensembles – State of the Art," Proc. 14th IEEE Int'l Symp. on Object/Component/Service-oriented Real-time Distributed Computing Workshops (ISCORCW'11), IEEE, 2011, pp.110–119.

Introduction

Ideally, autonomous systems are intelligent systems employing knowledge to become aware of situations, recognize changes and eventually respond to changing conditions. Knowledge is the key to such autonomous behavior. The fundamental questions are how to represent knowledge in such systems and how to make them use and manage that knowledge.

Current and ongoing research at Lero – the Irish Software Engineering Research Centre, is focused on the problem of knowledge representation for autonomous systems formed as ensembles of special autonomous service components. Such components encapsulate rules, constraints and mechanisms for self-adaptation and acquire and process knowledge about themselves, other service components and their environment. One of the expected major scientific contributions of this research is a formal approach to knowledge representation and reasoning mechanisms that help autonomous components acquire and structure comprehensive knowledge in such a way that it can be effectively and efficiently processed, so the system becomes aware of itself and its environment.

Autonomic Service-Component Ensembles (ASCENS) (see http://www.ascens-ist.eu/) is a class of multi-agent systems formed as mobile, intelligent and open-ended swarms of special autonomic service components capable of local and distributed reasoning. Such service components encapsulate rules, constraints and mechanisms for self- adaptation and acquire and process knowledge about themselves, other service components and their environment. ASCENS systems pose distinct challenges for knowledge representation languages.

Kinds of Knowledge for ASCENS

There have been determined four basic knowledge domains (kinds of knowledge) for ASCENS systems:

- the individual component structure and behavior;
- the system structure and behavior;
- the environment structure and behavior;
- situations where the system might end up in.

These knowledge domains have been used to derive distinct knowledge models (each representing a distinct knowledge domain) for ASCENS forming a high-level knowledge structure that is to be maintained by any service component (SC) member of a service-component ensemble (SCE):

- SC knowledge model – knowledge about internal configuration, resource usage, content, behavior, services, goals, communication ports, actions, events, metrics, etc.;
- SCE knowledge model – knowledge about the whole system, e.g., architecture topology, structure, system-level goals and services, behavior, communication links, public interfaces, etc.;

- environment knowledge model – parameters and properties of the operational environment, e.g., external systems, external communication interfaces, integration with other systems, etc.;
- situational knowledge patterns – specific situations, involving one or more SCs and eventually the environment.

By representing the knowledge in such models, an individual SC shall be able to query information about both the SC itself and the SCE, by considering the environment's parameters and properties. Moreover, this helps SCs understand and reason about themselves and discover situations through the use of probabilistic methods working over the knowledge modeled as situational knowledge patterns.

3. Structure of the Knowledge Representation The four knowledge domains for ASCENS are represented by four distinct knowledge corpuses – SC Knowledge Corpus, SCE Knowledge Corpus, Environment Knowledge Corpus and Situational Knowledge Corpus. Each knowledge corpus is structured into a special domain-specific ontology and a logical framework. The domain-specific ontology gives a formal and declarative representation of the knowledge domain in terms of explicitly described domain concepts, individuals (or objects) and the relationships between those concepts/individuals. The logical framework helps to realize the explicit representation of particular and general factual knowledge, in terms of predicates, names, connectives, quantifiers, and identity. Thus, the logical framework provides additional to the domain ontology computational structures that basically determine the logical foundations helping a SC reason and infer knowledge.

All the four ASCENS knowledge corpuses form together the ASCENS Knowledge Base (AKB). The AKB is a sort of knowledge database where knowledge is stored, retrieved and updated. Therefore, in addition to the knowledge corpuses, the AKB implies a knowledge-operating mechanism providing for knowledge storing, updating and retrieval/querying. Ideally, we can think of an AKB as a black box whose interface consists of two methods called TELL and ASK. TELL is used to add new sentences to the knowledge base and ASK can be used to query information. Both methods may involve knowledge inference and therefore, an AKB should be equipped with a special Inference Engine that reasons about the information in the knowledge base for the ultimate purpose of formulating new conclusions, i.e., inferring new knowledge.

3.23 Slow Feature Analysis: Learning with the Slowness Principle

Laurenz Wiskott (Ruhr-Universität Bochum, DE)

Joint work of Wiskott, Laurenz; Berkes, Pietro; Franzius, Mathias; Sprekeler, Henning; Wilbert, Niko; Zito, Tiziano
Main reference Wiskott, L.; Berkes, P.; Franzius, M.; Sprekeler, H.; Wilbert, N. (2011) Slow feature analysis. Scholarpedia, 6, 5282
URL http://www.scholarpedia.org/article/Slow_feature_analysis/

Slow feature analysis (SFA) is a biologically motivated algorithm for extracting slowly varying features from a quickly varying signal and has proven to be a powerfull general-purpose prepocessing method for spatio-temporal data. We have applied SFA to the learning of complex cell receptive fields, visual invariances for whole objects, and place cells in the hippocampus. On the technical side SFA can be used to extract slowly varying driving forces of dynamical systems and to perform nonlinear blind source separation. Here I will introduce the SFA algorithm and give an overview over these different applications.

3.24 Learning to see and understand

Rolf P. Wuertz (Ruhr-Universität Bochum, DE)

Joint work of Wuertz, Rolf P.; Müller, Marco K.; Walther, T.

Sensor interpretation is one of the central requirements for computing systems to interact sensibly with their environment. Moreover, the transformation of sensor data into semantically meaningful representations is necessary to exploit the types of computation in which machines are much better than humans.

The most important among the sensory modalities is visual processing. In this talk I reviewed self-organizing methods we have developed for the recognition of human faces and human bodies.

The first system [1] is based on similarity rank lists and can learn the transformation from a frontal face to a different pose strictly on the basis of examples. This can be carried out by a neural network in a natural way. It also presents a way of controlling the generalization of the recognitition system.

The second system [2] uses several state-of-the-art methods from computer vision to learn a model of the appearance and possible kinematics of the upper body of humans from videos in an unsuupervised way.

The resulting model is able to generalize over individuals, clothing, and a wide range of backgrounds, and is robust enough for still image interpretation.

References

1 Marco K. Müller and Rolf P. Würtz. Learning from examples to recognize faces in different situations. *Neural Networks*. Submitted.
2 Thomas Walther and Rolf P. Würtz. Autonomous learning of a human body model. In *Proc. IJCNN*. IEEE, 2011. In press.

3.25 Self-adaptive workload management on MPSoC

Johannes Zeppenfeld (TU München, DE)

Joint work of Zeppenfeld, Johannes; Bouajila, Abdelmajid; Stechele, Walter; Herkersdorf, Andreas
Main reference J. Zeppenfeld, A. Herkersdorf, "Applying Autonomic Principles for Workload Management in Multi-Core Systems on Chip," Int'l Conf. Autonomic Computing (ICAC'11), Karlsruhe, Germany, June 14–18, 2011

Due to the difficulty involved in designing ever more complex systems, a trend is developing to simply replicate existing hardware components, rather than extending the functional capabilities of existing components. While this simplifies the work of the hardware designers, it simply offloads the difficulty of designing complex systems to the software developers, who must somehow make use of all the resulting parallel processing units.

The Autonomic System on Chip paradigm suggests an alternative approach, namely by utilizing a portion of the gained resources for the addition of bio- inspired enhancements, which can autonomously accomplish at run time many of the tasks previously performed by the designer at design time. Not only does this reduce the burden on the designer, it also allows the system to adapt to unforseeable environment- or system-states. The presented work shows how such autonomic principles, applied to a general purpose multi-core system

on chip, can autonomically adjust the frequencies and distribute tasks across the available processing cores.

3.26 A glimpse of signaling pathways in the synapse

Junmei Zhu (EBI – Cambridge, GB)

Learning is a key to run-time design. One of the most studied learning elements is the synapse. Computationally, Hebbian plasticity can be conveniently described by one equation. In biology, however, thousands of genes have been identified in a typical synapse. The underlying molecular mechanism for plasticity seems to be escaping our grasp, with an ever-expanding list of involved molecules. I will introduce this picture and our effort to construct the complete signalling pathways. The complexity could be what is needed for the robustness and flexibility of organic systems, and thus OC design is not to be intimidated by large systems.

4 Working Groups

4.1 Group I Report – Model-based Self-Adaptation

4.1.1 Challenges

We identified five challenges related to model-based self-adaptation:

1. Defining what a model is and establishing a consistent, clear terminology. Models can be of the system itself and of the environment.
2. Selecting the right model: abstraction level, objective function, time scales.
3. Understanding control-loop design and making control-loop development explicit during the design process.
4. Identifying those environments that can benefit most from self-adaptation (as well as those that cannot benefit or will present significant challenges). One aspect of these environments is exploration vs. exploitation and competition for resources.
5. Reducing the need for the design-time. This challenge aims to develop some framework that will let the developer specify a high-level system goal and then have the system design itself (perhaps evolve) into one that satisfies that goal. We believe that breaking up the target system into a set of tiers may be a fruitful approach. Tiers are made up of building blocks. Each tier can "rearrange" its building blocks in order to achieve it's goal. Each tier also sets the goals for its building blocks, which are themselves the next tier. For example, a 3-tier traffic light system may consist of:
   - Neighborhood (top tier): many adjacent intersections
   - Single intersection (middle tier): four traffic lights at one intersection
   - Single traffic light (bottom tier).

The developer would specify the high-level goal to the top tier. For example, "minimize the average travel time through the neighborhood while avoiding accidents." That tier would then attempt to set a goal for the middle tier. e.g., "minimize the maximum waiting time for a car at the intersection while ensuring no intersecting car paths have the green light at

the same time." The top tier can (1) manipulate the intersections and (2) set the goals for the intersections. The intersections then do the same for the bottom tier: the traffic lights. The goals for the bottom tier may be "only one light on at a time, green → yellow → red → yellow → green ...".

We believe that each tier must operate at a slower time scale than the tier below it. Possible techniques for manipulating the lower tier to accomplish a goal include: simulated annealing, artificial evolution, genetic algorithms. One concern is whether such exploration is likely to get stuck in local minima. Background exploration with one of the above techniques can help get out of such situations, but it is unlikely to consistently find the optimum solution. The design of the tiers constrains the search space. This is one way to constrain that search space, and we are not sure that better ways do not exist. We hope there can be tools made to automate, semi-automate, or otherwise help developers design these tiers.

4.2 Group II Report – Learning and Context-dependent Dynamic Knowledge Representation

4.2.1 Challenges

4.2.1.1 How can we LEARN emergent features on several layers / hierarchies?

(Gabriele Peters, Laurenz Wiskott) As one open question we identified the problem of getting knowledge into a system. If learning takes place in a hierarchical manner, then probably different methods for knowledge acquisition are needed for different levels of learning. Whereas on a lower level a statistical approach for learning features may be sufficient, this probably does not hold true for higher levels. On a low level features are relatively simple and low-dimensional and plenty of samples are available for training. Examples are small image patches or sound snippets. This makes purely statistical methods feasible. On a high level, features become more complex and high dimensional. Examples are whole visual objects or scenes. It is therefore not possible anymore to learn them by brute force statistical methods. Another reason why it might be good to have different mechanisms on a high level is that the system needs to do reasoning about things, which requires a more symbolic representation, at least in addition to a feature based representation (where a feature is understood here as something that has a graded value, while a symbol is binary in nature).

From, e.g., feature learning on the one end until scene understanding / interpretation on the other end there has to be one point in the hierarchy where the acquired knowledge has to be transferred from an implicit / numerical form to a more explicit / symbolic representation that allows for interpretation and reasoning. One possibility would be a symbolic representation in the form of rules. Gabriele Peters has shown in her talk that it is possible that such rules can be learned "organically" during runtime [1]. On a lower level an implicit representation of appropriate behavior is learned by reinforcement learning; on a higher level rules are generated and checked against the reinforcement model to formalize and make explicit what has been learned implicitly on the lower level. The symbolic and rule based representation on the higher levels permits reasoning, planing, and anticipation as needed for the WHAT IF strategy, see below.

We also discussed the question whether one could replace the structured learning on the higher levels by pure memory. Besides the fact that any system has a limited memory capacity that would quickly be filled with sensory data, are there any other advantages for not storing everything and then working on the memories. It seems clear that working

on the collection of all memories not only raises the problem of limited memory capacity but also poses a problem for other resources, e.g. computational resources. For instance, a nearest neighbor classifier becomes quite inefficient, if too many samples are stored. Thus, it is also computational limitations that prevent a system from working on complete memory collections. Another problem is the lack of interpolation and regularization. If data is noisy and/or sparse the system needs to perform some regularization in order to learn appropriate functionality that generalizes well. This could, of course be done on the fly each time the functionality is needed, but this would be computationally even more expensive.

Thus, the need for a condensed and explicit/symbolic representation of the world on higher layers, i.e. the need for an efficient model of the world, results from memory limitation, computational limitations and the need for planing and reasoning. Quasi symbolic representations can also be derived from single exemplars given the right representation, see Section "How can we learn from few examples".

(Phyllis R. Nelson) There is also the question of how to start the system. What knowledge, models, methods and procedures should the designers put in? Should this set of initial resources be privileged in some way to ensure that the system can always be restarted? For any complex system, the designers don't know enough about some aspects of the system (otherwise we wouldn't need OC), and may also know so much from their own experience that the initial knowledge must be restructured by the system.

References

1 Thomas Leopold, Gabriele Kern-Isberner, and Gabriele Peters, Combining Reinforcement Learning and Belief Revision – A Learning System for Active Vision. 19th British Machine Vision Conference (BMVC 2008), Vol. 1, pp. 473–482, 2008

4.2.1.2 Knowlege organization: How to not get stuck in too many interconnections between elements?

(Christoph Landauer) We have proved [LB98] that any system of knowledge with any construction mechanism will "get stuck"; eventually, each new item needs to be connected with so many other items that it cannot be effectively entered, even by humans. Circumventing this problem is an important issue in knowledge representation, and we think that it will require a very different approach to representation of knowledge.

[LB98] Christopher Landauer, Kirstie L. Bellman, "Situation Assessment via Computational Semiotics," pp. 712-717 in Proceedings of ISAS'98: The 1998 International MultiDisciplinary Conference on Intelligent Systems and Semiotics, 14–17 September 1998, NIST, Gaithersburg, Maryland (1998)

(Emil Vassev) Knowledge should be presented at different levels of abstraction and conceptually organized. Different abstract levels shall form levels of generality (introducing certain degree of uncertainty though) that can be applied to reasoning algorithms to release heavy computations by moving from a low-level knowledge to a more generic one.

The concept-based representation is necessary, because this is probably one of the most efficient ways to give both meaning and semantics to the raw data. Ontologies and description logics (as one of the most prominent ways to write ontologies) may help not to have unnecessary interconnections between the knowledge elements. An ontology is open-ended and should not include all the possible conceptual descriptions of the context, but only the relevant ones, i.e., giving sufficient level of description taking into consideration the operational domain of the system in question. Moreover, techniques like subsumtion shall help to make new generalizations over the concepts, thus reducing the amount of concepts

that must be processes when the system reasons. For example, subsumtion may help to exploit the taxonomy structure of concepts that are defined in the ontology and compute a new taxonomy for a set of concepts.

4.2.1.3 Finding the relevant knowledge in a given context (contect might change on-line)

(Christoph Landauer) This issue is about how a system uses knowledge that it has. Whatever knowledge scheme is used, only a tiny fraction of that knowledge is appropriate for any given situation, and we need to have systems that can separate and isolate that relevant knowledge in real time.

(Wolfgang Konen) Another way to phrase it is that in a future OC system there might be a large feature base where certain sets of features are activated / deactivated with a mechanism yet-to-be-found depending on the context / the environment situation. The challenge is here to find mechanisms which are able to work in many relevant situations.

(Phyllis R. Nelson) The WIM and play (exploration) are a tool for finding out about those special contexts that are at the boundaries of the capabilities of the system. For a system with physical parts, playing games provides a structure for examining potentially "dangerous" behaviors and learning better models of those actions before they are invoked in "real" responses. In a game, the system can start and and then back away from these regions of operation since they are not directly part of accomplishing its main goals and purposes. This type of exploration is particularly important because the limits of physical subsystems are often nonlinear and depend on the history of the components (wear, old batteries, etc.).

4.2.1.4 Parallel knowledge representations (playing piano): how to merge / how to make that they support each other

(Christoph Landauer) The question here is how to merge suggestions from different kinds of knowledge in a seamless way. For the piano example, it is a combination of knowledge of the way the music looks on the score, knowledge of the way the music should sound, muscle knowledge of the way to move to make the sounds (this is why you practice), and knowledge of the way the music sounds just now. These different sources of very different kinds of knowledge are combined to make a sequence of actions that play the piece. In the piano example the different representations may support each other to perform the complex motor actions and memory retrievals. The hard part in general is blending the different knowledge sources, deciding if there is a conflict among the sources, resolving that conflict in real time, using whatever source is available to determine the next actions, and making all the relevant decisions very quickly.

(Phyllis R. Nelson) There is also the question of selecting the variables or features that will be measured in order to "close the loop" when the system acts. Producing the desired behaviors requires knowledge and models of the available sensing and actuation and how those variables that can be measured relate to the desired behavior. Play (exploration) is a way to build these connections.

4.2.2 Other points

4.2.2.1 How to learn from few examples?

(Rolf P. Würtz, Gabriele Peters, Laurenz Wiskott) Gradient-based learning methods usually require the iterated presentation of a large number of training examples, which makes it

comparably slow. Given an appropriate format, however, a single example may be sufficient to "know" the visual example. This works perfectly in human vision and nicely in face recognition systems. On a symbolic level, single examples can easily prove (exist) or disprove (forall) rules. If rules have a measure of evidence, this can be updated by every example. In G. Peter's system, each new example is able to qualitatively modify the behavior. The hierarchical SFA system by Laurenz Wiskott learns invariant representations from the presentation of few examples. Thus, learning from few examples is both feasible and desirable, especially for online real-time learning. *Christopher Landauer)* agree, when we decide that parallel performance improvement methods must be used, we can take into account the local structure of problems. An early example of learning symmetries from only one example was given by Konen and von der Malsburg in *Konen, W., von der Malsburg, C., "Learning to generalize from single examples in the dynamic link architecture," Neural Computation, 5, 1993, p. 719-735.*

4.2.2.2 Humans have a very elaborate What-If-mechanism (WIM)

(Wolfgang Konen, Laurenz Wiskott) Humans are good at combining prior experience and facts from new a environment in "What-If-simulations". The mechanism on how we can set up and utilize this very flexible mechanism is not yet fully understood and poses a challenge to be captured by future OC systems. But it is clear that "What-If-simulations" are a powerful tool to extend and recombine existing knowledge on the one hand and to save resources and avoid risks on the other hand.

(Laurenz Wiskott) However, What-If simulations have at least two types of limitations. If the world model is inaccurate, What-If simulations are bound to become unrealistic and therefore useless with increasing depth. This makes What-If simulations useful only up to a certain point in the future and then one has to take the real action in order to verify the predictions of the model. If the world model is perfect, What-If simulations can in principle substitute completely for the real exploration. Learning how to play a game by simulated self-play is a good example of WIM, but also of exploration. Perhaps this approach is useful in order to help limit the choices enough so that useful and relevant hypotheses can be inferred. However, in many cases this might be limited by memory capacity. For a human, for instance, it would be impossible (with a few exceptions, maybe) to learn to play chess by just imagining games against oneself, simply because it is so hard to keep the positions in mind accurately during an imagined game.

(Laurenz Wiskott) Thus, real exploration is needed if the model is inaccurate or if memory limitations prevent the system from What-If simulations of sufficient depth. A third reason for real exploration is, of course, if there is no adequate model for that particular domain, which however could be considered an extreme case of model inaccuracy. In [1] for instance a system is proposed where a model is learned from scratch by exploration. It is clear that the What-If simulations give strong hints as to where real exploration might pay off most. For instance, a good model should also maintain an estimate for the accuracy of its predictions. If predictions become uncertain, then exploring this part to improve the model is probably particularly advantageous.

References

1 Thomas Leopold, Gabriele Kern-Isberner, and Gabriele Peters. Belief Revision with Reinforcement Learning for Interactive Object Recognition. 18th European Conference on Artificial Intelligence (ECAI 2008), Vol. 178, pp. 65-69, 2008

4.2.2.3 How to cope with the combinatorial complexity of world in our brain model

(Laurenz Wiskott) The world is tremendously complex overall, but it is not possible to integrate all that in a model to begin with. Thus, one has to employ different mechanisms to start simple and grow complex gradually. During this process the degrees of freedom of a model should always be smaller than the available data that constrain it, in order to avoid overfitting and guarantee good generalization. To do so, it is good the start with a simple model, to which more complexity can be added as more data becomes available and more complexity is needed to explain it. When starting with such a simple model, it might be advantageous to first confront it only with part of the data, which allows for simplified modeling. When learning a language, for instance, start first with short sentences. Once theses are mastered by the model proceed to longer and more complex sentences, see [1]. This can be done because the world is highly structured and not just complex. There are simple parts that can be modeled largely independently of the rest and still one can use it to generalize to more complex domains. Or one can approximate complex domains by simple models and still get useful predictions out of it.

(Wolfgang Konen) An open issue is how to cope with the combinatorial complexity of the environment and it might call for knowledge representation models which incorporate also a similar way of combinatorial complexity.

References

1 Jeffery L. Elman. Learning and development in neural networks: the importance of starting small. *Cognition*, Vol. 48, pp. 71–99, 1993.

4.3 Group III Report – Dealing with Uncertainty in Organic Computing Systems

Being able to deal with uncertainty is a central concept that permeates many – if not most – aspects of Organic Computing systems, including run-time adaptation, knowledge representation, modeling, learning or feature extraction. In fact, one of the major promises of Organic Computing Systems – and one that potentially distinguishes them from traditionally engineered systems – is that they shall be able to deal with unanticipated situations, surprises and rare events. Furthermore, Organic Computing systems need to adapt and evolve at run-time, thus potentially introducing additional uncertainty with respect to their run-time behavior. In our group discussion we first tried to characterize different sources of uncertainty in Organic Computing Systems.

4.3.1 Environmental Uncertainty

One of the primary sources of uncertainty is the environment into which our computing systems are embedded. Here, uncertainty may result for example from technical failures, attacks, changing environmental conditions or human behavior. In most cases, one cannot actively control these environmental conditions, one can merely try to characterize the resulting uncertainties and reason about them either in a qualitative or in a quantitative fashion. And even this constitutes a considerable challenge. When trying to reason for example about the probabilities of certain events, one must deal with the fact that we typically don't even know the probability space to begin with, that is we don't know in advance which events may possibly occur, let alone their probabilities. In traditionally engineered systems, we are frequently reminded of this fact by the occurrence of unanticipated events. When

dealing with rare events – like for instance catastrophic earthquakes – we also need to face the challenge that there are – if any – only few prior instances from which we could learn about the probability of these events. Here, a modeling in terms of extreme value theory seems to be crucial. Also, second order probability techniques can often be reasonably applied to characterize and measure a system's uncertainty about its environment.

In the discussion above, we inherently assumed that environmental uncertainty results from a lack of knowledge which can be reduced by observation, experimentation and probabilistic reasoning. Whether the fundamental reason for uncertainty is a mere lack of knowledge about the system's details or whether there are processes that are inherently random is a question that is still being discussed in fundamental physics and philosophy. In fact, we can even find different interpretations of uncertainty when considering frequentistic or Bayesian notions of probability in mathematics. Independent of this philosophical question, in practice there clearly is a limit to how much we can reduce environmental uncertainty by means of experimentation and probabilistic reasoning. An important limiting factor is that in distributed systems our knowledge about the current state of a system is necessarily imperfect and incomplete and that many systems exhibit a sensitivity to initial conditions that tightly limits their predictability.

While we typically embrace self-organization as a crucial concept in the design of Organic Computing systems, self-organization processes taking place in the environment are a frequent source of uncertainty and surprising effects. In fact, simple probabilistic models about the behavior of human, social and technical systems often do not account for complex collective behavior and correlations that result from the complex and often very subtle interactions between individual elements. This often spontaneously occurs at certain critical points in a system's parameter space, thus considerably hindering a sound reasoning about a system's behavior.

4.3.2 Operational Uncertainty

A different kind of uncertainty in Organic Computing systems may be called operational uncertainty. In contrast to environmental uncertainty, here we refer to the use of probabilistic approaches that deliberately introduce an uncertainty about the system's exact behavior that can nevertheless be controlled, adapted and reasoned about. In particular, the active and controlled use of probabilistic schemes is an important approach in the design of systems that incorporate a degree of variation and permissiveness, thus introducing a range of variability that is the source of self-adaptation and self-optimization in Organic Computing Systems. The importance of randomness and variability for is particularly visible in evolutionary algorithms, particle swarm optimization and simulated annealing techniques that are being regularly applied in the design of self-adaptive systems.

Two further aspects of a deliberate and meaningful introduction of randomness into a system can be related to the order-from-noise principle that has been proposed by Heinz von Foerster. First of all, the sensible introduction of randomness or noise into a system can paradoxically result in self-organization processes producing patterns and structures that are more stable and predictable. However, for the degree of noise there typically is a critical point above which noise hinders the self-organization of structures and patterns. As such, the deliberate and meaningful introduction of randomness into a system can be a powerful tool both to foster beneficial self-organization as well as fight the self-organized formation of unwanted patterns and correlations. So one might actually be tempted to say that it is often reasonable to fight uncontrolled uncertainty by deliberately introducing uncertainty in a controlled fashion, thus facilitating a sound stochastic reasoning about systems. A

particularly simple illustrating example for this kind of approach in traditional algorithm engineering is the randomization of input sequences in order to facilitate a sound reasoning about the performance of sort algorithms like QuickSort (for instance in order to get rid of correlations in the input that might result in worst-case performance). Similarly, in Organic Computing systems randomization can be a very useful technique to get rid of unwanted correlations that potentially threaten the system's functioning and performance.

As suggested by the idea of fighting uncontrolled uncertainty by introducing controlled uncertainty, there is an intimate relationship between environmental and operational uncertainty. In fact, we may view the run-time adaptation of systems as as process that matches these two.

4.3.3 Designing Systems for Uncertainty

Having characterized different kinds of uncertainty as well as their importance for Organic Computing systems, we discussed different approaches that promise to improve the handling of unforeseen and rare events in practical systems.

4.3.3.1 Coping with everyday operations

First of all, simply due to the rarity and magnitude of extremal events, it often seems to be a best-practice approach to explicitly distinguish between everyday operations and rare, critical situations. One important reason for this distinction are for instance the often different requirements we may have. In everyday operations we are interested in the system's performance, cost-effectiveness and – at least to a certain degree – its optimality. In critical situations however, we are rather interested in the fact that the system maintains certain basic characteristics that are usually much more modest and relaxed. Under many circumstances it seems reasonable for systems to actively switch between two modes of operation.

4.3.3.2 Building fail-safe systems

In order to build fail-safe systems that are able to survive critical situations, we discussed the approach of goal-shedding. This is intimately related to the switch in the operational mode discussed above. The idea is to formulate a set of prioritized goals which the systems tried to achieve depending on the situation. In critical situations the system should be allowed to throw certain of these goals overboard in order to be enabled to meet at least the more important ones. Here, we discussed a practical example of a robotic submarine which is allowed to sacrifice crucial mission goals in order to at least safely return home based on the remaining power supply. Sometimes pursuing such goal-shedding strategies can require seemingly unintuitive behavior. Here we discussed the example of controlling cascading failures in complex networked systems. In certain classes of network topologies, it can actually be shown that intentionally killing almost all nodes in the system at the first indication of cascading faults is a good strategy to fight critical systemic faults. This admittedly drastic approach sacrifices most of the system's functionality but at the same time guarantees that the system can at least maintain a very basic service, while it would fail completely if no measures were taken. Preventing such errors can also facilitate the system to restart gracefully from a controlled state rather than being brought to a situation where nothing can be done anymore without outer intervention. Such approaches resemble passivity-based control strategies.

4.3.3.3 Reasoning about Organic Computing Systems

The existence of both environmental and operational uncertainty poses severe challenges when it comes to reasoning about the behavior of Organic Computing systems. In classically designed systems, testing is the main method being used to reach a certain level of confidence. In order to successfully deploy Organic Computing systems, we need to be able to derive at least an equal level of confidence about the performance of the system. As these systems are typically reconfigurable at run-time and often deliberate involve operational uncertainty, this necessarily involves stochastic reasoning. Such a stochastic reasoning can provide a number of benefits and can deliver strong guarantees. In traditionally designed systems, guarantees about a system's behavior often tend to become weaker as the system grows in size and complexity. In contrast, in systems consisting of probabilistically behaving and interacting elements, we are often able to extract stochastic guarantees about the system's aggregate behavior that tend to become stronger as the size of the systems grows. This closely resembles guarantees on bulk material properties we are quite used to in thermodynamics and material science. In this particular case, the decrease of uncertainty at the aggregate level is due to scaling effects and based on certain mathematical prerequisites like the enforcement of truly random and uncorrelated individual behavior. However, in many other cases it is less clear how the many small uncertainties related to a system's individual components can be composed to an aggregate picture.

In our discussion we also briefly highlighted potentially interesting further approaches that might be useful to make sound statements about Organic Computing systems, including reachability set analysis, stability notions in dynamical systems and control theory as well as verifiable probabilistic assumption guarantees (like e.g. the probabilistic model checker PRISM from the University of Oxford).

4.3.3.4 Conclusion and Challenges

Uncertainty is a concept of primary importance in the design of Organic Computing (OC) Systems. OC systems should be able to discover and modify models for environmental uncertainty and adapt and manipulate their own probabilistic behavior in accordance. For environmental uncertainty that is due to a lack of knowledge about the environmental conditions, it seems reasonable to employ approaches like e.g. self experimentation, learning and second-order probability techniques in order to proactively reduce uncertainty at least up to a reasonable level. In a sense, acknowledging the importance of uncertainty and probabilistic techniques in Organic Computing Systems retraces the findings of quantum mechanics which embraces probability theory and randomness as fundamental concepts that are necessary to adequately model our reality. Similarly, we now begin to understand the importance of uncertainty for the design and operation of complex technical systems. As Organic Computing community, we need to investigate what techniques from other disciplines can be used to handle the different kinds of uncertainty that are present in our technical systems. Furthermore, the Organic Computing perspective on complex engineered systems might be able to contribute new ideas and abstractions that can be used in other contexts. As a conclusion, it is justified to say that uncertainty is, at the same time, an important motivation, a tough challenge as well as a valuable tool in Organic Computing related research.

Participants

- Michael Beigl
KIT – Karlsruhe Institute of Technology, DE
- Jürgen Branke
University of Warwick, GB
- Uwe Brinkschulte
Univ. Frankfurt am Main, DE
- Sven A. Brueckner
Jacobs Technology Inc. – Ann Arbor, US
- Yuriy Brun
University of Washington – Seattle, US
- Jörg Hähner
Leibniz Univ. Hannover, DE
- Andreas Herkersdorf
TU München, DE
- Michael G. Hinchey
University of Limerick, IE
- Wolfgang Konen
FH Köln, DE
- Robert Laddaga
Doll Inc. – MA, US
- Chris Landauer
Aerospace Corp. – Los Angeles, US
- Falk Langhammer
Living Pages Research GmbH – München, DE
- Erik Maehle
Universität Lübeck, DE
- Christian Müller-Schloer
Leibniz Univ. Hannover, DE
- Phyllis R. Nelson
Cal Poly – Pomona, US
- Gabriele Peters
FernUniversität in Hagen, DE
- Marco Platzner
Universität Paderborn, DE
- Franz J. Rammig
C-LAB – Paderborn, DE
- Paul Robertson
Doll Inc. – MA, US
- Hartmut Schmeck
KIT - Karlsruhe Institute of Technology, DE
- Ingo Scholtes
Universität Trier, DE
- Bernhard Sick
Universität Kassel, DE
- Jürgen Teich
Univ. Erlangen-Nürnberg, DE
- Emil Vassev
University of Limerick, IE
- Laurenz Wiskott
Ruhr-Universität Bochum, DE
- Rolf P. Würtz
Ruhr-Universität Bochum, DE
- Johannes Zeppenfeld
TU München, DE
- Junmei Zhu
EBI – Cambridge, GB

Report from Dagstuhl Seminar 11182

Exploiting graph structure to cope with hard problems

Edited by

Andreas Brandstädt[1], Martin Charles Golumbic[2], Pinar Heggernes[3], and Ross McConnell[4]

1 **Universität Rostock, DE**
2 **Haifa University, IL,** golumbic@cs.haifa.ac.il
3 **University of Bergen, NO,** pinar@ii.uib.no
4 **Colorado State University, US,** rmm@cs.colostate.edu

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 11182 "Exploiting graph structure to cope with hard problems" which has been held in Schloss Dagstuhl – Leibniz Center for Informatics from May 1st, 2011 to May 6th, 2011. During the seminar experts with a common focus on graph algorithms presented various new results in how to attack NP-hard graph problems by using the structure of the input graph. Moreover, in the afternoon of each seminar's day new problems have been posed and discussed.

Seminar 01.–06. May, 2011 – www.dagstuhl.de/11182
1998 ACM Subject Classification G.2.1 Combinatorics, G.2.2 Graph Theory, G.2.3 Applications, F.2.2 Nonnumerical Algorithms and Problems, F.2.3 Tradeoffs among Complexity Measures
Keywords and phrases Graph Classes, Graph Algorithms, NP-completeness, Width Parameters, Approximation Algorithms, Parameterized Complexity
Digital Object Identifier 10.4230/DagRep.1.5.29
Edited in cooperation with Christian Hundt

1 Executive Summary

Andreas Brandstädt
Martin Charles Golumbic
Pinar Heggernes
Ross McConnell

The organizers of the seminar were Andreas Brandstädt, University of Rostock, Germany; Martin C. Golumbic, University of Haifa, Israel; Pinar Heggernes, University of Bergen, Norway; and Ross McConnell, University of Colorado, USA. The collector of the presented seminar material and the editor of the final report was Christian Hundt, University of Rostock, Germany.

One of the main goals of this Dagstuhl seminar was to gather experts with a common focus on graph algorithms but with various specializations to attack NP-hard graph problems using the structure of the input graph. This goal was achieved to a great extent, as the number of participants of our seminar was above the limit that was given beforehand. The seminar was granted space for 30 participants, and we had 35 participants on site. Still there were many experts and young researchers in the field that would like to come but that could

Exploiting graph structure to cope with hard problems, *Dagstuhl Reports*, Vol. 1, Issue 5, pp. 29–46
Editors: Andreas Brandstädt, Martin Charles Golumbic, Pinar Heggernes, and Ross McConnell
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

not be invited due to lack of space. The participants that came to our seminar were from the following countries: 9 from Germany, 7 from USA, 5 from France, 4 from Israel, 4 from Norway, 1 from Canada, 1 from Korea, 1 from Taiwan, 1 from Turkey, 1 from Greece, and 1 from Great Britain.

By bringing together experts with backgrounds in graph classes, optimization, width parameters, and parameterized and exact computing, our aim was that several of the hard problems arising in real applications would eventually find practical solutions. The seminar program was divided into two parts each day of the seminar: (1) presentation of new results, and (2) posing and discussions on new problems. Each of the presented new results and the discussed problems will be explained in detail in the sections below. In this section, we briefly summarize the program of the seminar.

On the first day of the seminar, Jesper Nederlof presented new results on solving connectivity problems parameterized by treewidth in single exponential time, Charis Papadopoulos presented new results on characterizing the linear clique-width of a class of graphs by forbidden induced subgraphs, Yngve Villanger presented new kernelization and approximation results on minimum fill-in of sparse graphs, and Martin Milanič presented new results on hereditary efficiently dominatable graphs. During the problem solving session of the first day, Pinar Heggernes posed a problem on dense subgraphs on proper interval graphs, Andreas Brandstädt and Christian Hundt posed problems on k-leaf powers, Dieter Rautenbach posed a problem on 2-domination on strongly chordal graphs, Pavol Hell asked which problems that are hard on general digraphs become efficiently solvable on interval digraphs, and Feodor Dragan introduced and asked questions about the short fill-in problem.

On the second day of the seminar, Ann Trenk presented a survey on the total linear discrepancy of a poset, Elad Cohen presented new results on vertex intersection graphs of paths on a grid, Yahav Nussbaum presented new results on the recognition of probe proper interval graphs, and Tinaz Ekim presented a survey on polar graphs. During the problem solving session of the second day, Daniel Lokshtanov posed a problem on map graphs, Christian Hundt gave more details on his problem on k-leaf powers, Martin Golumbic posed a problem on the recognition of one-bend EPG and VPG graphs, and Sang-il Oum posed a problem on Bott equivalence between directed acyclic graphs.

On the third day of the seminar, Dieter Rautenbach presented new results on unit interval graphs, R. Sritharan presented new results on finding a sun in graphs, Bernard Ries presented new vertices on coloring vertices of triangle-free graphs, and Wen-Lian Hsu presented new results on PC-trees and planar graphs. In the afternoon of the third day, there was an excursion to Trier and a hike in the close by surroundings of Dagstuhl area.

On the fourth day of the seminar, Pavol Hell presented new results on partitioning chordal graphs, Pim van 't Hof presented new results on contracting graphs to paths and trees, Daniel Lokshtanov presented new results on contracting graphs to bipartite graphs, and Frédéric Maffray presented new results on 3-colorable P_5-free graphs. During the problem solving session of the fourth day, Daniel Lokshtanov posed a problem on the parameterized complexity of protein folding, Yngve Villanger posed a question about forbidden induced subgraphs of circular arc graphs, Christophe Paul asked a question on recognizing circle graphs, and Van Bang Le posed a problem on modified circle graphs.

On the fifth day of the seminar, Feodor Dragan presented new results on graph classes, tree decomposition and approximation algorithms, Elias Dahlhaus presented new results on minimal fill-in ordering of planar graphs in linear time, and Christian Hundt presented new results on the dominating induced matching problem for hole-free graphs. The seminar ended with a brief discussion on all the presented results.

2 Table of Contents

3 Overview of Talks

Monday, May 2nd, 2011

3.1 Solving Connectivity Problems Parameterized by Treewidth in Single Exponential Time

Jesper Nederlof (University of Bergen, NO)

Joint work of Cygan, Marek; Nederlof, Jesper; Pilipczuk, Marcin; Pilipczuk, Michał; van Rooijk, Johan M. M.; Wojtaszczyk, Jakub Onufry
Main reference Marek Cygan, Jesper Nederlof, Marcin Pilipczuk, MichałPilipczuk, Johan van Rooij, Jakub Onufry Wojtaszczyk, "Solving connectivity problems parameterized by treewidth in single exponential time", arXiv:1103.0534v1 [cs.DS]
URL http://arxiv.org/abs/1103.0534v1

For the vast majority of local graph problems (local in the sense that a solution can be verified by checking separately the neighborhood of each vertex) on graphs of small treewidth, standard dynamic programming techniques give $c^{tw(G)}|V|^{O(1)}$ time algorithms, where $tw(G)$ is the treewidth of the input graph $G = (V, E)$ and c is a constant. On the other hand, for problems with a global requirement (usually connectivity) the best known algorithms were naive dynamic programming schemes running in at least $tw(G)^{tw(G)}$ time.

We breach this gap by introducing a technique we named Cut & Count that allows to produce $c^{tw(G)}|V|^{O(1)}$ time Monte Carlo algorithms for most connectivity-type problems, including DIRECTED HAMILTONIAN PATH, STEINER TREE, FEEDBACK VERTEX SET and CONNECTED DOMINATING SET.

This talk will be entirely self-contained. For convenience I will introduce a simplification of the treewidth concept, called $t*n$-graphs. Then I will introduce and proof the Isolation Lemma due to Mulmuley et al. (STOC 1987), and finally I will use this lemma by demonstrating the Cut & Count technique and showing how to solve the STEINER TREE problem in $t * n$-graphs in $c^t n^{O(1)}$ time. If time allows I will also discuss some other results from the paper.

3.2 Characterizing the Linear Clique-Width of a Class of Graphs by Forbidden Induced Subgraphs

Charis Papadopoulos (University of Ioannina, GR)

Joint work of Papadopoulos, Charis; Pinar Heggernes; Daniel Meister

We study the linear clique-width of graphs that are obtained from paths by disjoint union and adding true twins. We show that these graphs have linear clique-width at most 4, and we give a complete characterization of their linear clique-width by forbidden induced subgraphs. As a consequence, we obtain a linear-time algorithm for computing the linear clique-width of the considered graphs.

3.3 Minimum Fill-in of Sparse Graphs: Kernelization and Approximation

Yngve Villanger (University of Bergen, NO)

Joint work of Fomin, Fedor V; Philip, Geevarghese; Villanger, Yngve

The Minimum Fill-in problem is to decide if a graph can be triangulated by adding at most k edges. The problem has important applications in numerical algebra, in particular in sparse matrix computations. We develop kernelization algorithms for the problem on several classes of sparse graphs. We obtain linear kernels on planar graphs, and kernels of size $O(k^{3/2})$ in graphs excluding some fixed graph as a minor and in graphs of bounded degeneracy. As a byproduct of our results, we obtain approximation algorithms with approximation ratios $O(\log k)$ on planar graphs and $O(k^{1/2} \log k)$ on H-minor-free graphs. These results significantly improve the previously known kernelization and approximation results for Minimum Fill-in on sparse graphs.

3.4 On Hereditary Efficiently Dominatable Graphs

Martin Milanič (University of Primorska, SI)

Main reference Martin Milanič, "A hereditary view on efficient domination,' to appear in the Proceedings of the 10th Cologne-Twente Workshop on Graphs and Combinatorial Optimization (CTW 2011).

An efficient dominating set (or perfect code) in a graph is a set of vertices the closed neighborhoods of which partition the graph's vertex set. Determining whether a given graph contains an efficient dominating set is NP-complete in general. We study the class of hereditary efficiently dominatable graphs, that is, graphs every induced subgraph of which contains an efficient dominating set.

Based on a decomposition theorem for (bull, fork, C_4)-free graphs, we derive the forbidden induced subgraph characterization of hereditary efficiently dominatable graphs. We also present a polynomial time algorithm to find an efficient dominating set in a hereditary efficiently dominatable graph.

Tuesday, May 3rd, 2011

3.5 The Total Linear Discrepancy of a Poset

Ann Trenk (Wellesley College, US)

Main reference Howard, David; Shull, Randy; Streib, Noah; Trenk, Ann; "The total linear discrepancy of an ordered set," Discrete Math. 310(2010) 1022-1025.

In this talk we discuss the total linear discrepancy of a poset. If L is a linear extension of a poset P, and x, y is an incomparable pair in P, the height difference between x and y in L is $|L(x) - L(y)|$. The total linear discrepancy of P in L is the sum over all incomparable pairs

of these height differences. The total linear discrepancy of P is the minimum of this sum taken over all linear extensions L of P. While the decision problem of determining whether the (ordinary) linear discrepancy of a poset is at most k is NP-complete, the total linear discrepancy can be computed in polynomial time. In this talk we characterize those linear extensions that are optimal for total linear discrepancy. The characterization provides an easy way to count the number of optimal linear extensions.

3.6 Intersection Graphs of Paths on a Grid

Elad Cohen (University of Haifa, IL)

Joint work of Asinowski, Andrei; Golumbic, Martin C.; Limouzy, Vincent; Lipshteyn, Marina; Stern, Michal

We investigate the class of vertex intersection graphs of paths on a grid, and specifically consider the subclasses that are obtained when each path in the representation has at most k bends (turns). We call such a subclass the B_k-VPG graphs, $k \geq 0$.

We present a complete hierarchy of VPG graphs relating them to other known families of graphs. String graphs are equivalent to VPG graphs. The grid intersection graphs are shown to be equivalent to the bipartite B_0-VPG graphs.

Chordal B_0-VPG graphs are shown to be exactly Strongly Chordal B_0-VPG graphs. We prove the strict containment of B_0-VPG and circle graphs into B_1-VPG. Planar graphs are known to be in the class of string graphs, and we prove here that planar graphs are B_3-VPG graphs.

In the case of B_0-VPG graphs, we observe that a set of horizontal and vertical segments have strong Helly number 2. We show that the coloring problem for B_k-VPG graphs, for $k \geq 0$, is NP-complete and give a 2-approximation algorithm for coloring B_0-VPG graphs. Furthermore, we prove that triangle-free B_0-VPG graphs are 4-colorable, and this is best possible.

This work was presented at LAGOS 2011.

3.7 Recognition of Probe Proper Interval Graphs

Yahav Nussbaum (Tel Aviv University, IL)

In a probe graph the vertex set is partitioned into probes and non-probes. A probe proper interval graph is an intersection graph of intervals on the line such that every vertex is mapped to an interval, no interval contains another, and two vertices are adjacent if and only if the corresponding intervals intersect and at least one of them is a probe.

In this talk I will present the first linear-time algorithm that determines whether the input graph is a probe proper interval graph, and if the answer is positive then the algorithm constructs a corresponding set of intervals.

3.8 Polar Graphs

Tinaz Ekim (Bogaziçi University – Istanbul, TR)

G is polar if its vertex set admits a partition (A, B) where $G[A]$ is complete multipartite and $G[B]$ is a disjoint union of cliques. We distinguish a special case, called monopolar, when A is a stable set. The recognition of polar and monopolar graphs are NP-complete in general, but they become polynomially solvable in some classes of graphs. In this talk, we survey the recognition and the forbidden subgraph characterization of polar and monopolar graphs in cographs, in chordal graphs, in permutation graphs, in line graphs and in claw-free graphs. We also discuss some related research topics.

Wednesday, May 4th, 2011

3.9 Unit Interval Graphs — A Story with Open Ends

Dieter B. Rautenbach (Universität Ulm, DE)

Joint work of Rautenbach, Dieter B.; Szwarcfiter, Jayme L.

We give two structural characterizations of the class of finite intersection graphs of the open and closed real intervals of unit length. This class is a proper superclass of the well known unit interval graphs.

3.10 Finding a Sun in Graphs

Sritharan, R (University of Dayton, US)

Joint work of Eschen, Elaine; Hoang, Chinh; Spinrad, Jeremy; Sritharan, R

A sun is a graph with a Hamiltonian cycle $(x_1, y_1, ..., x_n, y_n), n >= 3$, where each x_i has degree two and the y_i vertices form a clique. Deciding whether an arbitrary graph contains a sun is NP-complete, while for some graph classes (e.g. chordal, hhd-free) the problem is solvable in polynomial time. We give a polynomial-time algorithm to test for suns in building-free graphs. Building-free graphs generalize Meyniel graphs (and hence, hhd-free, i-triangulated, and parity graphs). We also discuss an elimination scheme of vertices for graphs that do not contain any gem or a building.

3.11 Coloring Vertices of Triangle-free Graphs

Bernard Ries (Université Paris-Dauphine, FR)

The vertex coloring problem is known to be NP-complete in the class of triangle-free graphs. Moreover, it remains NP-complete even if we additionally exclude a graph F which is not a forest. We study the computational complexity of the problem in (K_3, F)-free graphs with F being a forest. From known results it follows that for any forest F on 5 vertices, the vertex coloring problem is polynomial-time solvable in the class of (K_3, F)-free graphs. In the present paper, we study the problem for (K_3, F)-free graphs with F being a forest on 6 vertices. It is known that in the case when F is the star $K_{1,5}$, the problem is NP-complete. It turns out that this is the only hard case.

3.12 PC-Trees and Planar Graphs

Wen-Lian Hsu (Academica Sinica – Taipei, TW)

Linear time planarity test was first established by Hopcroft and Tarjan in 1974 based on a path addition approach. A vertex addition approach, originally developed by Lempel, Even and Cederbaum, was later improved by Booth and Lueker in 1976 to run in linear time using a data structure called "PQ-tree". PQ-tree can also be used to test the consecutive ones property and to recognize interval graphs. Shih and Hsu developed a linear time planarity test based on PC-trees. PC-tree, a generalization of PQ-tree [W.L. Hsu and R. McConnell, PC-trees and circular-ones arrangement, Theoretical Computer Science 296(1), 2003, pp. 99-116], is more natural in representing the relationships between biconnected components and nodes in planar graphs. An earlier version of Shih and Hsu has been referred to as the simplest linear time planarity test by Thomas in his lecture notes. In this talk we shall describe an ultimate version of planarity test based on PC-trees, which is much simpler than any previous version. Moreover, we shall describe how to extend this algorithm naturally to find maximal planar subgraphs in linear time for arbitrary graphs.

Thursday, May 5th, 2011

3.13 Graph Partitions

Pavol Hell (Simon Fraser University – Burnaby, CA)

Some partition problems for chordal graphs have forbidden induced subgraph characterizations and others don't. I will describe some examples and classify all partitions with up to three parts. The main problem of classification is open. The talk will include joint work with Tomas Feder, Shekoofeh Nekooei-Rizi, Juraj Stacho, Geordie Schell, and Wing Xie.

3.14 Contracting Graphs to Paths and Trees

Pim van 't Hof (University of Bergen, NO)

The problems Path-Contractibility and Tree-Contractibility take as input an undirected graph G and an integer k, and ask whether we can obtain a path or a tree, respectively, by contracting at most k edges of G. Both problems are NP-complete, and fixed parameter tractability follows from Courcelle's Theorem. We present algorithms with running time $c^k n^{O(1)}$ with small constants $c < 5$ for both problems. Furthermore, we show that Path-Contractibility has a kernel with at most $5k + 3$ vertices, while Tree-Contractibility does not have a polynomial kernel unless coNP is a subset of NP/poly. Interestingly, Feedback Vertex Set, which can be seen as the vertex deletion variant of Tree-Contractibility, is known to have a kernel with $O(k^2)$ vertices.

3.15 Contraction to Bipartite Graphs is Fixed Parameter Tractable

Daniel Lokshtanov (University of California – San Diego, US)

Main reference Pinar Heggernes, Pim van 't Hof, Benjamin Lévêque, Daniel Lokshtanov, Christophe Paul, "Contracting Graphs to Paths and Trees," arXiv:1104.3677v1 [cs.DS]
URL http://arxiv.org/abs/1104.3677v1

We initiate the study of the Bipartite Contraction problem from the perspective of parameterized complexity. In this problem we are given a graph G and integer k, and the task is to determine whether we can obtain a bipartite graph from G by a sequence of at most k edge contractions. Our main result is that Bipartite Contraction admits a $f(k)n^{O(1)}$ time algorithm.

Despite a strong resemblance between Bipartite Contraction and the classical Odd Cycle Transversal (OCT) problem, the methods developed to tackle OCT do not seem to be directly applicable to Bipartite Contraction. Our algorithm is based on a novel combination of the irrelevant vertex technique introduced by Robertson and Seymour and the concept of important separators. Both techniques have previously been used as key components of cornerstone theorems in parameterized complexity.

However, to the best of our knowledge, this is the first time the two techniques are applied in unison.

3.16 On the Structure of 3-colorable P_5-free Graphs

Frédéric Maffray (CNRS – Grenoble, FR)

Joint work of Maffray, Frédéric; Morel, Grégory
Main reference Maffray Frédéric, Morel Grégory, "On 3-colorable P5-free graphs," Cahiers Leibniz Research Report 191 (Laboratoire G-SCOP, Grenoble, France), March, 2011.
URL http://www.g-scop.inpg.fr/CahiersLeibniz/

We consider the class of 3-colorable P_5-free graphs. We give a complete description of the structure of the graphs in that class.

Our main result is as follows. If G is a connected 3-colorable P_5-free graph, then either G has a homogeneous set (that induces a bipartite subgraph), or G is the complement of a chordal graph, or G has one among four possible types of well-defined structures.

From this structural description of 3-colorable P_5-free graphs, we can derive a linear time algorithm that tests membership in the class, and a linear-time algorithm that finds a maximum weight stable set.

Friday, May 6th, 2011

3.17 An Approximation Algorithm for the Tree t-Spanner Problem on Unweighted Graphs via Generalized Chordal Graphs

Feodor F. Dragan (Kent State University, US)

A spanning tree T of a graph G is called a tree t-spanner of G if the distance between every pair of vertices in T is at most t times their distance in G. In this talk, we present an algorithm which constructs for an n-vertex m-edge unweighted graph G:

- a tree $(2\lfloor \log_2 n \rfloor)$-spanner in $O(m \log n)$ time, if G is a chordal graph (i.e., every induced cycle of G has length 3);
- a tree $(2\rho\lfloor \log_2 n \rfloor)$-spanner in $O(mn \log^2 n)$ time or a tree $(12\rho\lfloor \log_2 n \rfloor)$-spanner in $O(m \log n)$ time, if G is a graph admitting a Robertson-Seymour's tree-decomposition with bags of radius at most ρ in G; and
- a tree $(2\lceil t/2 \rceil \lfloor \log_2 n \rfloor)$-spanner in $O(mn \log^2 n)$ time or a tree $(6t\lfloor \log_2 n \rfloor)$-spanner in $O(m \log n)$ time, if G is an arbitrary graph admitting a tree t-spanner.

For the latter result we use a new necessary condition for a graph to have a tree t-spanner: if a graph G has a tree t-spanner, then G admits a Robertson-Seymour's tree-decomposition with bags of radius at most $\lceil t/2 \rceil$ and diameter at most t in G.

3.18 Minimal Fill-in Ordering of Planar Graphs in Linear Time

Elias Dahlhaus (DB Systems GmbH – Frankfurt, DE)

We present an alternative linear time algorithm for ordering the vertices of a planar graph such that the set of fill-in edges is minimal with respect to the subset relation. The algorithm is simpler than the algorithm in [E. Dahlhaus, Minimal Elimination of Planar Graphs, Algorithm Theory – SWAT'98, LLNCS 1432 (1998), pp. 210–221] and is easily parallelizable, as it does not rely on the computation of a breadth-first search tree.

3.19 Efficient Edge Domination on Hole-free Graphs

Christian Hundt (Universität Rostock, DE)

Joint work of Brandstädt, Andreas; Hundt, Christian; Nevries, Ragnar
Main reference Andreas Brandstädt, Christian Hundt and Ragnar Nevries. Efficient Edge Domination on Hole-free Graphs in Polynomial Time. In Proc. 9th Latin American Theoretical Informatics Symposium, LATIN 2010, LNCS, Springer-Verlag 2010, pp. 650-661.
URL http://dx.doi.org/10.1007/978-3-642-12200-2_56

This talk deals with the *Efficient Edge Domination* Problem (*EED*, for short), also known as *Dominating Induced Matching Problem*, which asks, given an undirected graph $G = (V, E)$, for an induced matching $M \subseteq E$ that simultaneously dominates all edges of G. Thus, the distance between edges of M is at least two and every edge in E is adjacent to an edge of M. EED is related to parallel resource allocation problems, encoding theory and network routing. The problem is NP-complete even for restricted classes like planar bipartite and bipartite graphs with maximum degree three.

However, the complexity has been open for chordal bipartite graphs. We show that EED can be solved in polynomial time on hole-free graphs. Moreover, we provide even linear time for chordal bipartite graphs. Finally, we strengthen the NP-completeness result to planar bipartite graphs of maximum degree three with girth k for every fixed k.

4 Open Problems

Monday, May 2nd, 2011

4.1 Dense Subgraphs in a Proper Interval Graph

Pinar Heggernes (University of Bergen, NO)

Dense Subgraph in a Proper Interval Graph
Given: a proper interval graph $G = (V, E)$ and an integer $k \in \mathbb{N}$
Find: an induced subgraph of k vertices in G that has the largest possible number of edges

Is this problem polynomial-time solvable or NP-hard?

4.2 Structure and Recognition of k-Leaf Powers

Andreas Brandstädt (Universität Rostock, DE)

A graph $G = (V, E)$ is a *k-leaf power* for a $k \in \mathbb{N}$ if there is a tree T such that V is exactly the set of leaves of T and $uv \in E$ if and only if $d_T(u, v) \leq k$. A graph G is a *leaf power* if it is a k-leaf power for some $k \in \mathbb{N}$. For instance, the set of 3-leaf powers forms precisely the class of (bull, dart, gem)-free chordal graphs. Moreover, some results are known for $k = 4$ and $k = 5$.

What is the structure of leaf powers, respectively of k-leaf powers for $k \geq 6$? How complex are the corresponding recognition problems?

4.3 The Complexity of 2-Domination on Strongly Chordal Graphs

Dieter B. Rautenbach (Universität Ulm, DE)

The *dominating set problem* on chordal graphs is NP-hard. But it is computable in polynomial time on strongly chordal graphs.

2-Domination
Given: a graph $G = (V, E)$
Find: a smallest set $D \subseteq V$ such that the neighborhood $N(v)$ of all $v \in V \setminus D$ fulfills $|N(v) \cap D| \geq 2$

Is this problem solvable in polynomial time on strongly chordal graphs, too? Is the problem's complexity known for interval graphs? (R. Sritharan)

4.4 Problems Becoming Feasible on Classes of Interval Digraphs

Pavol Hell (Simon Fraser University – Burnaby, CA)

Directed interval graphs $G = (V, E)$ are the graphs that have an interval model where every vertex v is represented by two intervals (I_v, J_v) and there is a directed edge $uv \in E$ if and only if I_u intersects J_v. What problems that are intractable on general directed graphs become tractable on interval digraphs?

The same question is also interesting for the following subclasses of interval digraphs:

Interval catch digraphs, the digraphs $G = (V, E)$ that have a model where for all $v \in V$ the intervals J_v consist of just one point contained in the interval I_v. They have a forbidden subgraph characterization that resembles AT-freeness and they also exhibit an ordering characterization.

Adjusted interval digraphs, the digraphs $G = (V, E)$ that have a model where for all $v \in V$ it is true that the intervals I_v and J_v have the same left endpoint. Such graphs have another ordering characterization.

Adjusted interval catch digraphs, the intersection of the previous two classes. The graphs in this class can be recognized in linear time. Moreover, they can be characterized in the following way: In an adjusted interval catch digraph $G = (V, E)$ it is true for all $u, v \in V$ that there is an edge uv if and only if J_u and I_v intersect and I_u starts earlier than I_v.

4.5 Structure and Complexity of Short Fill-in for Strongly Chordal Graphs and Tree Powers

Feodor F. Dragan (Kent State University, US)

Short Fill-in with Property Π
Given: a graph $G = (V, E)$
Find: the smallest number $d \in \mathbb{N}$ that permits an edge set E' between vertices of G such that $G' = (V, E \cup E')$ has property Π and for all $uv \in E'$ it is true $d_G(u, v) \leq d$

If Π is the property of G' being chordal, then the problem is equivalent to finding the tree length of G. What is the structure and complexity of the problem, if Π is the property of G' being strongly chordal or the power of a tree?

The decision version of the problem is related to sandwich problems, because, if d is given, then G' is a graph between G and the d'th power G^d of G.

Tuesday, May 3rd, 2011

4.6 Fast and Simple Recognition of Map Graphs

Daniel Lokshtanov (University of California – San Diego, US)

Map graphs are a generalization of planar graphs that can be defined as the graphs $G = (V, E)$ that permit a planar bipartite graph $H = (V \cup R, E')$ having a bipartition into the vertex set V of G and an additional vertex set R such that two vertices $u, v \in V$ are adjacent in G, i. e., $uv \in E$, if and only if u and v have a common neighbor in H, hence, if there is a vertex $w \in R$ with $uw, vw \in E'$. This means that $G = H^2[V]$.

The recognition of map graphs can be done in $O(|V|^{120})$ time by an algorithm of Mikkel Thorup in 1998. Is it possible to realize a faster recognition of map graphs with a nice and simple algorithm?

4.7 The Leaf Rank of Proper Interval Graphs

Christian Hundt (Universität Rostock, DE)

Leaf Rank of Proper Interval Graphs
Given: a proper interval graph $G = (V, E)$
Find: the smallest possible number $k \in \mathbb{N}$ such that G is a k-leaf power

Proper interval graphs $G = (V, E)$ without true twins have a unique (up to reversal) ordering of their vertices V. In particular, such an ordering $(v_1, \ldots, v_n)$, $n = |V|$ fulfills for all $i < j \in \{1, \ldots, n\}$ that $v_i v_j \in E$ implies $v_i v_r \in E$ and $v_r v_j \in E$ for all $i < r < j$. Based on this ordering the above problem is equivalent to the following one:

Weighted Linear Steiner Root of a Proper Interval Graph
Given: a proper interval graph $G = (V, E)$ without true twins where $n = |V|$
Find: the smallest integer s such that there exist integers $x_1, \ldots, x_n, w_1, \ldots, w_{n-1} \in \mathbb{N}_0$ which fulfill for all $i < j \in \{1, \ldots, n\}$ that $v_i v_j \in E$ if and only if $x_i + w_i + \ldots + w_{j-1} + x_j \leq s$ where $(v_1, v_2, \ldots, v_n)$ is the unique ordering of the vertices V

In particular, the leaf rank of a given proper interval graph G is k if and only if for $s = k - 2$ there is a s-weighted linear Steiner root for G', the graph that emerges from G by removing all true twins. How complex are the two problems above?

4.8 The Recognition Problem of B_1-EPG and B_1-VPG Graphs

Martin Charles Golumbic (Haifa University, IL)

The class *EPG* consists of all graphs $G = (V, E)$ where every vertex $v \in V$ corresponds to a path $P(v)$ in a grid $\mathcal{G}$ and for all $u, v \in V$ the edge uv belongs to E if and only if the paths $P(u)$ and $P(v)$ intersect in at least one grid edge. Analogously, the graphs $G = (V, E)$ in *VPG* correspond to collections of paths in a grid where $uv \in E$ if and only if $P(u)$ and $P(v)$ intersect in at least one grid point. If all paths may have at most one bend, i. e., direction change, then the described graph classes are referred to as B_1-*EPG* and B_1-*VPG*, respectively.

No characterization is known for B_1-EPG graphs or for B_1-VPG graphs, and the recognition problem for both classes is open. Is it NP-hard to recognize B_1-EPG graphs, respectively B_1-VPG graphs, or are there polynomial time algorithm for these problems?

The same question is also interesting for some subclasses. We call a set of single bend paths *boundary generated* if the grid has a rectangular boundary and the endpoints of all paths are on the boundary. If for instance paths are restricted to start from the left y-axis and bend down to the x-axis, then B_1-VPG becomes exactly the well known class of permutation graphs. Under the same restriction, the B_1-EPG graphs can be shown to be the graphs whose edge set can be covered by two clique partitions of G. In this case it is known that these graphs are recognizable in polynomial time. In general, what happens to the complexity of the recognition problems for boundary generated B_1-EPG graphs, and respectively B_1-VPG graphs, for each subset of the four possible direction changes?

4.9 Bott-Equivalent and Diffeomorphilic Directed Acyclic Graphs

Sang-il Oum (KAIST – Daejeon, KR)

Real Bott manifolds can be described in terms of directed acyclic graphs. Given two directed acyclic graphs H and G, the graph H is *Bott-equivalent* to G if H can be obtained from G by the operations of (1) local complementation and (2) slide. Is it possible to decide in polynomial time whether two given directed acyclic graphs H and G are Bott-equivalent?

Two real Bott manifolds are *diffeomorphilic* if and only if the corresponding graphs H and G fulfill that H is isomorphic to a graph that is Bott-equivalent to G. How complex is the decision of this problem?

Thursday, May 5th, 2011

4.10 Fixed Parameter Tractability of Protein Folding

Daniel Lokshtanov (University of California – San Diego, US)

How complex is it to decide if a given graph G is a subgraph of a grid? Is there an FPT algorithm with parameter k, if G is a path with k additional edges?

This problem is related to the *protein folding problem*. Hence, is protein folding in FPT when parameterized by the number of points?

4.11 Forbidden Subgraph Characterization of Circular-Arc Bipartite Graphs

Yngve Villanger (University of Bergen, NO)

Circular arc graphs can be recognized in $O(n+m)$ time due to an algorithm by R. McConnell. What is the forbidden induced subgraph characterization of circular arc graphs?

A bipartite graph H, with a fixed bipartition (X, Y), is an interval bigraph if the vertices of H can be represented by a family of intervals $I_v, v \in X \cup Y$, so that, for $x \in X$ and $y \in Y$, x and y are adjacent in H if and only if I_x and I_y intersect. A circular-arc bigraphs is, the generalization of interval bigraphs, where the intervals of the models are replaced by arcs of a circle. What is the forbidden induced subgraph characterization of circular-arc bigraphs graphs?

4.12 Linear Time Recognition of Circle Graphs

Christophe Paul (CNRS, UniversitéMontpellier II, FR)

Recognizing circle graphs works in $O(n^2)$ time via split decomposition. Split decomposition itself can be realized in linear time $O(n+m)$ but in this way it cannot be used to recognize circle graphs. Recent split decomposition algorithms that work in an incremental fashion, run in $O((n+m)\alpha(n,m))$ time where α is the inverse of the Ackermann function, and they can be used to recognize circle graphs in $O((n+m)\alpha(n,m))$ time. Is it possible to get rid of the $\alpha(m,n)$ factor in the running time, i. e., can circle graphs be recognized in linear time?

4.13 Characterization and Recognition of a Circle Graph Generalization

Van Bang Le (Universität Rostock, DE)

Consider the generalization of circle graphs that have intersection models consisting of (not necessarily straight) lines where at least one point (but not necessarily two) is situated on a circle. Can these graphs be recognized in polynomial time? Particularly, are there graphs that do not belong to this class? (E. Köhler) Is it possible to find the independence number for such graphs if the intersection model is given?

Participants

- Andreas Brandstädt
Universität Rostock, DE
- Elad Cohen
University of Haifa, IL
- Elias Dahlhaus
DB Systems GmbH – Frankfurt, DE
- Feodor F. Dragan
Kent State University, US
- Tinaz Ekim
Bogaziçi Univ. – Istanbul, TR
- Elaine M. Eschen
West Virginia University – Morgantown, US
- Martin Charles Golumbic
Haifa University, IL
- Pinar Heggernes
University of Bergen, NO
- Pavol Hell
Simon Fraser University – Burnaby, CA
- Wen-Lian Hsu
Academica Sinica – Taipei, TW
- Christian Hundt
Universität Rostock, DE
- Ekkehard Köhler
BTU Cottbus, DE
- Dieter Kratsch
Univ. Paul Verlaine – Metz, FR
- Van Bang Le
Universität Rostock, DE
- Daniel Lokshtanov
University of California – San Diego, US
- Frédéric Maffray
CNRS – Grenoble, FR
- Ross McConnell
Colorado State University, US
- Daniel Meister
Universität Trier, DE
- Martin Milanič
University of Primorska, SI
- Haiko Müller
University of Leeds, GB
- Jesper Nederlof
University of Bergen, NO
- Ragnar Nevries
Universität Rostock, DE
- Yahav Nussbaum
Tel Aviv University, IL
- Sang-il Oum
KAIST – Daejeon, KR
- Charis Papadopoulos
University of Ioannina, GR
- Christophe Paul
CNRS, Univ. Montpellier II, FR
- Andrzej Proskurowski
University of Oregon, US
- Dieter B. Rautenbach
Universität Ulm, DE
- Bernard Ries
Université Paris-Dauphine, FR
- R. Sritharan
University of Dayton, US
- Ioan Todinca
Université d'Orleans, FR
- Ann Trenk
Wellesley College, US
- Pim van 't Hof
University of Bergen, NO
- Yngve Villanger
University of Bergen, NO
- Egon Wanke
Universität Düsseldorf, DE

Report from Dagstuhl Seminar 11191

Graph Drawing with Algorithm Engineering Methods

Edited by

Camil Demetrescu[1], Michael Kaufmann[2], Stephen Kobourov[3], and Petra Mutzel[4]

1 **Sapienza University of Rome, IT,** demetres@dis.uniroma1.it
2 **Universität Tübingen, DE,** mk@informatik.uni-tuebingen.de
3 **University of Arizona – Tucson, US,** kobourov@cs.arizona.edu
4 **TU Dortmund, DE,** petra.mutzel@cs.tu-dortmund.de

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 11191 "Graph Drawing with Algorithm Engineering Methods". We summarize the talks, open problems, and working group discussions.

Seminar 8.–13. May, 2011 – www.dagstuhl.de/11191
1998 ACM Subject Classification E.1 Data structures, F.2 Analysis of algorithms and problem complexity, G.1.6 Optimization, G.2.1 Combinatorics, G.2.2 Graph Theory
Keywords and phrases Algorithm Engineering, Graph Drawing
Digital Object Identifier 10.4230/DagRep.1.5.47
Edited in cooperation with Carsten Gutwenger and Karsten Klein

1 Executive Summary

Camil Demetrescu
Michael Kaufmann
Stephen Kobourov
Petra Mutzel

Automated graph drawing deals with the layout of relational data arising from computer science (database design, data mining, software engineering), and other sciences such as bioinformatics and sociology (social networks). The relational data are typically modeled as graphs, which can be visualized through diagrams drawn in the plane. The main objective is to display the data in a meaningful fashion, (i.e., in a way that shows well the underlying structures) and that often depends on the application domain. Although high quality algorithms exist for many optimization problems that arise in graph drawing, they are often complex and difficult to implement, and theoretically efficient algorithms may have unacceptable runtime behavior even for small-to-medium sized real-world instances. Also large graphs like, e.g., molecular interaction networks, may render exact but complex algorithms infeasible and require approximate or heuristic solutions.

Integrating automated graph drawing techniques into real-world software systems poses several algorithm engineering challenges. To achieve effective implementations, algorithms and data structures designed and analyzed on abstract machine models must be carefully tuned for performance on real hardware platforms. This task is becoming increasingly more difficult due to the impressive growth of data to be visualized in modern applications, as

well as their highly dynamic and data-intensive nature. Developers can no longer ignore architectural aspects such as the presence of complex memory hierarchies and multiple cores, which are likely to shape the design of novel algorithmic techniques and the way they will be implemented and engineered in the future.

The aim of this seminar was to bring together researchers from the algorithm engineering and graph drawing communities in order to strengthen and foster collaborations in this area and to identify key research directions for the future.

The seminar was attended by 48 participants from both academia and industry. Much was accomplished, fostered by the productive atmosphere of the Dagstuhl Center. Here we describe some of the more important achievements.

The program consisted of a wide variety of presentations, working group sessions and discussion sessions. The presentations included several survey lectures:

- Beppe Liotta provided a survey on graph visualization paradigms, and discussed general design principles for the realization of effective graph visualization systems.
- Emden Gansner suggested rules in order to get efficient and accurate graph drawing algorithms.
- Ulrik Brandes discussed experimental algorithmics and the relationship between graph drawing algorithms and algorithm engineering.
- Rudolf Fleischer's talk about algorithm engineering and his statement (taken from the definition in the German priority program SPP 1307 *Algorithm Engineering*) that the algorithm engineering cycle should be driven by falsifiable hypotheses, started a lively discussion among the participants.
- Rico Jakob provided a talk on engineering architecture aware algorithms and provided some thoughts about hardware sensitive algorithms. He convinced us that the new computer architectures will strongly influence future algorithmic research.
- Kurt Mehlhorn introduced us into the new and exciting area of slime mould that solves shortest path and network design problems. He would be interested in seeing if slime mould could possible solve graph drawing problems.

Beyond the survey lectures, highlights of the seminar included the two introductory sessions, the open problem sessions, and the working groups.

In two sessions, we have identified over two dozen open problems, which later crystallized into about a dozen well-defined problems, each of which were of interest to several participants. We had working groups on the following topics: Rotating binary trees, feedback arc set convergence, edge bundling models, co-occurence in bipartite graphs, RAC drawings, BRAC drawings, minimum branch spanning tree, cluster tree embedding, point set embeddings, parallel graph drawing, and library of graphs. Participants shared ideas and material using the online seminar Wiki.

The dissemination sessions at the end of the workshop showed that many of the working groups have achieved initial results, which may lead to future publications.

Arguably the most-appreciated features of the Seminar were the lively open discussion sessions, which led to several concrete proposals for the future of the field which, as a result of the workshop, are now being actively pursued.

A big step forward has been done concerning an *online library of graphs*. The graph drawing community would like to have such a library, however, there was no consensus about the requirements on such a graph archive. The working group conducted a survey on requirements for a graph archive during the Dagstuhl seminar. Two groups (Dortmund and Tübingen) presented their ideas and prototypes of such an archive. In order to foster

future work and to encourage participation and contributions, it was suggested that the GD proceedings should offer the opportunity to publish papers concerning the library. Moreover, the collection of many benchmark graphs has already begun.

We used the opportunity to bring together experts in algorithm engineering for multi-core algorithms with graph drawing researchers in order to discuss how graph drawing algorithms can be re-engineered to better take advantage of modern computer architecture into account. This working group was inspired by the many different backgrounds of group members. They have discussed how to improve data locality, or exploit multi-core processors, in particular for the widely used Sugiyama drawing method.

Subjectively (from interacting with the attendees) and objectively (from the official feedback data) we believe that the participants enjoyed the great scientific atmosphere offered by Schloß Dagstuhl, and profited from the scientific program and the fruitful discussions. We are grateful for having had the opportunity to organize this seminar. Special thanks are due to Carsten Gutwenger and Karsten Klein for their invaluable assistance in the organization and the running of the seminar.

2 Table of Contents

3 Overview of Talks

In this section, we report the abstracts of five survey talks on algorithm engineering aspects arising in graph drawing applications. The presentations provided a starting point for the seminar and covered foundational aspects, allowing participants with different grounds of expertise to share common methodologies and goals.

3.1 Experimental Algorithmics

Ulrik Brandes (Universität Konstanz, DE)

Joint work of Brandes, Ulrik; Karrenbauer, Andreas

The algorithm engineering cycle is said to be driven by falsifiable hypotheses that are validated by experiment. There is little evidence, however, that this is indeed common practice. We sketch the concepts of formal experimentation as established in other disciplines and propose a mapping to experimental algorithmics. After reviewing experimental work in graph drawing, we ask whether more formal experimentation is needed.

3.2 AE meets GD

Rudolf Fleischer (Fudan University – Shanghai, CN)

The scientific method postulates a certain framework for doing experiments and how to interpret experimental results theoretically. Experimentors in computer science and graph drawing seem sometimes to be unaware of these principles.

3.3 Notes on Practical Graph Drawing

Emden R. Gansner (AT&T Research – Florham Park, US)

To be useful, graph drawing algorithms need to be engineered to be efficient and accurate. To be used, it helps if the graph drawing algorithms are implemented and run using some practical rules of thumb. These include:

- Follow general software engineering principles
- Use optimizations where possible; use heuristics when necessary
- Leverage the geometry
- Provide the user with a rich set of drawing features
- Construct simple, flexible, reusable interfaces
- Stress robustness, especially as regards scalability

In this talk, we discuss these rules, providing motivations and examples.

3.4 Engineering Architecture Aware Algorithms – two case studies and some thoughts on Algorithms Engineering

Riko Jacob (TU München, DE)

I will present preliminary results from two ongoing projects, both regarding memory intensive tasks in (huge) internal memory. One is asking the question if I/O-efficient (sorting based) algorithms can outperform the direct algorithm when permuting several gigabytes of data in internal memory. The other one reports on an efficient implementation for numerical computations in high dimensional settings using so called sparse grids. Again, the improved implementation is inspired by I/O-efficient algorithms. Finally, I will try to show how these two examples fit into a more general theme of using and combining theoretical models and experiments to find the algorithm and implementation that performs best on a given hardware.

3.5 The Anatomy of a graph visualization system

Giuseppe Liotta (University of Perugia, IT)

Graph visualization addresses the problem of efficiently conveying the structure of relational information, which is typically modeled by networks. Therefore, graph visualization systems are largely used for information exploration and knowledge discovery, particularly in those applications that need to manage, process and analyze large sets of data. The design of a graph visualization systems typically addresses questions that belong to the intersection of different disciplines, such as graph algorithms, data mining, software engineering, algorithm engineering, and visual analytics. In this talk I will shortly review some common and emerging graph visualization paradigms, discuss general design principles present application examples, and compare different models for the realization of effective graph visualization systems.

4 Working Groups

4.1 Graph Archive

Christian Bachmaier, Franz J. Brandenburg, Philip Effinger, Carsten Gutwenger, Jyrki Katajainen, Karsten Klein, Miro Spönemann, and Michael Wybrow

In order to evaluate, compare, and tune graph (drawing) algorithms, experiments on well designed benchmark sets have to be performed. This, together with the goal of reproducibility of experimental results, creates a demand for an adequate archive to gather and store graph instances. Such an archive would ideally allow annotation of instances or sets of graphs with

additional information like graph properties and references to the respective experiments and results. We examine the requirements and formulate the next steps needed to produce an easily accessible library of graphs that provides the required services. Through successful community involvement, it is expected that the archive will contain a representative selection of both real-world and generated graph instances, covering significant application areas as well as relevant classes of graphs.

4.2 Spanning Trees with Few Branches

Markus Chimani, Aparna Das, and Joachim Spoerhase

Given a graph G finding the spanning tree T of G with the smallest number of *branch nodes* (nodes of degree at least 3) is known as the *minimum branch node spanning tree* problem (MBST). The problem is motivated by optical networks where switches of degree greater than two are expensive as they require sophisticated hardware to split light. Unfortunately [1] shows that the problem is not only NP-hard but also nonapproximable.

Observe that a spanning tree that minimizes the number of branchings also maximizes the number of low-degree nodes (leaves and degree-two nodes). Although this maximization version of MBST leads to the same set optimum solutions it has better approximability properties. In fact any spanning tree is already a 2-approximation since at least half the nodes have degree 1 or 2.

We investigate the following local search algorithm. A *legal k-flip* replaces at most k edges of T with the same number of edges in $E(G) - E(T)$ such that the resulting graph T' is a spanning tree of G with strictly more low-degree nodes than T. (In a more refined algorithm a legal k-flip increases a suitably defined potential function.) Starting with an arbitrary spanning tree we perform legal 2-flips until we obtain a spanning tree (local optimum) for which no legal 2-flips can be performed anymore.

We conjecture that the performance guarantee of this algorithm is strictly better than 2. Our current proof sketch suggests a factor of at most 1.8. Our analysis is based on the observation that the performance ratio of 2 can only be achieved by spanning trees in which (almost) all branch nodes have degree 3. Our approach consists in showing that any node u that has degree 3 in a local optimum either has degree 3 in a global optimum, too, or is "associated" with certain nodes of degree $\neq 3$ in T.

References

1 L. Gargano, P. Hell, L. Stacho and U. Vaccaro, Spanning trees with bounded number of branch vertices, Proc. of ICALP 02, LNCS vol. 2380, Springer (2002), pp. 355–365.

4.3 Parallel Graph Drawing

Deepak Ajwani, Camil Demetrescu, Carsten Gutwenger, Robert Krug, Henning Meyerhenke,Petra Mutzel, Stefan Näher, Georg Sander, and Matthias Stallmann

In recent years hardware architects have hit the power wall. Increasing the processor clock speeds further would yield an overproportial power consumption and heat production. That is why smaller circuit designs are exploited by adding more computing cores instead. Nowadays commodity processors already have four or six cores, a number expected to rise significantly in the next decade. Special processors already support hundreds of hardware threads today. To make efficient use of the available hardware for an algorithmic problem, it is necessary to expose the inherent parallelism in a problem at hand. Recently it became popular to use GPUs to accelerate computations. So it seems natural to parallelize graph drawing algorithms. Few papers consider algorithms based on force-directed methods on GPUs and similar architectures.

One of the most popular algorithms for drawing hierarchical layouts of directed graphs is the Sugiyama algorithm [1]. We are not aware of any parallel Sugiyama-type algorithms. The main phases are ranking, crossing minimization and coordinate assignment. The time-critical task is crossing minimization, whose main procedure sweeps repeatedly up and down the hierarchy. A sweep starts at the second level and computes the barycenter or barycenter values, respectively, for each node on that level. These values are normally based on the position of the neighbors in the previous level. Then the vertices are sorted according to these values and the sweep continues with the next level. The dependence on results from a previous step limits the available parallelism. In our group we came up with new ideas for modifying Sugiyama's method in order to introduce a higher amount of parallelism.

Partitioning-based Approaches.

First ideas of distributing work to processors were based on partitioning. One option is to cut the layers into horizontal slabs and assign each slab to one processing element. Here the border layers may experience conflicted orderings.

We had several ideas to resolve these conflicts, but it remains to be investigated in experiments which of these techniques are successful. The second possibility is the partitioning into vertical strips. The objective is to partition such that the number of edges running between different vertical slabs is small. Although graph partitioning is an NP-hard problem in general, there are efficient tools that yield good solutions. After that each vertical strip is processed by one processing element in parallel. The hope is that the number of bad orderings close to the strip boundaries is small due to the small number of edges between strips.

Odd-even Layer Processing.

The next idea was to reduce the sequential dependicies to exploit more parallelism. If the tasks to be performed have no sequential dependencies, any mapping to processing elements would allow for concurrent and conflict-free computations. In the traditional approach layers must be treated one after each other. Therefore, we tried to decouple the dependency of

adjacent layers. For this purpose we have developed the *odd-even barycentric algorithm* (OEBA).

Instead of up- or downward sweeps, the algorithm proceeds in passes. Similar to odd-even sort, in each pass one processes either the odd or the even layers. All odd layers are mutually independent, likewise all even layers. Hence, in each pass the currently visited layers can be processed independently in parallel.

Within each layer L_i, for each node $v \in L_i$ the new position is computed as the both-sided barycenter, i.e. taking predecessors and successors into account at the same time. Two alternatives have been considered, either taking the average over all neighbors or by taking the average of the barycenter of the neighbors in the layer above and of the barycenter of the neighbors in the below.

Next, the new positions of the nodes within the layer need to be determined. This is either done in the classical way by sorting the nodes according to their barycentric values directly after the barycentric computations of layer L_i have been finished. Alternatively, one computes the barycentric values of all layers first and then sorts all layers at once. These options allow for different types of parallelism and also alter the control flow and thus the output of the algorithm. Experiments will have to show which method yields the best performance.

The experiments we have conducted so far are not entirely conclusive yet, but certainly promising. For dense random graphs OEBA seems to yield crossing number results similar to the traditional barycentric method. However, for certain sparse graphs (Rome graphs) OEBA does reduce the crossings significantly, but not quite as much as the sequential approach.

Parallel Implementation Techniques.

We aim at shared-memory parallelism for multicore processors. Therefore, we anticipate two main parallel implementation techniques. One is OpenMP, a user-friendly standardized runtime system offered by most current compilers. The second one is a task-parallel job queue runtime system on top of the operating system's native threads. While the latter requires the implementor to break down the work into small independent pieces, it is expected to give the programmer more opportunities for low-level optimizations.

References

1 K. Sugiyama and S. Tagawa and M. Toda, Methods for Visual Understanding of Hierarchical System Structures, *IEEE Transactions on Systems, Man, and Cybernetics*, 11(2):109–125, 1981.

4.4 Circular-Arc Drawings with Right-Angle Crossings

Muhammad Jawaherul Alam, Martin Nöllenburg, Sankar Veeramoni, and Kevin Verbeek

We studied a problem arising from the combination of two recent topics in graph drawing: circular-arc or Lombardi drawings [3] and right-angle crossing (RAC) drawings [2]. More precisely, we looked for drawings, where each edge is drawn as a single circular arc and where the arc tangents at each crossing of two edges form right angles. Such a drawing is called a *circular-arc RAC* drawing or *CRAC* drawing.

During the seminar we obtained the some initial results for characterizing graphs that have CRAC drawings. The drawing cannot have four mutually crossing arcs such that all crossings form right angles, i.e., the crossing graph of the drawing cannot contain K_4 as a subgraph. Furthermore it seems that the upper bound of Arikushi et al. [1] that a graph admitting a 1-bend polyline RAC drawing has at most $6.5n - 13$ edges still holds for CRAC drawings.

We also looked at some specific examples of graphs and constructed CRAC drawings for $K_{3,5}, K_6$, and $K_{4,4}$; however we did not yet succeed in drawing K_7, which does have a 1-bend RAC drawing. We suspect that neither CRAC nor 1-bend RAC drawings are proper subclasses of each other.

We can use the same construction showing that every graph has a 3-bend RAC drawing [2] to show that every graph has a CRAC drawing where each edge is a differentiable poly-arc consisting of five circular arcs. What happens, if we reduce the complexity of the poly-arcs to three or four arcs? What is the relationship between the classes of k-bend RAC drawings and k-bend smooth CRAC drawings?

References

1 K. Arikushi, R. Fulek, B. Keszegh, F. Moric, and C. D. Tóth. Graphs that admit right angle crossing drawings. In D. M. Thilikos, editor, *Proc. 36th International Workshop on Graph Theoretic Concepts in Computer Science*, volume 6410, pages 135–146, 2010. http://dx.doi.org/10.1007/978-3-642-16926-7_14.

2 W. Didimo, P. Eades, and G. Liotta. Drawing graphs with right angle crossings. In F. Dehne, M. Gavrilova, J.-R. Sack, and C. D. Tóth, editors, *Proc. 11th International Symposium on Algorithms and Data Structures (WADS'09)*, volume 5664, pages 206–217, 2009. http://dx.doi.org/10.1007/978-3-642-03367-4_19.

3 C. A. Duncan, D. Eppstein, M. T. Goodrich, S. G. Kobourov, and M. Nöllenburg. Lombardi drawings of graphs. In U. Brandes and S. Cornelsen, editors, *Proc. 18th International Symposium on Graph Drawing (GD'10)*, volume 6502, pages 195–207, 2011. http://dx.doi.org/10.1007/978-3-642-18469-7_18.

4.5 Edge Bundling

David Auber, Stefan Diehl, Christian Duncan, Cesim Erten, Rudolf Fleischer, Emden Gansner, Michael Kaufmann, Lev Nachmanson, and Michael Wybrow

The background of the participants covered a wide range of fields information visualization, algorithm engineering, algorithms, graph theory. After having considered the existing practical approaches (e.g., [6, 4, 3, 1, 7, 5, 2]), the group tried to formulate a possibly unified mathematical model for the edge bundling problem. In several iterations, this model has been designed, checked on validity and reformulated. After that, the group including several subgroups considered the reformulation of the most important approaches within the proposed model. This led to interesting and stimulating insights about the differences of those approaches, and how to measure the differences.

A possible survey paper on the topic of this working group has been prospected and sketched.

References

1 Weiwei Cui, Hong Zhou, Huamin Qu, Pak Chung Wong, and Xiaoming Li. Geometry-based edge clustering for graph visualization. *IEEE Trans. Vis. Comput. Graph.*, 14(6):1277–1284, 2008.

2 E.R. Gansner, Yifan Hu, S. North, and C. Scheidegger. Multilevel agglomerative edge bundling for visualizing large graphs. In *Pacific Visualization Symposium (PacificVis), 2011 IEEE*, pages 187 –194, march 2011.

3 Danny Holten. Hierarchical edge bundles: Visualization of adjacency relations in hierarchical data. *IEEE Trans. Vis. Comput. Graph.*, 12(5):741–748, 2006.

4 Danny Holten and Jarke J. van Wijk. Force-directed edge bundling for graph visualization. *Comput. Graph. Forum*, 28(3):983–990, 2009.

5 Antoine Lambert, Romain Bourqui, and David Auber. Winding roads: Routing edges into bundles. *Comput. Graph. Forum*, 29(3):853–862, 2010.

6 Doantam Phan, Ling Xiao, Ron B. Yeh, Pat Hanrahan, and Terry Winograd. Flow map layout. In *INFOVIS*, page 29. IEEE Computer Society, 2005.

7 Sergey Pupyrev, Lev Nachmanson, and Michael Kaufmann. Improving layered graph layouts with edge bundling. In *Proceedings of the 18th international conference on Graph drawing*, GD'10, pages 329–340, Berlin, Heidelberg, 2011. Springer-Verlag.

4.6 On the Rotation Distance of Rooted Binary Trees

Ulrik Brandes, Rudolf Fleischer, Seok-Hee Hong, Tamara Mchedlidze, Ignaz Rutter, and Alexander Wolff

Given a triangulation of a regular n-gon with $n \geq 4$, a new triangulation can be obtained by *flipping* any internal edge. The triangulation resulting from flipping an edge e is obtained by first removing the edge e, and then inserting the diagonal of the resulting quadrilateral that is different from e; see Fig. 1.

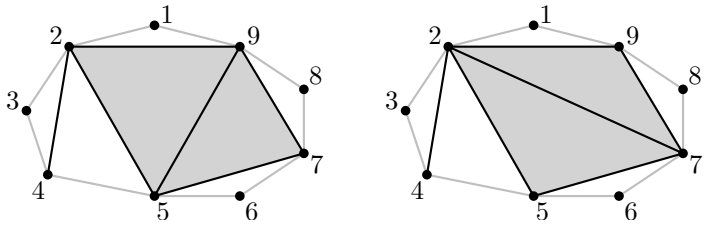

Figure 1 An example of a flip in a triangulation.

For any $n \geq 4$, there is a one-to-one correspondence between $(n-2)$-node ordered binary trees and triangulations of the regular (or any other convex) n-gon (with a designated "root" edge). The correspondence is such that rotations between pairs of trees are translated into flips between triangulations. This was observed by Sleator et al. [4]. Since flips are simpler to visualize and understand, we will use the flip-and-triangulation language in the following.

Let $\mathcal{P}_n$ be the regular n-gon, with vertices labeled $1, 2, \ldots, n$ in counterclockwise order. The *flip graph* $\mathcal{F}_n$ is the undirected graph whose vertices correspond to the triangulations

of $\mathcal{P}_n$; the graph has an edge between two triangulations if they can be obtained from each other by flipping a single edge.

The *flip distance* of two triangulations $\mathcal{T}_1$ and $\mathcal{T}_2$ of $\mathcal{P}_n$ is the minimum number of flips needed to transform $\mathcal{T}_1$ into $\mathcal{T}_2$. This is the length of a shortest path from $\mathcal{T}_1$ to $\mathcal{T}_2$ in $\mathcal{F}_n$. Sleator et al. [4] showed that the diameter of $\mathcal{F}_n$ is bounded by $2n-6$. Using involved arguments from hyperbolic geometry, they proved that this bound is tight for large values of n.

The complexity of the problem of computing the flip distance between two given triangulations of the same convex polygon (with a designated root edge) is unknown; there is neither an efficient algorithm nor an NP-hardness proof. Cleary and St. John [1] showed, however, that the problem is fixed-parameter tractable (FPT) with respect to the parameter flip distance. Their FPT algorithm runs in $O(n+4^{10k})$ time, where k is the flip distance between the given pair of triangulations of the regular n-gon. There are also three approximation algorithms. The first algorithm, by Li and Zhang [3], depends on the maximum number of diagonals, Δ, incident to a vertex in either of the given triangulations. Their algorithm has an approximation factor of $(2-2/(4(\Delta-1)(\Delta+6)+1))$ and runs in cubic time. Note that this approximation factor is always less than 2 and tends to 2 when Δ grows. The second, by the same authors [3], yields a 1.97-approximation if none of the triangulations contains an inner triangle. The third, by Cleary and St. John [2], yields a 2-approximation in linear time.

We developed a simpler and much faster FPT algorithm than that of Cleary and St. John. Our algorithm (see below) runs in $O(n+4^k/\sqrt{k})$ time.

We want to solve the following problem, which we call k-FlipDistance. Given a pair $\langle\mathcal{T}_1,\mathcal{T}_2\rangle$ of triangulations of $\mathcal{P}_n$ and a positive integer k, determine whether the flip distance of $\mathcal{T}_1$ and $\mathcal{T}_2$ is at most k and, if so, compute a shortest $\mathcal{T}_1$–$\mathcal{T}_2$ path in $\mathcal{F}_n$.

Sleator et al. [4] observed that, on a shortest path connecting $\mathcal{T}_1$ and $\mathcal{T}_2$ in $\mathcal{F}_n$, diagonals common to $\mathcal{T}_1$ and $\mathcal{T}_2$ are never flipped. Moreover, if $\mathcal{T}_1$ admits a flip that increases the number of diagonals common to $\mathcal{T}_1$ and $\mathcal{T}_2$, then there is a shortest $\mathcal{T}_1$–$\mathcal{T}_2$ path in $\mathcal{F}_n$ starting with that flip.

Let c be the number of diagonals common to $\mathcal{T}_1$ and $\mathcal{T}_2$, and let d be the number of diagonals that are in $\mathcal{T}_1$ but not in $\mathcal{T}_2$. Clearly, $c+d=n-3$. Since the problem k-FlipDistance allows us to flip at most k diagonals, we have $d\le k$ for all Yes-instances.

Our very simple algorithm is as follows. If $d>k$, we return "no". Otherwise, we split the problem into $c'\le c+1$ independent subproblems of sizes $d_1,\dots,d_{c'}$ with $\sum_{i=1}^{c'} d_i = d$. For subproblem $i=1,\dots,c'$, we do a breadth-first search in $\mathcal{F}_{d_i}$ starting with the appropriate part of $\mathcal{T}_1$. We stop as soon as we have reached the corresponding part of $\mathcal{T}_2$. Let ℓ_i be the length of the path between the two parts in $\mathcal{F}_{d_i}$. We return "no" if $\sum_{i=1}^{c'}\ell_i > k$, "yes" otherwise.

Obviously, the overall running time of our algorithm is bounded by the size of $\mathcal{F}_d$. It is well-known that the number of vertices of $\mathcal{F}_d$ equals the $(d-2)$-th Catalan number C_{d-2}. The number of edges of $\mathcal{F}_d$ is $(d-3)C_{d-2}/2$ since $\mathcal{P}_d$ has $d-3$ diagonals. Recall that $C_d=\binom{2d}{d}/(d+1)\approx 4^d/(d^{3/2}\sqrt{\pi})$. Hence, our algorithm runs in $O(n+4^d/\sqrt{d})$ time. We summarize our result with the following theorem.

▶ **Theorem 1.** *The problem k-FlipDistance can be decided in time $O(n+4^k/\sqrt{k})$, where n is the size of the polygon. For Yes-instances, a corresponding flip sequence can be computed within the same time bound.*

References

1 Sean Cleary and Katherine St. John. Rotation distance is fixed-parameter tractable. *Information Processing Letters*, 109:918–922, 2009.

2 Sean Cleary and Katherine St. John. A linear-time approximation for rotation distance. *J. Graph Algorithms Appl.*, 14(2):385–390, 2010.

3 Ming Li and Louxin Zhang. Better approximation of diagonal-flip transformation and rotation transformation. In Wen-Lian Hsu and Ming-Yang Kao, editors, *Proc. 4th Annu. Int. Conf. Comput. Combin. (COCOON'98)*, volume 1449, pages 85–94, 1998.

4 Daniel D. Sleator, Robert E. Tarjan, and William P. Thurston. Rotation distance, triangulations and hyperbolic geometry. *J. Amer. Math. Soc.*, 1(3):647–681, 1988.

Participants

- Deepak Ajwani
University College Cork, IE
- Muhammad Jawaherul Alam
Univ. of Arizona – Tucson, US
- David Auber
Université Bordeaux, FR
- Christian Bachmaier
Universität Passau, DE
- Melanie Badent
Universität Konstanz, DE
- Carla Binucci
University of Perugia, IT
- Franz J. Brandenburg
Universität Passau, DE
- Ulrik Brandes
Universität Konstanz, DE
- Markus Chimani
Universität Jena, DE
- Aparna Das
Univ. of Arizona – Tucson, US
- Camil Demetrescu
University of Rome "La Sapienza", IT
- Walter Didimo
University of Perugia, IT
- Stephan Diehl
Universität Trier, DE
- Christian Duncan
Louisiana Tech University, US
- Philip Effinger
Universität Tübingen, DE
- Cesim Erten
Kadir Has Univ. Istanbul, TR
- Rudolf Fleischer
Fudan University – Shanghai, CN
- Emden R. Gansner
AT&T Res. – Florham Park, US
- Carsten Gutwenger
TU Dortmund, DE
- Seok-Hee Hong
The University of Sydney, AU
- Riko Jacob
TU München, DE
- Jyrki Katajainen
University of Copenhagen, DK
- Michael Kaufmann
Universität Tübingen, DE
- Karsten Klein
TU Dortmund, DE
- Stephen Kobourov
Univ. of Arizona – Tucson, US
- Marcus Krug
KIT – Karlsruhe Institute of Technology, DE
- Robert Krug
Universität Tübingen, DE
- Giuseppe Liotta
University of Perugia, IT
- Tamara Mchedlidze
National TU – Athens, GR
- Kurt Mehlhorn
MPI für Informatik – Saarbrücken, DE
- Henning Meyerhenke
Georgia Inst. of Technology, US
- Petra Mutzel
TU Dortmund, DE
- Lev Nachmanson
Microsoft Corp. – Redmond, US
- Stefan Näher
Universität Trier, DE
- Martin Nöllenburg
KIT – Karlsruhe Institute of Technology, DE
- Maurizio Patrignani
Roma TRE University, IT
- Md. Saidur Rahman
Bangladesh University of Eng.& Technology, BD
- Ignaz Rutter
KIT – Karlsruhe Institute of Technology, DE
- Georg Sander
IBM – Frankfurt, DE
- Miro Spönemann
Universität Kiel, DE
- Joachim Spoerhase
Universität Würzburg, DE
- Matt Stallmann
North Carolina State Univ., US
- Sankar Veeramoni
Univ. of Arizona – Tucson, US
- Kevin Verbeek
TU Eindhoven, NL
- Stephen Wismath
University of Lethbridge, CA
- Alexander Wolff
Universität Würzburg, DE
- Michael Wybrow
Monash Univ. – Clayton, AU
- Katharina A. Zweig
Universität Heidelberg, DE

Report from Dagstuhl Seminar 11201

Constraint Programming meets Machine Learning and Data Mining

Edited by

Luc De Raedt[1], Siegfried Nijssen[2], Barry O'Sullivan[3], and Pascal Van Hentenryck[4]

1 **K.U. Leuven, BE,** `luc.deraedt@cs.kuleuven.be`
2 **K.U. Leuven, BE,** `siegfried.nijssen@cs.kuleuven.be`
3 **University College Cork, IE,** `b.osullivan@cs.ucc.ie`
4 **Brown University – Providence, US,** `pvh@cs.brown.edu`

Abstract

This report documents the programme and the outcomes of Dagstuhl Seminar 11201 *Constraint Programming meets Machine Learning and Data Mining*. Our starting point in this seminar was that machine learning and data mining have developed largely independently from constraint programming till now, but that it is increasingly becoming clear that there are many opportunities for interactions between these areas: on the one hand, data mining and machine learning can be used to improve constraint solving; on the other hand, constraint solving can be used in data mining in machine learning. This seminar brought together prominent researchers from both communities to discuss these opportunities.

Seminar 15.–20. May, 2011 – www.dagstuhl.de/11201
1998 ACM Subject Classification F.4.1. [Mathematical Logic: Logic and constraint programming]; H.2.8. [Database applications: Data mining]; I.2.6. [Learning]
Keywords and phrases Machine learning, data mining, constraint programming, constraints
Digital Object Identifier 10.4230/DagRep.1.5.61

1 Executive Summary

Luc De Raedt
Siegfried Nijssen
Barry O'Sullivan
Pascal Van Hentenryck

Goal of the Seminar

Over the past two decades the fields of constraint programming, machine learning and data mining have become well-established research fields within computer science. They have contributed many foundational techniques that are routinely applied in real-life scientific and industrial applications. At the same time, awareness has grown that constraints can be very useful during mining and learning, and also that machine learning and data mining may allow one to automatically acquire constraints from data.

Both the data mining and machine learning communities have been interested in *constraint-based mining and learning*, that is, the use of constraints to formalize mining and learning problems. Examples are the specification of desirable properties of patterns to be mined,

or clusters to be found. The task of the data mining or machine learning system is to generate all patterns or to compute the optimal clustering satisfying the constraints. A wide variety of constraints for local pattern mining, clustering and other machine learning problems exist and they have been implemented in an even wider range of specific data mining and machine learning systems for supporting such constraints. Some of these methods are based on mathematical programming techniques, such as linear programming or quadratic programming; other problems, however, cannot be modeled using these techniques. So far, the machine learning and data mining communities have been unable to develop general solvers that are applicable to a wide range of machine learning and data mining problems.

On the other hand, the artificial intelligence community has studied several types of constraint-satisfaction solvers. The most general systems are now gathered in the area of *constraint programming*. In constraint programming, the user specifies the model, that is, the set of constraints to be satisfied and constraint solvers generate solutions. Thus, the goals of constraint programming and constraint based mining and learning are similar; it is only that constraint programming targets *any* type of constraint satisfaction problem, whereas constraint-based mining and learning *specifically* targets data mining and machine learning applications. Therefore, it is surprising that despite the similarities between these two endeauvours, the two fields have evolved independently of one another, and also, that – with a few recent exceptions – constraint programming tools and techniques are not yet applied to data mining and machine learning, and, vice versa, that problems and challenges from data mining and machine learning have not yet been taken up by the constraint programming community. Exploring the possibilities for exploiting constraint programming in data mining and machine learning was one goal of this seminar.

The second goal was to study the use of machine learning and data mining in constraint programming. Practitioners of constraint programming have to formulate explicitly the constraints that underly their application. This is often a difficult task. Even when the right constraints are known, it can be challenging to formalize them in such a way that the constraint programming system can use them efficiently. This raises the question as to whether it is possible to (semi)- automatically learn such constraints or their formulations from data and experience. Again, some initial results in this direction exist, but we are away from a complete understanding of the potential of this approach.

In this seminar, we aimed at bridging the gap between these two fields by investigating, on the one hand, how standard constraint-programming techniques can be used in data mining and machine learning, and on the other hand, how machine learning and data mining can contribute to constraint programming. Therefore, this workshop brought together researchers in the areas of constraint programming, machine learning and data mining to discuss these issues, to identify interesting opportunities and challenges for research, and to consolidate and strengthen a promising line of research.

Seminar Organization

Given the various backgrounds of the participants, the seminar started with several introductory presentations. The focus was first on constraint programming and the use of machine learning and data mining in constraint programming; Helmut Simonis provided an introduction of constraint programming, while Barry O'Sullivan presented the possibilities for using machine learning and data mining in constraint programming.

Subsequently, the focused moved to data mining and machine learning and the possibilities

for using constraint solvers in data mining and machine learning. Bart Goethals provided an introduction to data mining; subsequently, Ian Davidson discussed uses of constraint solving in clustering, Dan Roth uses of solvers in inference problems, Siegfried Nijssen uses of solvers in pattern mining, and Tijl De Bie uses of solvers in statistical machine learning.

The remainder of the program consisted of a mix of technical presentations as well as meetings of discussion groups, each of which focused in more detail on the various possibilities for combining machine learning, data mining and constraint programming.

The topics of the discussion groups were determined after discussion at the first day of the seminar and were the following:

- declarative data mining, to discuss questions regarding the use of constraint programming solvers and declarative modelling to solve data mining problems;
- programming languages for machine learning, to discuss questions regarding the development of languages specifically for machine learning;
- applications, to discuss applications in which machine learning, data mining and constraint programming can be used;
- challenges, to discuss the possibilities for setting up benchmark problems –both real-life and artificial– that can be used to determine the success of combining machine learning, data mining and constraint programming;
- learning to solve, to discuss the use of machine learning and data mining to improve the efficiency and quality of constraint solving;
- learning constraints, to discuss the use of machine learning and data mining to learn appropriate models from examples.

The seminar was concluded with plenary presentations in which the results of the discussion groups were summarized.

The intention is to publish the results of the seminar in an edited book.

2 Table of Contents

3 Overview of Talks

3.1 Fast Algorithms for Finding Extremal Sets

Roberto Bayardo (Google Inc. – Mountain View, US)

Joint work of Bayardo, Roberto; Panda, Biswanath
Main reference Roberto J. Bayardo, Biswanath Panda, "Fast Algorithms for Finding Extremal Sets," Proc. of 11th SIAM International Conference on Data Mining (SDM'11), pp. 25–34.
URL http://siam.omnibooksonline.com/data/papers/089.pdf#page=1

Identifying the extremal (minimal and maximal) sets from a collection of sets is an important subproblem in the areas of data-mining and satisfiability checking. For example, extremal set finding algorithms are used in the context of mining maximal frequent itemsets, and for simplifying large propositional satisfiability instances derived from real world tasks such as circuit routing and verification. We describe two new algorithms for the task and detail their performance on real and synthetic data.

Each algorithm leverages an entirely different principle one primarily exploits set cardinality constraints, the other lexicographic constraints. Despite the inherent difficulty of this problem (the best known worst-case bounds are nearly quadratic), we show that both these algorithms provide excellent performance in practice, and can identify all extremal sets from multi-gigabyte itemset data using only a single processor core. Both algorithms are concise and can be implemented in no more than a few hundred lines of code.

3.2 Declarative Experimentation

Hendrik Blockeel (K.U. Leuven, BE)

Main reference Hendrik Blockeel, Joaquin Vanschoren, "Experiment Databases: Towards an Improved Experimental Methodology," in 11th European Conf. on Principles and Practice of Knowledge Discovery in Databases (PKDD'07), pp. 6–17, 2007
URL http://dx.doi.org/10.1007/978-3-540-74976-9_5

Data mining and machine learning are highly experimental scientific areas: many results are backed up with empirical data. With the current state of the art, to find the answer to certain experimental questions or testing experimental hypotheses, one needs to specify the experiments in a completely procedural way.

This makes the process of setting up experiments and interpreting the results very error-prone. We believe it should be possible to describe questions or hypotheses declaratively, and let the computer set up the right experiments. To make this possible, a formal language for experimental questions or hypotheses is required. We expect that elements from constraint programming will play a major role in such a language. In this talk we will discuss a number of issues that arise in languages for declarative experimentation, and point out the links with constraint-based data mining.

3.3 Discovering knowledge by combining primitives in the constraint programming framework

Bruno Cremilleux (Université de Caen, FR)

Joint work of Boizumault, Patrice; Crémilleux, Bruno; Khiari, Mehdi; Loudni, Samir; Métivier, Jean-Philippe
Main reference Patrice Boizumault, Bruno Crémilleux, Mehdi Khiari, Samir Loudni, Jean-Philippe Métivier, "Discovering Knowledge using a Constraint-based Language," Dagstuhl Preprint Archive, arXiv:1107.3407v1 [cs.LG]
URL http://arxiv.org/abs/1107.3407v1

Discovering pattern sets or global patterns is an attractive issue from the pattern mining community in order to provide useful information. By combining local patterns satisfying a joint meaning, this approach produces patterns of higher level and thus more useful for the end-user than the usual and well-known local patterns, while reducing the number of patterns. In parallel, recent works investigating relationships between data mining and constraint programming (CP) show that the CP paradigm is a nice framework to model and mine such patterns in a declarative and generic way. In this talk, we present a set of primitives which enables us to define queries addressing patterns sets and global patterns.

We provide several examples of such queries. The usefulness to propose a declarative approach is highlighted by the example of the clustering based on associations. We express the primitives in the CP framework by using numeric constraints and set constraints. Then a CP solver performs a sound and complete resolution of the queries. Finally, we discuss of the implementation of our approach.

References

1 L. De Raedt, T. Guns, and S. Nijssen. Constraint Programming for Itemset Mining. In *ACM SIGKDD Int. Conf. KDD'08*, Las Vegas, Nevada, USA, 2008.
2 L. De Raedt, T. Guns, and S. Nijssen. Constraint programming for data mining and machine learning. In *Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI-10)*, pages 1671–1675, 2010.
3 M. Khiari, P. Boizumault, and B. Crémilleux. Local constraint-based mining and set constraint programming for pattern discovery. In *From Local Patterns to Global Models (LeGo-09), ECML/PKDD-09 Workshop*, pages 61–76, Bled, Slovenia, 2009.
4 M. Khiari, P. Boizumault, and B. Crémilleux. Constraint programming for mining n-ary patterns. In *16th Int. Conf. on Principles and Practice of Constraint Programming (CP'10)*, volume 6308 of *LNCS*, pages 552–567. Springer, 2010.

3.4 Bayesian network learning with cutting planes

James Cussens (University of York, GB)

Main reference James Cussens, "Bayesian network learning with cutting planes," Proc. 27th Conference on Uncertainty in Artificial Intelligence (UAI 2011).
URL http://www-users.cs.york.ac.uk/~jc/research/uai11

The problem of learning the structure of Bayesian networks from complete discrete data with a limit on parent set size is considered. Learning is cast explicitly as an optimisation problem where the goal is to find a BN structure which maximises log marginal likelihood (BDe score). Integer programming, specifically the SCIP framework, is used to solve this

optimisation problem. Acyclicity constraints are added to the integer program (IP) during solving in the form of *cutting planes*. Finding good cutting planes is the key to the success of the approach—the search for such cutting planes is effected using a sub-IP. Results show that this is a particularly fast method for exact BN learning.

3.5 Continuous Optimization – Problems and successes

Tijl De Bie (University of Bristol, GB)

URL http://www.tijldebie.net

Continuous optimization (or mathematical programming), and convex optimization in particular [1], has had a profound impact on certain areas of machine learning research in the past decade [2, 3]. This talk is intended to provide an overview of the reasons for these successes, as well as some of the limitations.

The focus of the talk is on similarities with constraint programming as a largely declarative way of formulating search problems through large parameter spaces.

References

1 Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press (2004).
2 Kristin Bennet and Emilio Parrado-Hernandez. The interplay of optimization and machine learning research. *Journal of Machine Learning Research* 7 (2006).
3 Tijl De Bie. Deploying SDP for machine learning. *Proceedings of the 15th European Symposium on Artificial Neural Networks (ESANN 2007)*, Bruges, April 2007.

3.6 Reinforced Adaptive Large Neighborhood Search

Yves Deville (Université Catholique de Louvain, BE)

Joint work of Mairy, Jean-Baptiste; Deville, Yves; Van Hentenryck, Pascal

Large Neighborhood Search (LNS) is a hybrid search paradigm that combines local search (LS) and constraint programming (CP). At each step of LNS, a subset of variables of the current solution are relaxed. The potential solutions of the relaxed problems form the neighborhood. CP is used to solve the relaxed problem and tries to find a neighbor improving the current solution.

The CP part has a search limit. If no improving solution is found, the current solution is unchanged. In LNS, one has to fix the size of the relaxed fragment and the search limit. These two parameters are crucial in LNS. One has also to decide how the relaxed variables are chosen. Usually, a random selection is performed.

In this talk, we report preliminary results on applying reinforcement learning in LNS. Reinforcement learning is also used to guide the selection of the variables to relax at each LNS step.

References

1 Laurent Perron, Paul Shaw and Vincent Furnon. *Propagation Guided Large Neighborhood Search*. CP 2004, LNCS 3258, (2004) 468–481.
2 Laurent Perron and Paul Shaw. *Combining Forces to Solve the Car Sequencing Problem*. CPAIOR 2004, LNCS 3011, (2004) 225–239.
3 Richard S. Sutton, Andrew G. Barto. *Reinforcement learning: an introduction*. Cambridge, MA: MIT Press, 1998.

3.7 Two combinatorial optimization problems in Machine Learning

Pierre Dupont (UC Louvain-la-Neuve, BE)

In this talk we present two classical ML problems which can be formulated as combinatorial optimization problems. The open question is to which extent constraint programming could help to better address them.

Regular grammar induction aims at learning a formal regular language from sets of positive and negative strings. Learning a regular grammar is classically restated as the minimal DFA consistency problem: finding a minimal deterministic finite state machine accepting all positive strings and rejecting all negative ones.

We briefly describe the search space of this problem, present standard state-merging techniques and sketch the key properties of the winning algorithm of the recent Stamina competition.

Many supervised classification problems are characterized by a very large number (p) of features, but only a limited number (n) of samples. Transcriptome analysis from DNA microarray data is a typical example.

In such a small n big p setting, the risk of overfitting a classification model to the training sample is particularly present.

Feature selection aims at reducing, sometimes drastically, the number of actual features on which the classifier is effectively built.

The simplest approach selects top-ranked features according to a relevance index, such as the mutual information between each input feature and the response variable.

Such an approach is particularly efficient from a computational viewpoint but it is univariate and hence lacks to model the dependence between various features. A more advanced alternative aims at selecting a subset of maximally relevant but minimally redundant features.

Finding the optimal feature subset is computationally hard and typical approximations relies on a greedy search with no global optimum guarantee.

References

1 D. Angluin, *On the complexity of minimum inference of regular sets*, Information and Control **39** (1978), 337–350.
2 F. Coste and J. Nicolas, *How considering incompatible state mergings may reduce the DFA induction search tree*, Grammatical Inference, ICGI'98 (Ames, Iowa), Lecture Notes in Artificial Intelligence, no. 1433, Springer Verlag, 1998, pp. 199–210.
3 P. Dupont, B. Lambeau, C. Damas, and A. van Lamsweerde, *The QSM algorithm and its application to software behavior model induction*, Applied Artificial Intelligence **22** (2008), 77–115.

4 P. Dupont, L. Miclet, and E. Vidal, *What is the search space of the regular inference ?*, Grammatical Inference and Applications, ICGI'94 (Alicante, Spain), Lecture Notes in Artificial Intelligence, no. 862, Springer Verlag, 1994, pp. 25–37.

5 E.M. Gold, *Complexity of automaton identification from given data*, Information and Control **37** (1978), 302–320.

6 I. Guyon, S. Gunn, M. Nikravesh, and L. Zadeh (eds.), *Feature extraction, foundations and applications*, Series Studies in Fuzziness and Soft Computing, vol. 207, Springer, 2006.

7 T. Helleputte and P. Dupont, *Feature selection by transfer learning with linear regularized models*, European Conference on Machine Learning, Lecture Notes in Artificial Intelligence, no. 5781, 2009, pp. 533–547.

8 T. Helleputte and P. Dupont, *Partially supervised feature selection with regularized linear models*, International Conference on Machine Learning, 2009.

9 M. Heule and S. Verwer, *Exact DFA identification using SAT solvers*, International Colloquium on Grammatical Inference, Lecture Notes in Artificial Intelligence, vol. 6339, Springer Verlag, 2010, pp. 66–79.

10 B. Lambeau, C. Damas, and P. Dupont, *State-merging DFA induction algorithms with mandatory merge constraints*, Lecture Notes in Artificial Intelligence, vol. 5278, 2008, pp. 139–153.

11 G. Obozinski, B. Taskar, and M.I. Jordan, *Joint covariate selection and joint subspace selection for multiple classification problems*, Statistics and Computing (2009), 1–22.

12 H. Peng, F. Long, and C. Ding, *Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy*, IEEE Transactions on Pattern Analysis and Machine Intelligence **27** (2005), no. 8, 1226–1238.

13 Y. Saeys, I. Inza, and P. Larranaga, *A review of feature selection techniques in bioinformatics*, Bioinformatics **23** (2007), no. 19, 2507–2517.

14 N. Walkinshaw, K. Bogdanov, C. Damas, B. Lambeau, and P. Dupont, *A framework for the competitive evaluation of model inference techniques*, 1st International workshop on Model Inference In Testing, 2010.

3.8 Data Uncertainty and Constraint Programming Models

Carmen Gervet (German University in Cairo, EG)

Joint work of Gervet, Carmen; Saad, Aya; Abdennadher, Slim

Main reference Aya Saad, Carmen Gervet, Slim Abdennadher, "Constraint Reasoning with Uncertain Data Using CDF-Intervals," Proc. 7th Int'l Conf. on Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems (CPAIOR'10), pp. 292–306, 2010.

URL http://dx.doi.org/10.1007/978-3-642-13520-0_32

The classical CP framework has been extended in the past 10 years to address ill-defined, uncertain real-world optimisation problems. Extensions include probabilistic models, quantified models, interval data models. Each tackles different aspects of data uncertainty, due to data errors, incompleteness, forecasting and aims at ensuring that the problem is faithfully represented with what is known for sure about the data. The solution sought vary from robust solutions, to reliable solutions and covering set solutions. In all cases, the model does have a great impact on the algorithm complexity, and new approaches have to be implemented.

We summarize existing approaches and introduce a novel approach based on a new definition of the confidence interval .

3.9 Sequence Classification in High Dimensional Predictor Space

Georgiana Ifrim (University College Cork, IE)

Joint work of Ifrim, Georgiana; Wiuf, Carsten;
Main reference G. Ifrim, C. Wiuf, "Bounded Coordinate-Descent for Biological Sequence Classification in High Dimensional Predictor Space," Proc. 17th ACM SIGKDD Int'l Conf. on Knowledge Discovery and Data Mining (SIGKDD'11), pp. 708–716, 2011.
URL http://doi.acm.org/10.1145/2020408.2020519

In this talk I present a framework for sequence classification where we build linear classifiers on a rich but very high dimensional feature space. I present an optimisation algorithm based on coordinate gradient-descent coupled with branch-and-bound to efficiently build such classifiers.

Furthermore, I discuss accuracy, interpretability and scalability of this learning framework. This technique can be applied to any classification problem where the input data can be represented in symbolic form as a sequence of tokens (e.g., text, biological sequences, time series).

References

1 G. Ifrim. Gradient-bounded coordinate-descent for sequence regression. http://www.4c.ucc.ie/ ifrim/seql/data, 2010.
2 G. Ifrim, C. Wiuf. Bounded Coordinate-Descent for Biological Sequence Classification in High Dimensional Predictor Space In *KDD*, USA, 2011.
3 G. Ifrim, G. Bakir, and G. Weikum. Fast logistic regression for text categorization with variable-length n-grams. In *KDD*, pages 354–362, USA, 2008.
4 T. Kudo, E. Maeda and Y. Matsumoto. An Application of Boosting to Graph Classification. In *NIPS*, 2004.

3.10 Using and Learning Constraints in Inductive Process Modeling

Pat Langley (ASU – Tempe, US)

URL http://www.isle.org/process/

Most research on computational knowledge discovery has focused on descriptive models that only summarize data and utilized formalisms developed in AI or statistics. In contrast, scientists typically aim to develop explanatory models that make contact with background knowledge and use established scientific notations. In this talk, I present an approach to computational discovery that encodes scientific models as sets of processes that incorporate differential equations, simulates these models' behavior over time, uses background knowledge to guide construction of model structures, and fits the model structures to time-series data. I illustrate this framework on data and models from a number of fields but focusing on aquatic ecosystems. One important form of background knowledge – constraints – is used to rule out implausible model structures. I also report an extension that learns new constraints based on the success and failure of candidate models. In closing, I discuss intellectual influences on the work and directions for future research.

This talk describes joint work at Stanford University and ISLE with Kevin Arrigo, Stuart Borrett, Matt Bravo, Will Bridewell, and Ljupco Todorovski under funding from the National

Science Foundation. The URL http://www.isle.org/process/ includes a list of publications on inductive process modeling.

References

1 Chunki Park, Will Bridewell, Pat Langley: *Integrated Systems for Inducing Spatio-Temporal Process Models.* AAAI 2010
2 Will Bridewell, Ljupco Todorovski: *The Induction and Transfer of Declarative Bias.* AAAI 2010
3 Will Bridewell, Pat Langley, Ljupco Todorovski, Saso Dzeroski: *Inductive process modeling.* Machine Learning 71(1): 1-32 (2008)

3.11 Search and Black-Box Constraint-Programming Solvers

Laurent Michel (University of Connecticut – Storrs, US)

Joint work of Michel, Laurent; Van Hentenryck, Pascal
Main reference P. Van Hentenryck, L. Michel. *Constraint-Based Local Search.* MIT Press, Cambridge, London, 2009.
URL http://mitpress.mit.edu/catalog/item/default.asp?ttype=2&tid=11963

Constraint Programming Tools traditionally rely on the combination of a modeling and a search component to solve challenging combinatorial optimization problems. Advances in modeling are often aimed at delivering stronger propagation and filtering algorithms while progress on the search components focus on declarative control abstractions that deliver clear, modular and reusable search procedures.

The purpose of this talk is to review the state of the art and recent advances in search. In particular, it illustrates how control abstractions like non- deterministic tryall, selectors, explorers, controllers and labeling heuristics found in the COMET language are instrumental in delivering simple, elegant and fully reusable search procedures.

The talk reviews how such abstractions can be used to build classic generic search procedures found in black-box constraint-programming solvers. It discusses in details a recent addition: activity-based search, the idea of using the activity of variables during propagation to guide the search. ABS is compared experimentally to impact-based search and the WDeg heuristics on a variety of benchmarks to illustrate its robustness and performance.

3.12 MIME: Discovering Interesting Patterns Visually and Interactively

Sandy Moens (University of Antwerp, BE)

Joint work of Goethals, Bart; Moens, Sandy; Vreeken, Jilles
Main reference Goethals, B., Moens, S., and Vreeken, J., "MIME: A Framework for Interactive Visual Pattern Mining," In Proc. 17th ACM SIGKDD Int'l Conf. on Knowledge Discovery and Data Mining (SIGKDD'11), pp. 757–760, 2011
URL http://doi.acm.org/10.1145/2020408.2020529

Pattern mining is concerned with discovering interesting patterns from data. Formalizing interestingness, however, is notoriously difficult, as it is inherently subjective. We propose to discover subjectively interesting patterns. In our system, users can visually browse

through data and patterns. To assist users in identifying interesting patterns, a toolbox of interestingness measures and algorithms for mining and post-processing are provided. All statistics can be combined in interconnected views, and results of actions in one view can be directly used in others. All information is computed on-the-fly, where priority is given to what is currently shown to the user, while off-screen results are calculated in the background.

By making the process interactive, we enable the user to combine their subjective notion of interestingness, and background knowledge, with a wide variety of objective measures in order to easily mine the most interesting patterns. Basically, we enable the user to become an essential part of the mining algorithm, by allowing the construction of patterns, deciding what algorithms to run and how to analyze what patterns in detail.

References

1 Goethals, B., Moens, S. and Vreeken, J. *MIME: A Framework for Interactive Visual Pattern Mining*. In Proc. KDD, ACM, 2011.
2 Goethals, B., Moens, S. and Vreeken, J. *MIME: A Framework for Interactive Visual Pattern Mining*. In Proc. ECML PKDD, Springer, 2011.

3.13 Constraint programming for Itemset Mining

Siegfried Nijssen (K.U. Leuven, BE)

Joint work of Nijssen, Siegfried; Guns, Tias; De Raedt, Luc
Main reference Tias Guns, Siegfried Nijssen, Luc De Raedt, "Itemset mining: A constraint programming perspective," *Artif. Intell.*, 175(12–13): 1951–1983 (2011)
URL http://dx.doi.org/10.1016/j.artint.2011.05.002

The field of data mining has become accustomed to specifying constraints on patterns of interest. A large number of systems and techniques has been developed for solving such constraint-based mining problems, especially for mining itemsets. The approach taken in the field of data mining contrasts with the constraint programming principles developed within the artificial intelligence community. While most data mining research focuses on algorithmic issues and aims at developing highly optimized and scalable implementations that are tailored towards specific tasks, constraint programming employs a more declarative approach. The emphasis lies on developing high-level modeling languages and general solvers that specify what the problem is, rather than outlining how a solution should be computed, yet are powerful enough to be used across a wide variety of applications and application domains.

Here we present a declarative constraint programming approach to data mining.

More specifically, we show that it is possible to employ off-the-shelf constraint programming techniques for modeling and solving a wide variety of constraint-based itemset mining tasks, such as frequent, closed, discriminative, and cost-based itemset mining. In particular, we develop a basic constraint programming model for specifying frequent itemsets and show that this model can easily be extended to realize the other settings. This contrasts with typical procedural data mining systems where the underlying procedures need to be modified in order to accommodate new types of constraint, or novel combinations thereof. Even though the performance of state-of-the-art data mining systems outperforms that of the out-of-the-box constraint programming approach on some standard tasks, we show that by developing specialized CP systems, it is possible to overcome most of the differences in efficiency. Furthermore, we show that there exist problems where the out-of-the-box

constraint programming approach already leads to significant performance improvements over state-of-the-art methods in data mining and leads to new insights into the underlying data mining problems. Many such insights can be obtained by relating the underlying search algorithms of data mining and constraint programming systems to one another. We provide an illustration of the approach on a problem in bioinformatics and finally discuss a number of interesting new research questions and challenges raised by the declarative constraint programming approach to data mining.

References

1 Luc De Raedt, Tias Guns, Siegfried Nijssen: Constraint programming for itemset mining. KDD 2008: 204-212
2 Tias Guns, Hong Sun, Kathleen Marchal, Siegfried Nijssen: Cis-regulatory module detection using constraint programming. BIBM 2010: 363–368
3 Tias Guns, Siegfried Nijssen, Luc De Raedt: Itemset mining: A constraint programming perspective. Artif. Intell. 175(12–13): 1951–1983 (2011)
4 Siegfried Nijssen, Tias Guns: Integrating Constraint Programming and Itemset Mining. ECML/PKDD (2) 2010: 467–482
5 Siegfried Nijssen, Tias Guns, Luc De Raedt: Correlated itemset mining in ROC space: a constraint programming approach. KDD 2009: 647–656

3.14 Integrations of Machine Learning and Data Mining in Constraint Satisfaction

Barry O'Sullivan (University College Cork, IE)

Main reference Barry O'Sullivan, "Automated Modelling and Solving in Constraint Programming," Proc. 24th AAAI Conf. on Artificial Intelligence (AAAI'10), pp. 1493–1497, 2010.
URL http://www.aaai.org/ocs/index.php/AAAI/AAAI10/paper/view/1899

Constraint programming can be divided very crudely into modeling and solving. Modeling defines the problem, in terms of variables that can take on different values, subject to restrictions (constraints) on which combinations of variables are allowed. Solving finds values for all the variables that simultaneously satisfy all the constraints. However, the impact of constraint programming has been constrained by a lack of "user-friendliness". Constraint programming has a major "declarative" aspect, in that a problem model can be handed off for solution to a variety of standard solving methods. These methods are embedded in algorithms, libraries, or specialized constraint programming languages. To fully exploit this declarative opportunity however, we must provide more assistance and automation in the modeling process, as well as in the design of application-specific problem solvers. Automated modelling and solving in constraint programming presents a major challenge for the artificial intelligence community. Artificial intelligence, and in particular machine learning, is a natural field in which to explore opportunities for moving more of the burden of constraint programming from the user to the machine. This talk presents technical challenges in the areas of constraint model acquisition, formulation and reformulation, synthesis of filtering algorithms for global constraints, and automated solving. We also present the metrics by which success and progress can be measured.

References

1 Barry O'Sullivan: Automated Modelling and Solving in Constraint Programming. AAAI 2010
2 Hadrien Cambazard, Tarik Hadzic, Barry O'Sullivan: Knowledge Compilation for Itemset Mining. ECAI 2010: 1109-1110
3 Christian Bessière, Emmanuel Hebrard, Barry O'Sullivan: Minimising Decision Tree Size as Combinatorial Optimisation. CP 2009: 173-187
4 Xuan-Ha Vu, Barry O'Sullivan: A Unifying Framework for Generalized Constraint Acquisition. International Journal on Artificial Intelligence Tools 17(5): 803-833 (2008)
5 Christian Bessière, Remi Coletta, Frédéric Koriche, Barry O'Sullivan: A SAT-Based Version Space Algorithm for Acquiring Constraint Satisfaction Problems. ECML 2005: 23-34

3.15 Constrained Conditional Models: Integer Linear Programming Formulations for Natural Language Understanding

Dan Roth (University of Illinois at Urbana-Champaign, USA)

URL www.aaai.org/Papers/AAAI/2008/AAAI08-252
Main reference M. Chang, L. Ratinov, N. Rizzolo and D. Roth, "Learning and Inference with Constraints," Proc. of the National Conference on Artificial Intelligence (AAAI), 2008.

Intelligent Information Access and Extraction suggest significant challenges for Natural Language analysis. Tasks of interest include semantic role labeling (determining who did what to whom, when and where), information extraction (identifying events, entities and relations), transliteration of names, and textual entailment (determining whether one utterance is a likely consequence of another). A computational approached to these challenges often involves assigning values to sets of interdependent variables and thus frequently necessitate performing global inference that accounts for these interdependencies. This talk presented research on Constrained Conditional Models (CCMs), a framework that augments probabilistic models with declarative constraints as a way to support such decisions. We presented a framework we introduced a few years ago, formulating decision problems in NLP as Integer Linear Programming problems, but focused on new algorithms for training these global models using indirect supervision signals. Learning models for structured tasks is difficult partly since generating supervision signals is costly. We showed that it is often easy to obtain a related indirect supervision signal, and discussed several options for deriving this supervision signal, including inducing it from the world's response to the model's actions. Our learning framework is "Constraints Driven" in the sense that it allows and even gains from global inference that combines statistical models with expressive declarative knowledge (encoded as constraints), modeled via Constrained Conditional Models. Experimental results showed the significant contribution of easy-to-get indirect supervision on several NLP tasks including information extraction, Transliteration and Textual Entailment.
See also the tutorial from NAACL'10:
http://l2r.cs.uiuc.edu/%7Edanr/Talks/ILP-CCM-Tutorial-NAACL10.ppt.

3.16 Gecode – an open constraint solving library

Christian Schulte (KTH – Kista, SE)

Joint work of Gecode team
URL http://www.gecode.org

Gecode is a widely used toolkit for developing constraint-based systems and applications. Gecode provides a constraint solver with state-of-the-art performance while being modular and extensible. Gecode is: open (documented interfaces support tasks from modeling to implementing new variables, constraints, search engines, ...), free (MIT license), portable (standard C++), accessible (extensive tutorial and reference documentation, efficient (excellent performance, has won the 2008-2010 MiniZinc challenges in all categories), and parallel (exploits today's multi- core hardware).

The talk provides an overview of what Gecode is, what one can do with it, and what are the basic ideas behind its architecture.

3.17 Lifelong learning in Constraint Solving

Michele Sebag (LIMSI – CNRS, FR)

Joint work of Arbelaez, Alejandro; Hamadi, Youssef; Sebag, Michele
Main reference Alejandro Arbelaez, Youssef Hamadi, Michèle Sebag, "Continuous Search in Constraint Programming," Proc. 22nd IEEE Int'l Conf. on Tools with Artificial Intelligence (ICTAI'10), pp. 53–60, 2010.
URL http://dx.doi.org/10.1109/ICTAI.2010.17

The goal is to allow any user to eventually get top performance from his constraint solver. This goal is reached by customizing the solver to the user's problem instance distribution, using the computer idle time to launch exploratory experiments on the user's problem, observing the results and learning the most appropriate tunings of the heuristic portfolio.

This approach can be viewed as an instance of Meta-Learning problem, with the difference that many descriptive features have been proposed in the CP literature to characterize the problem to solve and the current search state [4]. The novelty compared to the state of the art (e.g.

CPHydra, [5]) is that the computer does not require a set of problems, representative of the user's activity, to be available beforehand.

Instead, the computer uses an exploration/exploitation approach (lifelong learning) to gradually become an expert into the user's problem distribution.

Experimental validation suggests that Continuous Search can design efficient mixed strategies after considering a moderate number of problem instances.

References

1 Alejandro Arbelaez. *Learning During Search*. PhD thesis, Université Paris-Sud, France, 2011.
2 Alejandro Arbelaez, Youssef Hamadi, Michèle Sebag: *Continuous Search in Constraint Programming*. ICTAI (1) 2010: 53–60.

3 Alejandro Arbelaez and Youssef Hamadi and Michele Sebag: *Autonomous Search*. In: Youssef Hamadi and Eric Monfroy and Frederic Saubion (editors): Continuous Search in Constraint Programming, Springer, 2011.
4 Frank Hutter, Holger H. Hoos, Thomas Stützle: *Automatic Algorithm Configuration Based on Local Search*. AAAI 2007: 1152–1157.
5 Eoin O'Mahony, Emmanuel Hebrard, Alan Holland, Conor Nugent and Barry O'Sullivan: *Using Case-based Reasoning in an Algorithm Portfolio for Constraint Solving*. Proceedings of the 19th Irish Conference on Artificial Intelligence and Cognitive Science, 2008.

3.18 A Constraint Seeker: Finding and Ranking Global Constraints from Examples

Helmut Simonis (University College Cork, IE)

Joint work of Beldiceanu, Nicolas; Simonis, Helmut
Main reference N. Beldiceanu and H. Simonis, "A Constraint Seeker: Finding and Ranking Global Constraints from Examples," In Principles and Practice of Constraint Programming (CP2011), Perugia, Italy, September 2011.

The global constraint catalog provides a valuable repository of global constraint information for both researchers in the constraint field and application developers searching for the right modelling abstraction. The catalog currently describes over 350 constraints on more than 2800 pages. This wealth of information is also its main weakness.

For a novice (and even for an experienced) user it can be quite challenging to find a particular constraint, unless the name in the catalog is known. As a consistent naming scheme (like the naming scheme for chemical compounds http://en.wikipedia.org/wiki/IUPAC_nomenclature, for example) does not exist in the constraint field, and different constraint systems often use incompatible names and argument orders for their constraints, there is no easy way to deduce the name of a constraint and the way its arguments are organized from its properties. The catalog already provides search by name, or by keywords, and provides extensive cross-references between constraints, as well as a classification by the required argument type. All of these access methods can be helpful, but are of limited use if one does not know too much about the constraint one is looking for.

The Constraint Seeker (http://seeker.mines-nantes.fr) provides another way of finding constraint candidates, sorted by potential relevance, in the catalog.

The user provides positive and/or negative ground instances for a constraint, and the tool provides a ranked list of possible candidates which match the given examples.

Besides introducing the Constraint Seeker as an application, we want to illustrate the power of meta data and meta programming in the context of future constraints platforms. We use meta data for explicitly describing different aspects of global constraints such as argument types, restrictions on the arguments, typical use of a constraint, symmetries w.r.t. the arguments of a constraint, links between constraints (e.g. implication, generalisation). The electronic version of the global constraint catalog provides such information in a systematic way for the different global constraints. The Constraint Seeker illustrates the fact that the same meta data can be used for different purposes, unlike ad-hoc code in a procedural language which is designed for a specific (and unique) usage and a single system.

3.19 APOPCALEAPS: Automatic Pop Composer And Learner of Parameters

Jon Sneyers (K.U. Leuven, BE)

Joint work of Sneyers, Jon; De Schreye, Danny
Main reference Jon Sneyers and Danny De Schreye, "APOPCALEAPS: Automatic Music Generation with CHRiSM," 22nd Benelux Conf. on Artificial Intelligence (BNAIC'10), Luxembourg, October 2010.
URL http://people.cs.kuleuven.be/ jon.sneyers/papers/chrism_music.pdf

APOPCALEAPS is a system for automatic pop music generation and learning, implemented in CHRiSM [1]. CHRiSM is a new programming language, combining features of the existing languages CHR [2] and PRISM [3]. It is a high-level rule-based formalism for probabilistic-logic learning, allowing a powerful mix of domain-specific constraints and probabilistic inference. The current implementation of CHRiSM is based on the K.U.Leuven CHR system in B-Prolog. The user interface of APOPCALEAPS was made using fltk; its output is rendered using LilyPond.

The goal of the APOPCALEAPS project is to create a personal automatic music generator. The user creates music, selects the examples that are good according to his or her taste, and then these selected examples are used as a training set for learning. After several iterations, the result is a music generator that produces personalized music.

References

1 Jon Sneyers, Wannes Meert, Joost Vennekens, Yoshitaka Kameya and Taisuke Sato. *CHR(PRISM)-based Probabilistic Logic Learning.* 26th International Conference on Logic Programming (ICLP'10), Edinburgh, UK, July 2010.
2 Jon Sneyers, Peter Van Weert, Tom Schrijvers, and Leslie De Koninck. *As Time Goes By: Constraint Handling Rules – A Survey of CHR Research between 1998 and 2007.* Theory and Practice of Logic Programming, 10(1):1-47, 2010.
3 Taisuke Sato. *A glimpse of symbolic-statistical modeling by PRISM.* Journal of Intelligent Information Systems, Vol.31, No.2, pp.161-176, 2008.

3.20 MiniZinc – Towards a standard CP modelling language

Guido Tack (K.U. Leuven, BE)

Joint work of Nethercote, Nicholas; Stuckey, Peter J.; Becket, Ralph; Brand, Sebastian; Duck, Gregory J.; Tack, Guido
Main reference N. Nethercote, P. J. Stuckey, R. Becket, S. Brand, G. J. Duck, and G. Tack, "MiniZinc: Towards a standard CP modelling language," . Proc. 13th Int'l Conf. on Principles and Practice of Constraint Programming (CP'07), pp. 529–543, 2007.
URL http://dx.doi.org/10.1007/978-3-540-74970-7_38

MiniZinc arose as a response to the extended discussion at CP'06 of the need for a standard modelling language for CP. Since 2006, MiniZinc and its ecosystem have grown and matured considerably. Today, in addition to several CP solvers that handle MiniZinc (through the FlatZinc solver interface language), there is support for SMT, SAT, and MIP solvers as well. This makes MiniZinc an ideal candidate language for experimentation and model exchange.

The talk consists of an introduction to the MiniZinc language using example models, a quick overview of the compilation and execution approach, and the current plans for MiniZinc's future.

3.21 On Learning Constraint Problems

Christel Vrain (Université d'Orleans, FR)

Joint work of Lallouet, Arnaud; Lopez, Matthieu; Martin, Lionel; Vrain, Christel
Main reference Arnaud Lallouet, Matthieu Lopez, Lionel Martin, Christel Vrain, "On Learning Constraint Problems," Proc. 22nd IEEE Int'l Conf. on Tools with Artificial Intelligence (ICTAI'10), pp. 45–52, 2010.
URL http://dx.doi.org/10.1109/ICTAI.2010.16

We address the problem of learning constraint models expressed in a higher level language. For learning it, we use Inductive Logic Programming and we develop a new search strategy to tackle the problem.

We start from the observation that some novice users have difficulties to model with constraints and it would be desirable to provide them tools that help them to come with a constraint model by using only the informal knowledge they have on their problem.

Often users can provide examples and counter-examples for the problem to solve or answer questions about it. But often the problems they tackled in the past were not exactly the same but only similar. For example, it may happen that the user has solved the 4-queens and 5-queens problem but wants to solve n-queens.

How to reuse this knowledge to provide a general model?

In this work, we learn a high level model in a first-order logic language retaining some of the features of middle-level constraint modeling languages like Minizinc. The user provides some historical data and the model is learned by ILP techniques. We found that existing ILP learners were subject to blind search and we had to provide a new bidirectional learning algorithm to find solutions.

Then the model is transformed into a CSP through a rewriting phase which also takes as input the actual variables of the new problem to solve. Experiments have been conducted on graph coloring, timetable and jobshop scheduling and n-queens.

References

1 Arnaud Lallouet, Matthieu Lopez, Lionel Martin, Christel Vrain: On Learning Constraint Problems. ICTAI (1) 2010: 45-52
2 Matthieu Lopez, Lionel Martin, Christel Vrain: Learning Discriminant Rules as a Minimal Saturation Search. ILP 2010: 146-157

4 Working Groups

4.1 Declarative Data Mining

Siegfried Nijssen (K.U. Leuven, BE)

The aim of this working group was to consider how declarative approaches to problem solving, such as constraint programming, may be used to improve the principles and practice of data mining. The discussion started by compiling a list of questions:

- What is a declarative approach to data mining?
- What is currently missing in declarative solvers if we wish to use them in data mining?

- Which data mining problems are good candidates for a declarative approach to data mining?
- Is there a catalog of data mining problems that could give constraint programmers a better idea what the challenges are?
- What are expected benefits of a declarative approach data mining approach?

Subsequently, these questions were addressed in the further discussion.

Improved flexibility and easiness were identified as main reasons for declarative data mining. Declarative data mining should allow to implement new data mining tasks with only a small number of lines of code, hence making the implementation of new data mining tasks easier. Furthermore, declarative data mining system should ideally be compositional. This would make it possible to combine data mining tasks in novel ways, making data mining more flexible.

A declarative data mining language would be a language that allows to express *what* a data mining task is, and would decouple this from *how* the data mining task is solved. Hence, it was argued that such a language could be solver independent, where the Zinc language was mentioned as an example from the constraint programming community. A solver would ideally process a model in the language automatically, but the model may also be used a starting point for manual tuning.

A topic that attracted extensive discussion was how easy would be to make such an approach sufficiently efficient. It was argued that declarative approaches have been successful in the database community as any statement written down in a language such as (traditional) SQL will always be executed in polynomial time. This is not the case for statements written in a constraint programming language. It was argued that a library should be constructed that makes clear the complexity of the primitives in a language, such that users have a clear idea how to model problems efficiently. Convex optimization in machine learning was used as an example: certain mathematical operations are known to be more efficient than others, and hence machine learning researchers focus on using these.

The working group was concluded with a discussion on problems involving probabilistic and/or soft constraints. Many probabilistic inference problems in machine learning are hard to solve exactly and would require inexact solvers or solvers based on sampling. From a constraint programming perspective addressing such problems is still a challenge.

4.2 Learning Constraint Satisfaction Problems

Luc De Raedt (K.U. Leuven, BE)

This working group discussed how machine learning techniques could help to improve constraint programming. A constraint satisfaction problem is usually defined as a triplet (V, D, C), with V the variables, D the domain of the variables and C the constraints. In addition, there can be a preference function f that specifies which solutions are to be preferred.

The working group concluded that there were at least two possible roles for machine learning. The first role is for learning the model of the problem, that is, the triple (V, D, C) and possible the preference function f from data. The most direct task is that of learning the constraints C from data, that is, from possible assignments to the variables V that

satisfy the constraints. This is essentially a task that is akin to that of concept-learning. Learning constraints has been addressed within the database and inductive logic programming communities. Several variations on the task can be imagined: learning from positives only, learning from positive as well as negative examples, active learning (where the learner may ask questions), and learning *soft constraints*. Learning soft constraints would typically involves the learning of the weights or parameters of the constraints, though also the structure might be learned. This setting is akin to that of learning probabilistic models. As an alternative to learning soft constraints, one might learn the preference function f from examples of the form assignment x is better than assignment y. This is known in the machine learning literature as preference learning.

The second role for machine learning is to learn how to improve the performance of the solvers using experience, that is, speed-up learning. In this case it is usually assumed that the model (V, D, C) is given and passed runs of the solver are used to learn additional constraints or heuristics. This role for machine learning was extensively discussed in the talk by Michele Sebag. One can basically use deductive as well as inductive techniques for speed-up learning. Inductive techniques would gather data from previous test-runs and learn constraints or heuristics in a similar way as above, the key difference being that the learned hypotheses have to be tested against the model rather than against the data. On the other hand, deductive techniques would employ methods of explanation based learning to analyse traces and proofs and deductively obtain constraints that are implied by the given model.

4.3 Learning in Constraint Solvers

Barry O'Sullivan (University College Cork, IE)

This working group discussed how machine learning techniques could be used to improve the constraint solving process and, in particular, how machine learning techniques could be employed in the design and development of constraint solvers. The group discussed the opportunities for applying machine learning at a variety of levels, such as:

- solving specific problem instances, e.g. how can machine learning be used to solve a specific instance of a class of CSPs.
- designing portfolios of constraint solvers, building on existing approaches such as SATzilla (runtime prediction), CPhydra (case-based reasoning), ACME (solver selection as classification).
- lifelong learning, e.g. how can constraint solving experience be generalised and applied in an automated way?
- using learning to exploit problem structure, e.g. neighborhood identification for large-neighbourhood search methods.
- solver configuration, e.g. inspired by the MultiTAC system which parameterizes constraints solvers in terms of search heuristics, propagators, etc. Which technology should be used and when?

The key challenges for applying learning to solver design/development are: to identify the appropriate learning task that defines an appropriate success metric; to define an appropriate set of features to represent the training data; to select an appropriate distribution of scenarios/problem instances to use as training data; and, to select an appropriate learning

method. As evidenced in these working notes, there has been some progress in these areas. However, many rather mundane, but complex, challenges remain, e.g. on what basis should we claim that one solving procedure is better than another? Clearly solving time is important, but so is robustness if that solver must solve many instances within a problem class, or across several classes.

Many interesting opportunities for classification were discussed. For example, one can frame a classification task around the identification of "good" vs "bad" search trees to inform a randomized search algorithm which uses restarts. If a search tree is considered "bad" the algorithm should restart. A related discussion considered the opportunities for reinforcement learning to react to the context defined at a node during search. For example, should the problem be decomposed, should we propagate the state? These are all complex challenges which experienced constraint programmers deal with, and which could be aided by the appropriate use of machine learning.

Many researchers were interested in the use of learning for designing new search heuristics. There are many opportunities to gather statistics during search which can be used as training data for credit assignment methods and rule learners (both symbolic and numeric) upon which new heuristics can be learned.

A key concern for constraint programmers is to learn how to store and add constraints from dead-ends, often referred to as nogood learning. Challenges here relate to the identification, generalization, and forgetting of nogoods.

Finally, there was a long discussion about the use of learning for problem decomposition. An interesting question is how to generalize branching on assignments to branching on constraints. Machine learning can help us learn from past experience about the utility of solving particular problems using stream-lining constraints whereby a decision is made to impose a constraint that solutions should have specific structure. Similarly, machine learning can help learn grammars for composing interesting constraints.

In conclusion there was agreement that research at the interface between machine learning and constraint solving offered considerable promise and had the potential to yield many high-impact and valuable techniques.

5 Acknowledgements

We would like to thank the staff of Schloss Dagstuhl for their excellent help in organising this seminar.

Participants

- Rolf Backofen
Universität Freiburg, DE
- Roberto Bayardo
Google Inc. – Mountain View, US
- Michele Berlingerio
CNR – Pisa, IT
- Hendrik Blockeel
K.U. Leuven, BE
- Jean-François Boulicaut
INSA – Lyon, FR
- Maurice Bruynooghe
K.U. Leuven, BE
- Toon Calders
TU Eindhoven, NL
- Bruno Crémilleux
Université de Caen, FR
- James Cussens
University of York, GB
- Ian Davidson
Univ. of California – Davis, US
- Tijl De Bie
University of Bristol, GB
- Luc De Raedt
K.U. Leuven, BE
- Yves Deville
UC Louvain-la-Neuve, BE
- Pierre Dupont
UC Louvain-la-Neuve, BE
- Pierre Flener
Uppsala University, SE
- Paolo Frasconi
University of Firenze, IT
- Alan Frisch
University of York, GB
- Carmen Gervet
German University in Cairo, EG
- Bart Goethals
University of Antwerp, BE
- Tias Guns
K.U. Leuven, BE
- Georgiana Ifrim
University College Cork, IE
- Kristian Kersting
Fraunhofer IAIS – St. Augustin, DE
- Arnaud Lallouet
GREYC – Caen, FR
- Pat Langley
ASU – Tempe, US
- Laurent Michel
University of Connecticut – Storrs, US
- Sandy Moens
University of Antwerp, BE
- Barry O'Sullivan
University College Cork, IE
- Dan Roth
Univ. of Illinois – Urbana, US
- Salvatore Ruggieri
University of Pisa, IT
- Christian Schulte
KTH – Kista, SE
- Michele Sebag
LIMSI – CNRS, FR
- Arno Siebes
Utrecht University, NL
- Helmut Simonis
University College Cork, IE
- Jon Sneyers
K.U. Leuven, BE
- Guido Tack
K.U. Leuven, BE
- Pascal Van Hentenryck
Brown Univ. – Providence, US
- Peter Van Roy
UC Louvain-la-Neuve, BE
- Christel Vrain
Université d'Orleans, FR
- Toby Walsh
The University of New South Wales, Australia, AU

Report from Dagstuhl Seminar 11211

Geometric Modeling

Edited by

Thomas Grandine[1], Stefanie Hahmann[2], Jörg Peters[3], and Wenping Wang[4]

1 **The Boeing Company – Seattle, US,** thomas.a.grandine@boeing.com
2 **University of Grenoble, FR,** stefanie.hahmann@imag.fr
3 **University of Florida, Gainesville, US,** jorg@cise.ufl.edu
4 **University of Hong Kong, HK,** wenping@cs.hku.hk

Abstract

This report documents the program and the results of Dagstuhl Seminar 11211 "Geometric Modeling", taking place May 22-27 2011. The focus of the seminar was to discuss modern and emerging topics in Geometric Modeling by researchers and industrial scientists from all over the world.

Seminar 22.–27. May, 2011 – www.dagstuhl.de/11211
1998 ACM Subject Classification I.3.5 Computational Geometry and Object Modeling
Keywords and phrases Geometric modeling, design, curve and surface modeling, splines, CAD, B-splines, reconstruction, subdivision methods, multiresolution, parameterization
Digital Object Identifier 10.4230/DagRep.1.5.84

1 Executive Summary

Thomas Grandine
Stefanie Hahmann
Jörg Peters
Wenping Wang

The 8th Dagstuhl seminar was attended by 51 leading researchers coming from 4 continents and 21 countries (!). This high number shows the strong interest of the community for this event. The feedback from participants was very positive. A total of 45 presentations were grouped together into 12 sessions and let a lot of room for stimulating and fruitful discussions in the harmonic Dagstuhl atmosphere.

Geometric Modeling is the branch of Computer Science concerned with the acquisition, representation, modeling and analysis of 3-dimensional geometry. While its combination of technically complex and often interdisciplinary approaches is grounded both in applied mathematics and computer science data structures and theory, applications of the field therefore cover a wide collection of areas from classical product design, virtual prototyping and simulation to computer graphics, scientific visualization, medical imaging, multimedia and entertainment. It is therefore fitting that the seminar was attended by 7 leading scientists and engineers from industry.

The presentations ranged from surface reconstruction, GPU programming, to curve and surface modeling with classical splines, surface meshes and new subdivision methods based

on algebraic and differential geometry methods with applications in medical or architectural modeling and in videao and gaming industry. As with all previous Dagstuhl Seminars on Geometric Modeling, the conference proceedings will be published as a special issue in an international journal.

A special event during the conference was the John Gregory Memorial Award honoring Carl de Boor, Malcolm Sabin and Gershon Elber. This award is presented every three years at Dagstuhl and honors fundamental contributions to the field of geometric modeling.

The organizers thank all the attendees for their contributions and extend special thanks to the team of Schloss Dagstuhl for helping to make this workshop a success. As always, we enjoyed the warm atmosphere of the Schloss, which supports both formal presentations as well as informal exchanges of ideas.

2 Table of Contents

3 Overview of Talks

3.1 An intuitive way for constructing parametric quadric triangles

Gudrun Albrecht (University of Valenciennes, FR)

Joint work of G. Albrecht; M. Paluszny; M. Lentini

We present an intuitive algorithm for constructing rational triangular quadratic patches which lie on quadrics. The input of the algorithm is three vertex data points in 3D and normals at these points. It emanates from two existing methods thanks to an interesting geometric relation connecting them.

The sufficient condition for a configuration of vertex and normal data to allow for the existence of a rational triangular quadratic patch lying on a quadric whose tangent planes at the vertices are those prescribed by the given normals is the concurrence of certain cevians. When these conditions are not met we offer an optimization procedure to tweak the normals, without varying the vertex data, so that for the new normals there is a rational triangular quadratic patch that lies on a quadric.

3.2 Geometric Modeling and Analysis of Protein Complexes using algebraic splines

Chandrajit Bajaj (Univ. of Texas at Austin, US)

The rate limiting step of atomistic molecular dynamics (and binding affinity estimations in disease therapeutics), is the accurate evaluation of the molecular configurational potential energy E, and the vector atomic derivatives of E (a.k.a. force fields). In this talk I shall first introduce an implicit solvent free-energy model of protein complexes based on either Generalized Born, and Poisson-Boltzmann formulations, and present analysis algorithms for the rapid computation of bound and unbound free-energies. Our new algorithms rely on a novel algebraic surface spline element and the use of irregular fast Fourier methods, for capturing both the geometry of molecular solvation interfaces, as well as the efficient computation of polarization energies/forces.

3.3 A Volume Approach to Model Repair and Smoothing

Pere Brunet (Universitat Politècnica de Catalunya – Barcelona, ES)

Joint work of Pere Brunet, Antoni Chica, Eva Monclus, Isabel Navazo, Alvar Vinacua

In this talk we propose a novel strategy to automatically segment volume medical data using a high-quality mesh segmentation of a "example" model as a guiding example. The example mesh is deformed until it matches the relevant volume features. The algorithm starts from a volume model (scalar field of densities) to be segmented, together with an already existing

segmentation (polygonal mesh) of the same organ, usually from a different person. The preprocess step computes a suitable atracting scalar field in the volume model. After an approximate 3D registration between the example mesh and the volume (this is the only step requiring user intervention), the algorithm works by minimizing an energy and adapts the shape of the polygonal mesh to the volume features in order to segment the target organ.

3.4 Snake based segmentation of teeth from digitized dental casts

Guido Brunnett (TU Chemnitz, DE)

Joint work of Guido Brunnett, Thomas Kronfeld, David Brunner

During the past years several innovative and technological developments have been made in oral surgery. Today, digitized dental casts are common for simulation and planning of orthodontic interventions. In order to work with these models, knowing the exact position of the teeth is of high importance. In this paper, we present a new method for tooth segmentation with minimal user interaction. At the beginning, an initial estimate for the separating curve between the teeth and the gum is computed and optimized by use of an active contour. The second step calculates the dental arch and the interstices between the teeth. In order to detect each tooth surface exactly, we finally position a snake around the cusp of each tooth.

3.5 Adaptive Finite Element Methods for Elliptic Equations using PHT-Splines

Falai Chen (Univ. of Science & Technology of China – Anhui, CN)

Joint work of Falai Chen, Li Tian, Qiang Du

In this talk, we propse to use the polynomial splines over hierarchical T-meshes (PHT-splines) to solve elliptic equations with adaptive finite element methods. We develop a residual-based a posteriori error estimate for the finite element discretization of elliptic equations and establish their approximation properties. In addition, we conduct numerical experiments to verify the theoretical results and demonstrate the robustness of the error estimate and the effective approximations provided by the new spline space.

3.6 Computing with A B-spline like basis for the Powell Sabin 12-split

Elaine Cohen (University of Utah – Salt Lake City, US)

The Powell Sabin 12-split has the attribute that it can be used as a macro-element over arbitrary triangulations to create C1 smooth surfaces. We present attributes of a new B-spline like basis based on simplex splines and demonstrate applications.

3.7 Representing Non-Manifold Simplicial Shapes in Arbitrary Dimensions

Leila De Floriani (University of Genova, IT)

Joint work of Leila De Floriani, David Canino

Simplicial complexes are commonly used to describe non-manifold shapes in several applications, including finite element analysis, solid modeling, animation, terrain modeling and visualization of scalar and vector fields. In our work, we have developed several data structures for non-manifold simplicial shapes discretized as d-dimensional simplicial Euclidean complexes, both dimension-specific, and, more recently, dimension-independent. We will present the dimension-independent data structures, analyze and compare them, also with other representations, based on expressive power, storage requirements, and effectiveness in supporting navigation and updates. We will present our dimension-independent library for encoding and manipulating Euclidean simplicial complexes based on such data structures, which will soon be released in the public domain. We will finally discuss issues on topological modeling of non-manifold shapes.

3.8 A Sketching Interface for Feature Curve Recovery of Free-Form Surfaces

Ellen Dekkers (RWTH Aachen, DE)

Joint work of Ellen Dekkers, Leif Kobbelt, Richard Pawlicki, Randall C. Smith

In this talk, we present a semi-automatic approach to efficiently and robustly recover the characteristic feature curves of a given free-form surface where we do not have to assume that the input is a proper manifold. The technique supports a sketch-based interface where the user just has to roughly sketch the location of a feature directly on the input mesh. The system then snaps this initial curve to the correct position based on a graph-cut optimization scheme that takes various surface properties into account. Additional position constraints can be placed and modified manually which allows for an interactive feature curve editing functionality. We apply our technique to two practical scenarios. At first, feature curves can be used as handles for intuitive deformation of non-manifold surfaces. Secondly, we consider a practical problem in reverse engineering where we generate a statistical (PCA) shape model for car bodies.

3.9 Locally Refined B-splines

Tor Dokken (SINTEF – Oslo, NO)

Isogeometric analysis (IGA), introduced in 2005 by T.J. R Hughes, replaces traditional Finite Elements with element descriptions using tensor product 3-varaite NonUniform Rational

B-splines (NURBS), thus combining the exact topology description in Finite Element Analysis with the accurate shape representation in Computer Aided Design. NURBS based IGA shows great promise; however, NURBS lack the necessary local refinement needed. Locally Refined splines (LR B-splines), such as T-splines and the novel Locally Refined B-splines (LR B-splines) show great promise with respect to the needed local refinement. The talk will address IGA, and LR B-splines, and report on the work on the theoretical basis of LR B-splines.

3.10 Reconstruction of 3D objects from 2D cross-sections by the 4-point subdivision schemes

Nira Dyn (Tel Aviv University, IL)

We adapt the 4-point subdivision schemes to refine 2D sets, by expressing it in terms of repeated binary averages, also with negative small weights, and by using a newly designed affine combination of sets. The resulting set refining subdivision scheme is convergent and the limits generated by it approximates objects when the scheme is applied to their cross-sections. The quality of the approximation depends on the continuity properties of the sampled object. Several examples of reconstructed objects by this method will be demonstrated.

3.11 Class-A surfacing in practice

Matthias Eck (Dassault Systémes Deutschland GmbH – Hannover, DE)

Class-A surfacing characterizes since decades now the creation of aesthetic surfaces with highest visual expectations. We briefly sketch the main industrial application areas, some simple but important rules for the creation, and how software tools can support that. After outlining some exemplary actual requests, we consider next-generation possibilities including associative Class-A surfacing and tighter process integration with Concept Modeling.

3.12 Optimal Ruled Surface Fitting

Gershon Elber (Technion – Haifa, IL)

Ruled surfaces play a major role in many manufacturing processes, such as wire EDM and side CNC milling. In this work, we explore an optimal ruled surface fitting scheme to a given general freeform rational surface, S.

The solution is divided into two parts. In the first and given a line segment that spans two (boundary) points on S, we derive a scheme that efficiently computes the maximal distance from the line segment to S. Then, and using a discrete sampled set of boundary points on S,

all the possible line segments between point-pairs are considered and the best ruling fit is derived using dynamic-programming.

The result is an optimal ruling fit for general hyperbolic regions that is less effective on convex domains, a concern we also address in the examples section.

3.13 Agnostic G1 Gregory Surface Fitting

Gerald Farin (ASU – Tempe, US)

Joint work of Gerald Farin, Ashish Amresh

We utilize a geometric formulation of G1 conditions between adjacent Gregory patches. These conditions work equally well for three- or four- sided patches: they are "agnostic" to the patch type. We thus arrive at a simple and geometric framework for generating mixed patch type surfaces.

3.14 Spatial Pythagorean hodographs, quaternions, and rotations in $\mathbb{R}^3$ and $\mathbb{R}^4$

Rida T. Farouki (Univ. of California – Davis, US)

Quaternions, the first example of a non-commutative algebra, arose as a by-product of Hamilton's failed attempt to construct an "algebra of triples". Hamilton envisaged the quaternions as the "new language" of science and technology, but their place was usurped by vector analysis, an algebraically crude and overtly pragmatic subset of the quaternion algebra. A simple quaternion expression automatically generates Pythagorean quartuples of polynomials, thus yielding an elegant rotation invariant characterization of Pythagorean hodographs in $\mathbb{R}^3$. Quaternions provide compact and intuitive descriptions for rotations in $\mathbb{R}^3$, a fact that has lead to a renewed interest in them for robotics, computer graphics, animation, and related fields. Quaternions also describe rotations in $\mathbb{R}^4$, whose strange properties provide a cautionary tale against extrapolating our geometric intuition from $\mathbb{R}^2$ and $\mathbb{R}^3$ to Euclidean spaces of higher dimension.

3.15 Quantum Bezier Curves

Ron Goldman (Rice University – Houston, US)

Joint work of Ron Goldman, Plamen Simeonov, Vasilis Zafiris

Quantum Bezier curves are an extension of standard Bezier curves based on the quantum calculus. To study these new quantum Bezier curve, we extend standard blossoming (polar forms) to quantum blossoming. We apply quantum blossoming to derive extensions of

well known algorithms for standard Bezier curves, including de Casteljau's algorithms for recursive evaluation and subdivision, to quantum Bezier curves. We also derive algorithms for quantum differentiation of quantum Bezier curves as well as new quantum versions of Marsden's identity.

3.16 The Convolution of two B-splines

Thomas Grandine (The Boeing Company – Seattle, US)

An important application has arisen at The Boeing Company for which spline convolution appears to be a remarkably effective solution. Although the literature on this topic is fairly extensive, most of it is devoted to very special cases, e.g. approximation with cardinal splines, which is not directly applicable to The Boeing problem. One old result, however, did prove to be unexpectedly relevant to this problem. This talk will review that result and show some numerical results of that old technique applied to a spline convolution problem related to solar energy production.

3.17 Parametrisation in Isogeometric Analysis

Jens Gravesen (Technical University of Denmark, DK)

One of the attractive features of isogeometric analysis is the exact representation of the geometry. The geometry is furthermore given by a relative low number of control points and this makes isogeometric analysis an ideal basis for shape optimisation.

One problem is that the geometry of the shape is given by the boundary alone. And, it is the boundary and its parametrisation that is changed by the optimisation procedure. But isogeometric analysis requires a parametrisation of the whole domain. So in every optimisation cycle we need to extend a parametrisation of the boundary of a domain to the whole domain. It has to be fast in order not to slow the optimisation down but it also has to be robust and give a parametrisation of high quality. These are conflicting requirements so we propose the following approach. During the optimisation a fast linear method is used, but if the parametrisation becomes singular or close to singular then the optimisation is stopped and the parametrisation is improved using a nonlinear method. The optimisation then continues using a linear method.

We will explain how the validity of a parametrisation can be checked and we will describe various ways to parametrise a domain. We will in particular study the Winslow functional which turns out to have some desirable properties.

3.18 Realistic-Looking Wrinkles added to Coarse Cloth Simulation

Stefanie Hahmann (University of Grenoble – LJK-INRIA, FR)

Joint work of Stefanie Hahmann, D. Rohmer, M.-P. Cani, A. Sheffer, T. Poppa

Moving garments and other cloth objects exhibit dynamic, complex wrinkles. Generating such wrinkles in a virtual environment currently requires either a time-consuming manual design process, or a computationally expensive simulation, often combined with accurate parameter-tuning requiring specialized animator skills. Our work presents an alternative approach for wrinkle generation which combines coarse cloth animation with a post-processing step for efficient generation of realistic-looking fine dynamic wrinkles. Our method uses the stretch tensor of the coarse animation output as a guide for wrinkle placement. To ensure temporal coherence, the placement mechanism uses a space-time approach allowing not only for smooth wrinkle appearance and disappearance, but also for wrinkle motion, splitting, and merging over time. Our method generates believable wrinkle geometry using specialized curve-based implicit deformers. The method is fully automatic and has a single user control parameter that enables the user to mimic different fabrics.

3.19 Generalization of the incenter subdivision scheme

Victoria Hernandez-Mederos (ICIMAF – La Habana, CU)

In this talk we present an interpolatory Hermite subdivision scheme depending on a free parameter, which generalizes in certain way Deng's incenter subdivision scheme (CAGD, 2010). We prove that for any value of the free parameter the limit curve is G1 continuous. Moreover, if the vertices of the initial polygon and the tangent vectors are sampled from a circle with any arbitrary spacing, then the subdivision curve is the circle. The proposed scheme is shape preserving, avoids undesirable oscillations of the subdivision curve and introduces inflection points only in those regions of the curve where the control polygon suggests a change of convexity. Several examples are presented to demonstrate the performance of the scheme and we formulate some conjectures supported by numerical experiments.

3.20 Geometric Subdivision Curves

Kai Hormann (Universität Lugano, CH)

Joint work of Kai Hormann, Nira Dyn

While linear subdivision schemes for curves and surfaces are well-understood by now, we still have but a few tools for analyzing the non-linear setting, although non-linear schemes often yield better results than their linear counterparts. In this talk we discuss a geometric condition which guarantees that a non-linear interpolating subdivision scheme produces

tangent continuous limit curves. As an example, we present two non-linear variants of the classical 4-point scheme which satisfy this condition.

3.21 Topological Classification of Intersections of Two Ring Tori

Xiaohong Jia (University of Hong Kong, HK)

We provide an enumeration for the topology of the intersection curve of two ring tori in generic positions and orientations. Based on the enumeration, a simple algebraic approach of deciding the topology of the intersection curve of two ring tori is presented, which requires only counting the real root number of two univariate quartic polynomials constructed from the implicit equations of these two ring tori.

3.22 Parameterization of Contractible Domains Using Sequences of Harmonic Maps

Bert Jüttler (University of Linz, AT)

Joint work of Bert Jüttler, Thien-Tuan Nguyen

We propose a new method for parameterizing a contractible domain (called the computational domain) which is defined by its boundary. Using a sequence of harmonic maps, we first build a bijective mapping from the computational domain to the parameter domain, i.e., the unit square or unit cube. Then we parameterize the original domain by a spline approximation of the inverse mapping. Numerical simulations of our method were performed with several shapes in 2D and 3D to demonstrate that our method is suitable for various shapes. The method is particular useful for isogeometric analysis because it provides an extension from a boundary representation of a model to a volume representation.

3.23 Constructing Smooth Branching Surfaces from Parallel Cross Sections

Panagiotis Kaklis (National TU – Athens, GR)

Joint work of P. Kaklis, M. Karavelas, A.-A. Ginnis

We deal with the problem of constructing G^1 surfaces that interpolate data points on parallel cross sections, consisting of simple disjoined and non-nested contours, the number of which varies from plane to plane. We shall present preliminary results towards a more stable and general approach, versus that developed in [1] and [2]. In our novel approach we modify the way the "surrounding curve" is constructed, and we generalize our approach so as to be able to incorporate in the construction contours that are away from the convex hull of contours

(which was not handled in [1] and [2]). These modifications/generalizations affect a crucial step in our smooth branching surface framework, namely the treatment of the "one-to-many" configuration. The approach is, so far, limited to convex contours and is based on the Voronoi diagram / Delaunay graph of each coplanar contour-set. The Voronoi diagram/Delaunay graph of the contours is used for deciding:

1. membership and order of the contours touching the surrounding curve (via proximity to the convex-hull boundary),
2. (combinatorial) description of the surface patches to be created away from the surrounding curve (via neighboring information provided by the Voronoi diagram), and
3. location of the surface saddle areas (via immersing Voronoi vertices). The approach is completed by introducing a methodology for the construction of:
   - the (G^1) surrounding curve,
   - the (G^1) guiding curves along with tangent ribbons defined on them and finally,
   - a patchwork of locally defined quadrilateral Gordon-Coons patches that define the final surface.

References

1 N.C. Gabrielides, A.-A.I. Ginnis and P.D. Kaklis, "Constructing Smooth Branching Surfaces from Cross Sections", in Proceedings of the ACM Symposium on Solid and Physical Modeling 2006, June 6-8, 2006, Cardiff, UK, Leif Kobbelt and Wenping Wang (eds.), ACM, pp. 161-170.

2 N.C. Gabrielides, A.-A.I. Ginnis, P.D. Kaklis and M.I. Karavelas, "G^1-smooth Branching Surface Construction from Cross Sections", CAD, vol. 39, pp. 639-651 (2007).

3.24 Rational splines for free-form design with basic shapes

Rimvydas Krasauskas (Vilnius University, LT)

Joint work of Rimvydas Krasauskas, Jorg Peters

We present a rational G^2 (curvature continuous) analogue of non-uniform C^2 cubic B-spline paradigm. These rational splines can exactly reproduce the parts of multiple basic shapes, such as cyclides and quadrics in one smoothly connected structure.

3.25 Construction of spline basis over quad-linear sub-domains with T-junctions

Tae-Wan Kim (Seoul Nat. University, KR)

Joint work of Tae-Wan Kim, Yeong-hwa Seo, Bernard Mourrain

We give a talk about the existence of a partition of unity basis and the method of getting that basis having local support property in the case of local domain subdivision with T-junctions through linear algebraic analysis. From the subspace in which B-splines exists, we construct a matrix of which the null space is equivalent to the spline space. The shape of domain

is generalized from rectangular to the shape that can be grouped to several rectangles. Furthermore, the regularity imposed on each boundary between neighboring cells even can be n-1 if the degree of the spline space is n. The domain can be extended to one of k-parametric dimensions. From that null space, we get a basis and from the basis we get a partition of unity basis having local support property by multiplying change of basis matrices on it or solving a constrained quadratic model. In addition to this, we give a talk about the level by level local refinement for some specific domain and regularity conditions, especially one of PHT- splines(Polynomial Splines over Hierarchical T-meshes), that guarantees the increment of the dimension whenever the domain is crossly subdivided.

3.26 Sampling conditions and topological guarantees for shape reconstruction algorithms

Andre Lieutier (Dassault Systeme, FR)

In this talk, I shall first recall some results on shape reconstruction algorithms from sampled points coming with topological guarantees and, for that, recall the definition of some simplicial complexes built on finite set of points in Euclidean space: alpha-shapes, Cech and Rips complexes. Of course, no topological guarantee can be given without associated sampling condition. I shall then gives some properties of the Euclidean distance function to compact sets and associated results for sampling conditions on non smooth sets, mentioning the relation with topological persistence. I shall then quickly introduce sampling conditions and associated results based on the recent notion of defect of convexity.

3.27 Sketching Mesh Cutting

Ligang Liu (Zheiiang University – Hangzhou, CN)

Joint work of Ligang Liu, Min Meng, Lubin Fan, Zhongping Ji

Shape decomposition is an important ingredient in geometry processing and shape analysis. However, it remains a challenge to develop automatic mesh segmentation algorithms due to the complicated human perception. Therefore, user interaction is generally needed to obtain the desired meaningful parts. We will presenting a series of our works on sketch-based interactive shape decomposition. We will illustrate that users can easily express their intentions to obtain the desired semantic subparts by simply drawing some rough sketches over the shapes in an intuitive manner. The users do not care much about how precisely they draw the sketches and the algorithms can obtain the expected decompositions.

3.28 Surface meshing with planar quadrilateral faces

Yang Liu (Microsoft Research – Beijing, CN)

Joint work of Yang Liu, Weiwei Xua, Jun Wang, Lifeng Zhu

Planar quadrilateral (PQ) meshes are essential in architectural geometry for discretizing a freeform architectural structure with planar quad faces due to the economy and easy fabrication. An intuitive tool for designing and generating PQ meshes is much demanded by architects. However the existing methods do not provide a general solution which maximizes degrees of design freedom. To accomplish this task, we present a novel method based on conjugate direction fields (CDF) which allow k/4 ($k \in Z$) singularities. We characterize the smoothness of a CDF by the signed permutation and convert it to a simple summation of trigonometric functions for measurement. Starting with a triangle discretization of a freeform shape, we first compute an as smooth as possible CDF satisfying the user's directional and angular constraints, then apply global parameterization and planarization techniques to generate a PQ mesh which approximates the input shape and follows the CDF well. The effectiveness and robustness of our method are demonstrated on various 3D models.

3.29 Feature Adaptive GPU Rendering of Catmull-Clark Subdivision Surfaces

Charles Loop (Microsoft Research – Redmond, US)

Catmull-Clark subdivision surfaces. Unlike previous methods, our algorithm computes the true limit surface up to machine precision, and is capable of rendering surfaces that conform to the full RenderMan specification for Catmull-Clark surfaces. Specifically, our algorithm can accommodate base meshes consisting of arbitrary valence vertices and faces, and the surface can contain any number and arrangement of semi-sharp creases and hierarchically defined detail. We also present a variant of the algorithm which guarantees watertight positions and normals, meaning that even displaced surfaces can be rendered in a crack-free manner. Finally, we describe a view dependent level-of-detail scheme which adapts to both the depth of subdivision and the patch tessellation density. Though considerably more general, the performance of our algorithm is comparable to the best approximating method, and is considerably faster than Stam's exact method.

3.30 Realistic plant modeling

Geraldine Morin (ENSEEIHT – IRIT Toulouse, FR)

Joint work of Geraldine Morin, Sebastien Mondet, Jérome Guénard

Just as in the real world, plants are important objects in virtual worlds for creating pleasant and realistic environments, especially those involving natural scenes. As such, much effort

has been made in realistic modeling of plants. As the trend moves towards networked and distributed virtual environments, however, the current models are inadequate as they are not designed for progressive transmissions. We fill in this gap by proposing a progressive representation for plants based on generalized cylinders. We model the shape and thickness of branches in a plant as Bézier curves, group the curves according to the similarity, and differentially code the curves to represent the plant in a compact and progressive manner. Then, we derive a fully automatic, analysis-by-synthesis method to generate realistic vine plant models from easy-to-get images or video streams.

3.31 Uniformly stable linear 'wavelets' on general triangulations

Knut M. Morken (University of Oslo, NO)

Joint work of Knut M. Morken, Solveig Bruvoll, Tom Lyche

In this talk we construct wavelet-like functions on arbitrary triangulations. Nested triangulations are obtained through refinement by two standard strategies, in which no geometric regularity is enforced. One strategy inserts a new node anywhere inside a triangle and splits it into three smaller triangles by connecting the node to the three vertices of the triangle. The other strategy splits two neighbouring triangles into four smaller triangles by inserting a new node anywhere on the edge between the triangles and connecting this node to the two opposite vertices. The refinement results in nested spaces of piecewise linear functions. These functions are made to satisfy certain orthogonality conditions, which locally correspond to vanishing linear moments. We show that this construction is uniformly stable in the uniform norm, independently of the geometry of the original triangulation and the refinements, and illustrate with an example.

3.32 Rational linear re-parameterizations

Jörg Peters (University of Florida, US)

There is a unique family of rational linear re-parameterizations and associated surface partitions that allow three- or four-sided patches to join with geometric continuity. How can we take advantage of the fact that these re-pararameterization preserve the total degree of rational patches?

3.33 CubeCover - Cubical grids for bounded volumes

Konrad Polthier (FU Berlin, DE)

Joint work of Konrad Polthier, Matthias Nieser, Ulrich Reitebuch

We discuss novel techniques to fill a bounded volumetric shape with a (preferably coarse) cubical voxel structure. Among the optimization goals are alignment of the voxels with the bounding surface as well as simplicity of the voxel grid. Mathematical analysis of the possible singularities is given.

The algorithm is uses a tetrahedral volume mesh plus a user given guiding frame field as input. Then it constructs an atlas of chart functions, i.e. the parameterization function of the volume, such that the images of the coordinate lines align with the given frame field. Formally the function is given by a first-order PDE, namely the gradients of the coordinate functions are the vectors of the frame. In a first step, the algorithm uses a discrete Hodge decomposition to assure local integrability of the frame field. A subsequent step assures global integrability along generators of the first homology group and alignment a face of the boundary cube with the original surface boundary. All steps can be merged into solving linear equations.

Conceptually the presented CubeCover-algorithm extends the known QuadCover-algorithm from surface meshes to volumes meshes.

3.34 Circular Arc Structures

Helmut Pottmann (KAUST – Jeddah, SA)

Joint work of Helmut Pottmann, Pengbo Bo, Martin Kilian, Wenping Wang, Johannes Wallner

The most important guiding principle in computational methods for freeform architecture is the balance between cost efficiency on the one hand, and adherence to the design intent on the other. Key issues are the simplicity of supporting and connecting elements as well as repetition of costly parts. We propose so-called circular arc structures as a means to faithfully realize freeform designs without giving up smooth appearance. In contrast to non-smooth meshes with straight edges where geometric complexity is concentrated in the nodes, we stay with smooth surfaces and rather distribute complexity in a uniform way by allowing edges in the shape of circular arcs. We are able to achieve the simplest possible shape of nodes without interfering with known panel optimization algorithms. We study remarkable special cases of circular arc structures which possess simple supporting elements or repetitive edges, we present the first global approximation method for principal patches, and we show an extension to volumetric structures for truly three-dimensional designs.

3.35 The construction(s) of Ck subdivision schemes

Hartmut Prautzsch (Universität Karlsruhe, DE)

Subdivision algorithms for arbitrary meshes generating Ck surfaces that are also Ck at their extraordinary points have first been described 16 years ago. Several mprovements were developed and published until recently. I show that all possible constructions obey the same simple framework, which is based on the necessary and sufficient Ck conditions for subdivision schemes. I discuss the free parameters in this framework and show their limits and point to parameters that have not been used so far.

3.36 EC funding in the PEOPLE programme – a short tour

Ewald Quak (Technical University – Tallinn, EE)

Based on my experiences as an evaluator for the European framework programmes, I would like to give some pointers concerning funding opportunities, especially in the PEOPLE programme, including also those which concern cooperation with non-European research partners.

3.37 Ambient B-Splines

Ulrich Reif (TU Darmstadt, DE)

We propose a new spline representation of C^k-surfaces of arbitrary topology: To construct such a surface S from given data, say a triangular mesh S', we employ a smooth reference manifold M of equal genus and start with determining a parametrization $\Phi : M \to S'$ of S' over M. Let $v : M \to \mathbb{R}^3$ be a smooth, non-tangential vector field on M, e.g., the Gauss map. Then there exists a neighborhood Ω of M such that any point $y \in \Omega$ can be written as $y = x + tv(x)$ for unique $x \in M$ and $t \in \mathbb{R}$. Now, a trivariate function $F : \Omega \to \mathbb{R}^3$ is constructed by letting

$$F(y) := \Phi(x).$$

Then, the function F is approximated by some trivariate tensor product spline $G : \Omega \to \mathbb{R}^3$ of arbitrary degree and smoothness. Finally, G is restricted to M to obtain the desired representation $S : M \to \mathbb{R}^3$ of the surface.

The method is easy to implement and possesses many variants and generalizations. In particular, tensor product B-splines may be replaced by any other set of trivariate basis functions.

3.38 Interpolatory blending net subdivision schemes of Dubuc-Deslauriers type

Lucia Romani (Universita Bicocca – Milan, IT)

Joint work of Lucia Romani, Costanza Conti, Nira Dyn

In this talk we deal with interpolatory net subdivision schemes, i.e. iterative procedures which repeatedly refine nets of univariate functions and converge to continuous bivariate functions interpolating the initial net. We extend the family of Dubuc-Deslauriers interpolatory point subdivision schemes and construct an interpolatory blending net subdivision scheme. To this purpose, at each recursion step we use a net refinement operator based on the evaluation of an interpolating Gordon surface. Doing so the limit surface is not only interpolating the initial net of univariate functions, but also all the nets generated by the iterative procedure. Convergence and smoothness properties of these blending net subdivision schemes are proved in relation to the properties of the cardinal blending function used to define the Gordon surface, and by exploiting the notion of proximity with the tensor-product Dubuc-Deslauriers schemes for points. We conclude by presenting an example of a family of interpolatory blending net subdivision schemes whose first two members can be used to design C^1 and C^2 surfaces from given nets of 3D curves.

3.39 A superficial survey of approaches to CAD/CAE interfacing

Malcolm A. Sabin (Numerical Geometry Ltd., GB)

The interfacing of CAD models to analysis is a matter of current concern. There are a dozen or so approaches either in use or proposed. The talk outlines these and identifies some issues which cause each of them problems, and then goes on to suggest that we should regard the approaches as sources of good ideas rather than as competing technologies.

3.40 Powell-Sabin B-splines for adaptive local refinement in Isogeometric Analysis

Maria Lucia Sampoli (University of Siena, IT)

In this talk we present a novel approach, within the new paradigm of Isogeometric Analysis introduced by Hughes et al. (2005), to deal with numerical solution of partial differential equations. The new method is based on Powell-Sabin B-splines which, along with the ability of accurately describe the approximate solution, allow also local refinement strategies.

3.41 Creation and Manipulation of Discrete Indicator Functions

Scott Schaefer (Texas A&M University, US)

Joint work of Scott Schaefer, Josiah Manson, Jason Smith

In this talk, we will explore the creation and extraction of surfaces from Discrete Indicator Functions. These functions are box-filtered versions of exact Indicator functions and have a variety of applications in Computer Graphics and Geometric Modeling. To construct these functions from boundary representations, we will derive exact integrals of a wavelet decomposition of the function. This method yields an efficient algorithm for rasterizing polygons, fonts and implicitizing 3D surfaces. Next we consider the inverse problem of constructing a surface from a discrete indicator function and show that applying Marching Cubes with linear interpolation performs poorly. We provide a simple replacement for linear interpolation that dramatically improves the quality of the results with almost no computational overhead.

3.42 Possible Use of Subdivision Surfaces in Industrial Modelling

Timothy L. Strotman (Siemens – Milford, US)

Industrial modelling consists of two stages, design and analysis. Currently subdivision surfaces are rarely used in the design stage. Many analysis applications, such as finite element modeling, toolpath generation, and visualization, work with different faceted versions of the design part. In this talk we will explore how a single multi-resolution subdivision surface model could replace these facetted models in the analysis applications. We will present the strengths and weaknesses of this approach, hoping to generate a discussion for possible improvements to this approach.

3.43 Smoothing Operation on Spherical Covering of Partition of Unity Implicits

Hiromasa Suzuki (University of Tokyo, JP)

Partition of unity implicit surfaces (PU implicits) is a method proposed in [1] to reconstruct a triangular mesh from a set of scanned points. As with other implicit function based reconstruction approaches, an implicit function is first defined to roughly represent signed distance from the surface underlying the set of points. Then a triangular mesh is extracted as the zero level set of the implicit function by iso-surfacing. In the PU implcits an implicit function is defined by blending linear implicit functions each of which locally approximates the surface within a spherical support. These spherical supports all together form a covering of the set of points.

For dealing with noisy scanned points, it is necessary to apply a smoothing operation on the points or on a reconstructed triangular mesh. We rather derive new smoothing operation acting on the spherical covering of PU implicits [2]. New differential operators are proposed to compose a smoothing operation to the spherical covering. This process achieves noise robust surface reconstruction from scattered points.

References

1 Y. Ohtake, A. Belyaev, M. Alexa, G. Turk, and H.-P. Seidel. Multi-level partition of unity implicits. ACM Transaction of Graphics, 22(3):463–470, 2003.
2 Y. Nagai, Y. Ohtake, and H. Suzuki. Smoothing of partition of unity implicit surfaces for noise robust surface reconstruction. Computer Graphics Forum, 28(5):1339–1348, 2009. Proceedings of Symposium on Geometry Processing 2009.

3.44 Iso-geometric Finite Element Analysis Based on Catmull-Clark Subdivision Solids

Georg Umlauf (HTWG Konstanz, DE)

We present a volumetric iso-geometric finite element analysis based on Catmull-Clark solids. This concept allows to use the same representation for the modeling, the physical simulation, and the visualization, optimizing the design process reducing the gap between CAD and CAE. The underlying geometric models are Catmull-Clark surfaces with optional corners and creases. Hence, these models can easily be refined to increase the accuracy of the simulation. The crucial point in the simulation phase is the evaluation at arbitrary parameter values for the assembly of the stiffness matrices. We propose a method similar to the standard subdivision surface evaluation technique, such that numerical quadrature can be used.

Experiments for linear elastic materials show that our approach converges faster than methods based on tri-linear and tri-quadratic elements and is numerically more stable than tri-quadratic finite elements. However, the topological structure of Catmull-Clark elements is as simple as the structure of linear elements. Furthermore, the Catmull-Clark elements we use are C^2-continuous at the boundary and in the interior, except at extraordinary points.

We also use this technique to solve the incompressible steady-state Navier-Stokes equations for flow simulations. The non-linear system of equations which arises during the simulation is solved by Newton's method. Our experiments show that the approach requires approximately the same number of iterations as linear Lagrangian hexahedral finite elements for the same number of degrees of freedom. As Catmull-Clark finite elements are tri-cubic we can achieve higher accuracy with the same number of degrees of freedom and almost the same computational cost.

3.45 N-sided Transfinite Surface Interpolation Revisited

Tamas Varady (Technical University of Budapest, HU)

Joint work of Tamas Varady, Peter Salvi and Alyn Rockwood

Transfinite surface interpolation is a classic topic of CAGD, and various representational schemes to interpolate boundary and cross-derivative functions have been published for non-quadrilateral domains. Discussion starts with Coons patches that can be generalized in three possible ways combining side or corner interpolants. The direct generalization of Coons' scheme apparently has not been explored until now. In order to minimize shape artifacts and to provide a more natural patch interior for uneven and strongly curved boundary configurations former approaches are enhanced and extended. The proposed representation is based on irregular convex domains, special blending functions and balanced sweep-line parameterizations that avoid parametric shearing. Sweep-line constructions to create one- or two-sided patches are also presented.

3.46 Surface fitting with cyclide splines

Wenping Wang (University of Hong Kong, HK)

Joint work of Wenping Wang, Pengbo Bo and Yang Liu

We consider the problem of fitting a cyclide spline surface to a free-form target surface. It is shown that, by using a global optimization scheme with relaxed corner points and corner frames of cyclide patches, we can fit a free form surface by a G^1 spline surface composed of cyclide patches. T- joints are allowed for adaptive approximation and spherical patches are included in the cyclide spline to fill in singular regions around the umblic points of the underlying target surface.

Participants

- Gudrun Albrecht
University of Valenciennes, FR
- Chandrajit Bajaj
Univ. of Texas at Austin, US
- Robert Ellis Barnhill
ASU – Mesa, US
- Pere Brunet
Universitat Politécnica de Catalunya – Barcelona, ES
- Guido Brunnett
TU Chemnitz, DE
- Falai Chen
Univ. of Science & Technology of China – Anhui, CN
- Elaine Cohen
University of Utah – Salt Lake City, US
- Leila De Floriani
University of Genova, IT
- Ellen Dekkers
RWTH Aachen, DE
- Tor Dokken
SINTEF – Oslo, NO
- Nira Dyn
Tel Aviv University, IL
- Matthias Eck
Dassault Systémes Deutschland GmbH – Hannover, DE
- Gershon Elber
Technion – Haifa, IL
- Gerald Farin
ASU – Tempe, US
- Rida T. Farouki
Univ. of California – Davis, US
- Ron Goldman
Rice University – Houston, US
- Thomas A. Grandine
The Boeing Comp. – Seattle, US
- Jens Gravesen
Technical Univ. of Denmark – Lyngby, DK
- Hans Hagen
TU Kaiserslautern, DE
- Stefanie Hahmann
Inria Rhône-Alpes, FR
- Dianne Hansford
Paradise Valley, US
- Victoria Hernandez-Mederos
ICIMAF – La Habana, CU
- Kai Hormann
Universität Lugano, CH
- Xiaohong Jia
University of Hong Kong, HK
- Bert Jüttler
University of Linz, AT
- Panagiotis Kaklis
National TU – Athens, GR
- Kestutis Karciauskas
Vilnius University, LT
- Tae-Wan Kim
Seoul National University, KR
- André Lieutier
Dassault Systeme, FR
- Ligang Liu
Zheiiang Univ. – Hangzhou, CN
- Yang Liu
Microsoft Research – Beijing, CN
- Charles Loop
Microsoft Res. – Redmond, US
- Geraldine Morin
Université de Toulouse, FR
- Knut M. Morken
University of Oslo, NO
- Jörg Peters
University of Florida – Gainesville, US
- Konrad Polthier
FU Berlin, DE
- Helmut Pottmann
KAUST – Jeddah, SA
- Hartmut Prautzsch
KIT – Karlsruhe Institute of Technology, DE
- Ewald Quak
Technical Univ. – Tallinn, EE
- Ulrich Reif
TU Darmstadt, DE
- Richard F. Riesenfeld
University of Utah – Salt Lake City, US
- Lucia Romani
Universita Bicocca – Milan, IT
- Malcolm A. Sabin
Numerical Geometry Ltd. – Cambridge, GB
- Maria Lucia Sampoli
University of Siena, IT
- Scott Schaefer
Texas A&M University, US
- Timothy L. Strotman
Siemens – Milford, US
- Hiromasa Suzuki
University of Tokyo, JP
- Georg Umlauf
HTWG Konstanz, DE
- Tamas Varady
Technical Univ. of Budapest, HU
- Wenping Wang
University of Hong Kong, HK