DAGSTUHL
REPORTS

Volume 1, Issue 2, February 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
July, 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.1.2.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 Perspectives Workshop 11061

Online Privacy: Towards Informational Self-Determination on the Internet

Edited by

Simone Fischer-Hübner[1], Chris Hoofnagle[2], Kai Rannenberg[3], Michael Waidner[4], Ioannis Krontiris[5], and Michael Marhöfer[6]

1 **Karlstad University, Sweden,** simone.fischer-huebner@kau.se
2 **UC Berkeley, USA,** choofnagle@law.berkeley.edu
3 **Goethe University Frankfurt, Germany,** Kai.Rannenberg@m-chair.net
4 **TU Darmstadt, Germany,** michael.waidner@sit.fraunhofer.de
5 **Goethe University Frankfurt, Germany,** ioannis.krontiris@m-chair.net
6 **Nokia Siemens Networks – München, Germany,** michael.marhoefer@nsn.com

Abstract

The Dagstuhl Perspectives Workshop "Online Privacy: Towards Informational Self-Determination on the Internet" (11061) has been held in February 6-11, 2011 at Schloss Dagstuhl. 30 participants from academia, public sector, and industry have identified the current status-of-the-art of and challenges for online privacy as well as derived recommendations for improving online privacy. Whereas the Dagstuhl Manifesto of this workshop concludes the results of the working groups and panel discussions, this article presents the talks of this workshop by their abstracts.

Seminar 6.–11. February, 2011 – www.dagstuhl.de/11061
1998 ACM Subject Classification D.4.1 [Computers and Society: Public Policy Issues: Privacy]; H.4 [Information Systems Applications: Miscellaneous: Personalized Services, Business Processes, Web 2.0, (mobile) Internet]; K.6.5 [Computing Milieux: Management of Computing and Information Systems: Security and Protection]
Keywords and phrases Online privacy, Data protection, Data security, Data loss prevention, Informational self-determination, Web 2.0, (mobile) Internet
Digital Object Identifier 10.4230/DagRep.1.2.1
Edited in cooperation with Sven Wohlgemuth

1 Executive Summary

Simone Fischer-Hübner
Chris Hoofnagle
Kai Rannenberg
Michael Waidner

While the collection and monetization of user data has become a main source for funding "free" services like search engines, on-line social networks, news sites and blogs, neither privacy-enhancing technologies nor its regulations have kept up with user needs and privacy preferences.

The aim of this Dagstuhl Perspectives Workshop is to raise awareness for the actual state of the art of on-line privacy, especially in the international research community and

in ongoing efforts to improve the respective legal frameworks, and to deliver soon after the workshop a Dagstuhl Manifesto providing recommendations to industry, regulators, and research agencies for improving on-line privacy. In particular we have examined how the basic principle of informational self-determination, as promoted by European legal doctrines, could be applied to infrastructures like the Internet, Web 2.0 and mobile telecommunication networks.

It was deemed necessary and timely to bring together a *broad spectrum of key contributors* in order to promote both legally and commercially viable foundations for a balanced on-line privacy:

- *Academia* (specifically data security, privacy, cyber-law, and privacy-influential technologies & services),
- *Public sector* (data protection officers, organizers of relevant research programs, relevant civil rights organizations), and
- *Industry* (providers of communication solutions, browsers and apps; data aggregation and web analytics companies; providers of major Internet and mobile Internet services)

This workshop and its planned Dagstuhl Manifesto have four goals, aside from galvanizing an emerging research community:

1. *Provide a big picture of on-line privacy, which can be understood widely*
   Because of swift progress in the mobile Internet, on-line social networks, and on-line advertisements, it is a challenge for non-experts (and perhaps even experts themselves) to understand the current state of on-line privacy including the technologies and systems to collect personal information on-line.
2. *Compile the industry and engineering options to improve on-line privacy*
   On-line privacy depends on the technologies and systems used to access Internet/Web 2.0 services as well as on the services provided to users. Therefore industry has a strong influence.
3. *Update the respective legislative and regulative authorities on their options for enforcing practical, commercially viable informational self-determination of users in global infrastructures (e.g. EU's Privacy Directive to be revised in 2011)*
   Access to personal information is critical to self-determination; it is also seen as a right that serves a policing function among information-intensive firms. However, legal and business structures have often foreclosed rights of access, or made them impracticable for consumers to exercise.
4. *Foster industry's and academia's research for creating effective on-line privacy technologies, components, and systems that promote informational self-determination*
   Corresponding to additional risks for on-line privacy, new approaches are required in research to again establish adequate levels of on-line privacy.

This workshop has been structured into four parts, for each part, a topic responsible has been assigned:

- Part 1: Current S-o-A of on-line privacy w.r.t. to informational self-determination
  Responsible: Alma Whitten, Google Research, Great Britain
- Part 2: Industry & engineering options to improve on-line privacy
  Responsible: Michael Waidner, ex-IBM CTO Security, then TU Darmstadt, Germany
- Part 3: Recommendations for improving regulations of online privacy
  Responsible: Caspar Bowden, Microsoft WW Technology Office, Great Britain

- Part 4: Recommendations for research to improve the S-o-A of online privacy
  Responsible: Kai Rannenberg, Goethe University Frankfurt, Germany

A Dagstuhl Manifesto will conclude this workshop according to http://www.dagstuhl.de/en/program/dagstuhl-perspectives/.

2 Table of Contents

3 Overview of Talks

3.1 Shining Light on Leakage of Private Information via the Web

Craig E. Wills (Worcester Polytechnic Institute, USA)

This talk seeks to shine light on the leakage of private information via the Web. We first examine longitudinal results showing the size of users' privacy footprints continues to grow as presence of third-party trackers increases on first-party sites. We then examine the leakage of private information about users to these third parties via traditional and mobile social networking sites. We conclude with directions of current and future work.

3.2 Online Privacy – "The Mobile Aspect" Privacy in Mobile Applications and Beyond

Kai Rannenberg (Goethe University Frankfurt, Germany)

Location Information enables or at least supports many mobile applications such as pollen warning, children tracking, location based advertising, and mobile communities. At the same time processing of location information either by providers or by community peers creates sensitive issues, such profiling and dangers on personal safety. As mobile applications usually involve consortia with at least two providers, privacy and information flow issues are relevant and a sensitive matter. Often mobile telecommunications providers are in a key position and exposed to the privacy issues, as they maintain the customer relationship and their mobile communications systems (e.g. GSM, UMTS) hold the location information. However with the development of mobile sensors such as GPS receivers location information can be sensed widely and is available to more players in the value chain. Enabling privacy without disabling essential parts of the applications requires the users to make decisions on information flows.

This presentation reports on the data gathered in mobile communication systems and the activities of mobile phones to this regard, e.g. collecting data, reporting data to 3rd parties, and leaving traces. Solution approaches from projects such as the PRIME and PICOS are introduced, e.g. the PRIME LBS application prototype and the PICOS mobile angling and gaming community applications to demonstrate how users can be enabled to protect their privacy considering the tension between restricting information flows and their respective application interests.

3.3 Trust and Privacy: What is missing?

Claire Vishik (Intel – London, Great Britain)

In the last 10-15 years, significant advances were achieved in the area of trust and privacy. In the area of trust, understood as obtaining proof of expected behavior, new metrics and attestation protocols as well as technical trust elements in other technologies were added to the list of available approaches. The area of online privacy is much harder to define, but those engaged in privacy by design processes adopt specific parameters that reflect level of privacy in various types of technologies.

The presentation covers advances in bringing more privacy to data handling processes, hardware and software design as well as advances in building legal framework in regulatory frameworks. But the progress made thus far is not sufficient for modern computing environments. The study of levels of privacy **across domains** remains an emerging area at the time when most electronic processes and data sharing cut across domains. The **evidence of trust and privacy** that could work in cross-domain environments is in the very early stages of definition. **Policy enforcement** as opposed to policy interpretation is still in its infancy. **Truly multidisciplinary studies** are needed in trust and privacy where technical solutions are necessary, but not sufficient for progress. Greater pragmatism is also required to develop deployable and adoptable approaches to online privacy. A lot of work needs to be done, and a multidisciplinary group like the one that has gathered in Dagstuhl are necessary to make rapid and lasting progress.

3.4 What can Engineers and Industry do to improve Online Privacy?

Alma Whitten (Google Research, Great Britain)

Engineers and industry have a great deal of valuable work to do to improve online privacy. There is much progress still to be made on offering better transparency and clarity in our products and in our communications, and on employing innovative techniques to enhance understanding. Similarly, there is still much that can be done to offer people improved choice and control that better aligns with their needs and concerns. Finally, progress is steadily continuing on strengthening the safety of online systems through cryptography, sandboxing, more efficient patching, and more.

3.5 Privacy in Online Social Networks – Past Experiences, Future Challenges

Andreas Poller (Fraunhofer SIT – Darmstadt, Germany)

In their study in 2008, Fraunhofer SIT evaluated seven online social networks for measures to protect the privacy of their users [1]. The analysis and evaluation was based on a criteria catalogue which considers the very risks of social network platforms and state of the art privacy concepts. None of the tested platforms fully convinced the testers. In many cases Fraunhofer SIT dissuaded from using several platform functions.

Since then, the platforms solved most of their teething problems like missing TLS encryption for whole user sessions. However, several issues remain: Up to now, there exists no convincing business model which can respect the users' privacy, external audits take place rarely, and access control concepts are difficulty to use. In addition, the platforms are becoming more complex by integrating third-party applications. Particularly biometric and augmented reality functions foster new privacy threats.

To meet the further challenges, it is necessary to identify the several stakeholders like the individual data subject, the other platform users, non-members, platform provider and third-party application provider. It is required to analyze their relationships, the data or information flow among them, and their individual privacy needs. For example, the users' privacy concerns about the flow of their personal information to other users differ from privacy concerns towards the service provider as a data collector. Further research shall distinguish these problems and propose pertinent solution, be it new regulations or new usable privacy mechanism.

References

1 A. Poller. Privatsphärenschutz in Soziale-Netzwerke-Plattformen. Fraunhofer Institute for Secure Information Technology. Technical Report. 2008. http://www.sit.fraunhofer.de/Images/SocNetStudie_Deu_Final_tcm501-35966.pdf

3.6 Technology and Privacy: A lost Battle!?

Jan Camenisch (IBM Research – Zürich, Switzerland)

Engineers have and are building lots of devices and tool for people to communicate with each other and for tapping into the digital world. The way these have been built makes them leaving lots of traces that endanger the users' privacy. This is despite there being lots of technologies available that would allow one to build such tools and devices in a privacy respective and enhancing way. Of course, doing so will come at some cost in performance similarly as when building-in security. Thus: we need to consider and find an answer to why are engineers are today not doing privacy by design although they could?

3.7 Online Privacy: Reflections on the Regulatory Aspects

Jos Dumortier (K.U. Leuven, Belgium)

The provision of personalised services is an essential element of the Internet business model. Personalised service offering is not possible without processing personal information. Some people estimate that the best solution for protecting the individual in this context consists in asking this individual's consent before registering his personal data and using him for the provision of personalised online services. Obtaining such consent in an online environment is usually very easy. Consequently so-called "informational self-determination" is very popular in the commercial profiling and direct marketing business.

In Europe, however, the law doesn't consider privacy exclusively as an individual's business but rather as a societal good. Privacy is in the first place necessary as a condition for maintaining democracy. This viewpoint is clearly reflected in the European Convention of Human Rights (ECHR) and in the jurisprudence of the European Court of Human Rights. Privacy is closely connected to diversity since it is the contrary of societal control and conformity. Privacy protection is mainly necessary to guarantee free self-expression, which is a condition sine qua non for a democratic society. This is the reason why we consider privacy as a fundamental right, a right which cannot be given away by the individual. Regulation to protect online privacy should therefore not primarily focus on informational self-determination but on the prevention of the societal risks connected to the large "oceans of data" that are created in the context of the Internet business model.

3.8 On Regulations of Online Privacy

Caspar Bowden (Microsoft WW Technology Office, Great Britain)

The EU Data Protection Directive 95/46/EC addresses personal data as information relating to an identified or identifiable natural person (data subject). The principles of protection shall not apply to data rendered anonymous in such a way that the data subject is no longer identifiable. However, scientific discoveries about data privacy and de-anonymization attacks, e.g. k-anonymity, show that data subjects can be re-identified even though their data have been anonymized. Currently, data is considered "atomically" and there is no proportionality according to scale. However, systems increasingly collect identifiable transactional data with the "side-effect" that a database of all transactions is retained. This talk addresses the questions by whom data subjects are identifiable and how to define the concept of data to be regulated. This talk stresses the importance to consider data sets and not atomic data. A proposal is to establish "red line" limits (absolute rules) against new threats, e.g. storage of e-mails and "life logs" as well as to eliminate consent as an "escape clause" towards a "right to lie". It also addresses the question how a regulator can carry out a meaningful inspection of , e.g., cloud computing and how does one certify a privacy system.

3.9 Regulating Online Privacy: Why, What, and Where

Omer Tene (Israeli College of Management School of Law, Israel)

When assessing the regulation of online privacy we must answer three fundamental questions [10], namely **why** should we regulate cyberspace? **What** exactly should be regulated? And **where** will privacy regulation apply geographically?

Why Privacy regulation can be justified by one of two basic hypotheses: First, from a law and economic perspective, regulation (any regulation) is justified where there is a **market failure** [5]. Arguably, this is the case for online privacy, given consumers' relative ignorance of privacy policies and weak bargaining position *vis-á-vis* online service providers. Conversely, if the online market is sufficiently competitive (as it is widely considered to be), it can be expected to clear any informational and bargaining discrepancies to obtain an efficient equilibrium. The second basis for regulation in this sphere is the conception of **privacy as a fundamental human right**, tightly linked to human dignity and autonomy, and not subject to market forces [9]. Under this view, privacy regulation is justified regardless of the market equilibrium, and may be effected by paternalistic decisions concerning individuals' welfare.

What The two thorny issues for online privacy regulation are the definition of personal data and the scope of consent. First, the **definition of personal data**, the basic building block of any privacy regime, has come under stress recently based on researchers' demonstrations of the ability to re-identify or de-anonymize the people hidden in anonymized data sets. "Re-identification science disrupts the privacy policy landscape by undermining the faith that we have placed in anonymization." [4] Second, **consent** has proven to be a weak basis for processing data in an online environment which is increasingly complex, involves multiple parties (many of which are invisible to the consumer), and is largely based on the American "notice and choice" model of regulation, which has largely failed. Indeed, in its recent Preliminary Staff Report, Protecting Consumer Privacy in an Era of Rapid Change, the Federal Trade Commission states: "the notice-and-choice model, as implemented, has led to long, incomprehensible privacy policies that consumers typically do not read, let alone understand." [2] This view is echoed in the Department of Commerce "Green Paper" on privacy and innovation in the Internet economy: "From the consumer perspective, the current system of notice-and-choice does not appear to provide adequately transparent descriptions of personal data use, which may leave consumers with doubts (or even misunderstandings) about how companies handle personal data and inhibit their exercise of informed choices." [3] Yet consent cannot be entirely done away with as it is inexorably linked with the definition of privacy itself. We must therefore find a way to reinvigorate transparency and allow consumers to make meaningful choices with respect to the collection and use of their personal data.

Where Choice of law and jurisdiction (which law applies and who is to apply it) have always raised dense problems in the online ecosystem. This is due to the fact that choice of law and jurisdiction are typically determined according to geographical markers, whereas cyberspace transcends national borders [7]. In addition, the paradigm shift to cloud computing and storage of personal data in the cloud pose risks to privacy, as data changes hands, crosses borders, and may be accessed and used without the knowledge and meaningful consent of individuals [1, 8]. The European Union Justice Commissioner Viviane Reding recently announced that legislation proposed next summer will call for "four pillars", including the

extraterritorial application of the EU Data Protection Directive to entities in the United States collecting information online from European data subjects [6]. This solution (namely, a "targeting" test initially introduced in the United States in the *Zippo case* [11]) has benefits and costs, given that increased scope may add pressure on enforcement resources which are already scarce and yield suboptimal results.

References

1 A. Cavoukian. Privacy in the Clouds: A White Paper on Privacy and Digital Identity – Implications for the Internet. May 28, 2008. http://www.ipc.on.ca/images/Resources/privacyintheclouds.pdf

2 Preliminary FTC Staff Report. Protecting Consumer Privacy in an Era of Rapid Change: A Proposed Framework for Businesses and Policymakers. December 2010. http://www.ftc.gov/os/2010/12/101201privacyreport.pdf

3 The Department of Commerce Internet Policy Task Force. Commercial Data Privacy and Innovation in the Internet Economy: A Dynamic Policy Framework. At page 22. December 2010. http://www.ntia.doc.gov/reports/2010/iptf_privacy_greenpaper_12162010.pdf

4 P. Ohm. Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization. 57 UCLA Law Review 1701. 2010.

5 R. A. Posner. The Right to Privacy. 12 Ga. L. Rev. 393. 1978.

6 V. Reding, Vice-President of the European Commission EU Justice Commissioner. Your data, your rights: Safeguarding your privacy in a connected world Privacy Platform. The Review of the EU Data Protection Framework. Brussels. March 16, 2011. http://europa.eu/rapid/pressReleasesAction.do?reference=SPEECH/11/183

7 J. Reidenberg. Technology and Internet Jurisdiction. 153 Penn. L. Rev. 1951. 2005.

8 W. Robison. Free at What Cost?: Cloud Computing Privacy Under the Stored Communications Act. 98 Geo. L. J. 1195. 2010.

9 S. Simitis. Reviewing Privacy in the Information Society. 135 Penn. L. Rev. 707. 1987.

10 O. Tene. Privacy: The New Generation, 1 International Data Privacy Law 15. url-http://idpl.oxfordjournals.org/content/1/1/15.full. 2011.

11 Zippo Manufacturing Co. v. Zippo Dot Com, Inc., 952 F. Supp. 1119 (W.D. Pa. 1997). 1997.

3.10 Online Privacy – a European Commission Perspective

Jesus Villasante (European Commission – Brussels, Belgium)

The European Commission's initiatives in the field of online privacy consist of several strands including policy and regulatory initiatives, the involvement of end-users and society as well as support for research and innovation. The objective of the talk by Jesus Villasante is to put the technological innovation in the field of online privacy into context with the EU's privacy policies and its research activities.

The Digital Agenda for Europe of May 2010 summarises the European Commission actions in the area of ICT. One of its pillars is dedicated to "Trust and Security" focusing in particular on the safety and privacy of online content and services. The actions foresee among others the implementation of privacy and personal data protection, where research results and innovative solutions could provide crucial support to tackle the burning issues of online privacy.

Due to the dynamic changes of digital society, privacy issues gain in importance and policy must keep up to date with emerging technological challenges. In order to enable the user to control his privacy online, the current open issues include privacy by design, the right to be forgotten and emerging privacy issues in cloud computing and the Internet of Things.

Research and Development is one way of the European Commission to address these open issues and this summer the opening of FP7-ICT Call 8 will provide an excellent occasion for researchers to receive substantial funding for projects in the field of Trust, eID, and Privacy Management Infrastructures.

At the same time, 2011 will see intensive discussions on the future European Research Framework Programme "FP8". Consultations will try identifying the remaining technological challenges for Privacy, ID management, and trustworthy ICT, which need to be prioritised in the coming years to enable the application of European principles of privacy in the Future Internet.

3.11 Recommendation on Structure and Form of Manifesto

Jacques Bus (Digitrust EU – Brussels, Belgium)

This presentation does not intend to give proposals for a Manifesto text. That would be the task of the participants in the discussions to follow. I will only raise questions and make suggestions on issues that are in my view important when writing a Manifesto that intends to give recommendations to politicians and researchers about the problems and possible research agenda for solutions in the field of privacy in the digital environment. Following the organizational structure we expect the following parts in the Manifesto, where of course parts three and four depend on the results of the first two parts:

- Part 1: Current S-o-A of online privacy w.r.t. to informational self-determination, including background and relevance
- Part 2: Industry & engineering options to improve online privacy, including the existing challenges
- Part 3: Recommendations for improving regulations of online privacy
- Part 4: Recommendations for research to improve the S-o-A of online privacy

Who do we want to address and influence?

When writing the Manifesto, the main question is who we want to address and influence. If we aim at politicians, privacy commissioners, lawyers, regulators etc., then we must ask the question how can technology play its role in creating transparency, privacy assurance, auditing. What is the societal and industrial motivation for protecting and strengthening privacy and the arguments for doing research in this field. How do governments and citizens react to the digital world?

If we aim mostly on those who will have to fund the research in government and industry we must think about political and societal arguments, as well as arguments of industrial competitiveness and innovation. This holds at the EU level as well as at the Member State level. If we aim at the researchers in academia it is mainly about interesting research and potential publications, patents, and generally recognition.

A terminological minefield

Let us just address a few terms in the field that we discuss in this workshop.

Security: We may have many different things in mind when we use the word "security", even if we would only restrict its use to information. We can mean the protection of the secrecy of information, by hiding it away or encrypt it. We can also mean the safety and protection of people and relate it to data protection and informational privacy of citizens, but also to secrecy of information avoiding that sensitive information gets in the hands of criminals or terrorists. We may think of the protection of critical infrastructures and the control structures with data that can be infiltrated. And we can think about national security as protected by intelligence agencies, armies, police, within the state, at its borders and beyond. All these aspects lead to different solutions and the debate on the perceived balance between privacy and public security through surveillance is only one example of the difficulties we get involved in.

Identity: Davis [2] distinguishes three concepts:

1. Metaphysical identity: what are the essential qualities of a person that makes him unique
2. Physical identity: the carrier in flesh and blood of all the roles and qualities
3. Epistemological identity: created by relations to institutions; or existing because of various practices connected to our culture, language, ...

We can also talk about multiple (partial) identities if we consider every creation of a relation or an existence of practice that form together the epistemological identity, as one (partial) identity. In general we can say that an "identity" in a certain context is a particular set of credentials (attributes), called a partial identity. FIDIS [3] distinguishes (1) the structural perspective (ID as set of attributes) and the (2) process perspective (ID as set of processes of disclosure and usage of ID data, i.e. authentication). Many more perspectives are given in literature, demonstrating the complexity and fuzziness when we use the term "identity".

Privacy: Maybe it started with Warren and Brandeis [7] in 1890 and their plea for privacy as the "right to be left alone". Allen [1] considers:

1. Physical privacy (seclusion, solitude)
2. Informational privacy (confidentiality, secrecy, data protection and control over personal information)
3. Proprietary privacy (control over names, likeliness and repositories of personal information)

Helen Nissenbaum [6] gives an excellent account and framework of contextual informational integrity, demonstrating the dependence of privacy per/Users/sven/Desktop/Dagstuhl - abstract Jesus Villasante.docception on context and social norms. Privacy is laid down as a human right and the Data Protection Regulation of the EU as well as the so-called Privacy Regulation (on data protection in digital services) are reflections of that. Privacy and Identity are closely related subjects and proper identity management is a pre-requisite for privacy, but not sufficient.

Confidence, Trust, and Trustworthiness: Confidence can be had in institutions, organizations, technology to do what it is expected to do, although we often say we trust (or not) the government, a company, etc. However, trust has a positive connotation and technology can do what is expected, which might be negative. For example we can be confident that viruses are harmful to our system. Hardin [5] uses therefore "confidence" instead of "trust" in institutions (the latter he reserves for interpersonal relations). But Fukuyama [4] talks about trust in government, society (societal trust as a measure of opinion). Trust can be seen as a context-dependent (also culture, character or psychology) – relation between entities (often reserved for persons) to have a certain benign behavior or acting.

Trustworthiness is the quality of an entity, as believed by the truster, to behave in a certain way (One can trust an entity without the entity being trustworthy!)

It is clear that we must in this workshop and in the Manifesto it produces be careful in the use of the terms mentioned above. What type of entities do we consider? How do these terms relate? What terminology do we use in particular in the context of technology?

The choices depend on the audience (policy, industry, researchers), and in general we must avoid abstract and rigid use of language (unless it is meant solely for researchers). It must be understood that people want to recognize their thinking and preaching and be able to integrate new ideas in their normal talking. In general it is difficult for politicians to change language once they have presented their basic vision and policy documents for their job period. It is often better if we address a larger public to use various words and meanings and explain them by metaphors. Finally, in general, research program language is vague and abstract to avoid strong prescription, potential errors and the risk of being already out-of-date when it is published. It should also leave creativity to the proposers (some years later !!).

Confidence in Technology

The main requirements for users to get confidence in the technology they are offered are:

1. Technology providers must be open and transparent about how it works, how they make profit from it, and how they provide redress in case of harm done.
2. Government must develop effective regulation and law, which is as much as possible technology neutral and enforceable (also globally).
3. The technology application must give users the feeling that its use is compliant with their norms, that they understand the general picture and dangers and that they have ways of controlling such dangers.

Research Directions in Privacy?

When proposing research directions in privacy we must take account of:

1. It is about informational privacy and takes account of the essential factors: context; social norms; potential of data inference; and the need for data security.
2. The developments in industry and society (ad-nets, targeted advertisement, profiling, location data collection, data in the cloud, social networks).
3. It takes account of developments in the regulatory environment (focus on Privacy by Design, privacy assurance methodology, auditing, reproducibility).
4. Take account of societal developments: increasing general worries with seemingly little relation to the actual behavior.
5. Ensuring attention for confidence building during the whole product life-cycle from design till customer service.
6. The need for real multi-disciplinary research.

And in doing so we must consider the timing. What need to be done in the short (1-2 yr) term, what in 3-5 years and what beyond 5 years. For example, the expected revision of the EU Data Protection Framework might in particular need research on Privacy by Design, assurance and certification, modular and transparent data management processes and auditing.

Conclusions

Summarizing, in writing the Manifesto we must:

1. Think from the world of the audience
2. Be tolerant with their language and understanding

3. Accept their worries, understand their goals
4. Be rational with timing

But the real barriers to make long term progress in online privacy technology are:

1. Including the dynamicity, diversity and cultural and normative essence of life.
2. To achieve essential multi-disciplinarity in all future work.

References

1 A. Allen. Uneasy Access, Totowa, NJ, Rowman & Littlefield. 1998.
2 S. Davis. A conceptual analysis of identity. In Kerr, Steeves and Lucock (eds) Lessons from the identity trail. Anonymity, Privacy and Identity in a Networked Society, Oxford. 2009.
3 FIDIS. Future of Identity in the Information Society. 2009. http://www.fidis.net
4 F. Fukuyama. Trust: the social virtues and the creation of prosperity. Free press, NY. 1995.
5 R. Hardin. Trust & Trustworthiness. Russel Sage Foundation, NY. 2002.
6 H. Nissenbaum. Privacy in Context: Technology, Policy and the Integrity of Social Life. Stanford Univ Press. 2010.
7 S. Warren & L. Brandeis. The right to Privacy. In F.D. Schoeman (ed) Philosophical dimensions of Privacy: an Anthology, Cambridge, MA. 1890.

Participants

- Andreas Albers
Goethe University Frankfurt, Germany
- Caspar Bowden
Microsoft WW Technology Office, Great Britain
- Sonja Buchegger
KTH Stockholm, Sweden
- Johannes A. Buchmann
TU Darmstadt, Germany
- Jacques Bus
Digitrust EU – Brussels, Belgium
- Jan Camenisch
IBM Research – Zürich, Switzerland
- Fred Carter
IPC – Toronto, Canada
- Ingo Dahm
Deutsche Telekom AG, Germany
- Claudia Diaz
K.U. Leuven, Belgium
- Jos Dumortier
K.U. Leuven, Belgium
- Simone Fischer-Hübner
Karlstad University, Sweden
- Dieter Gollmann
TU Hamburg-Harburg, Germany
- Marit Hansen
ULD SH – Kiel, Germany
- Jörg Heuer
Deutsche Telekom AG Laboratories, Germany
- Stefan Köpsell
TU Dresden, Germany
- Ioannis Krontiris
Goethe University Frankfurt, Germany
- Michael Marhöfer
Nokia Siemens Networks – München, Germany
- Andreas Poller
Fraunhofer SIT – Darmstadt, Germany
- Kai Rannenberg
Goethe University Frankfurt, Germany
- Thomas L. Roessler
W3C, France
- Kazue Sako
NEC, Japan
- Omer Tene
Israeli College of Management School of Law, Israel
- Hannes Tschofenig
Nokia Siemens Networks – Espoo, Finnland
- Claire Vishik
Intel – London, Great Britain
- Michael Waidner
TU Darmstadt, Germany
- Rigo Wenning
W3C / ERCIM, France
- Alma Whitten
Google Research, Great Britain
- Craig E. Wills
Worcester Polytechnic Institute, USA
- Jesus Villasante (Observer)
European Commission – Brussels, Belgium
- Sven Wohlgemuth
National Institute of Informatics – Tokyo, Japan

Report from Dagstuhl Seminar 11062

Self-Repairing Programs

Edited by

Mauro Pezzè[1], Martin C. Rinard[2], Westley Weimer[3], and Andreas Zeller[4]

1 **University of Lugano, CH,** `mauro.pezze@usi.ch`
2 **MIT – Cambridge, US,** `rinard@lcs.mit.edu`
3 **University of Virginia, US,** `weimer@cs.virginia.edu`
4 **Saarland University, DE,** `zeller@cs.uni-saarland.de`

Abstract

Dagstuhl seminar 11062 "Self-Repairing Programs" included 23 participants and organizers from research and industrial communities. Self-Repairing Programs are a new and emerging area, and many participants reported that they initially felt their first research home to be in another area, such as testing, program synthesis, debugging, self-healing systems, or security. Over the course of the seminar, the participants found common ground in discussions of concerns, challenges, and the state of the art.

Seminar 06.–11. February, 2011 – www.dagstuhl.de/11062

1998 ACM Subject Classification D.2.1 Requirements/Specifications, D.2.4 Software/Program Verification, D.2.5 Testing and Debugging, D.2.7 Distribution, Maintenance, and Enhancement, F.3.1 Specifying and Verifying and Reasoning about Programs, I.2.2 Automatic Programming

Keywords and phrases automated program repair, contract, debugging, fault, patch, self-healing

Digital Object Identifier 10.4230/DagRep.1.2.16

1 Executive Summary

Mauro Pezzè
Westley Weimer
Andreas Zeller

Dagstuhl seminar 11062 "Self-Repairing Programs" included 23 participants and organizers from research and industrial communities. Self-Repairing Programs are a new and emerging area, and many participants reported that they initially felt their first research home to be in another area, such as testing, program synthesis, debugging, self-healing systems, or security. Over the course of the seminar, the participants found common ground in discussions of concerns, challenges, and the state of the art.

Why Self-Repairing Programs?

Recent years have seen considerable advances in automated debugging. Today, we have techniques that automatically determine problem causes — in the source code, in program input, in the change history, or in internal data structures. While these approaches make it

considerably easier to find the causes of defects, their precision is still insufficient to suggest a single concrete course of action — a human in the loop is still required to design and apply the patch. At the same time, there is an ongoing need for self-healing systems — systems that can recover from failures and even reconfigure themselves such that the failure no longer occurs. Most research efforts in this direction, though, assume planned recovery — that is, well-defined recovery strategies for anticipated failures.

An alternative is to explore self-repairing systems from a generic perspective — that is, to develop techniques that repair systems that are as generic, unassuming, and non-intrusive as program analysis and debugging. The idea is to determine actual fixes to state, to configuration, or to code. These fixes can be seen as guidance for the developer on how to fix the problem and evolve the software. However, fixes can also be deployed automatically, and effectively lead to programs that fix themselves. Such techniques may be particularly useful for orphaned systems that are no longer maintained or for critical software for which downtime is extremely expensive or even unacceptable. In these situations, there is no time to wait for a human developer to find and fix the bug. A synthesized patch can form a first line of defense against failures and attacks, a "first aid" approach to buy time while more expensive or manual methods are deployed. At the same time, automatically generated fixes provide a much richer diagnostic quality then simple fault localization, and thus may dramatically reduce the time it takes to debug a problem.

Goals of the Seminar

This main goals of this seminar was to provide knowledge exchange, mutual inspiration, and opportunities for collaborations for a rapidly developing field. The seminar aimed to bring together researchers in dynamic program analysis, automated debugging, specification mining, software survival techniques, and autonomic computing to increase awareness of these issues and techniques across relevant disciplines (program analysis, debugging and self-adaptive systems), and to discuss:

- how to monitor systems to detect abnormal state and behavior
- how to generate fixes and how to choose the best fixes
- how to deploy them in real-life systems and how to deal with the issues that arise when automatically correcting errors in software systems

Format and Presentations

The seminar started with summary presentations to bring all participants up to the speed on the state of the art and establish a common terminology. Subsequent activities alternated between technical presentations and plenary discussion sections. The seminar participants also split into two groups based on the self-identified focus areas of "The Architecture of Self-Repairing Systems" and "Validating Automated Repairs via Testing and Specification". Some evenings featured demonstrations or special-interest talks.

Common Concerns and Insights

As a whole, the group identified four challenge areas and opportunities for self-repairing programs: Architecture, Redundancy, Efficiency and Trust. In terms of Architecture, there was an acknowledgment that overall progress could be made by tackling particular problems or subdomains (e.g., fixing only atomicity violations or fixing only web applications, etc.) and potentially combining solutions later. There was a broad realization that redundancy is important on many levels: as a source of comparison for finding bugs or specifications; as a source of repair components; and as a main component of self-healing or self-adaptive systems at the architectural level. In terms of efficiency, the speed of the repair process — including the time required to validate a candidate repair — was of some concern, although many current techniques take minutes rather than hours to produce repairs. Trust was perhaps the most universally accepted issue: a notion that it is the responsibility of the repair process to provide an assurance argument, backed up by evidence, that would give a user or developer confidence that a repair can be applied safely. As general guidelines, we felt that an automatically-generated repair should not (or should at most minimally) regress the program by impairing functionality, and that applying such a repair should not be worse than doing nothing.

The group also identified two cross-cutting concerns related to correctness and evidence. The first was a notion that the evidence used to produce a repair (e.g., a few test cases or a partial specifications used for fault localization or repair construction) might be different from the evidence used to validate a final candidate repair (e.g., a larger test suite or a more complete specification). In addition, emphasis was placed on a clear characterization of common versus anomalous (or incorrect) behavior, possibly via a learned specification.

Challenge Areas Identified

The participants also identified a number of challenge areas or difficult tasks. By far the most popular was a notion of benchmarking. While the group acknowledged that the field is still quite new, and that formal benchmarks may not be appropriate, there was a desire for representative instances of programs with defects, tests that demonstrate those defects, normal regression tests, and indications of how humans fixed those defects.

The second challenge identified was the need for low-overhead, continuous monitoring to learn formal specifications for correct behavior, detect anomalies, and validate a system after repair deployment. The third challenge was to provide "just-in-time" repairs that were as quick as the auto-correction in Word or Eclipse. A fourth challenge related to documenting repairs or otherwise equipping them with evidence and arguments that would give confidence that they fix the system without causing additional harm.

A number of additional concerns were identified but were supported by a smaller segment of the participants. These included focusing on the economic value of repairs (e.g., targeting high severity defects or measuring the effort saved), the desire to repair programs even if a regression test suite is note available, the desire to have tools that succeed or fail with certainty (i.e., rather than producing incorrect repairs), some notion of automated repair techniques fixing 10% (or 50%, or 70%) of all reported bugs with some level of confidence, and the desire to improve automated fault localization techniques and allow them to report causes, not just locations.

2 Table of Contents

3 Overview of Talks

During the seminar, we alternated plenary and subgroup discussions and presentations. Here we summarize the main presentation.

3.1 Self-supervising BPEL processes

Luciano Baresi (Politecnico di Milano, IT)

Service compositions suffer changes in their partner services. Even if the composition does not change, its behavior may evolve over time and become incorrect. Such changes cannot be fully foreseen through pre-release validation, but impose a shift in the quality assessment activities. Provided functionality and quality of service must be continuously probed while the application executes, and the application itself must be able to take corrective actions to preserve its dependability and robustness. The talk proposes the idea of self-supervising BPEL processes, that is, special-purpose compositions that assess their behavior and react through user-defined rules. Supervision consists of monitoring and recovery. The former checks the system's execution to see whether everything is proceeding as planned, while the latter attempts to fix any anomalies. The talk introduces two languages for defining monitoring and recovery and explains how to use them to enrich BPEL processes with self-supervision capabilities. Supervision is treated as a cross-cutting concern that is only blended at runtime, allowing different stakeholders to adopt different strategies with no impact on the actual business logic. The talk also presents a supervision-aware run-time framework for executing the enriched processes.

3.2 Angelic Debugging

Satish Chandra (IBM TJ Watson Research Center – Hawthorne, US)

Software ships with known bugs because it is expensive to pinpoint and fix the bug exposed by a failing test. To reduce the cost of bug identification, we compute expressions that are likely causes of bugs and thus candidates for repair. Our symbolic method closely approximates an ideal approach to fixing bugs, which is to explore the space of all edits to the program, searching for one that repairs the failing test without breaking any passing test. We approximate this expensive ideal by computing not syntactic edits to an expression but instead the set of values whose substitution for the expression results in a correct execution. We observe that an expression is a repair candidate if it can be replaced with a value that fixes a failing test and, crucially, in each passing test, its value can be changed to another value without breaking the test. Such an expression is flexible because it permits multiple values; therefore, the repair of the expression is less likely to break a passing test. The method is called angelic debugging because the values are computed by angelically non-deterministic statements. We implemented the method on top of the Java PathFinder model

checker. Our experiments show that angelic debugging can pinpoint the source of the bug in both synthetic and realistic programs.

Based on joint work with Emina Torlak (formerly, IBM Research), Shaon Barman (Berkeley) and Ras Bodik (Berkeley).

3.3 What should we repair and how

Brian Demsky (University of California – Irvine, US)

I presented a summarizing talk entitled "What should we repair and how?" that covered previous work on repair. My talk covered work on repairing data structures, program values, program environment, and source code. The talk then extracted common themes. One theme is the problem of selecting repairs — avoiding trivial and undesirable themes. The next theme is a tradeoff between providing strong guarantees and the expressiveness of the system. Another theme is where the repair actions come from. Finally the question of whether the human is in the loop.

3.4 Introductory Rabble Rousing Talk

Stephanie Forrest (University of New Mexico – Albuquerque, US)

This introductory talk attempted to lay out some of the big questions for the field of self-repairing programs, including: How we know what a program should be doing; "How we know that a program is behaving incorrectly;" and "How do we find the bug"? The talk briefly highlighted common approaches to these problems, emphasizing anomaly detection approaches.

3.5 Mutational robustness

Stephanie Forrest (University of New Mexico – Albuquerque, US)

One form of redundancy in software is those statements or instructions that have no discernible effect on the execution of the program. We have measured this effect at both the Abstract Syntax Tree (using CIL) and at the assembly code level using the following procedure:

1. Start with an unmutated working program
2. Generate a random mutation using the mutation mechanisms described in our ICSE 09 and ASE 10 papers.
3. Run the mutated program on the test cases for the program.
4. Call a mutation that does not change test case behavior "neutral."
5. Repeat Steps 1–4.

This experiment produces a rate of neutral mutations ranging from 20 to 60%. Subsequent experiments showed that these startling results are not due to expected sources (e.g., code coverage of test suites, insertion of dead code, etc.).

We believe that this may be an important source of free redundancy for automated program repair.

3.6 Automatic Workarounds for Web Applications

Alessandra Gorla (University of Lugano, CH)

Faults in Web APIs can escape the testing process, and consequently applications relying on these libraries may fail. Reporting an issue and waiting until developers fix faults in failing Web APIs is a time consuming activity, and in this time frame many users may be affected by the same issue.

In this talk I will present a technique that finds and executes workarounds for faulty Web applications automatically and at runtime. Automatic workarounds exploit the inherent redundancy of Web applications, whereby a functionality of the application can be obtained through different sequences of invocations of Web APIs. In general, runtime workarounds are applied in response to a failure, and require that the application remain in a consistent state before and after the execution of a workaround. Therefore, they are ideally suited for interactive Web applications, since those allow the user to act as a failure detector with minimal effort, and also either use read-only state or manage their state through a transactional data store. This work focuses on faults found in the access libraries of widely used Web applications such as Google Maps. It starts with a classification of a number of reported faults of the Google Maps and YouTube APIs that have known workarounds. From those we derive a number of general and API-specific program-rewriting rules, which we then apply to other faults for which no workaround is known. Our experiments show that workarounds can be readily deployed within Web applications, through a simple client-side plug-in, and that program-rewriting rules derived from elementary properties of a common library can be effective in finding valid and previously unknown workarounds.

3.7 Improving population-based automated patch generation

Dongsun Kim (The Hong Kong University of Science & Technology, HK)

Generating patches is one of key activities in software maintenance. Once a buggy code is located, developers try to add, remove, or change source code in order to fix the bug. This bug resolution work is tedious and time consuming. Recent few work attempted to automate bug resolution using population-based approaches. However, only few cases can be resolved by these techniques. Our goal is to improve the current state of the art. A novel technique includes fix patterns and similarity measures to enhance the current practice. Preliminary experimental results show our technique expands the space of automated patch generation.

3.8 Automated Atomicity-Violation Fixing

Ben Liblit (University of Wisconsin – Madison, US)

Fixing software bugs has always been an important and time-consuming process in software development. Fixing concurrency bugs has become especially critical in the multicore era. However, fixing concurrency bugs is challenging, in part due to non-deterministic failures and tricky parallel reasoning. Beyond correctly fixing the original problem in the software, a good patch should also avoid introducing new bugs, degrading performance unnecessarily, or damaging software readability. Existing tools cannot automate the whole fixing process and provide good-quality patches.

I will present AFix, a tool that automates the whole process of fixing one common type of concurrency bug: single-variable atomicity violations. AFix starts from the bug reports of existing bug-detection tools. It augments these with static analysis to construct a suitable patch for each bug report. It further tries to combine the patches of multiple bugs for better performance and code readability. Finally, AFix's run-time component provides testing customized for each patch. Experimental evaluation shows that patches automatically generated by AFix correctly eliminate six out of eight real-world bugs and significantly decrease the failure probability in the other two cases. AFix patches never introduce new bugs and have similar performance to manually-designed patches.

3.9 Dynamic Analysis for Diagnosing Integration Faults

Leonardo Mariani (Università di Milano–Bicocca, IT)

In this talk we present the BCT analysis technique. BCT uses dynamic analysis to automatically identify the causes of failures and locate the related faults. BCT augments dynamic analysis techniques with model-based monitoring. In this way, BCT identifies a structured set of interactions and data values that are likely related to failures (failure causes), and indicates the components and the operations that are likely responsible for failures (fault locations).

3.10 Automated Regression Testing of Modified Software

Alessandro Orso (Georgia Institute of Technology, US)

Throughout its lifetime, software is modified to enhance its functionality, repair it, eliminate faults, and adapt it to new platforms. One common way to ensure that changes made on the program behave as intended and did not introduce unintended side effects is to run the new version of the program against a set of test cases (i.e., a test suite). Unfortunately, the existing test suite for the program may be inadequate for this task. First, the test suite

may contain too many test cases that do not test the modified parts of the program, and thus waste testing resources if run. Second the test suite may not contain test cases needed to adequately exercise the changes in the code. To address these issues, we present two approaches: the first approach analyzes the changes between two versions of a program and identifies the test cases in an existing test suite that do not need to be rerun; the second approach identifies behavioral differences between the two versions through test generation and differential dynamical analysis, and suitably presents them to the developers. In this talk, we present the two techniques, discuss their applicability in the context of self-repairing programs, and sketch possible future research directions.

3.11 Evidence-based automated program fixing

Yu Pei (ETH Zürich, CH)

Many programmers, when they encounter an error, would like to have the benefit of automatic fix suggestions—as long as they are, most of the time, adequate. Initial research in this direction has generally limited itself to specific areas, such as data structure classes with carefully designed interfaces, and relied on simple approaches.

To provide high-quality fix suggestions in a broad area of applicability, the present work relies on the presence of contracts in the code, and on the availability of dynamic analysis to gather evidence on the values taken by expressions derived from the program text.

The ideas have been built into the AutoFix-E2 automatic fix generator. Applications of AutoFix-E2 to general-purpose software, such as a library to manipulate documents, show that the approach provides an improvement over previous techniques, in particular purely model-based approaches.

3.12 Self healing lessons that may be learned from concurrency testing to self healing

Shmuel Ur (University of Bristol, GB)

I started by discussing how healing is done in the context of deadlock. One very interesting point about concurrent software is that bugs cause failure for specific inputs only some times. This means that for the same input there are interleaving that fail and some that succeed. The healing can be done by reducing the interleaving to those that do not have the bug without modifying the source code. Care must be taken not to reduce interleaving in such a way that a new interleaving is created. I showed how gate locks can be added to the code to protect from deadlock resulting from violation of lock discipline. It is interesting that in order to heal deadlocks we do not need one to occur first but the possibility of one is enough. I explained that care must be taken as the healing itself may cause new deadlocks and explained how to avoid such results.

The second topic I discussed was how to pinpoint the location of concurrent bugs. We use instrumentation to evaluate how timing change at each point are likely to expose the

bug. This give each point in the program a bug finding score. We refine the score by looking at the delta between close points in the control flow. The idea is that if one point is good at finding bugs and the next one is not as good, than this is a good location to show the programmer.

I briefly described additional relevant concurrent research. I showed how concurrent coverage is used. I explained that the noise that is used to reveal bugs must be carefully applied as too much will actually not be good at finding bugs and we discussed performance healing by slowing down some requests.

3.13 How should repairs be validated and deployed?

Westley Weimer (University of Virginia, US)

The validation and deployment of repairs is of critical concern to self-repairing systems. This summary presentation discusses possible settings (e.g., human-in-the-loop, short-term fix, long-term repair) as well as issues of trust. One key goal is that automated repairs must provide an assurance argument, based on evidence, that they are safe to apply. Most research thus far has focused on safety properties, but liveness properties and notions like dependability and reliability are also important. In addition, it is possible that insights from formal verification can be used to aid automated repair. Some researchers have taken advantage of the special structure or rich semantics of languages to provide additional information to the repair process. Others obtain help from humans or automatically mine partial specifications. By far the most common approach, however, is to use test cases to validate repairs.

3.14 Automated Fixing of Programs with Contracts

Yi Wei (ETH Zürich, CH)

In program debugging, finding a failing run is only the first step; what about correcting the fault? Can we automate the second task as well as the first? The AutoFix-E tool automatically generates and validates fixes for software faults. The key insights behind AutoFix-E are to rely on contracts present in the software to ensure that the proposed fixes are semantically sound, and on state diagrams using an abstract notion of state based on the boolean queries of a class.

Out of 42 faults found by an automatic testing tool in two widely used Eiffel libraries, AutoFix-E proposes successful fixes for 16 faults. Submitting some of these faults to experts shows that several of the proposed fixes are identical or close to fixes proposed by humans.

3.15 First Step Towards Automatic Correction of Firewall Policy Faults

Tao Xie (North Carolina State University, US)

Firewalls are critical components of network security and have been widely deployed for protecting private networks. A firewall determines whether to accept or discard a packet that passes through it based on its policy. However, most real-life firewalls have been plagued with policy faults, which either allow malicious traffic or block legitimate traffic. Due to the complexity of firewall policies, manually locating the faults of a firewall policy and further correcting them are difficult. Automatically correcting the faults of a firewall policy is an important and challenging problem. In this paper, we make three major contributions. First, we propose the first comprehensive fault model for firewall policies including five types of faults. For each type of fault, we present an automatic correction technique. Second, we propose the first systematic approach that employs these five techniques to automatically correct all or part of the misclassified packets of a faulty firewall policy. Third, we conducted extensive experiments to evaluate the effectiveness of our approach. Experimental results show that our approach is effective to correct a faulty firewall policy with three of these types of faults.

3.16 Pex for Fun: Tool Support for Human to Repair Programs for Fun and Learning

Tao Xie (North Carolina State University, US)

Although there are various emerging serious games developed for education and training purposes, there exist few serious games for practitioners or students to improve their programming or problem-solving skills in the computer science domain. To provide an open platform for creating serious games in learning computer science, in 2010 summer, Microsoft Research released a web-based serious gaming environment called Pex for Fun, in short as Pex4Fun (http://www.pexforfun.com/), for learning critical computer science skills such as problem solving skills and abstraction skills.

Within Pex4Fun, coding duels are interactive puzzles to offer both fun and learning. In a coding duel, a player's task is to implement the Puzzle method to have exactly the same behavior as another secret Puzzle method (which is never shown to the player), based on feedback on what selected values the player's current version of the Puzzle method behaves differently and the same way, respectively. Pex4Fun uses Pex, a white-box test generation tool for .NET based on dynamic symbolic execution, to automatically generate such feedback.

3.17 Programming with Delegation

Jean Yang (MIT – Cambridge, US)

Access control and information flow are outside core functionality but critical to program correctness. Such issues are at odds with innovation because they are global concerns: securely adding a program feature requires reasoning about interaction with existing functionality. Supporting the separation core functionality and security concerns would facilitate rapid development. I describe programming with delegation, a programming model and execution strategy that allows a program to be run according to a security policy that the programmer provides but that the system automatically enforces. With delegation, the programmer can associate values with policies and the runtime system is responsible for ensuring these policies are satisfied. The programmer gives the system flexibility to do so by introducing nondeterminism; the programmer governs the nondeterminism using constraints. The system executes such programs using symbolic execution and constraint propagation. In this talk, I describe the Jeeves programming language for programming with delegation, implementation of the Jeeves interpreter, and performance results that suggest the feasibility of this approach.

3.18 Dynamic Generation of Processes

Rogerio de Lemos (University of Kent, GB)

In this talk we present the development of a framework for the dynamic generation of processes that factors out common process generation mechanisms and provides explicit customisation points to tailor process generation capabilities to different application domains. The framework encompasses a reference process for managing the dynamic generation of processes, a reusable infrastructure for generating processes and a methodology for its instantiation in different application domains. The framework explores model driven technology for simplifying the generation of processes in different domains, and includes fault-tolerance mechanisms for dealing with faults during generation and execution of processes.

3.19 How should repairs be validated and deployed?

Rogerio de Lemos (University of Kent, GB)

This talk has presented an overview of validation and deployment of self-repair of software from two perspectives: feedback control loop (MAPE or CADA) and fault tolerance. The objective is to scope the issues that could be covered in an overview kind of talk. Based on a brief introduction several other points were raised. What to validate? Whether the actual system or a model of the system should be validated. What kind of evidence can be obtained when the validation produces inconclusive results? How to combine development- and run-time evidence? The talk concluded with the presentation of some challenges.

Participants

- Luciano Baresi
Politecnico di Milano, IT
- Michael Carbin
MIT – Cambridge, US
- Antonio Carzaniga
Universitt Lugano, CH)
- Satish Chandra
IBM TJ Watson Research Center – Hawthorne, US
- Rogerio de Lemos
University of Kent, GB
- Brian Demsky
Univ. of California – Irvine, US
- Stephanie Forrest
University of New Mexico – Albuquerque, US
- Alessandra Gorla
University of Lugano, CH
- Patrick Hurley
AFRL/RIGA - New York, US
- Dongsun Kim
The Hong Kong University of Science & Technology, HK
- Ben Liblit
University of Wisconsin – Madison, US
- Leonardo Mariani
Universita Bicocca–Milano, IT
- Alessandro Orso
Georgia Institute of Tech., US
- Yu Pei
ETH Zürich, CH
- Mauro Pezzè
University of Lugano, CH
- Stelios Sidiroglou-Douskos
MIT – Cambridge, US
- Armando Solar-Lezama
MIT – Cambridge, US
- Shmuel Ur
University of Bristol, GB
- Yi Wei
ETH Zürich, CH
- Westley Weimer
University of Virginia, US
- Tao Xie
North Carolina State Univ., US
- Jean Yang
MIT – Cambridge, US
- Andreas Zeller
Universität des Saarlandes, DE

Report from Dagstuhl Seminar 11071

Theory and Applications of Graph Searching Problems (GRASTA 2011)

Edited by

Fedor V. Fomin[1], Pierre Fraigniaud[2], Stephan Kreutzer[3], and Dimitrios M. Thilikos[4]

1 **University of Bergen, NO,** fomin@ii.uib.no
2 **Université Paris Sud, FR,** pierre.fraigniaud@liafa.jussieu.fr
3 **University of Oxford, GB,** kreutzer@comlab.ox.ac.uk
4 **National and Kapodistrian University of Athens, GR,** sedthilk@math.uoa.gr

Abstract

From February 14, 2012 to February 18, 2012, the Dagstuhl Seminar 11071 "Theory and Applications of Graph Searching Problems (GRASTA 2011)" was held in Schloss Dagstuhl – Leibniz Center for Informatics. During the seminar, participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and open problems are put together in this paper. The first section describes the seminar topics and goals in general. The second section contains the abstracts of the talks and the third section includes the open problems presented during the seminar.

Seminar 14.–18. February, 2011 – www.dagstuhl.de/11071
1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems, G.2.1 Combinatorics, G.2.2 Graph Theory, G.2.3 Applications, I.2.9 Robotics
Keywords and phrases Graph Searching, Pursuit Evasion Games, Cop and Robers Games, Fugitive Search Games
Digital Object Identifier 10.4230/DagRep.1.2.30

1 Executive Summary

Fedor V. Fomin
Pierre Fraigniaud
Stephan Kreutzer
Dimitrios M. Thilikos

Graph searching is often referred to, in a more playful language, as a pursuit-evasion game (or, alternatively, cops and robbers game). This is a kind of game where one part is a set of escaping mobile entities, called evaders (or fugitives), that hide in a graph representing a network, and the other part is a number of chasing agents, called searchers (or pursuers), that move systematically in the graph. The game may vary significantly according to the capabilities of the evaders and the pursuers in terms of relative speed, sensor capabilities, visibility, etc. The objective of the game is to capture the evaders in an optimal way, where the notion of optimality itself admits several interpretations.

Theory and Applications of Graph Searching Problems, *Dagstuhl Reports*, Vol. 1, Issue 2, pp. 30–46
Editors: Fedor V. Fomin, Pierre Fraigniaud, Stephan Kreutzer, and Dimitrios M. Thilikos
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Graph searching revealed the need to express in a formal mathematical way intuitive concepts such as avoidance, surrounding, sense of direction, hiding, persecution, and threatening. There are many variants of graph searching studied in the literature, which are either application driven, i.e. motivated by problems in practice, or are inspired by foundational issues in Computer Science, Discrete Mathematics, and Artificial Intelligence including

- Information Seeking
- Robot motion planning
- Graph Theory
- Database Theory and Robber and Marshals Games
- Logic
- Distributed Computing
- Models of computation
- Network security

The objective of the seminar was to bring researchers from the widest possible variety of disciplines related to graph searching and we will especially encourage the maximum interplay between theory and applications. The meeting initiated the exchange of research results, ideas, open problems and discussion about future avenues in Graph Searching. As a fruit of this encounter new research results, open problems, and methodologies will appeared, especially those of interdisciplinary character.

2 Table of Contents

3 Overview of Talks

3.1 Cops and Robbers played on random graphs

Pawel Pralat (West Virginia Univ. – Morgantown, US)

We study the vertex pursuit game of Cops and Robbers, in which cops try to capture a robber on the vertices of the graph. The minimum number of cops required to win on a given graph G is called the cop number of G. We present asymptotic results for the game of Cops and Robbers played on random graph. In particular we show that:

- the Meyniel's conjecture holds a.a.s. for a random d-regular graph $G(n,d)$ as well as a binomial random graph $G(n,p)$ – joint work with Wormald,
- the cop number of $G(n,p)$ as a function of an average degree forms an intriguing zigzag shape – joint work with Luczak,
- almost all cop-win graphs contain a universal vertex – joint work with Bonato and Kemkes.

Other related problems will be mentioned as well.

3.2 Complexity of Cops and Robber Game

Petr Golovach (University of Durham, GB)

The Cops and Robbers game was defined independently by Winkler-Nowakowski and Quilliot in the 1980s and since that time has been studied intensively. Despite of such a study of the combinatorial properties of the game, almost no algorithmic results on this game are known. Perhaps the main algorithmic result known about Cops and Robbers game is the observation that determining whether the cop number of a graph on n vertices is at most k can be done by a backtracking algorithm which runs in time $n^{O(k)}$ (thus polynomial for fixed k). From the hardness side, Goldstein and Reingold in 1995 proved that the version of the Cops and Robbers game on *directed* graphs is EXPTIME-complete. Also, they have shown that the version of the game on undirected graphs when the cops and the robber are given their initial positions is also EXPTIME-complete. They also conjectured that the game on undirected graphs is also EXPTIME-complete. However, even NP-hardness of the problem was proved only in 2008 by Fomin, Golovach and Kratochvíl. We survey the known complexity results about the Cops and Robbers game and its variants and give a list of open problems.

3.3 Robotic Pursuit Evasion and Graph Search

Athanasios Kehagias (Aristotle University of Thessaloniki GR)

Robotic Pursuit Evasion (PE) is a hot research area in the robotics community. Among the various mathematical tools roboticists use to model the PE problem, the Graph Search (GS) theory is a prominent (but not the only) example. In this talk I will present and compare several GS-based approaches to robotic PE. I will point out similarities but also differences between robotic PE and graph search. In particular, I will compare the goals, methodology and outlook of roboticists, pure mathematicians and applied mathematicians who have attacked the problem. I will also present some robotic PE problems which require extensions of the "classical" GS setup and I will briefly discuss models of robotic PE which use graphs but *not* the graph search setup.

3.4 Polygon reconstruction from local observations

Peter Widmayer (ETH Zürich, CH)

We study the problem of reconstructing an unknown simple polygon from a series of certain local observations, similar in spirit to the reconstruction of an unknown network by exploring it. For mobile agents that move in simple ways inside a polygon, we are interested in understanding what types of local observations carry enough information to allow polygon reconstruction. This is part of a more general effort to understand when and how simple primitives allow mobile agents to draw global conclusions about the environment from local observations.

3.5 The price of connectivity in graph searching games

Dariusz Dereniowski (Gdansk University of Technology, PL)

In the edge searching problem the goal is to clear a simple graph that is initially entirely contaminated. The task is performed by a team of searchers that are allowed to make three types of moves: a searcher is placed on a vertex, a searcher is removed from a vertex, and a searcher slides from a vertex to one of its neighbors. The fugitive is invisible, fast, and has complete knowledge about the graph and the strategy used by the searchers. The fugitive is considered captured if a searcher reaches his location. We are interested in determining the minimum number of searchers (i.e. the search number) required to search a given graph. In the connected graph searching problem we have an additional restriction: the subgraph that is free of the fugitive is always connected. In this talk we discuss the connection between the search number and the connected search number, including an algorithm that converts a given search strategy using k searchers into a connected one using at most $2k + 3$ searchers.

3.6 On the Fast Edge Searching Problem

Boting Yang (University of Regina, CA)

In this talk, we consider the problem of finding the minimum number of steps to capture the fugitive. We introduce the fast edge searching problem in the edge search model, we describe relations between the fast edge searching and other searching problems, such as the fast searching and the node searching problems, and we present some recent progress on lower bounds and upper bounds of fast search numbers.

3.7 Algorithms for solving infinite games on graphs

Marcin Jurdzinski (University of Warwick, GB)

This talk is a selective survey of algorithms for solving a number of infinite-path-following games on graphs, such as parity, mean-payoff, and discounted games. The games considered are zero-sum, perfect-information, and non-stochastic. Several state-of-the-art algorithms for solving infinite games on graphs are presented, exhibiting disparate algorithmic techniques, such as divide-and-conquer, dynamic programming/value iteration, local search/strategy improvement, and mathematical programming, as well as hybrid approaches that dovetail some of the former. While the problems of solving infinite games on graphs are in NP and co-NP, and also in PLS and PPAD, and hence unlikely to be complete for any of the four complexity classes, no polynomial-time algorithms are known for solving them.

3.8 On the complexity of CSP decompositions

Zoltán Miklos (EPFL – Lausanne, CH)

We give a short overview of the relation of certain graph and hypergraph games (robbers and cops/marshals) and CSP decompositions. We discuss the complexity of these problems, in particular the case of tree decompositions. Finally, we report on some progress about the analogous hypergraph problems.

3.9 Searching Games

Maria Serna (UPC – Barcelona, ES)

We consider a general multi-agent framework in which a set of n agents are roaming a network (such as internet or social networks) where m valuable and sharable goods (or resources or

services) are hidden in m different vertices of the network. We analyze several strategic situations that arise in this setting by means of game theory. To do so we introduce search games, in those games agents have to select a simple path from a predetermined set of initial vertices. Depending on how the goods are splitted among the agents we consider two search game types: finders-share in which the agents that find a good split among them the corresponding benefit and first-share in which only the agents that first find a good share the corresponding benefit. We show that finders-share games always have pure Nash equilibria (PNE). For obtaining this result we introduce the notion of Nash preserving reduction between strategic games. We show that finders-share search games are Nash reducible to single-source network congestion games. This is done through a series of Nash preserving reductions. For first-share search games we show the existence of games with and without PNE. Furthermore we identify some graph families in which the first-share search game has always a PNE that is computable in polynomial time. We discuss also some variants of searching games and the associated graph parameters.

3.10 An overview of The Firefighter Problem

Margaret-Ellen Messinger (Mount Allison University – Sackville, CA)

The Firefighter Problem is a simplified model for the spread of a fire (or disease or computer virus) in a network. Initially, a fire breaks out at a vertex in a connected graph. At each subsequent time step, firefighters *protect* a fixed number of unburned vertices and then the fire spreads to all *unprotected* neighbors. Since its introduction in 1995, there has been a steady growth of both structural and algorithmic results. One possible objective is to maximize the number of saved vertices: this generally requires a strategy on the part of the firefighters, while the fire itself spreads without any strategy. Another possible objective is to find the number of firefighters needed to save a particular number of, or fraction of, or subset of the vertices. (These objectives are sometimes in conflict.) I will discuss some interesting results as well as variants and open problems.

3.11 Graphs with average degree smaller than $\frac{30}{11}$ are burning slowly

Pawel Pralat (West Virginia Univ. – Morgantown, US)

We consider the following *firefighter problem* on a finite graph $G = (V, E)$. Suppose that a fire breaks out at a given vertex $v \in V$. In each subsequent time unit, a firefighter protects one vertex which is not yet on fire, and then the fire spreads to all unprotected neighbours of the vertices on fire. Since the graph is finite, at some point each vertex is either on fire or is protected by the firefighter, and the process is finished. The objective of the firefighter is to save as many vertices as possible. Let $\mathbf{sn}(G, v)$ denote the number of vertices in G the firefighter can save when a fire breaks out at vertex $v \in V$, assuming the best strategy is used. The surviving rate $\rho(G)$ of G is defined as the expected percentage of vertices that can be saved when a fire breaks out at a random vertex of G, that is, $\rho(G) = \frac{1}{n^2} \sum_{v \in V} \mathbf{sn}(G, v)$. The main focus of the talk is on sparse graphs. Let $\epsilon > 0$. We show that any graph G on n vertices with at most $(\frac{15}{11} - \epsilon)n$ edges can be well protected, that is, $\rho(G) > \frac{\epsilon}{60} > 0$.

Moreover, a construction of a random graph is proposed to show that the constant $\frac{15}{11}$ cannot be improved.

3.12 Cops and Robbers on Directed Graphs

Jan Obdrzalek (Masaryk University, PL)

We survey the current status of cops and robber games on directed graphs. After presenting the different variants of the game for the most common digraph measures we ask the following question: Is there a digraph width measure which is powerful (i.e. a big class of problems is decidable in linear/polynomial time if this measure is bounded), significantly different from tree-width and yet, at the same time, characterizable by a variant of the cops-and-robber game for tree-width? We show that, under some standard complexity theory assumption, this is not so. We also show a new improvement of this result: That we do not need the measure to be efficiently orientable for our theorem to hold.

3.13 Cop and robber games when the robber can hide and ride

Nicolas Nisse (INRIA Sophia Antipolis, FR)

In the classical cop and robber game, two players, the cop C and the robber R, move alternatively along edges of a finite graph $G = (V, E)$. The cop captures the robber if both players are on the same vertex at the same moment of time. A graph G is called cop win if the cop always captures the robber after a finite number of steps. Nowakowski, Winkler (1983) and Quilliot (1983) characterized the cop-win graphs as graphs admitting a dismantling scheme. In this talk, we characterize in a similar way the cop-win graphs in the game in which the cop and the robber move at different speeds s' and $s, s' \leq s$. We also investigate several dismantling schemes necessary or sufficient for the cop-win graphs in the game in which the robber is visible only every k moves for a fixed integer $k > 1$. We characterize the graphs which are cop-win for any value of k.

3.14 Complexity of the cop and robber guarding game

Tomas Valla (Charles University – Prague, CZ)

The guarding game is a game in which several cops has to guard a region in a (directed or undirected) graph against a robber. The robber and the cops are placed on vertices of the graph; they take turns in moving to adjacent vertices (or staying), cops inside the guarded region, the robber on the remaining vertices (the robber-region). The goal of the robber is to enter the guarded region at a vertex with no cop on it. The problem is to determine whether

for a given graph and given number of cops the cops are able to prevent the robber from entering the guarded region. The problem is highly nontrivial even for very simple graphs. It is know that when the robber-region is a tree, the problem is NP-complete, and if the robber-region is a directed acyclic graph, the problem becomes PSPACE-complete [Fomin, Golovach, Hall, Mihalák, Vicari, Widmayer: How to Guard a Graph? Algorithmica, DOI: 10.1007/s00453-009-9382-4]. We solve the question asked by Fomin et al. and we show that if the graph is arbitrary (directed or undirected), the problem becomes ETIME-complete.

3.15 Multi-target ray searching problems

Spyros Angelopoulos (CNRS – Paris, FR)

We consider the problem of exploring m concurrent rays (i.e., branches) using a single searcher. The rays are disjoint with the exception of a single common point, and in each ray a potential target may be located. The objective is to design efficient search strategies for locating t targets (with $t \leq m$). This setting generalizes the extensively studied ray search (or star search) problem, in which the searcher seeks a single target. In addition, it is motivated by applications such as the interleaved execution of heuristic algorithms, when it is required that a certain number of heuristics have to successfully terminate. We study the problem under two different cost measures, and show how to derive optimal search strategies for each measure.

3.16 Characterizations of k-cop win graphs

Nancy Clarke (Acadia University – Wolfville, CA)

We give two characterizations of the graphs on which k cops have a winning strategy in the game of Cops and Robber. These generalize the corresponding characterizations that are known in the one cop case. In particular, we give a relational characterization of k-copwin graphs, for all finite k, and then use this characterization to obtain a vertex elimination order characterization of such graphs. Instead of the elimination order being of the vertices of the given graph G as in the one cop case, it is an ordering of the vertices of the $(k+1)$-fold categorical product of G with itself. Most of our results hold for variations of the game and some of them extend to infinite graphs.

3.17 Some thoughts on constrained cops-and-robbers

Gena Hahn (Université de Montréal, CA)

This talk is essentially about open questions. First we propose a more general setting for cops-and-robbers games on graphs. Next, we suggest a way to model position and move constraints for the games and observe that there is a partially ordered set of constraints. We then ask what the structure of the poset might be, having observed that the theorem of Nowakowski and Winkler that characterizes cop-win graphs via a binary relation on the set of vertices carries over to the general setting. We close by suggesting that graphs that have some, but not all, loops should be studied, as well as tournaments, and propose a few problems.

3.18 Hypergraph searching as notion justification

Andrei Krokhin (University of Durham, GB)

We discuss a class of hypergraphs that appeared recently in the study of the constraint satisfaction problem. We show that this class can be described by a natural variant of the hypergraph searching game.

3.19 Monitoring on a Grid

Dieter Mische (UPC – Barcelona, ES)

We consider a set of g walkers W moving on the $n \times n$ integer grid. Initially, each walker chooses a vertex u.a.r., and in each step, each walker chooses u.a.r. and independently from the other walkers, one neighboring vertex. Moreover, we are given a set D of fixed devices, which are also placed on the integer grid. The devices are used to read data from walkers, and a device can read data of a walker if the walker is within a certain grid distance. We give bounds on the expected number of steps it takes to read data from all walkers for the case where all devices are put onto the halving line of the grid and for the case where all devices are regularly spread on the grid (in a grid-like way).

3.20 LIFO-search

Dimitrios M. Thilikos (National and Kapodistrian University of Athens, GR)

We study a variant of classic fugitive search game called LIFO-search where searchers are ranked, i.e. are assigned different numbers. The additional rule is that a searcher can be removed only if no searchers of lower rank are in the graph at that moment. We introduce the notion of shelters in graphs and we prove a min-max theorem implying their equivalence with the tree-depth parameter. As shelters provide escape strategies for the fugitive, this implies the the LIFO-search game is monotone and that the LIFO-search parameter is equivalent with the one of tree-depth.

4 Open problems

4.1 Cops and Robbers, parameterized algorithms

Fedor V. Fomin and Petr Golovach

By making use of backtracking algorithm, it is possible to decide if k cops can win on an n-vertex graph in time $n^{O(k)}$. It is easy to show that if the treewidth a graph is at most t, then the cop number of G is at most $t+1$. Thus on graphs of constant treewidth computing the minimum number of cops sufficient to win can be done in polynomial time. What is the parameterized complexity of the problem parameterized by the treewidth of the graph? Similar questions can be asked about the parameterization by the clique-width, the genus, and by the size of the excluded minor. The cop number of a graph is bounded by functions of these parameters.

4.2 Computing edge and nodes search numbers on special graph classes

Pinar Heggernes

Let $es(G)$ and $ns(G)$ be the edge and node search numbers of a graph G, respectively. These parameters are closely related: $ns(G)-1 \leq es(G) \leq ns(G)+1$. In general, both parameters are NP-hard to compute, but there are families of graphs, like interval graphs, split graphs, and cographs, on which both parameters can be computed in polynomial time. Is there there a class $\mathcal{F}$ of graphs such that $ns(G)$ can be computed in polynomial time for every graph $G \in \mathcal{F}$, whereas computing the edge search number is NP-hard on $\mathcal{F}$? Natural candidate classes to look at are those on which the computation of node search number, or equivalently pathwidth, can be done in polynomial time, but no results are known on the computation of their edge search number. As a first case to consider, $ns(G)$ can be computed in polynomial

(even linear) time if G is a permutation graph. Is the edge search number of permutation graphs computable in polynomial time?

4.3 Connected node search number

Dimitrios M. Thilikos

Let $\mathbf{n}s(G)$ and $\mathbf{cns}(G)$ be the node and connected node search numbers of graph G. We denote by $\mathbf{mcns}(G)$, the monotone connected search number It was recently shown that

$$\mathbf{ns}(G) \leq \mathbf{cns}(G) \leq \mathbf{mcsn}(G) \leq 2 \cdot \mathbf{ns}(G).$$

Since deciding if $\mathbf{ns}(G) \leq k$ is fixed parameter tractable parameterized by k, this gives an FPT approximation algorithm for connected search number. Is deciding $\mathbf{cns}(G) \leq k$ or $\mathbf{mcns}(G) \leq k$ in FPT?

It is known that numbers $\mathbf{cns}(G)$ and $\mathbf{mcns}(G)$ can be different. How much can they be different? Is it correct that for almost all graphs $\mathbf{cns}(G)/\mathbf{mcns}(G) \to 1$ as the size of G goes to infinite?

It is believed that the parameter $\mathbf{cns}(G)$ is closed under contractions, i.e., contractions of edges do not make the parameter increase. Is there a formal proof of this? Is deciding $\mathbf{cns}(G) \leq k$ in NP?

4.4 Span-width

Isolde Adler

We say that a graph G has a span-width at most k, if there is a tree decomposition of G of width k such that every vertex belongs to at most $k+1$ bags. What is the parameterized complexity of deciding if the span-width of a graph is at most k, parameterized by k? Similar question for tree-spanners.

4.5 Cop number of toroidal graphs

Gena Hahn

The long standing conjecture of Schroeder is that the cop number of a graph of genus g is at most $g+3$.

It is known that for toroidal graph this number is at most 4. Do toroidal graphs have cop number at most 3 as conjectured by Andreae in 1986? Or is there a toroidal graph that actually needs 4 cops?

4.6 Kelly-width

Paul Hunter

The Kelly-width of a digraph D is defined as the minimum number of searchers required to catch an invisible, inert fugitive with a (fugitive-)monotone strategy. "Inert" means the fugitive is unable to move from a vertex unless a searcher is about to land on that vertex. When the fugitive is able to move, he may move along any directed path not occupied by a searcher. For a more precise definition, see Hunter & Kreutzer [4]. Digraphs of Kelly-width 1 are precisely the acyclic digraphs, and there is a known polynomial time algorithm for deciding if a digraph has Kelly-width at most 2, see [6].

Is deciding if the Kelly-width of digraph D is at most k in PTIME for any fixed $k \geq 3$?

4.7 Ray searching

Spyros Angelopoulos

In the m-lane *ray search* problem we are given a set of m semi-infinite lanes with a common origin O. A target is placed at an unknown ray at distance d from the origin. We seek a strategy that minimizes the worst-case ratio $cost/d$, where $cost$ denotes the overall distance traversed by the searcher up to the point it locates the target.

This problem has been solved in its deterministic variant by Gal [3]. The question of finding randomized strategies that minimize the worst-case ratio $E[cost]/d$ is not quite settled. In [5] a randomized strategy is presented which however is optimal only in the class of round-robin strategies. Can we find optimal randomized strategies without any restrictive assumptions?

4.8 Ratio of monotonicity

Stephan Kreutzer

Consider two version of searching on a directed graph:

- Inert invisible fugitive game.
- Visible fugitive game.

Both problems are known to be non-monotone. Is there a number d such that the ratio between monotone and non-monotone versions of these games is at most d? More generally, is there an FPT approximation of non-monotone via monotone parameters?

4.9 Best strategy to catch the drunk robber

Dimitrios M. Thilikos

So far, in all cop and robber game settings, cops where considered to be omniscient and lucky in the sense that they will always take the best decision in order to avoid, or delay, capture. An interesting topic would be to study the setting where the robber moves randomly and the cops are clever. This induces a mix between classic graph searching and random walks. An example is given below:

Given a graph such as a line, a cycle, or a (torodial) grid, or a 3-regular graph, consider a robber that chooses randomly its first position and then moves randomly in neighbor nodes of the graph. Assume also that there are so many cops as the searching number of the graph. The cops play first, may move simultaneously. The two parts play in rounds. The objective here is to compute the minimum, over all cop strategies, expected time of arrest of the the (drunk) robber.

Question 1. Are all the optimal strategies monotone in the sense that the expected capture time does not change if cop strategies are restricted to those that do not visit again an already searched location? (The question has some meaning even when the number of cops is smaller than the search number of the graph.)

Question 2. How the expected capture time changes when there are less cops than the cop number?

Question 3. What is the ratio between the expected capture time for a drunk robber and the maximum capture time a "sober" robber (i.e., one that makes its best to avoid capture). Is this ratio common for many (or even all) graphs? Is it a constant such as 2?

4.10 Escaping from random cops

Pierre Fraigniaud

The paper [1] analyzes a randomized cop-and-robber game on graphs. The cop and the robber do not see each other, unless they are on the same node, in which case the robber is caught. They both move along the edges of the graph, one edge per round, playing in turn. Given a randomized cop strategy, the escape length for that strategy is the worst case expected number of rounds it takes the cop to catch the robber, where the worst case is with regards to all (possibly randomized) robber strategies. Adler et al. [1] proposes a cop strategy with an escape length of $O(n \log D)$ in n-node diameter-D graphs. On the other hand, there is a trivial $\Omega(n)$ lower bound on the escape length.

Open problem: close the gap between the two bounds.

One restricted case that may deserve attention is the case where the cop is bounded to apply simple random walk. In that case, is the best strategy for the robber the one consisting in placing itself at the node with lowest steady state probability, and stay idle? Or, if the initial position of the cop given, is the best strategy for the robber the one consisting in placing itself at the node with highest hitting time, and stay idle?

References

1 M. Adler, H. Räcke, N. Sivadasan, C. Sohler, B. Vöcking, *Randomized Pursuit-Evasion in Graphs*, ICALP 2002: 901-912

2 T. Andreae, *On a pursuit game played on graphs for which a minor is excluded*, J. Combin. Theory Ser. B, 41 (1986), pp. 37–47.

3 S. Gal, *Minimax solutions for linear search problems*, SIAM J. on Applied Math., 27 (1974), pp. 17–30.

4 P. Hunter and S. Kreutzer, *Digraph measures: Kelly decompositions, games, and orderings*. Theor. Comput. Sci. 399(3): 206-219 (2008)

5 M-Y. Kao and Y. Ma and M. Sipser and Y.L. Yin, *Optimal constructions of hybrid algorithms*, Journal of Algorithms, 29 (1998), pp. 142–164.

6 D. Meister, J. A. Telle, M. Vatshelle, *Recognizing digraphs of Kelly-width 2*. Discrete Applied Mathematics 158(7): 741-746 (2010)

Participants

- Isolde Adler
Univ. Frankfurt am Main, DE
- Carme Alvarez
UPC – Barcelona, ES
- Spyros Angelopoulos
CNRS - Paris, FR
- Dietmar Berwanger
ENS – Cachan, FR
- Lélia Blin
Université d'Evry, FR
- Nancy Clarke
Acadia Univ. – Wolfville, CA
- Dariusz Dereniowski
Gdansk Univ. of Technology, PL
- Josep Diaz
UPC – Barcelona, ES
- Amalia Duch Brown
UPC – Barcelona, ES
- Fedor V. Fomin
University of Bergen, NO
- Pierre Fraigniaud
Univ. Paris-Diderot, CNRS, FR
- Petr Golovach
University of Durham, GB
- Gena Hahn
Université de Montréal, CA
- Pinar Heggernes
University of Bergen, NO
- Paul Hunter
University of Oxford, GB
- David Ilcinkas
Université Bordeaux, FR
- Marcin Jurdzinski
University of Warwick, GB
- Marcin Kaminski
University of Brussels, BE
- Athanasios Kehagias
Aristotle University of Thessaloniki, GR
- Stephan Kreutzer
University of Oxford, GB
- Andrei Krokhin
University of Durham, GB
- Margaret-Ellen Messinger
Mount Allison University – Sackville, CA
- Zoltan Miklos
EPFL – Lausanne, CH
- Dieter Mitsche
UPC – Barcelona, ES
- Nicolas Nisse
INRIA Sophia Antipolis, FR
- Jan Obdrzálek
Masaryk University, CZ
- Xavier Perez Gimenez
MPI für Informatik – Saarbrücken, DE
- Pawel Pralat
West Virginia Univ. – Morgantown, US
- Maria Serna
UPC – Barcelona, ES
- Dimitrios M. Thilikos
National and Kapodistrian University of Athens, GR
- Tomas Valla
Charles University – Prague, CZ
- Erik Jan van Leeuwen
University of Bergen, NO
- Peter Widmayer
ETH Zürich, CH
- Boting Yang
University of Regina, CA

Report from Dagstuhl Seminar 11081

Combinatorial and Algorithmic Aspects of Sequence Processing

Edited by

Maxime Crochemore[1], Lila Kari[2], Mehryar Mohri[3], and Dirk Nowotka[4]

1 **King's College London, GB,** maxime.crochemore@kcl.ac.uk
2 **University of Western Ontario, London, CA,** lila@csd.uwo.ca
3 **New York University, US,** mohri@cs.nyu.edu
4 **University of Stuttgart, DE,** nowotka@fmi.uni-stuttgart.de

Abstract

Sequences form the most basic and natural data structure. They occur whenever information is electronically transmitted (as bit streams), when natural language text is spoken or written down (as words over, for example, the latin alphabet), in the process of heredity transmission in living cells (as DNA sequence) or the protein synthesis (as sequence of amino acids), and in many more different contexts. Given this universal form of representing information, the need to process strings is apparent and actually a core purpose of computer use. Algorithms to efficiently search through, analyze, (de-)compress, match, learn, and encode/decode strings are therefore of chief interest. Combinatorial problems about strings lie at the core of such algorithmic questions. Many such combinatorial problems are common in the string processing efforts in the different fields of application.

Scientists working in the fields of *Combinatorics on Words*, *Computational Biology*, *Stringology*, *Natural Computing*, and *Machine Learning* were invited to consider the seminar's topic from a wide range of perspectives. This report documents the program and the outcomes of Dagstuhl Seminar 11081 "Combinatorial and Algorithmic Aspects of Sequence Processing".

Seminar 21.–25. February, 2011 – www.dagstuhl.de/11081
1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems, F.4.3 Formal Languages, G.2.1 Combinatorics, I.2.6 Learning, J.3 Life and Medical Sciences
Keywords and phrases Combinatorics on words, computational biology, stringology, natural computing, machine learning
Digital Object Identifier 10.4230/DagRep.1.2.47

1 Executive Summary

Maxime Crochemore
Lila Kari
Mehryar Mohri
Dirk Nowotka

The object of concern of this seminar, *sequences*, implies a large degree of generality. It plays an essential rôle in many fields and constitutes a true cross section area. Hence, the seminar was designed to bring together researchers from different disciplines whose interest

are string processing algorithms and related combinatorial problems on words. Scientists working in the following fields were invited to consider the seminar's topic from a wide range of perspectives:

- Combinatorics on Words,
- Computational Biology,
- Stringology,
- Natural Computing,
- Machine Learning.

This Dagstuhl seminar was attended by 40 researchers from 13 countries. Everyone of the five topics above was about equally represented. Given the extremely interdisciplinary approach of this meeting it was an obvious necessity to hold a tutorial on each one of the participating research areas. These tutorials were held over the first and the morning of the second seminar day (see the scientific schedule below). They provided a good introduction for the non-specialists and triggered the first scientific discussions and exchanges.

A second (and standard) element of this seminar were regular talks, of course. A total of 15 talks were presented. It has to be noted that one could experience a very productive atmosphere during the whole seminar. All talks were well-attended and accompanied with interesting comments. Plenty of time was reserved for questions and discussions which was actively used by the participants.

The third element of the seminar were open problem sessions which did yield a larger attention to a range of problems, only some of them are included in this report. These open problem sessions provided the ideal ground for the ignition of new research lines and cooperations. Just to mention one example, the paper "On the regularity of iterated hairpin completion of a single word" (arXiv:1104.2385v1) resulted from the collaboration of Steffen Kopecki and Shinnosuke Seki initiated at this Dagstuhl seminar. In the light of such developments, it can be safely claimed that this seminar was a success.

Given the quality of presentations on this seminar and the constructive intensity of discussions between and after the talks, it is self-evident that follow-ups will be attempted. After this initial meeting of different communities, where common problems were identified, personal contacts established and first cooperations initiated, further events can be sharpened in focus and more on particular cross section topics regarding combinatorial and algorithmic problems in sequence processing.

Finally, we would like to say that the organization of a meeting for researchers of so unusually diverse fields bears a certain risk. However, it can be said that the event turned out better than expected. It was more than worthwhile to have taken that risk. We are grateful to all participants for their contributions to this successful seminar as well as to the staff of Schloss Dagstuhl for their perfect service.

2 Table of Contents

3 Overview of Tutorials

3.1 Data Structures for Text Indexing and String Algorithms

Roberto Grossi (University of Pisa, IT)

URL http://www.dagstuhl.de/mat/Files/11/11081/11081.GrossiRoberto.Slides.pdf

This is an introductive tutorial to the basic data structures employed in stringology: tries, compact tries, suffix trees, suffix arrays, and suffix automata. The tutorial considers also the case of large texts, discussing the external-memory model and the cache-oblivious model, with examples for suffix arrays, suffix trees, and string B-trees.

3.2 Natural Computing Tutorial

Hendrik Jan Hoogeboom (Leiden University, NL)

URL http://www.dagstuhl.de/mat/Files/11/11081/11081.HoogeboomHendrikJan.Slides.pdf

This is an overview of some concepts in the field of Molecular Computing (aka. DNA Computing): Adlemans experiment, TicTacToe computer (and beyond), evolutionary DNA computing, self assembly, bio-inspired formal models (splicing systems, new operations, membrane computing), nature as computer (gene assembly in ciliates).

3.3 Introduction to Sequence Learning

Mehryar Mohri (New York University, US)

Joint work of Corinna Cortes; Mehryar Mohri

This tutorial presents an introduction to sequence learning. This includes a brief presentation of binary classification problems and solutions based on large-margin hyperplanes and kernel methods, and a detailed discussion of sequence kernels. In particular, we describe a general framework based on rational kernels, give the proof of the positive-definiteness of a general class of rational kernels, show how general families of count-based kernels can be defined using rational kernels, and give a variety of examples of PDS rational kernels relevant to computational biology and text and speech processing.

We also present a number of general open problems related to a faster computation of sequence kernels and to the characterization of the class of languages learnable with rational kernels. Finally, we discuss the problem of learning sequence kernels and that of determining more efficient optimization solutions using sequence kernels.

3.4 Combinatorics on Words: An Introduction

Jeffrey Shallit (University of Waterloo, CA)

URL http://www.cs.uwaterloo.ca/ shallit/Talks/introcw.pdf

In this talk I surveyed some of the main themes in combinatorics on words: periodicity, patterns and pattern avoidance, equations in words, and infinite words and their properties. Among other things, I covered the Lyndon-Schutzenberger theorems, primitive words, conjugates, Lyndon words, fractional powers, unbordered words, Duval's conjecture, the Fine-Wilf theorem, Sturmian words, the Thue-Morse sequence, construction of a square-free infinite word, Dejean's conjecture, avoidance of abelian powers, Makanin's algorithm, Plandowski's PSPACE results, subword complexity, automatic sequences, and Christol's theorem.

4 Overview of Talks

4.1 Intelligent Strategies for Remote Homology Detection

Juliana Bernardes (UPMC – Paris, FR)

Main reference Bernardes J.S, Carbone A. and Zaverucha Gerson. A discriminative method for family-based protein remote homology detection that combines inductive logic programming and propositional models. BMC Bioinformatics 2011, 12:83
URL http://www.biomedcentral.com/1471-2105/12/83

Remote homology detection is a hard computational problem. Most approaches have trained computational models by using either full protein sequences or multiple sequence alignments (MSA), including all positions. However, when we deal with proteins in the "twilight zone" we can observe that only some segments of sequences (motifs) are conserved. We introduce a novel logical representation that allows us to represent physico-chemical properties of sequences, conserved amino acid positions and conserved physico-chemical positions in the MSA. From this, Inductive Logic Programming (ILP) finds the most frequent patterns (motifs) and uses them to train propositional models, such as decision trees and support vector machines (SVM). Our results show that our methodology when using SVM performs significantly better than some of the state of the art methods, and comparable to other. However, our method provides a comprehensible set of logical rules that can help to understand what determines a protein function.

The strategy of selecting only the most frequent patterns is effective for the remote homology detection. This is possible through asuitable first-order logical representation of homologous properties, and through a set of frequent patterns, found by an ILP system, that summarizes essential features of protein functions.

4.2 Simple Real-Time Constant-Space String-Matching

Dany Breslauer (University of Haifa, IL)

Joint work of Dany Breslauer; Roberto Grossi; Filippo Mignosi
Main reference D. Breslauer R. Grossi, and F. Mignosi. Simple Real-Time Constant-Space String-Matching. 22nd Annual Symposium on Combinatorial Pattern Matching (CPM), 2011.

We use a simple observation about the locations of critical factorizations to derive a real-time variation of the Crochemore-Perrin constant-space string-matching algorithm. The real-time variation has a simple and efficient control structure.

4.3 Combinatorial Measure of Co-evolving Blocks and theirEvolutionary Pressure

Linda Dib (UPMC – Paris, FR)

Co-evolution signals have been detected on a few divergent protein families while families of conserved protein sequences remain untractable by current methods. A large scale investigation of residue networks can only be made with the development of refined methods treating conserved sequences as well. We propose a new combinatorial approach to overcome this difficulty. Based on the observation that co-evolving positions are usually not isolated and that their co-evolving behaviour concerns adjacent positions as well, our combinatorial method, named Blocks In Sequences (BIS), studies co-evolution of blocks of contiguous positions in sequences, where a block is possibly constituted by a single position. BIS determines whether blocks of residues co- evolve or not and at which strength. BIS can be applied to sets of very conserved sequences, possibly made by a few sequences, and yet it is able to detect positional differences between these sequences and evaluate possi- ble signals of co-evolution. BIS captures important information on folding processes. It gives no hint on the kinetics but rather on the actors (that is, residues, parts of secondary structures, 3D interactions) of the kinetics process. The level of importance of these actors is encoded on the strength of the co-evolution signal. This strength is measured by a symmetric signal coming from residue pairs and by the ressemblance of residues in a network with their environment, but also by the combinatorics of the relationships possibly existing between networks that can be highlighted by the method. Network overlapping and connected components of the associated interval graph are used to bring up the intricated structure of co-evolving networks. Results obtained by BIS on the Protein A were compared to Φ-analysis and the outcomes are remarkably similar.

4.4 Non-Archimedian Words

Volker Diekert (University of Stuttgart, DE)

Non-Archimedean words have been introduced as a new type of infinite words which can be investigated through classical methods in combinatorics on words due to a length function. The length function, however, takes values in the additive group of polynomials $\mathbb{Z}[t]$ (and not, as traditionally, in $\mathbb{N}$), which yields various new properties. Non-Archimedean words allow to solve a number of interesting algorithmic problems in geometric and algorithmic group theory. There is also a connection to logic and the first-order theory in free groups (Tarski Problems).

In my lecture I report on a joint work with Alexei Miasnikov. We provide a general method to use infinite words over a discretely ordered abelian group as a tool to investigate certain group extensions for an arbitrary group G. The central object is a group $E(A, G)$, which is defined in terms of a non-terminating, but confluent rewriting system. The group G as well as some natural HNN-extensions of G embed into $E(A, G)$ (and still "behave like" G), which makes it interesting to study its algorithmic properties.

The main result characterizes exactly when the Word Problem is decidable in all finitely generated subgroups of $E(A, G)$. We show that this property holds if and only if the Cyclic Membership Problem "u in $\langle v \rangle$?" is decidable for all v in G.

The results combine methods from combinatorics on words, string rewriting and group theory.

4.5 Fixed Points of Nontrivial Morphisms

Štěpán Holub (Charles University – Prague, CZ)

Main reference Holub, Štěpán, Polynomial algorithm for fixed points of nontrivial morphisms, Discrete Mathematics 309: 5069-5076 (2009)
URL http://dx.doi.org/10.1016/j.disc.2009.03.019

In this talk I presented an algorithm that decides whether given finite word is a fixed point of a nontrivial morphism. Such words are also called *morphically imprimitive*.

These words were characterized in [1] and [2]. In [3], the question was raised about the time complexity of the decision problem as well as about the possibility to find the corresponding morphism.

The algorithm works in a subquadratic time and outputs a morphism that is in a good sense unique minimal witness of the word being imprimitive.

I also mentioned a related open problem, known as The Conjecture of Billaud. Let w be a finite word, and let δ_x denote the morphism canceling the letter x and otherwise acting as the identity. The conjecture states that if $\delta_x(w)$ is imprimitive for all x from the alphabet of w, then also w is imprimitive. Some partial results regarding this conjecture can be found in [1].

References

1 T. Head. *Fixed languages and the adult languages of OL schemes.* Int. J. Comput. Math. 10(2) (1981) 103-107

2 D. Hamm and J. Shallit. *Characterization of finite and one-sided infinite fixed points of morphisms on free monoids*. Technical Report CS-99-17, University of Waterloo, July 1999
3 D. Reidenbach and J. C. Schneider. *Morphically primitive words*. Theor. Comput. Sci. 410 (2009) 2148-2161
4 F. Levé and G. Richomme. *On a conjecture about finite fixed points of morphisms*. Theor. Comput. Sci. 339(1)(2005) 103-128

4.6 Observations and Problems on k-abelian Avoidability

Juhani Karhumäki (University of Turku, FI)

Joint work of Mari Huova; Juhani Karhumäki
Main reference M. Huova and J. Karhumäki, Observations and Problems on k-abelian Avoidability, Dagstuhl Preprint Archive
URL http://arxiv.org/abs/1104.4273v1

We introduce new avoidability problems for words by considering equivalence relations, k-abelian equivalences, which lie properly in between equality and commutative equality, i.e. abelian equality. For two k-abelian equivalent words the numbers of occurrences of different factors of length k coincide and the prefixies (resp. suffixies) of length $k-1$ are equal as well.

The size of the smallest alphabet avoiding 2-repetitions of words, i.e. squares, is three and for abelian squares it is four. It follows that for 2-abelian squares this size has to be three or four. Similarly, the size of the smallest alphabet where 2-abelian cubes, i.e. 3-repetitions, can be avoided is two or three, because cubes (resp. abelian cubes) are avoidable in binary (resp. ternary) alphabet.

We show that for 2-abelian squares the required size is four, as in the case of abelian squares. The longest 2-abelian square-free ternary word is of length 537. The question for 2-abelian cubes is open. Though, we have computational evidence that the size would be two, since there exists 2-abelian cube-free binary word of length 100 000, meaning that the 2-abelian case would behave like that of words.

4.7 Hairpin Completion versus Hairpin Lengthening

Steffen Kopecki (University of Stuttgart, DE)

URL http://www.dagstuhl.de/mat/Files/11/11081/11081.KopeckiSteffen.Slides.pdf

The hairpin completion and the hairpin lengthening are operations on formal languages that have been inspired by the hairpin formation in biochemistry. It is known that the hairpin completion (resp. hairpin lengthening) of a regular language is not in general regular but always linear context-free. As regularity of a (linear) context-free language is undecidable in general, we investigate the decidability problem whether the hairpin completion (resp. hairpin lengthening) of regular languages is regular again. For the hairpin completion we solved the problem positively in former papers. Even though both operations seem quite similar, we were not able to use the same approach for the hairpin lengthening. Here, we provide partial results on the decidability problem for the hairpin lengthening and discuss some differences between the two operations. To name one of them, the hairpin completion

of a regular language is an unambiguous (linear) context-free language, but the hairpin lengthening may lead to an inherent ambiguous (linear) context-free language.

References

1 D. Cheptea, C. Martın-Vide, and V. Mitrana. A new operation on words suggested by DNA biochemistry: Hairpincompletion. in Proc. of *Transgressive Computing*, 216–228 (2006)
2 F. Manea, C. Martın-Vide, and V. Mitrana. Hairpin lengthening. Proc. of *CiE*, LNCS 6158, 296–306 (2010)
3 V. Diekert and S. Kopecki. It is NL-complete to decide whether a hairpin completion of regular languages is regular. CoRR, abs/1101.4824, 2011.
4 S. Kopecki. On the iterated hairpin completion. CoRR, abs/1010.3640, 2011.

4.8 Sequence and Chromatin Signatures Predict Transcription Factor Binding in the Human Genome

Christina Leslie (Memorial Sloan-Kettering Cancer Center – New York, US)

Gene regulatory programs are orchestrated by proteins called transcription factors (TFs), which coordinate expression of target genes both through direct binding to genomic DNA and through interaction with cofactors. Accurately modeling the DNA sequence preferences of TFs and predicting their genomic binding sites are key problems in regulatory genomics. These efforts have long been frustrated by the limited availability and accuracy of TF binding site motifs. Today, protein binding microarray (PBM) experiments and chromatin immunoprecipitation followed by sequencing (ChIP-seq) experiments are generating unprecedented high-resolution data on in vitro and in vivo TF binding. Moreover, genome-wide data on the cell-type specific chromatin state, including ChIP-seq experiments that profile histone modifications associated with active or inactive transcriptional states, provide additional information for predicting the genomic binding locations of TFs.

We will present a flexible new discriminative framework for representing and learning TF binding preferences using these massive data sets. We will first describe in vitro models of TF-DNA sequence affinities, where we train support vector regression (SVR) models with a novel string kernel on PBM data to learn the mapping from probe sequences to binding intensities. In a large data set of over 180 yeast and mouse TF binding experiments, our SVR models better predicted in vitro binding than popular motif discovery approaches or methods based on enrichment of k-mer patterns.

We will then show how to train kernel-based SVM models directly on TF ChIP- seq data to learn in vivo TF sequence models and present results from a large- scale evaluation on 184 TF ChIP-seq experiments from ENCODE. We confirmed that our discriminative sequence models significantly outperform existing motif discovery algorithms, and we found that ChIP-trained models greatly improved TF occupancy prediction over PBM-trained models, suggesting distinct in vivo sequence information (e.g. binding sites of cofactors). Finally, we trained discriminative chromatin models using histone modification ChIP-seq data and found that models combining sequence and chromatin signatures strongly outperformed using either one alone. We found that relatively few TFs in our study had pronounced cell-type specific binding patterns, but in those that did, we identified cell-type dependent sequence information. This work establishes effective new techniques for analyzing next generation

sequencing data sets to study the interplay of chromatin and sequence in TF binding in the human genome.

4.9 Some algorithmic and combinatorial problems in the RNA and DNA world

Jan Maňuch (Simon Fraser University – Burnaby, CA)

I have presented 3 problems in computational biology. The first problem, the energy barrier problem for pseudoknot-free RNA structures, asks whether what's the minimal energy barrier needed to transform from the initial to the final structure of an RNA molecule. We consider the simplest energy model in which only the number of basepairs is taken into account and transformation sequence removes and adds the basepairs in some order. We showed that the problem is NP-complete [1], but can be solved in polynomial time if the barrier is assumed to be constant [2]. I have also introduced a string displacement system which can model multi-strand scenario and can be modeled by a simple rewriting system with two types of strings (signals and templates).

The second and third problems arise in the DNA synthesis. DNA strand needs to be assembled from a shorter factors of the strand which avoid many types of collisions. The simplest type of collision is equality, which leads to the following word problem. Given a word on alphabet Sigma and integer k, is it possible to partition it to distinct factors of length at most. We show that this is NP-complete for alphabet size 4 [3]. We also consider other conditions which factors need to satisfy, e.g., prefix-freeness, factor-freeness, etc, and show that in all cases the problem is NP-complete even for the binary alphabet.

The third problem goes one step back. Starting from a protein sequence (sequence of amino acids), the task is to find the DNA sequence which is mapped that protein sequence and satisfies 2 constraints. One possible algorithm based on acyclic DFA leads to polynomial algorithm, however, its complexity is $O(n^{42})$. Is there a more efficient algorithm?

References

1 Maňuch, J., Thachuk, C., Stacho, L., Condon, A., NP-completeness of the energy barrier problem without pseudoknots and temporary arcs, *Nat. Comput.* **10**, No. 1, 391–405 (2011)
2 Thachuk, C., Maňuch, J., Rafiey, A., Mathieson, L-A., Stacho, L., Condon, A., An algorithm for the energy barrier problem without pseudoknots and temporary arcs, Proc. of *Pacific Symposium on Biocomputing* (PSB, Hawaii, USA, 2010), World Scientific Publishing, 108–119 (2010)
3 Condon, A., Maňuch, J., Thachuk, C., Complexity of a collision-aware string partition problem and its relation to oligo design for gene synthesis, Proc. of *Annual International Computing and Combinatorics Conference* (COCOON, Dalian, China, 2008), LNCS **5092**, 265–275 (2008).

4.10 Exact ensemble properties in combinatorial dynamic programming schemes

Yann Ponty (Ecole Polytechnique – Palaiseau, FR)

Joint work of Yann Ponty; Cédric Saule
Main reference Y. Ponty and C. Saule. A combinatorial framework for the design of (pseudoknotted) RNA algorithms, *WABI*, 2011.
URL http://www.lix.polytechnique.fr/ ponty/docs/EnsembleHypergraphsDP.pdf

We extend an hypergraph representation, introduced by Finkelstein and Roytberg, to unifydynamic programming algorithms in the context of RNA folding with pseudoknots. Classic applicationsof RNA dynamic programming (Energy minimization, partition function, base-pair probabilities ...) are reformulated within this framework, giving rise to very simple algorithms. This reformulationallows one to conceptually detach the conformation space/energy model — captured by thehypergraph model — from the specific application, assuming unambiguity of the decomposition. Toensure the latter property, we propose a new combinatorial methodology based on generating functions.We extend the set of generic applications by proposing an exact algorithm for extracting generalizedmoments in weighted distribution, generalizing a prior contribution by Miklos and al. Finally,we illustrate our full-fledged programme on three exemplary conformation spaces (secondary structures,Akutsu's simple type pseudoknots and kissing hairpins). This readily gives sets of algorithmsthat are either novel or have complexity comparable to classic implementations for minimizationand Boltzmann ensemble applications of dynamic programming.

4.11 On the Structure of Compacted Subword Graphs of Thue-Morse Words and Their Applications

Wojciech Rytter (University of Warsaw, PL)

Thue-Morse words (*TM* words, in short) form a famous family of words, due to many interesting properties related not only to text algorithms and combinatorics on words, but also to other disciplines, see [1]. In particular they do not contain factors of type $axaxa$, where a is a single letter (overlaps), consequently they do not contain cubes. A very good source for properties of these words is for example the book [3]. We rediscover/discover several known/unknown properties of *TM* words in a novel way: analyzing the subword graphs of finite and infinite *TM* words. This approach was already successfully applied by one of the authors to another well-known family of words, namely the Fibonacci words [8]. We also study how the cdawg of the infinite *TM* word is related to an infinite graph with 2-counting property and a numeration system, similar analysis for Fibonacci words and, in general, Sturmian words can be found in [7].

The structure of cdawg of a word w is closely related to right special factors of w (defined later on in the text). Such factors of *TM* words were already studied thoroughly in relation to the subword complexity function of the infinite *TM* word (i.e., the number of distinct factors of the word of a given length), see [4, 6, 9].

On the other hand, the vertices of cdawg of w can be seen as bispecial factors of w; bispecial factors of the infinite *TM* word are characterized in [2, 5].

Using the special structure of cdawgs we present several unknown properties of Thue-Morse words as well as new (graph-based) proofs of some well-known properties. A slight modification of the compact dawg of the infinite Thue-Morse word yields an infinite graph with 2-counting property.

References

1 J.-P. Allouche and J. Shallit. The ubiquitous Prouhet-Thue-Morse sequence. *Springer Ser. Discrete Math. Theor. Comput. Sci.*, pages 1–16, 1999.

2 L. Balkova, E. Pelantova, and W. Steiner. Return words in the Thue-Morse and other sequences. *arxiv:math/0608603v2*, 2006.

3 J. Berstel, A. Lauve, C. Reutenauer, and F. V. Saliola. *Combinatorics on Words: Christoffel Words and Repetitions in Words*. Amer. Mathematical Society, 2009.

4 S. Brlek. Enumeration of factors in the Thue-Morse word. *Discrete Applied Mathematics*, 24(1-3):83–96, 1989.

5 A. de Luca and L. Mione. On bispecial factors of the Thue-Morse word. *Inf. Process. Lett.*, 49(4):179–183, 1994.

6 A. de Luca and S. Varricchio. Some combinatorial properties of the Thue-Morse sequence and a problem in semigroups. *Theor. Comput. Sci.*, 63(3):333–348, 1989.

7 C. Epifanio, F. Mignosi, J. Shallit, and I. Venturini. On Sturmian graphs. *Discrete Applied Mathematics*, 155(8):1014–1030, 2007.

8 W. Rytter. The structure of subword graphs and suffix trees of Fibonacci words. *Theor. Comput. Sci.*, 363(2):211–223, 2006.

9 J. Tromp and J. Shallit. Subword complexity of a generalized Thue-Morse word. *Inf. Process. Lett.*, 54(6):313–316, 1995.

4.12 Enumeration and Automatic Sequences

Jeffrey Shallit (University of Waterloo, CA)

Main reference E. Charlier, N. Rampersad, and J. Shallit, Enumeration and decidable properties of automatic sequences, preprint, available at `http://arxiv.org/abs/1102.3698`.

URL http://arxiv.org/abs/1102.3698

In this talk, I discussed some new results, obtained with E. Charlier and N. Rampersad, on the decidability of properties of automatic sequences.

Here is a brief summary of our results:

1. Given a k-automatic sequence $\mathbf{x} = a(0)a(1)a(2)\cdots$, the sequence $\mathbf{b} = b(0)b(1)b(2)\cdots$ defined by $b(n) = 1$ if $\mathbf{x}$ has an unbordered factor of length n and 0 otherwise, is also k-automatic.

2. The following questions are decidable:

(a) given a k-automatic sequence, does it contain powers of arbitrarily large exponent?

(b) given a k-automatic sequence, does it contain arbitrarily large unbordered factors?

(c) given a k-automatic sequence, is it recurrent? linearly recurrent?

3. Many sequences counting properties of k-automatic sequences are k-regular, and constructively so. These include

(a) the number of distinct factors of length n;

(b) the number of palindromic factors of length n;

(c) the number of unbordered factors of length n;

and many other examples.

References

1 E. Charlier, N. Rampersad, and J. Shallit, Enumeration and decidable properties of automatic sequences, preprint, available at `http://arxiv.org/abs/1102.3698`.

4.13 Context Equivalence Problem

Arseny M. Shur (Ural State Univ. – Ekatarinenburg, RU)

The problem we study is a natural decision problem on words w.r.t. an arbitrary fixed language L. The instance is a pair of words; the problem is to decide whether these two words are "equally placed" in L in the sense that they have exactly the same contexts in L. By a context of a word v w.r.t. L we mean a pair (x, z) of words such that the word xvz belongs to L. From an algebraic point of view, the context equivalence problem is exactly the word problem in the syntactic monoid of the language L. It is little known about decidability and complexity of this interesting problem. Solving it for a given language L, we can get a lot of information about the internal structure of L.

We briefly explain some cases when this problem can be easily solved either because the language L is simple (the case of regular languages) or because the solution is trivial (the case of uniformly recurrent languages). Then we present a sophisticated, but linear-time solution of the context equivalence problem for the language of binary overlap-free words. Finally, we shortly discuss this problem for other power-free languages.

References

1 A. V. Klepinin. *On syntactic congruences of uniformly recurrent languages*, Proc. Ural State Univ. Ser. Computer Science. 2006. Vol. 1 (43). P. 38–44. [Russian]

2 A. M. Shur. *Syntactic semigroups of avoidable languages*, Sibirskii Matematicheskii Zhurnal. 1998. Vol. 39(3). P. 683–702. [Russian; Engl. Transl. in Siberian Math. J. 1998. Vol. 39(3). P. 594–610.]

3 A. M. Shur. *Deciding context equivalence of binary overlap-free words in linear time*, Semigroup Forum. 2011. (Submitted)

4.14 Hashing for Strings

Alexander J. Smola (Yahoo! Research – Santa Clara, US)

In this talk I discuss applications of hashing to fast computation of string similarity measures. For this purpose I first give an overview over string kernels using Suffix Trees, then I will discuss how hashing can deal with the problem of an ever increasing memory footprint for

suffix trees, simply by allowing collisions between its vertices. Applications to personalized spam filtering and approximate matching are provided to show the feasibility of this approach in practice.

References

1 Q. Shi, J. Petterson, G. Dror, J. Langford, A. J. Smola, A. L. Strehl, and V. Vishwanathan. Hash kernels. *Journal of Machine Learning Research - Proceedings Track*, 5:496–503, 2009.
2 S. V. N. Vishwanathan and Alexander J. Smola. Fast kernels for string and tree matching. In S. Becker, S. Thrun, and K. Obermayer, editors, *NIPS*, pages 569–576. MIT Press, 2002.
3 K. Q. Weinberger, A. Dasgupta, J. Attenberg, J. Langford, and A. J. Smola. Feature hashing for large scale multitask learning. *CoRR*, abs/0902.2206, 2009.

5 Open Problems

5.1 Word Equations with Loops

Štěpán Holub (Charles University – Prague, CZ)

In this talk I gave a short introduction into the problem of independent equivalent subsystems of word equations.

It is known (see [1] and [2]) that each infinite system of equations over a finite set of unknowns has a finite equivalent subsystem (where "equivalent" means having the same set of solutions).

On the other hand, little is known about possible size of such equivalent finite subsystems. It is even not known whether the size is bounded by a function in the number of unknowns.

I listed several examples where the bounds are known, including equations with particular looping properties, see [3] or [4].

In this framework I presented an open problem asking whether the following system of equations has a nontrivial solution for some number n of unknowns:

$$(x_1 \cdot x_2 \cdots x_n)^2 = x_1^2 \cdot x_2^2 \cdots x_n^2$$
$$(x_1 \cdot x_2 \cdots x_n)^3 = x_1^3 \cdot x_2^3 \cdots x_n^3$$

It is known that the answer is negative if, in addition to previous two equalities, also

$$(u_1 \cdot u_2 \cdots u_n)^4 = u_1^4 \cdot u_2^4 \cdots u_n^4$$

is required. For more details and bibliography see [5].

References

1 V. S. Guba. *The equivalence of infinite systems of equations in free groups and semigroups with finite subsystems*. Mat. Zametki, 40 (1986) 321-324
2 M. H. Albert and J. Lawrence. *A proof of Ehrenfeucht conjecture*. Theoret. Comput. Sci., 41 (1985) 121–123
3 Š. Holub, *Local and global cyclicity in free semigroups*. Theoret. Comput. Sci., 262 (2001) 25-36
4 Š. Holub and J. Kortelainen. *On systems of word equations with simple loop sets*. Theoret. Comput. Sci., 380 (2007) 363-372
5 http://www.karlin.mff.cuni.cz/~holub/soubory/prizeproblem.pdf

5.2 Is morphic primitivity hereditary?

Štěpán Holub (Charles University – Prague, CZ)

A word w is said to be morphically imprimitive if there is a nontrivial morphism f such that $f(w) = w$. Let δ_x, where x is a letter occurring in w, denote the morphism canceling x and being the identity on all other letters.

Prove or disprove the following claim, known as The Conjecture of Billaud:

If $\delta_x(w)$ is morphically imprimitive for all x occurring in w, then also w is imprimitive.

For more information and some partial results see [1].

References

1 F. Levé and G. Richomme. *On a conjecture about finite fixed points of morphisms*. Theor. Comput. Sci. 339(1)(2005) 103-128

5.3 Asymptotic Number of Long-Armed Palindromes in a Word

Gregory Kucherov (Université de Marne-la-Vallée, FR)

Such a palindrome is called *long-armed* if $|v| \geq |u|$. Given a word w of length n, we are interested in all subwords of w that are long- armed palindromes. In [1], an algorithm has been proposed for computing all long-armed palindromes in time $O(n + S)$, where S is the size of the output, i.e. the number of long-armed palindromes found.

However, it is not known whether this number is linearly-bounded in n.

Trivially, for any fixed gap size ($|u|$) there can be a linear number of corresponding palindromes, as every position (or every letter) of w can be the center of only one palindrome. In a private communication after the Dagstuhl seminar, Jeffrey Shallit and Michael Domaratzki provided an example of a word containing order $3n$ long-armed palindromes. On the other hand, it is very easy to see that the number of long-armed palindromes if $O(n \log n)$. Proving (or refuting) the linear bound remains an open problem.

References

1 R. Kolpakov and G. Kucherov. Searching for gapped palindromes. *Theoretical Computer Science*, 410(51):5299–5382, 28 November 2009.

5.4 The Separating Words Problem

Jeffrey Shallit (University of Waterloo, CA)

Main reference E. D. Demaine, S. Eisenstat, J. Shallit, and D. A. Wilson, Remarks on Separating Words. *Preprint.*
URL http://arxiv.org/abs/1103.4513

In this talk I discussed the separating words problem, as introduced by Goralčík and Koubek in 1986.

In this problem we are given two words x and x of length $\leq n$, and we want a good bound for the size of the smallest DFA that accepts one of $\{w, x\}$ and rejects the other.

If $|w| \neq |x|$ then w and x can be separated by a DFA of $O(\log n)$ states, so the only interesting case is where $|w| = |x|$.

I mentioned two new results:

(1) If the Hamming distance between w and x is $< d$, then w and x can be separated using $O(d \log n)$ states.

(2) There exists a sequence of words w, x such that nondeterministic separation is arbitrarily better than deterministic separation.

5.5 Some open problems inspired by Dejean's conjecture

Arseny M. Shur (Ural State Univ. – Ekatarinenburg, RU)

Joint work of Irina A. Gorbunova; Alexei V. Samsonov; Arseny M. Shur

Among the repetitions in words, fractional powers constitute one of the most natural classes. Any word, in which some prefix coincides with some suffix, is a fractional power. Each such power is numerically expressed by its exponent: $\exp(w) = |w|/\pi(w)$, where $|w|$ is the length of the word w, and $\pi(w)$ is the smallest period of w. A word v is β-free if no one of its factors has the exponent $\geq \beta$, and is β^+-free if no one of its factors has the exponent $> \beta$. The repetition threshold on k letters, $RT(k)$, is the rational number γ_k such that the number of γ_k^+-free k-ary words is infinite, while the number of γ_k-free k-ary words is only finite. The values of $RT(k)$ were conjectured by Dejean in 1972:

$$RT(3) = 7/4,\ RT(4) = 7/5,\ RT(k) = k/(k-1) \text{ for } k = 2 \text{ and } k \geq 5.$$

The conjecture is now confirmed in all cases, due to Dejean, Pansiot, Moulen-Ollagnier, Currie, Mohammad-Noori, Carpi, Rampersad, and Rao (1972 to 2009). The proof stimulated further research in several directions.

We mention three such directions and an open problem in each direction.

However, we are not pretending to present an exhaustive list of Dejean- like problems.

1. Strengthening original conjecture.

Problem: estimate the growth of the language of the k-ary $RT(k)^+$- free language.

Exponential conjecture (folklore): all these languages have exponential growth. Confirmed by Ochem for $k = 3, 4$.

Growth rate conjecture (first stated by the author at JM 2008; see [2]): exponential growth rates of these languages tend to the limit $\alpha \approx 1.242$ as k approaches infinity.

2. Different notion of words.

Problem: find the analog(s) of repetition threshold for circular words.

There are three possible definitions of circular repetition threshold $CRT(k)$: weak (there are infinitely many circular $CRT(k)^+$-free words), intermediate (there are circular $CRT(k)^+$-free words of all but finitely many lengths), and strong (there are circular $CRT(k)^+$-free words of any length). For $k = 2$, these bounds are 2, 7/3, and 5/2, respectively (Aberkane, Currie).

Conjecture for $k \geq 3$: weak and intermediate thresholds both coincide with $RT(k)$ (we have a proof for $k = 3$). The strong threshold is strictly bigger (for $k \geq 9$, it is at least $(k-3)/(k-4)$, as follows from the results of [3]).

3. Different notion of powers.

Problem: find the analog(s) of repetition threshold for Abelian powers.

We mention only one of several possible definitions of Abelian fractional powers. This definition was first given by Cassaigne and Currie and suits well for the powers less than 2. According to it, a word is Abelian β- free ($\beta < 2$) if it has no factors of the form xyz such that x and z are Abelian equivalent and $|xyz|/|xy| \geq \beta$.

Conjecture (first stated at JM 2010; see [1]): the Abelian repetition threshold (for the above definition of Abelian β-freeness) equals 9/5 for $k = 4$ and $(k-2)/(k-3)$ for $k \geq 5$.

References

1 A. V. Samsonov, A. M. Shur. *On Abelian repetition threshold*, Proc. 13th Mons Days of Theoretical Computer Science. Univ. de Picardie Jules Verne, Amiens, 2010. P. 1–11.

2 A. M. Shur, I. A. Gorbunova. *On the growth rates of complexity of threshold languages*, RAIRO Inform. Theor. Appl. 2010. Vol. 44. P. 175– 192.

3 A. M. Shur. *On the existence of minimal β-powers*, Proc. 14th Int. Conf. on Developments in Language Theory. Berlin: Springer, 2010. P. 411–422. (LNCS Vol. 6224).

6 Scientific Schedule

Monday

09:00–10:15 Tutorial: Combinatorics on Words – *Jeffrey Shallit*
10:30–12:00 Tutorial: Machine Learning – *Mehryar Mohri*
14:00–14:50 Hashing for Strings – *Alexander Smola*
16:00–17:15 Tutorial: Natural Computing – *Hendrik Jan Hoogeboom*

Tuesday

09:00–10:30 Tutorial: Bioinformatics – *Rolf Backofen*
10:50–11:45 Tutorial: Stringology (part 1) – *Roberto Grossi*
11:45–12:30 Tutorial: Stringology (part 2) – *Alessandra Carbone*
14:00–14:45 Non-Archimedian Words – *Volker Diekert*
15:00–15:15 Open Problem – *Gad Landau*
15:45–16:30 Sequence and Chromatin Signatures Predict Transcription Factor Binding in the Human Genome – *Christina Leslie*
16:30–17:00 Combinatorial Measure of Co-evolving Blocks and theirEvolutionary Pressure – *Linda Dib*
19:30–20:00 Open Problems – *Volker Diekert, Štěpán Holub, Dirk Nowotka*

Wednesday

09:30–10:00 Real-Time, Constant Space String Matching – *Dani Breslauer*
10:15–11:00 K-abelian Equivalence – *Juhani Karhumäki*
11:15–12:00 Hairpin Completion versus Hairpin Lengthening – *Steffen Kopecki*
12:00–12:15 Open Problems – *Štěpán Holub, Alexander Smola*

Thursday

09:30 - 10:15 Words and Permutations – *Antonio Restivo*
10:30 - 11:15 Context Equivalence Problem – *Arseny Shur*
11:30 - 12:15 Intelligent Strategies for Remote Homology Detection – *Juliana Bernardes*
14:00 - 14:45 The Structure of Graphs Representing All Subwords of Thue-Morse Sequences – *Wojciech Rytter*
14:45 - 15:30 Open Problems – *Gregory Kucherov, Jeffrey Shallit, Aseny Shur*
15:45 - 16:30 Exact Ensemble Properties in Combinatorial Dynamic Programming Schemes – *Yann Ponty*

Friday

09:30 - 10:15 Energy Barrier Problem without Pseudo Knots – *Jan Maňuch*
10:30 - 11:15 Polynomial Algorithm for Fixed Points of Nontrivial Morphisms – *Štěpán Holub*
11:15 - 12:00 Some Decidable Properties of Automatic Sequences – *Jeffrey Shallit*

Participants

- Cyril Allauzen
Google – New York, US
- Rolf Backofen
University Freiburg, DE
- Marie-Pierre Beal
Univ. de Marne-la-Vallée, FR
- Juliana Bernardes
UPMC – Paris, FR
- Dany Breslauer
University of Haifa, IL
- Alessandra Carbone
UPMC – Paris, FR
- Corinna Cortes
Google – New York, US
- James Currie
University of Winnipeg, CA
- Alessandro De Luca
University of Napoli, IT
- Linda Dib
UPMC – Paris, FR
- Volker Diekert
University of Stuttgart, DE
- Mike Domaratzki
University of Manitoba, CA
- Roberto Grossi
University of Pisa, IT
- Stepan Holub
Charles University – Prague, CZ
- Hendrik Jan Hoogeboom
Leiden University, NL
- Costas S. Iliopoulos
King's College – London, GB
- Juhani Karhumäki
University of Turku, FI
- Juha Karkkainen
University of Helsinki, FI
- Steffen Kopecki
University of Stuttgart, DE
- Gregory Kucherov
Univ. de Marne-la-Vallée, FR
- Gad Landau
University of Haifa, IL
- Thierry Lecroq
Université de Rouen – Mont-Saint-Aignan Cedex, FR
- Christina Leslie
Memorial Sloan-Kettering Cancer Center – New York, US
- Jan Manuch
Simon Fraser University – Burnaby, CA
- Mehryar Mohri
New York University, US
- Dirk Nowotka
University of Stuttgart, DE
- Enno Ohlebusch
Universität Ulm, DE
- Yann Ponty
Ecole Polytech. – Palaiseau, FR
- Svetlana Puzynina
University of Turku, FI
- Gunnar Rätsch
MPI für biologische Kybernetik – Tübingen, DE
- Narad Rampersad
University of Liége, BE
- Antonio Restivo
Universitá di Palermo, IT
- Wojciech Rytter
University of Warsaw, PL
- Shinnosuke Seki
Univ. of Western Ontario, CA
- Jeffrey Shallit
University of Waterloo, CA
- Arseny M. Shur
Ural State Univ. – Ekatarinenburg, RU
- Alexander J. Smola
Yahoo! Res. – Santa Clara, US
- Sören Sonnenburg
TU Berlin, DE
- German Tischler
Universität Würzburg, DE
- Chris J. Watkins
RHUL – London, GB

Report from Dagstuhl Seminar 11091

Packing and Scheduling Algorithms for Information and Communication Services

Edited by

Klaus Jansen[1], Claire Mathieu[2], Hadas Shachnai[3], and Neal E. Young[4]

1 **Universität Kiel, DE,** kj@informatik.uni-kiel.de
2 **Brown University - Providence, US,** claire@cs.brown.edu
3 **Technion - Haifa, IL,** hadas@cs.technion.ac.il
4 **Univ. California - Riverside, US,** neal@cs.ucr.edu

Abstract

From 27.02.2011 to 4.03.2011, the Dagstuhl Seminar 11091 "Packing and Scheduling Algorithms for Information and Communication Services" was held in Schloss Dagstuhl Leibniz Center for Informatics. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

Seminar 27.02–4.03.2011 – www.dagstuhl.de/11091
1998 ACM Subject Classification F.2.2 Sequencing and scheduling, G.4 Algorithm Design and Analysis, G.2 Discrete Mathematics, C.2 Computer-Communication Networks
Keywords and phrases Packing, scheduling, information and communication services, combinatorial optimization, mathematical programming, parameterized complexity
Digital Object Identifier 10.4230/DagRep.1.2.67

1 Executive Summary

Klaus Jansen
Claire Mathieu
Hadas Shachnai
Neal E. Young

Packing and scheduling are one area where mathematics meets puzzles. While many of these problems stem from real-life applications, they have also been of fundamental theoretical importance. In a *packing* problem given is a set of items and one or more (multi-dimensional) bins. The objective is to maximize the profit from packing a subset of the items, or to minimize the cost of packing *all* items. In a *scheduling* problem, given are a set of jobs and a set of machines. One needs to schedule the jobs to run on the machines (under some constraints) so as to optimize an objective function that depends on the order of the jobs, on their completion times or on the machines by which they are processed.

Storage allocation in computer networks, cutting stock problems in various industries and production planning are only few of the applications of packing and scheduling. With the

growing impact of next generation technologies in information and communication services (some examples are Video-on-Demand systems, web applications and wireless networks), practitioners as well as theoreticians seek fast and efficient solutions for new variants of some classic packing and scheduling problems, which are crucial for optimizing the performance of these systems.

Since many of these problems are NP-hard, it is natural to seek efficient approximate solutions. Traditionally, such approximations are obtained by using fundamental tools from combinatorial optimization and mathematical programming. While for some of the problems there exist algorithms which achieve the best possible approximation ratio, one major effort of this community has been to close the gaps in running times between heuristic solutions, which perform well in practice, and algorithms which are provably efficient in terms of approximation ratio, but impractical in use. The large class of approximation schemes for packing and scheduling problems has been the recent target of this effort.

Parameterized complexity uses refined measures for the approximability of a given problem, by referring, e.g., to approximation with instance parameters, by defining performance functions (instead of performance ratios) and by defining the quality of approximation as parameter. Such measures provide further insight to the studied problems and lead to the design of algorithms that work efficiently if the parameters of the input instance are small (even if the size of the input is large). Efficient parameterization for packing and scheduling problems is a major challenge on the way to obtaining practical algorithms.

During the 5 days of the seminar, 24 talks were given by the participants. Five of these talks were two-hour tutorials and 60-minute survey talks on various topics:Kirk Pruhs gave an exciting tutorial on the challenges faced by designers of algorithms for green computing; Dániel Marx talked about several existing connections between approximation algorithms and fixed-parameter algorithms; Ola Svensson gave an overview of the implications and techniques of two fascinating hardness of approximation results for shops and precedence constraints scheduling; Neal Young talked about using lagrangian-relaxation algorithms to solve packing and covering problems, and Magnús Halldórsson gave an overview of recent analytic work on scheduling wireless links.

The seminar successfully brought together both experts and newcomers from the areas of packing and sequencing, combinatorial optimization, mathematical programming, and parameterized complexity, with many interesting interactions. The talks left plenty of time for discussion in the afternoon. An open problem session was held on Tuesday, and problems raised there were discussed by different groups throughout the seminar and in a research groups session on Friday. A session on current and future trends in scheduling was held on Thursday, and brought up some exciting issues relating to this area.

2 Table of Contents

3 Overview of Talks

3.1 On Packing Resizable Items and Covering by Holes

Sivan Albagli-Kim (Technion, IL)

Joint work of Albagli-Kim, Sivan; Shachnai, Hadas; Tamir, Tami

In next generation video services, movie files can be transmitted to the clients at different encodings, thus enabling the system to degrade Quality- of-Service for bounded time intervals, while increasing the number of serviced clients. Maximizing throughput in such systems can be modeled as the following problem of packing a set of items, whose sizes may change over time. Given is a set I of unit-sized items and a bin of capacity $B > 1$. The items need to be packed in the bin for a fixed time interval. Each item j can be compressed to the size $p_j \in (0, 1)$ for at most a fraction $q_j \in (0, 1]$ of its packing time. The goal is to pack in the bin, for the given time interval, a subset of the items of maximum cardinality. This problem of *packing resizable items (PRI)* is strongly NP-hard already for highly restricted instances.

In this paper we present approximation algorithms for two subclasses of instances of PRI which are of practical interest. For instances with uniform compression ratio, we develop an asymptotic fully polynomial time approximation scheme. For instances with uniform compression time, we give an almost optimal algorithm, which packs at least $OPT(I) - 1$ items, where $OPT(I)$ is the number of items packed by an optimal algorithm. We derive our results by using a non-standard transformation of PRI to the problem of covering a region by sliceable rectangles. The resulting problem, which finds numerous applications in computational geometry, is of independent interest.

3.2 Secretary Problems via Linear Programming

Niv Buchbinder (Open Univ., IL)

Joint work of Buchbinder, Niv; Jain, Kamal; Singh, Mohit

In the classical secretary problem an employer would like to choose the best candidate among n competing candidates that arrive in a random order. This basic concept of n elements arriving in a random order and irrevocable decisions made by an algorithm have been explored extensively over the years, and used for modeling the behavior of many processes. Our main contribution is a new linear programming technique that we introduce as a tool for obtaining and analyzing mechanisms for the secretary problem and its variants. Capturing the set of mechanisms as a linear polytope holds the following immediate advantages.

1. Computing the optimal mechanism reduces to solving a linear program. 2. Proving an upper bound on the performance of any mechanism reduces to finding a feasible solution to the dual program. 3. Exploring variants of the problem is as simple as adding new constraints, or manipulating the objective function of the linear program.

We demonstrate the applicability of these ideas in several settings including online auctions.

3.3 Approximating the Non-Contiguous Multiple Organization Packing Problem

Pierre-Francois Dutot (INRIA Rhôn-Alpes, FR)

We present in this paper a 5/2-approximation algorithm for scheduling rigid jobs on multi-organizations. For a given set of n jobs, the goal is to construct a schedule for N organizations (composed each of m identical processors) minimizing the maximum completion time (makespan). This algorithm runs in $O(n(N + \log(n)) \log(np_{\max}))$, where p max is the maximum processing time of the jobs. It improves the best existing low cost approximation algorithms. Moreover, the proposed analysis can be extended to a more generic approach which suggests different job partitions that could lead to low cost approximation algorithms of ratio better than 5/2.

3.4 Online Clustering with Variable Sized Clusters

Leah Epstein (Univ. of Haifa, IL)

Joint work of Csirik, János; Epstein, Leah; Imreh, Csanád; Levin, Asaf
Main reference J. Csirik, L. Epstein, C. Imreh, A. Levin, "Online Clustering with Variable Sized Clusters ", MFCS'10, pp.282–293, LNCS.
URL http://dx.doi.org/10.1007/978-3-642-15155-2_26

Online clustering problems are problems where the classification of points into sets (called clusters) is done in an online fashion. Points arrive at arbitrary locations, one by one, to be assigned to clusters at the time of arrival. A point can be either assigned to an existing cluster or a new cluster can be opened for it. We study a one dimensional variant on a line. Each cluster is a closed interval, and there is no restriction on the length of a cluster. The cost of a cluster is the sum of a fixed set-up cost and its diameter (or length). The goal is to minimize the sum of costs of the clusters used by the algorithm. We study several variants, all maintaining the essential property that a point which was assigned to a given cluster must remain assigned to this cluster, and clusters cannot be merged. In the strict variant, the diameter and the exact location of the cluster must be fixed when it is initialized. In the flexible variant, the algorithm can shift the cluster or expand it, as long as it contains all points assigned to it. In an intermediate model, the diameter is fixed in advance while the exact location can be modified. We give tight bounds on the competitive ratio of any online algorithm in each of these variants. In addition, for each one of the models, we also consider the semi-online case, where points are presented sorted by their location. The paper is joint work with J. Csirik, Cs. Imreh and A. Levin, and was presented in MFCS2010.

3.5 Competitive Strategies for Routing Flow Over Time

Lisa K. Fleischer (Dartmouth College, US)

Joint work of Bhaskar, Umang; Fleischer, Lisa K.; Anshelevich, Elliot

The study of routing games is motivated by the desire to understand the impact of individual user's decisions on network efficiency. To do this, prior work uses a simplified model of network flow where all flow exists simultaneously, and users route flow to either minimize their maximum delay or their total delay. Both of these measures are surrogates for measuring how long it takes to get all of your traffic through the network over time.

Instead of using these surrogates, we attempt a more direct study of how competition among users effects network efficiency by examining routing games in a flow-over-time model. We show that the network owner can reduce available capacity so that the competitive equilibrium in the reduced network is no worse than a small constant times the optimal solution in the original network using two natural measures of optimum: the time by which all flow reaches the destination, and the average amount of time it takes flow to reach the destination.

3.6 Potential Reduction Schemes in Structured Optimization

Michael D. Grigoriadis (Rutgers Univ., US)

Joint work of Grigoriadis,Michael D.; Khachiyan, Leonid G.; Villavicencio, J. U.

We study the performance of approximately computing a min-max [max- min] solution of a given set of M convex [concave], nonnegative-valued and block-separable coupling inequalities over the product of K convex compact "blocks". The generality of the model allows for a variety of specializations for applications in packing [covering] feasibility LPs, matrix games, block angular LPs, routing in multicommodity flows, and others. Optimization variants run within polylogarithmic factors. All of our FPTAS's include a quadratic term of 1/epsilon. Working within the well-known Lagrangian decomposition framework, we replace the underlying piecewise convex [concave] objective with its exponentially many breakpoints, by a smooth approximation, such as an exponential or logarithmic potential function, which is gradually improved. For implementations using the exponential potential, the original blocks are further restricted by their part of the coupling inequalities, adjustable by a restriction parameter. This helps in controlling the so-called "width". In contrast, logarithmic potential-based implementations are shown to be "width-free" and thus work with the unrestricted (original) blocks. We show that best coordination complexities obtain by using the logarithmic potential with unrestricted blocks for instances with roughly $M < K \log K$, but switching to the exponential potential with restricted blocks when $M > K \log K$. The exponential potential-based scheme solves (n,m)-matrix games A with elements in [-1,+1] to a prescribed relative error in quadratic $log(nm)$ time on an nm-processor EREW PRAM. In addition there is a parallel randomized approximation scheme for solving such games to within a given absolute accuracy, in expected quadratic $log(n+m)$ time on an $(n+m)/log(n+m)$-processor EREW PRAM, thus providing a sublinear support for such games. A roughly quadratic expected speedup is obtained relative to any deterministic approximation scheme.

Computational experiments show that optimal solutions of very large maximum concurrent flow problems are computed routinely to 4-digit accuracy, several orders faster than modern LP codes.

3.7 Wireless Scheduling in the Physical Model

Magnús M. Halldórsson (Reykjavik Univ., IS)

I will survey recent work on analytic work on scheduling wireless links in the SINR model. The first half will be focused on properties of the model, and on capacity (throughput) maximization in the case of uniform power. The second half will look at the problems involving power control, as well as other related issues, such as distributed algorithms.

3.8 A Polynomial Time OPT+1 Algorithm for the Cutting Stock Problem with a Constant Number of Object Lengths

Klaus Jansen (Universität Kiel, DE)

Joint work of Jansen, Klaus; Solis-Oba, Roberto

In the cutting stock problem we are given a set $T = T_1, ..., T_d$ of object types, where objects of type T_i have integer length $p_i > 0$. Given a set O of n objects containing n_i objects of type T_i, for each $i = 1, ..., d$, the problem is to pack O into the minimum number of bins of capacity beta. In this talk we consider the version of the problem in which the number d of different object types is constant and we present a polynomial time algorithm that computes a solution using at most $OPT + 1$ bins, where OPT is the value of an optimum solution.

3.9 The Cutting-Stock Approach to Bin Packing: Theory and Experiments

David S. Johnson(AT&T Research, US)

We report on results of an experimental study of the Gilmore-Gomory cutting-stock heuristic and related LP-based approaches to bin packing, as applied to instances generated according to discrete distributions. We examine the questions of how best to solve the knapsack problems used to generate columns in the Gilmore-Gomory approach, how the various algorithms' running times and solution qualities scale with key instance parameters, and how the algorithms compare to more traditional bin packing heuristics.

No polynomial running time bound is known to hold for the Gilmore-Gomory approach, and high-level empirical operation counts suggest that no straightforward implementation can have average running time $O(m^3)$, where m is the number of distinct item sizes. Our

experiments suggest that by using dynamic programming to solve the knapsack problems, one can robustly obtain average running times that are $o(m^4)$ and are feasible for m well in excess of 1,000. This makes a variant on the previously un-implemented asymptotic approximation scheme of Fernandez de la Vega and Lueker practical for arbitrarily large values of m and quite small values of ϵ.

In the process of performing these experiments we discovered two interesting anomalies: (1) running time *decreasing* as the number n of items increases and (2) solution quality improving as running time is reduced and an approximation guarantee is weakened. We provide explanations for these phenomena and characterize the situations in which they occur.

3.10 Disjoint-Path Facility Location: Theory and Practice

Howard Karloff (AT&T Research, US)

Internet service providers hope to provide their customers with superior Internet connectivity, but do they always do so? How can an ISP even know what quality of service it's providing to its customers? To this end, researchers recently proposed a new scheme an ISP could use in order to estimate the packet loss rates experienced by its customers.

To implement the new scheme, one has to approximately solve an interesting NP-Hard optimization problem on the ISP's network. Specifically, one must choose a small set of network nodes such that from each customer node there are arc-disjoint paths to *two* of the selected nodes. I will discuss recent work, mostly at ATT, attacking this problem and its surprisingly good results, in light of the problem's provable inapproximability in the worst case.

3.11 Procrastination Pays: Scheduling Jobs in Batches to Minimize Energy Usage

Samir Khuller (Univ. of Maryland, US)

Joint work of Chang, J.; Gabow, H.; Khuller, S.

We consider an elementary scheduling problem defined as follows. Given a collection of n jobs, where each job J_i has an integer length l_i as well as a set T_i of time intervals in which it can be feasibly scheduled. We are given a parallelism parameter P and can schedule up to P jobs at any time slot in which the machine is "active". The goal is to preemptively schedule all the jobs in the fewest number of active time slots.

The machine consumes a fixed amount of energy per time slot, regardless of the number of jobs scheduled at that slot (as long as the number of jobs is non-zero). In other words, subject to l_i units of each job i being scheduled in its feasible region and at each slot at most P jobs being scheduled, we are interested in minimizing the total time duration when the machine is active.

We present an O(n log n) algorithm for the case where jobs have unit length and T_i forms a single interval. For general T_i (and unit jobs), we show that the problem is NP-complete even for $P = 3$. However when $P = 2$, we show that it can be solved in polynomial time – we also present several extensions: for example when the jobs have non–unit requirements we can still solve this version in polynomial time.

No previous background knowledge on scheduling is expected. In addition, we will survey some recent work on bundling jobs in batches.

3.12 An AFPTAS for Variable Sized bin Packing with General bin Costs

Asaf Levin (Technion, IL)

Joint work of Epstein, Leah; Levin, Asaf

In variable sized bin packing problems, bins of different sizes are to be used for the packing of an input set of items. We consider variable sized bin packing with general costs. Each bin type has a cost associated with it, where the cost of a bin may be smaller or larger than its size, and the costs of different bin sizes are unrelated. For each bin type, this cost is to be paid for each instance which is used for the packing of input items. This generalized setting of the problem has numerous applications in storage and scheduling. We introduce new reduction methods and separation techniques, which allow us to design an AFPTAS for the problem.

3.13 Survey of connections between approximation algorithms and parameterized complexity

Dániel Marx (HU Berlin, DE)

Approximation algorithms and parameterized complexity are two well-studied approaches for attacking hard combinatorial problems. In my talk, I overview the ways approximation can be introduced into the framework of parameterized complexity, survey results in this direction, and show how parameterized hardness theory can be used to give lower bounds on the efficiency of approximation schemes.

3.14 Vertex Cover in Graphs with Locally Few Colors and Precedence Constrained Scheduling with Few Predecessors

Monaldo Mastrolilli (IDSIA - Lugano, CH)

In 1986 Erdös et. al. defined the local chromatic number of a graph as the minimum number of colors that must appear within distance 1 of a vertex. For any fixed $\Delta \geq 2$, they presented graphs with arbitrarily large chromatic number that can be colored so that: (i) no vertex neighborhood contains more than Δ different colors (bounded local colorability), and (*ii*) adjacent vertices from two color classes form an induced subgraph that is complete and bipartite (local completeness).

We investigate the weighted vertex cover problem in graphs when a locally bounded coloring is given as input. This generalizes in a very natural vein the vertex cover problem in bounded degree graphs to a class of graphs with arbitrarily large chromatic number. Assuming the Unique Game Conjecture, we provide a tight characterization. More precisely, we prove that it is UG-hard to improve the approximation ratio of $2 - 2/(\Delta + 1)$ if only the bounded local colorability, but not the local completeness condition holds for the given coloring. A matching upper bound is also provided. Vice versa, when both the above two properties (i) and (ii) hold, we present a randomized approximation algorithm with performance ratio of $2 - \Omega(1)\frac{\ln \ln \Delta}{\ln \Delta}$. This matches (up to the constant factor in the lower order term) known inapproximability results for the special case of bounded degree graphs.

Moreover, we show that when both the above two properties (i) and (ii) hold, the obtained result finds a natural application in a classical scheduling problem, namely the precedence constrained single machine scheduling problem to minimize the weighted sum of completion times. In a series of recent papers it was established that this scheduling problem is a special case of the minimum weighted vertex cover in graphs G of incomparable pairs defined in the dimension theory of partial orders. We show that G satisfies properties (i) and (ii) where $\Delta - 1$ is the maximum number of predecessors (or successors) of each job.

3.15 Min-Max Graph Partitioning and Small Set Expansion

Seffi Naor (Technion, IL)

We study graph partitioning problems from a min-max perspective, in which an input graph on n vertices should be partitioned into k parts, and the objective is to minimize the maximum number of edges leaving a single part. The two main versions we consider are where the k parts need to be of equal-size, and where they must separate a set of k given terminals. We consider a common generalization of these two problems, and design for it an approximation algorithm. This improves over an O(log2 n) approximation for the second version due to Svitkina and Tardos [ST04], and roughly O(k log n) approximation for the first version that follows from other previous work.

3.16 Green Computing Algorithmics

Kirk Pruhs (Univ. of Pittsburgh, US)

We are in the midst of a green computing revolution involving the redesign of information technology hardware and software at all levels of the information technology stack. Such a revolution spawns a multitude of technological challenges, many of which are algorithmic in nature. The most obvious type of algorithmic problem arising from this green computing revolution involves directly managing power, energy or temperature as a resource. Other algorithmic problems arise because the new technology, which was adopted for energy and power considerations, has different physical properties than previous technologies. I will try to give a feel for the current state of green computing algorithmics research, and provide some advice about how to contribute to this research.

3.17 Minimizing Busy Time in Multiple Machine Real-time Scheduling

Baruch Schieber (IBM TJ Watson Research Center, US)

Joint work of Khandekar, Rohit; Schieber, Baruch; Shachnai, Hadas; Tamir, Tami

We consider the following fundamental scheduling problem. The input consists of n jobs to be scheduled on a set of machines of bounded capacities. Each job is associated with a release time, a due date, a processing time and demand for machine capacity. The goal is to schedule all of the jobs non-preemptively in their release-time-deadline windows, subject to machine capacity constraints, such that the total busy time of the machines is minimized. Our problem has important applications in power-aware scheduling, optical network design and customer service systems. Scheduling to minimize busy times is APX-hard already in the special case where all jobs have the same (unit) processing times and can be scheduled in a fixed time interval.

Our main result is a 5-approximation algorithm for general instances. We extend this result to obtain an algorithm with the same approximation ratio for the problem of scheduling moldable jobs, that requires also to determine, for each job, one of several processing-time vs. demand configurations. Better bounds and exact algorithms are derived for several special cases, including proper interval graphs, intervals forming a clique and laminar families of intervals.

3.18 Bin Packing with Fixed Number of Bins Revisited

Ildiko Schlotter (Budapest Univ. of Technology & Economicsn, HU)

Joint work of Jansen, Klaus; Kratsch, Stefan; Marx, Dániel; Schlotter, Ildiko

As Bin Packing is NP-hard already for $k = 2$ bins, it is unlikely to be solvable in polynomial time even if the number of bins is a fixed constant. However, if the sizes of the items are

polynomially bounded integers, then the problem can be solved in time $n^{O(k)}$ for an input of length n by dynamic programming. We show, by proving the W[1]-hardness of Unary Bin Packing (where the sizes are given in unary encoding), that this running time cannot be improved to $f(k) \cdot n^{O(1)}$ for any function $f(k)$ (under standard complexity assumptions). On the other hand, we provide an algorithm for Bin Packing that obtains in time $2^{O(k \log^2 k)} + O(n)$ a solution with additive error at most 1, i.e., either finds a packing into $k+1$ bins or decides that k bins do not suffice.

3.19 Balanced Interval Coloring

Alexander Souza (HU Berlin, DE)

Joint work of Souza, Alexander; Antoniadis, Antonios; Hüffner, Falk; Lenzner, Pascal
URL http://arxiv.org/abs/1012.3932

We consider the discrepancy problem of coloring n intervals with k colors such that at each point on the line, the maximal difference between the number of intervals of any two colors is minimal. Somewhat surprisingly, a coloring with maximal difference at most one always exists. Furthermore, we give an algorithm with running time O(n log n + kn log k) for its construction. This is in particular interesting because many known results for discrepancy problems are non-constructive.

This problem naturally models a load balancing scenario, where n tasks with given start- and end-times have to be distributed among k servers.

Our results imply that this can be done ideally balanced.

When generalizing to d dimensional boxes (instead of intervals), a solution with difference at most one is not always possible. We show that for any d $>$ 1 and any k $>$ 1 it is NP-complete to decide if such a solution exists, which implies also NP-hardness of the respective minimization problem.

In an online scenario, where intervals arrive over time and the color has to be decided upon arrival, the maximal difference in the size of color classes can become arbitrarily high for any online algorithm.

3.20 Fast Separation Algorithms for Multidimensional Assignment Problems

Frits C.R. Spieksma (K.U. Leuven, BE)

In polyhedral combinatorics, the polytope corresponding to an integer programming formulation of a combinatorial optimization problem is examined in order to obtain families of valid inequalities. To incorporate such families of inequalities within a cutting plane algorithm requires an additional step: determining whether an inequality of a specific family is violated by a given vector x (the separation problem). The idea put forward in this work is to consider a compact representation of this given vector x, and to measure the complexity of a separation algorithm in terms of this compact representation.

We illustrate this idea on the separation problem of well-known families of inequalities associated to the (multi-index) assignment polytope, and we show that for these families of inequalities, better time-complexities than the current ones are possible.

3.21 Hardness of Shops and Optimality of List Scheduling

Ola Svensson (KTH - Stockholm, SE)

We give an overview of the implications and techniques of the following hardness of approximation results:

- Non-constant inapproximability results for various shop scheduling problems that essentially match the best known approximation algorithm for acyclic job shops and general flow shops.

- A 2 hardness assuming the unique games conjecture for the problem of scheduling jobs with precedence constraints on identical machines so as to minimize the makespan. This matches the classical 2-approximation algorithm by Graham from 66.

3.22 Scheduling with Bully Selfish Jobs

Tami Tamir (The Interdisciplinary Center - Herzliya, IL)

In job scheduling with precedence constraints, $i < j$ means that job j cannot start being processed before job i is completed. In this paper we consider selfish bully jobs who do not let other jobs start their processing if they are around. Formally, we define the selfish precedence–constraint where $i <_s j$ means that j cannot start being processed if i has not started its processing yet. Interestingly, as was detected by a devoted kindergarten teacher whose story is told below, this type of precedence constraints is very different from the traditional one, in a sense that problems that are known to be solvable efficiently become NP-hard and vice-versa. The work of our hero teacher, Ms. Schedule, was initiated due to an arrival of bully jobs to her kindergarten. Bully jobs bypass all other nice jobs, but respect each other. This natural environment corresponds to the case where the selfish precedence-constraints graph is a complete bipartite graph. Ms. Schedule analyzed the minimum makespan and the minimum total flow-time problems for this setting. She then extended her interest to other topologies of the precedence constraints graph and other special instances with uniform length jobs and/or release times. Finally, she defined a generalization of her problem, where the precedence constraints graph is weighted, and $w(i, j)$ specifies the minimal gap between the starting times of i and j. The paper was presented in FUN with Algorithms 2010.

3.23 How to use Lagrangian-Relaxation Algorithms to solve Packing and Covering Problems

Neal E. Young (Univ. of California - Riverside, US)

Following a brief review of the history of Lagrangian-relaxation algorithms, I will summarize recent results in the area in a concrete form that (hopefully) makes it easy to understand how to apply the results.

Given any linear program (LP) that includes some packing constraints and/or some covering constraints, the packing and/or covering constraints can be "dualized", replacing the packing constraints by a carefully chosen linear combination of the packing constraints, and likewise for the covering constraints. This replaces all m packing/covering constraints by just one or two constraints, and gives an LP relaxation LP' of the problem that is combinatorially simpler than the original problem.

Given any algorithm alg' for the simpler problem LP', there is a simple algorithm for the original problem that calls $alg'O(min(m, width) * log(m)/epsilon^2)$ times, then returns an epsilon-approximate solution to the original problem.

I will illustrate the ideas using zero-sum matrix games, the Held-Karp lower bound on TSP, the "configuration LP" for bin packing, and on multi-commodity flow problems.

3.24 A Truthful Constant Approximation for Maximizing the Minimum Load on Related Machines

Rob van Stee (MPI für Informatik - Saarbrücken, DE)

Joint work of Christodoulou, Giorgos; Kovacs, Annamaria; van Stee, Rob

Designing truthful mechanisms for scheduling on related machines is a very important problem in single-parameter mechanism design. We consider the covering objective, that is we are interested in maximizing the minimum completion time of a machine. This problem falls into the class of problems where the optimal allocation can be truthfully implemented. A major open issue for this class is whether truthfulness affects the polynomial-time implementation.

We provide the first constant factor approximation for deterministic truthful mechanisms. In particular we come up with a approximation guarantee of 2+eps, significantly improving on the previous upper bound of $min(m, (2 + eps)s_m/s_1)$.

4 Discussion notes

4.1 Current and Future Trends in Scheduling

Alexander Souza (HU Berlin, DE), souza@informatik.hu-berlin.de

Notes of the discussion on "Current and Future Trends in Scheduling" that took place at the Dagstuhl meeting 11091 on "Packing and Scheduling Algorithms for Information and Communication Services" from 27.2.2011 to 4.3.2011.

Theory and Applications

- *To what extent do scheduling problems in theory and practice relate?*
  1. "No relation", "There should be", "There are"
  2. Example of a project of a Steel company together with TU Berlin: Initiated at an OR conference in Germany; Company approached TU Berlin; Solution was implemented by TU Berlin, but no support was given; Approach was a dynamic programming algorithm combined with heuristics.
- *What does it bring you to be attached to reality?*
  More satisfying research; Algorithm engineering; Per-Instance-Guarantees; Modeling as an issue.

New Theoretical Promises and Challenges

- *Are there new theoretical problems that we need to work on?*
  (a) Inapproximability results: Long standing open questions; tight bounds
  (b) New variants of classical problems: Measures; Models; Green IT; Cloud Computing
  (c) Dynamic Aspects: Practically important; Maybe theoretically nice; Stochastic models; Observed distributions; Technically difficult; Modeling again an issue; Communication with other disciplines required (workshop); Availability of data; Storing solutions for later reuse; Markov chain models (for online scheduling)
  (d) Theoretical vs. practical results: Essentially same outline as with the Paging problem; "Why something does well"

Sustainability Domain

1. Example: Land-lot purchase; Can have scheduling components
2. Currently well funded; For example Carla Gomes

   `http://www.cs.cornell.edu/gomes/`

Per Instance Guarantees

1. Certificates of instance-wise approximation ratio
2. Without LP-bounds? Maybe by MILP solutions or lower bounds

Insights from Game Theory

- *What are the new insights gained from game theoretic approach to scheduling (Does it capture better than classic scheduling contemporary systems)?*

(a) "Next question"
(b) Canonical examples from Game Theory are scheduling and network design problems
(c) Behavioral economics: Model how people behave; Maybe not accurately reflected in scheduling (pain-scheduling at a dentist); Human aspect of scheduling; Interaction of schedules with people; Indirectly done already; Find out criteria and objective function is an issue; "User happiness" is the objective function (in order to have an impact)

Stochastic Scheduling and Robustness

- Is *stochastic scheduling the 'right' direction for future research?* (*Can this direction be fruitful in view of the experience of the 70's?*)

(a) Contacts with industry: Combining scheduling and routing; Transportation problems
(b) Robustness: Varying data (small perturbations); Stable schedules

Integer Programming Approach in Scheduling

- *What do you think about integer programming as an approach for solving scheduling problems?*

(a) CSP's are maybe better because more flexible
(b) Problem: CPLEX not available; "Black magic"; Free solvers available at TU Berlin; Practical algorithms for large-scale scheduling problems are available
(c) ILP research mostly in OR, but not so much in CS; More collaboration between OR and CS needed
(d) Formulations matter

Theoretical Knowledge in Companies

1. Theoretically good algorithms rarely implemented for critical systems; Mostly prototypes
2. Examples for benefit of theoretical knowledge in companies
   a. Algorithm used for something it was not designed for, but it worked
   b. Akamai; Theoretical knowledge went into applications
   c. Start-up companies sometimes initiated by CS PhD's (also theoretical)

Personal Motivation

- *What drives your interest in the area of scheduling?*

(a) Open fundamental problems; Optimization of resources (also in real life); New problems; Old problems; Get paid; Beauty
(b) Disconnect between "formal motivation" (the introduction of your paper) and "personal motivation" (why you really do it)
(c) Playground for new questions; Models; Techniques
(d Can be explained to people

5 Open Problems

Notes of the "Open problems" session that took place at the Dagstuhl meeting 11091 on "Packing and Scheduling Algorithms for Information and Communication Services" from 27.2.2011 to 4.3.2011.

5.1 Implementing the Sum-of-Squares Bin-Packing Algorithm

David Johnson (AT&T, US)

The Sum-of-Squares bin packing algorithm (SS) is designed for instances in which the bin size B is an integer, as are all the item sizes. It is an online algorithm that chooses the bin into which to place the next item as follows: Let $c[i]$ be the number of bins in the current packing whose gap (B minus the sum of the sizes of the items already packed in the bin) equals i. Initially all the $c[i]$'s are 0. It chooses a bin into which to pack the item so as to minimize

$$\sum_{i=1}^{B-1} c[i]^2,$$

where the choice can be either an existing bin or a new bin with initial gap B.

It is straightforward to implement this algorithm to run in time $O(nB)$, where n is the number of items, whereas the classical Best-Fit algorithm (place each item in a bin with the smallest gap that will contain it) can be implemented to run in time $O(n \log B)$ by maintaining a priority queue for the non-zero values of $c[i]$.

For instances, when the number of item sizes is bounded by some constant J, SS can be implemented in time $O(nJ \log B)$, by maintaining a priority queue for each item size. But what if there is no such bound, or if $J = \Omega(B)$?

Our question: Is there an implementation of SS that, for any fixed B and without restriction on J, runs in time $o(nB)$?

For a detailed discussion of the Sum-of-Squares algorithm and its performance, see [1].

References

1 J. Csirik, D. S. Johnson, C. Kenyon, J. B. Orlin, P. W. Shor and R. R. Weber. On the Sum-of-Squares Algorithm for Bin Packing. *J. ACM* 53, pp. 1–65, 2006).

5.2 Covering by Rectangles: Is Slicing Essential?

Sivan Albagli-Kim (Technion, IL)

The following open problem explores the relation between *packing resizable items (PRI)* and geometric covering. As shown in [1], PRI is equivalent to the problem of *covering with holes (CwH)*, defined as follows. Given is a set $H_I = \{h_1, \cdots, h_n\}$ of n holes; each hole h_j is associated with a length $0 < q_j < X$ and a width $0 < p_j < Y$. We want to determine whether it is possible to cover an $X \times Y$ rectangle by holes in H_I. A cover is a placement of

the holes. For each hole h_j, the solution specifies the x-interval $X_j = [x_{1j}, x_{2j}]$ in which h_j is spanned, such that $x_{2j} - x_{1j} = q_j$. An $X \times Y$ rectangle is covered if, for any $0 \leq t \leq X$, the total width of holes whose x-interval contains t is at least Y. Note that the holes need not be placed as rectangles and can be *sliced* along the y-axis. This type of cover models, e.g., applications in which the x-axis corresponds to time, and the y-axis corresponds to a resource whose allocation is not associated with specific location.

Figure 1 shows a cover of a 1×1 rectangle by 7 holes. Note, for example, that hole h_4 spans along $[0.6, 1]$ and its width is 0.5. Similarly, hole h_2 spans along $[0, 0.6]$ and its width is 0.3. We also note that it is possible to have overlapping holes, as well as holes whose intervals span beyond the covered area.

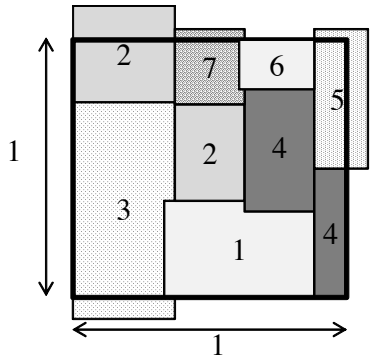

Figure 1 Covering a 1×1-rectangle by 7 holes.

In the *Covering with Rectangles (CwR)* problem, given is a set $R_I = \{r_1, \cdots, r_n\}$ of n rectangles, such that each rectangle r_j is associated with a length $0 < q_j < X$ and a hight $0 < p_j < Y$. We need to determine whether it is possible to cover an $X \times Y$ rectangle with rectangles in R_I. A cover is a placement of the rectangles. For each rectangle r_j, the solution specifies the x-interval $X_j = [x_{1j}, x_{2j}]$ such that $x_{2j} - x_{1j} = q_j$, and the y-interval $Y_j = [y_{1j}, y_{2j}]$ such that $y_{2j} - y_{1j} = p_j$. A solution covers an $X \times Y$ rectangle if, for any $0 \leq t \leq X$, the total hight of rectangles whose x-interval contains t is at least Y. Note that, unlike the CwH problem, in CwR the rectangles are rigid (and therefore, cannot be sliced).

Let $H_I = R_I$. Clearly, for all X, Y, a positive answer for CwR implies a positive answer for CwH. However, does the reverse hold, namely, does the existence of a solution for CwH imply the existence of a solution for CwR?

This open problem was settled during the seminar. We thank Jiří Sgall for the following example, which implies that the answer to the above is NO. The input consists of 8 holes: $(a)3 \times 2, (b)2 \times 1, (c)1 \times 4, (d)3 \times 2, (e)2 \times 1, (f)1 \times 4, (g)1 \times 3, (h)1 \times 1$. The holes need to cover a 4×7 rectangle. As shown in Figure 2, there exists a solution for CwH (in which g is sliced); however, there is no solution for CwR with this set of rectangles.

References

1 S. Albagli-Kim, H. Shachnai and T. Tamir. Approximation Algorithms for Packing Resizable Items and Covering by Holes. *Submitted*.

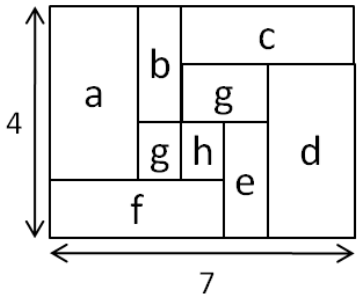

Figure 2 A solution for covering with holes in which slicing is essential.

5.3 Fixed-parameter Tractable Scheduling Problems

Dániel Marx (Humbold-Universität zu Berlin, DE)

Recall that a problem is *fixed-parameter tractable (FPT)* with some parameter k if it can be solved in time $f(k) \cdot n^{O(1)}$, where n is the input size and f is an arbitrary computable function depending only on k. Typically, if a problem is solvable in polynomial time for every fixed value of the parameter k (for example, there is an algorithm with running time $O(n^k)$), then it makes sense to ask if the problem is FPT, i.e., if we can remove k from the exponent of n and make it a multiplicative factor.

Compared to graph algorithms and other applications, there is surprisingly little work on the fixed-parameter tractability of scheduling problems (see [1,2] for a few examples). One problem is that it is not obvious how to choose relevant parameters that lead to interesting positive results. For example, the number k of processors is an obvious choice for the parameter, but a large fraction of the scheduling problems is NP-hard already for constant number of processors (and hence unlikely to be FPT with respect to this parameter).

A parameter which looks much more promising for obtaining fixed-parameter tractability results is the *number of rejected jobs.* Consider *any* scheduling problem that can be solved optimally in polynomial time. Then we can extend the problem by allowing rejections: the input contains an additional integer k, and the solution has to schedule all but k jobs. Assuming that the original problem is polynomial-time solvable, it is clear that the extended problem can be solved in $n^{O(k)}$ time: we first guess which k of the jobs to reject and solve the problem optimally for the remaining jobs. However, it is not obvious if the extended problem is fixed-parameter tractable parameterized by k. This question can be raised for any polynomial-time solvable scheduling problem and could be potentially interesting to explore. The open question is to find concrete scheduling problems, where the extended version with rejected jobs is NP-hard, but fixed-parameter tractable.

References

1 Fellows, M., and McCartin, C. On the parameterized complexity of minimizing tardy tasks. Theoretical Computer Science A 298 (2003), 317-324.
2 Bodlaender, H., and Fellows, M. On the complexity of k-processor scheduling. Operations Research Letters 18 (1995), 93-98.

5.4 Scheduling with Buffering on the Line

Adi Rosén (CNRS, FR)

We consider directed linear communication networks. The linear network consists of n nodes, $\{1,\ldots,n\}$, and $n-1$ directed edges, $(i,i+1)$, for $1 \leq i \leq n-1$. The system is synchronous, and at any time step, each edge can transmit one message. In one version of the problem, each node can store at any time an *infinite* number of messages. We are given a set $\mathcal{M}$, $|\mathcal{M}| = M$, of messages. Each message $m = (s_m, t_m, r_m, d_m) \in \mathcal{M}$ consists of a source node s_m, a target node t_m, a release time r_m, and a deadline d_m. For a message m, we define the slack of m, σ_m, to be $m = (d_m - r_m) - (t_m - s_m)$ (this is the number of steps the message can be idle and still make it to its destination by its deadline.).We define $\Sigma = \max_{m\in\mathcal{M}} \sigma_m$. We want to find a schedule for the messages that maximizes the number of messages that arrive to their destinations by their respective deadlines.

The open problem is whether there exists a polynomial-time algorithm with constant approximation ratio.

The problem is NP-hard [2]. A polynomial-time algorithm with approximation ratio $O(\min\{\log^* n, \log^* \Sigma, \log^* M\})$ is known [3].

References

1 Micah Adler, Sanjeev Khanna, Rajmohan Rajaraman, and Adi Rosén. Time-constrained scheduling of weighted packets on trees and meshes. *Algorithmica*, 36(2), pp. 123–152, 2003.
2 Micah Adler, Arnold L. Rosenberg, Ramesh K. Sitaraman, and Walter Unger. Scheduling time-constrained communication in linear networks. *Theory of Computing Systems* 35(6), pp. 599–623, 2002.
3 H. Räcke, A. Rosén, Approximation Algorithms for Time-Constrained Scheduling on Line Networks. In *Proc. of the 21st ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pp. 337–346, 2009.

5.5 Wireless Scheduling

Magnús Halldórsson (Reykjavik University, IS)

Let $p_1, p_2, \ldots, p_n$ be points on the real line with *capacities* $c_1, \ldots, c_n$. The problem is to partition $P = \{p_i\}$ into fewest sets $P_1, \ldots, P_t$, such that

$$\sum_{p'\in P_i, p'\neq p} |p - p'|^3 \leq c_i, \quad \text{for each } i \text{ and each } p \in P_i.$$

We seek an $O(1)$-approximation.

This problem statement captures the most basic open question in scheduling wireless links under the physical (or, SINR) model. Normally, links are given as sender-receiver pairs, but it is known that when messages are all transmitted with the same uniform power, we can blur the distinction between sender and receiver, by paying a constant factor. The problem is usually specified on the plane, or in a general distance metric, but results for the

one-dimensional case can typically be generalized relatively easily. The exponent "3", known as the path-loss constant, is situation dependent, and can be any number between 2 and 6.

An $O(1)$-approximation is known for the throughput problem of finding a single set P_1 of maximum cardinality within which all points satisfy the inequality above [1]. This immediately gives an $O(\log n)$-factor, but no better is known.

References

1 O. Goussevskaia, M. M. Halldorsson, R. Wattenhofer, and E. Welzl. Capacity of Arbitrary Wireless Networks. In *INFOCOM*, pages 1872–1880, April 2009.

5.6 Feedback Arc Set Problems with Near-metric Weights

Monaldo Mastrolilli (IDSIA Lugano, CH)

Introduction

The Minimum Feedback Arc Set problem (MinFAS) is a fundamental and classical combinatorial optimization problem that finds application in many different settings that span from circuit design, constraint satisfaction problems, artificial intelligence, scheduling, etc. (see e.g. Chapter 4 in [19] for a survey). For this reason it has been deeply studied since the late 60's (see, e.g., [17]).

Its input consists of a set of vertices V and nonnegative weights $\{w_{(i,j)}, w_{(j,i)} : \{i,j\} \subseteq V\}$ for every oriented pair of vertices. The goal is to find a permutation π that minimizes $\sum_{\pi(i)<\pi(j)} w_{(i,j)}$, i.e. the weight of pairs of vertices that comply with the permutation[1]. A *partially ordered set* (*poset*) $\mathbf{P} = (V, P)$, consists of a set V and a partial order P on V, i.e., a reflexive, antisymmetric, and transitive binary relation P on V, which indicates that, for certain pairs of elements in the set, one of the elements precedes the other. In the *constrained* MinFAS (see [23]) we are given a partially ordered set $\mathbf{P} = (V, P)$ and we want to find a linear extension of $\mathbf{P}$ of minimum weight.

MinFAS was contained in the famous list of 21 NP-complete problems by Karp [14]. Despite intensive research for almost four decades, the approximability of this problem remains very poorly understood due to the big gap between positive and negative results. It is known to be APX-hard [13], but no constant approximation ratio has been found yet. The best known approximation algorithm achieves a performance ratio $O(\log n \log\log n)$ [21, 10, 9], where n is the number of vertices of the digraph. Closing this approximability gap is a well-known major open problem in the field of approximation algorithms (see e.g. [25], p. 337). Very recently and conditioned on the Unique Games conjecture, it was shown [11] that for every constant $C > 0$, it is NP-hard to find a C-approximation to the MinFAS.

Important ordering problems can be seen as special cases of MinFAS with restrictions on the weighting function. Examples of this kind are provided by ranking problems related to the aggregation of inconsistent information, that have recently received a lot of attention [1, 2, 15, 24]. Several of these problems can be modeled as (constrained) MinFAS with weights satisfying either *triangle inequalities* (i.e., for any triple i, j, k, $w_{(i,j)} + w_{(j,k)} \geq w_{(i,k)}$), or *probability constraints* (i.e., for any pair i, j, $w_{(i,j)} + w_{(j,i)} = 1$). Ailon, Charikar and

[1] Different, but equivalent formulations are often given for the problem.

Newman [2] give the first constant-factor randomized approximation algorithm for the unconstrained MinFAS problem with weights that satisfy the triangle inequalities. When the probability constraints hold, Mathieu and Schudy [15] obtain a PTAS. The currently best known constant approximation algorithms for the (constrained) MinFAS with triangle inequalities on the weights can be found in [1, 24]. Another prominent example is given by a classical problem in scheduling, namely the precedence constrained single machine scheduling problem to minimize the weighted sum of completion times, denoted as $1|prec|\sum w_jC_j$ (see e.g. [16] and [12] for a 2-approximation algorithm). This problem can be seen as a constrained MinFAS where the weight of arc (i,j) is equal to the product of two numbers $p_i \cdot w_j$: p_i is the processing time of job i and w_j is a weight associated to job j (see [3, 4, 7, 8] for recent advances).

The (constrained) MinFAS can be described by the following natural (compact) ILP using linear ordering variables $\delta_{(i,j)}$ (see e.g. [24]): variable $\delta_{(i,j)}$ has value 1 if vertex i precedes vertex j in the corresponding permutation, and 0 otherwise.

$$\text{[FAS]} \qquad \min \sum_{i \neq j} \delta_{(i,j)} w_{(i,j)} \tag{1a}$$

$$\text{s.t.} \qquad \delta_{(i,j)} + \delta_{(j,i)} = 1, \qquad \text{for all distinct } i,j \tag{1b}$$

$$\delta_{(i,j)} + \delta_{(j,k)} + \delta_{(k,i)} \geq 1, \qquad \text{for all distinct } i,j,k \tag{1c}$$

$$\delta_{(i,j)} = 1, \qquad \text{for all } (i,j) \in P \tag{1d}$$

$$\delta_{(i,j)} \in \{0,1\}, \qquad \text{for all distinct } i,j \tag{1e}$$

Constraint (1b) ensures that in any feasible permutation either vertex i is before j or vice versa. The set of Constraints (1c) is used to capture the transitivity of the ordering relations (i.e., if i is ordered before j and j before k, then i is ordered before k, since otherwise by using (1b) we would have $\delta_{(j,i)} + \delta_{(i,k)} + \delta_{(k,j)} = 0$ violating (1c)). Constraints (1d) ensure that the returned permutation complies with the partial order P.

To some extent, one source of difficulty that makes the MinFAS hard to approximate within any constant is provided by Constraint (1b). To see this, consider, for the time being, the unconstrained MinFAS. The following covering relaxation obtained by relaxing Constraint (1b) behaves very differently with respect to approximation.

$$\min \sum_{i \neq j} \delta_{(i,j)} w_{(i,j)} \tag{2a}$$

$$\text{s.t.} \qquad \delta_{(i,j)} + \delta_{(j,i)} \geq 1, \qquad \text{for all distinct } i,j \tag{2b}$$

$$\delta_{(i,j)} + \delta_{(j,k)} + \delta_{(k,i)} \geq 1, \qquad \text{for all distinct } i,j,k \tag{2c}$$

$$\delta_{(i,j)} \in \{0,1\}, \qquad \text{for all distinct } i,j \tag{2d}$$

Problem (2) is a special case of the vertex cover problem in hypergraphs with edges of sizes at most 3. It admits "easy" constant approximate solutions, whereas problem (1) does not seem to have any constant approximation [11]. Moreover, the fractional relaxation of (2), obtained by dropping the integrality requirement, is a positive linear program and therefore fast NC approximation algorithms exists: Luby and Nissan's algorithm [18] computes a feasible $(1+\varepsilon)$-approximate solution in time polynomial in $1/\varepsilon$ and $\log N$, using $O(N)$ processors, where N is the size of the input (fast approximate solution can also be obtained through the methods of [20]). On the other side, the linear program relaxation of (1) is not positive.

In a recent (unpublished, but available upon request) paper we show that the covering relaxation (2) is an "optimal" relaxation, namely, a *proper* formulation, for the unconstrained

MinFAS when the weights satisfy the triangle inequalities. More precisely, we show that any α-approximate solution to (2) can be turned in polynomial time into an α-approximate solution to (1), for any $\alpha \geq 1$ and when the weights satisfy the triangle inequalities. The same claim applies to fractional solutions. (We also observe that the same result does not hold when the weights satisfy the probability constraints.)

Interestingly, a compact covering formulation can be also obtained for the more general setting with precedence constraints. In this case we need to consider the following covering relaxation[2] which generalizes (2) to partially ordered sets $\mathbf{P} = (V, P)$.

$$\min \sum_{i \neq j} \delta_{(i,j)} w_{(i,j)} \tag{3a}$$

$$\text{s.t. } \delta_{(x_1,y_1)} + \delta_{(x_2,y_2)} \geq 1, \qquad (x_2, y_1), (x_1, y_2) \in P \tag{3b}$$

$$\delta_{(x_1,y_1)} + \delta_{(x_2,y_2)} + \delta_{(x_3,y_3)} \geq 1, \qquad (x_2, y_1), (x_3, y_2), (x_1, y_3) \in P \tag{3c}$$

$$\delta_{(i,j)} \in \{0, 1\}, \qquad \text{for all distinct } i, j \tag{3d}$$

Open problems

The constrained MinFAS problem admits a natural covering formulation with an exponential number of constraints (see e.g. [5]):

$$\min \sum_{(i,j)} \delta_{(i,j)} w_{(i,j)} \tag{4a}$$

$$\text{s.t.} \sum_{i=1}^{c} \delta_{(x_i,y_i)} \geq 1, \qquad \text{for all } c \geq 2,\ (x_i, y_i)_{i=1}^{c} \text{ s.t. } (x_i, y_{i+1}) \in P \tag{4b}$$

$$\delta_{(i,j)} \in \{0, 1\}, \qquad \text{for all distinct } i, j \tag{4c}$$

The condition $(x_i, y_{i+1}) \in P$ in constraint (4b) is to be read cyclically, i.e. $(x_c, y_1) \in P$. The hyperedges in this vertex cover problem are exactly the alternating cycles of poset P (see e.g. [22]).

We know that when the weights satisfy the triangle inequality then we can drop from (4) all the constraints of size strictly larger than three. Generalizing, it would be nice to prove/disprove the following statement that we conjecture to be true.

▶ Hypothesis 1. When the weights satisfy the k-gonal inequalities[3], then there exists a constant $c(k)$, whose value depends on k, such that a proper formulation for the constrained MinFAS problem can be obtained by dropping from (4) all the constraints of size strictly larger than $c(k)$.

Moreover, it would be nice to use the large literature and techniques developed for covering problems to improve the best known ratios for MinFAS with (near-)metric weights. This was actually the case for the scheduling problem $1|prec|\sum w_j C_j$: in [3, 8] it was first shown that the structure of the weights for this problem allows for all the constraints of size strictly larger than two to be ignored, therefore the scheduling problem can be seen as a special case of the vertex cover problem. The established connection proved later to be very valuable both for positive and negative results: studying this graph yielded a framework that

[2] It is a relaxation to constrained MinFAS since if either Constraint (3b) or (3c) was violated then we would have a cycle.

[3] For all $a_1, \ldots, a_k \in V$ the following holds: $w_{(a_1,a_k)} \leq w_{(a_1,a_2)} + \ldots + w_{(a_{k-1},a_k)}$.

unified and improved upon previously best-known approximation algorithms [4, 6]; moreover, it helped to obtain the first inapproximability results for this old problem [7] by revealing more of its structure.

References

1. N. Ailon. Aggregation of partial rankings, -ratings and top- lists. *Algorithmica*, 57(2):284–300, 2010.
2. N. Ailon, M. Charikar, and A. Newman. Aggregating inconsistent information: Ranking and clustering. *J. ACM*, 55(5), 2008.
3. C. Ambühl and M. Mastrolilli. Single machine precedence constrained scheduling is a vertex cover problem. *Algorithmica*, 53(4):488–503, 2009.
4. C. Ambühl, M. Mastrolilli, N. Mutsanas, and O. Svensson. Scheduling with precedence constraints of low fractional dimension. In *Proceedings of IPCO 2007*, volume LNCS 4168, pages 28–39. Springer, 2007.
5. C. Ambühl, M. Mastrolilli, N. Mutsanas, and O. Svensson. Precedence constraint scheduling and connections to dimension theory of partial orders. *Bulletin of the European Association for Theoretical Computer Science (EATCS)*, 95:45–58, 2008.
6. C. Ambühl, M. Mastrolilli, and O. Svensson. Approximating precedence-constrained single machine scheduling by coloring. In *Proceedings of the APPROX + RANDOM*, volume LNCS 4110, pages 15–26. Springer, 2006.
7. C. Ambühl, M. Mastrolilli, and O. Svensson. Inapproximability results for sparsest cut, optimal linear arrangement, and precedence constraint scheduling. In *Proceedings of FOCS 2007*, pages 329–337, 2007.
8. J. R. Correa and A. S. Schulz. Single machine scheduling with precedence constraints. *Mathematics of Operations Research*, 30(4):1005–1021, 2005.
9. G. Even, J. Naor, S. Rao, and B. Schieber. Divide-and-conquer approximation algorithms via spreading metrics. *J. ACM*, 47(4):585–616, 2000.
10. G. Even, J. Naor, B. Schieber, and M. Sudan. Approximating minimum feedback sets and multicuts in directed graphs. *Algorithmica*, 20(2):151–174, 1998.
11. V. Guruswami, R. Manokaran, and P. Raghavendra. Beating the random ordering is hard: Inapproximability of maximum acyclic subgraph. In *FOCS*, pages 573–582, 2008.
12. L. A. Hall, A. S. Schulz, D. B. Shmoys, and J. Wein. Scheduling to minimize average completion time: off-line and on-line algorithms. *Mathematics of Operations Research*, 22:513–544, 1997.
13. V. Kann. *On the Approximability of NP-Complete Optimization Problems*. PhD thesis, Department of Numerical Analysis and Computing Science, Royal Institute of Technology, Stockholm, 1992.
14. R. Karp. *Reducibility Among Combinatorial Problems*, pages 85–103. Plenum Press, NY, 1972.
15. C. Kenyon-Mathieu and W. Schudy. How to rank with few errors. In *STOC*, pages 95–103, 2007.
16. E. L. Lawler, J. K. Lenstra, A. H. G. Rinnooy Kan, and D. B. Shmoys. Sequencing and scheduling: Algorithms and complexity. In S. C. Graves, A. H. G. Rinnooy Kan, and P. Zipkin, editors, *Handbooks in Operations Research and Management Science*, volume 4, pages 445–552. North-Holland, 1993.
17. A. Lempel and j. . I. n. . . v. . . y. . . p. . . I. Cederbaum, title = Minimum feedback arc and vertex sets of a directed graph.
18. M. Luby and N. Nisan. A parallel approximation algorithm for positive linear programming. In *STOC*, pages 448–457, 1993.
19. P. Pardalos and D. Du. *Handbook of Combinatorial Optimization: Supplement*, volume 1. Springer, 1999.

20 A. Plotkin, D. Shmoys, and E. Tardos. Fast Approximation Algorithms for Fractional Packing and Covering Problems. *Mathematics of Operation Research*, 20, 1995.

21 P. D. Seymour. Packing directed circuits fractionally. *Combinatorica*, 15(2):281–288, 1995.

22 W. T. Trotter. *Combinatorics and Partially Ordered Sets: Dimension Theory*. Johns Hopkins Series in the Mathematical Sciences. The Johns Hopkins University Press, 1992.

23 A. van Zuylen, R. Hegde, K. Jain, and D. P. Williamson. Deterministic pivoting algorithms for constrained ranking and clustering problems. In *SODA*, pages 405–414, 2007.

24 A. van Zuylen and D. P. Williamson. Deterministic pivoting algorithms for constrained ranking and clustering problems. *Math. Oper. Res.*, 34(3):594–620, 2009.

25 V. V. Vazirani. *Approximation Algorithms*. Springer, 2001.

Participants

- Sivan Albagli
Technion – Haifa, IL
- Evripidis Bampis
UPMC – Paris, FR
- Niv Buchbinder
Open University – Israel, IL
- Ed G. Coffman Jr.
Columbia University, US
- Pierre-Francois Dutot
INRIA Rhône-Alpes, FR
- Leah Epstein
University of Haifa, IL
- Lisa K. Fleischer
Dartmouth Coll. – Hanover, US
- Michael D. Grigoriadis
Rutgers Univ. – Piscataway, US
- Magnús M. Halldórsson
Reykjavik University, IS
- Klaus Jansen
Universität Kiel, DE
- David S. Johnson
AT&T Res. – Florham Park, US
- Howard Karloff
AT&T Res. – Florham Park, US
- Marek Karpinski
Universität Bonn, DE
- Samir Khuller
University of Maryland, US
- Asaf Levin
Technion – Haifa, IL
- Alejandro Lopez-Ortiz
University of Waterloo, CA
- Dániel Marx
HU Berlin, DE
- Monaldo Mastrolilli
IDSIA – Lugano, CH
- Claire Mathieu
Brown Univ. – Providence, US
- Ernst W. Mayr
TU München, DE
- Nicole Megow
MPI für Informatik – Saarbrücken, DE
- Rolf H. Möhring
TU Berlin, DE
- Seffi Naor
Technion – Haifa, IL
- Kirk Pruhs
University of Pittsburgh, US
- Christina Robenek
Universität Kiel, DE
- Adi Rosén
Univ. Paris-Diderot, CNRS, FR
- Nicolas Schabanel
Univ. Paris-Diderot, CNRS, FR
- Baruch Schieber
IBM TJ Watson Res. Center, US
- Ildiko Schlotter
Budapest Univ. of Technology & Economics, HU
- Ilka Schnoor
Universität Kiel, DE
- Jiri Sgall
Charles University – Prague, CZ
- Hadas Shachnai
Technion – Haifa, IL
- Alexander Souza
HU Berlin, DE
- Frits C.R. Spieksma
K.U. Leuven, BE
- Ola Svensson
KTH – Stockholm, SE
- Tami Tamir
The Interdisciplinary Center – Herzliya, IL
- Rob van Stee
MPI für Informatik - Saarbrücken, DE
- Gerhard Woeginger
TU Eindhoven, NL
- Prudence W. H. Wong
University of Liverpool, GB
- Neal E. Young
University of California – Riverside, US
- Shmuel Zaks
Technion – Haifa, IL