



DAGSTUHL REPORTS

Volume 3, Issue 3, March 2013

Computational Geometry (Dagstuhl Seminar 13101) <i>Otfried Cheong, Kurt Mehlhorn, and Monique Teillaud</i>	1
Scheduling (Dagstuhl Seminar 13111) <i>Susanne Albers, Onno J. Bozma, and Kirk Pruhs</i>	24
Bidimensional Structures: Algorithms, Combinatorics and Logic (Dagstuhl Seminar 13121) <i>Erik D. Demaine, Fedor V. Fomin, MohammadTaghi Hajiaghayi, and Dimitrios M. Thilikos</i>	51
Future Internet (Dagstuhl Seminar 13131) <i>Jon Crowcroft, Markus Fidler, Klara Nahrstedt, and Ralf Steinmetz</i>	75

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, 2013

Bibliographic information published by the Deutsche Nationalbibliothek

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

License

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



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

- Attribution: The work must be attributed to its authors.

The copyright is retained by the corresponding authors.

Aims and Scope

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

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

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

Editorial Board

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

Editorial Office

Marc Herbstritt (*Managing Editor*)

Jutka Gasirowski (*Editorial Assistance*)

Thomas Schillo (*Technical Assistance*)

Contact

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

Digital Object Identifier: 10.4230/DagRep.3.3.i

www.dagstuhl.de/dagrep

Computational Geometry

Edited by

Otfried Cheong¹, Kurt Mehlhorn², and Monique Teillaud³

1 KAIST – Daejeon, KR, otfried@kaist.edu

2 MPI für Informatik – Saarbrücken, DE, mehlhorn@mpi-inf.mpg.de

3 INRIA Sophia Antipolis – Méditerranée, FR, Monique.Teillaud@inria.fr

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 13101 “Computational Geometry”. The seminar was held from 3rd to 8th March 2013 and 47 senior and young researchers from various countries and continents attended it. Recent developments in the field were presented and new challenges in computational geometry were identified.

This report collects abstracts of the talks and a list of open problems.

Seminar 03.–08. March, 2013 – www.dagstuhl.de/13101

1998 ACM Subject Classification F.2 Analysis of Algorithms and Problem Complexity, G.2 Discrete Mathematics, G.4 Mathematical Software

Keywords and phrases Algorithms, geometry, theory, approximation, implementation, combinatorics, topology

Digital Object Identifier 10.4230/DagRep.3.3.1

Edited in cooperation with Lena Schlipf

1 Executive Summary

Otfried Cheong

Kurt Mehlhorn

Monique Teillaud

License © Creative Commons BY 3.0 Unported license
© Otfried Cheong, Kurt Mehlhorn, and Monique Teillaud

Computational Geometry and its Evolution

The field of computational geometry is concerned with the design, analysis, and implementation of algorithms for geometric and topological problems, which arise in a wide range of areas, including computer graphics, CAD, robotics computer vision, image processing, spatial databases, GIS, molecular biology, and sensor networks. Since the mid 1980s, computational geometry has arisen as an independent field, with its own international conferences and journals.

In the early years mostly theoretical foundations of geometric algorithms were laid and fundamental research remains an important issue in the field. Meanwhile, as the field matured, researchers have started paying close attention to applications and implementations of geometric and topological algorithms. Several software libraries for geometric computation (e.g. LEDA, CGAL, CORE) have been developed. Remarkably, this emphasis on applications and implementations has emerged from the originally theoretically oriented computational geometry community itself, so many researchers are concerned now with theoretical foundations as well as implementations.



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

Computational Geometry, *Dagstuhl Reports*, Vol. 3, Issue 3, pp. 1–23

Editors: Otfried Cheong, Kurt Mehlhorn, and Monique Teillaud



Dagstuhl Reports

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

Seminar Topics

The seminar presented recent developments in the field and identified new challenges for computational geometry. Below we list some of the most interesting subareas of the field at this stage, covering both theoretical and practical issues in computational geometry.

- *Theoretical foundations* of computational geometry lie in combinatorial geometry and its algorithmic aspects. They are of an enduring relevance for the field, particularly the design and the analysis of efficient algorithms require deep theoretical insights.
- *Geometric Computing* has become an integral part of the research in computational geometry. Besides general software design questions, especially *robustness* of geometric algorithms is important. Several methods have been suggested and investigated to make geometric algorithms numerically robust while keeping them efficient, which lead to interaction with the field of computer algebra, numerical analysis, and topology.
- *Computational topology* concentrates on the properties of geometric objects that go beyond metric representation: modeling and reconstruction of surfaces, shape similarity and classification, and persistence are key concepts with applications in molecular biology, computer vision, and geometric databases.
- In its early years, computational geometry concentrated on low dimensions. *High-dimensional data* has become very important recently, in particular, in work related to machine learning and data analysis. Standard solutions suffer from the curse of dimensionality. This has led to extensive work on dimension-reduction and embedding techniques.
- Various *applications* such as robotics, GIS, or CAD lead to interesting variants of the *classical topics* originally investigated, including convex hulls, Voronoi diagrams and Delaunay triangulations, and geometric data structures. For example, Voronoi diagrams and nearest-neighbor data structures under various metrics have turned out to be useful for many applications and are being investigated intensively.
- *Massive geometric data* sets are being generated by networks of sensors at unprecedented spatial and temporal scale. How to store, analyze, query, and visualize them has raised several algorithmic challenges. New computational models have been proposed to meet these challenges, e.g., streaming model, communication-efficient algorithms, and maintaining geometric summaries.

Participants

47 researchers from various countries and continents attended the seminar, showing the strong interest of the community for this event. The feedback from participants was very positive.

Dagstuhl seminars on computational geometry have been organized in a two year rhythm since a start in 1990. They have been extremely successful both in disseminating the knowledge and identifying new research thrusts. Many major results in computational geometry were first presented in Dagstuhl seminars, and interactions among the participants at these seminars have led to numerous new results in the field. These seminars have also played an important role in bringing researchers together, fostering collaboration, and exposing young talent to the seniors of the field. They have arguably been the most influential meetings in the field of computational geometry.

No other meeting in our field allows young researchers to meet with, get to know, and work with well-known and senior scholars to the extent possible at Dagstuhl. To accommodate new, younger researchers, the organizers held a *lottery* for the first time this year. From an initial list of selected researchers, we randomly selected a certain number of senior, young, and female participants. Researchers on the initial list who were not selected by the lottery were notified by us separately per email, so that they knew that they were not forgotten, and to reassure them that—with better luck—they will have another chance in future seminars.

We believe that the lottery created space to invite younger researchers, rejuvenating the seminar, while keeping a large group of senior and well-known scholars involved. The seminar was much “younger” than in the past, and certainly more “family-friendly.” Five young children roaming the premises created an even cosier atmosphere than we are used in Dagstuhl. Without decreasing the quality of the seminar, we had a more balanced attendance than in the past. Feedback from both seminar participants and from researchers who were not selected was uniformly positive.

Dagstuhl itself is a great strength of the seminar. Dagstuhl allows people to really meet and socialize, providing them with a wonderful atmosphere of a unique closed and pleasant environment, which is highly beneficial to interactions. Therefore, we warmly thank the scientific, administrative and technical staff at Schloss Dagstuhl!

2 Table of Contents**Executive Summary**

<i>Otfried Cheong, Kurt Mehlhorn, and Monique Teillaud</i>	1
--	---

Overview of Talks

Union of Random Minkowski Sums and Network Vulnerability Analysis <i>Pankaj Kumar Agarwal</i>	6
Fast Point Location for Easy Points <i>Boris Aronov</i>	6
Geometry-driven collapses for simplifying Cech complexes <i>Dominique Attali</i>	7
Exact Symbolic-Numeric Computation of Planar Algebraic Curves <i>Eric Berberich</i>	7
Four Soviets Walk the Dog – with an Application to Alt’s Conjecture <i>Kevin Buchin</i>	8
Trajectory Grouping Structures <i>Maike Buchin</i>	9
Approximate Shortest Descending Paths <i>Siu-Wing Cheng</i>	9
Graph Induced Complex on Point Data <i>Tamal K. Dey</i>	10
Geometric Input Models <i>Anne Driemel</i>	10
Efficiently hex-meshing things with topology <i>Jeff Erickson</i>	11
Theory Meets Practice: Two Videos <i>Sándor Fekete</i>	11
Random hypergraphs and small silhouettes <i>Marc Glisse</i>	12
Simplifying inclusion-exclusion formulas <i>Xavier Goaoc</i>	12
Fréchet Queries in Geometric Trees <i>Joachim Gudmundsson</i>	12
Geometric properties of space-filling curves: some results and open problems <i>Herman J. Haverkort</i>	13
Qualitative symbolic perturbations <i>Menelaos Karavelas</i>	13
Bottleneck Non-Crossing Matching in the Plane <i>Matthew J. Katz</i>	14
On the Complexity of Higher Order Abstract Voronoi Diagrams <i>Rolf Klein</i>	14

On numerical algorithms for the topology of curves with simple singularities <i>Guillaume Moroz</i>	15
Output-Sensitive Well-Separated Pair Decompositions for Dynamic Point Sets <i>David M. Mount</i>	15
Improved Approximation for Geometric Unique Coverage Problems <i>Yoshio Okamoto</i>	16
Kinetic data structures in the black-box model <i>Marcel J. M. Roeloffzen</i>	16
α -Visibility <i>Jörg-Rüdiger Sack</i>	17
Parallel computation of the Hausdorff distance between shapes <i>Ludmila Scharf</i>	17
Convex Transversals <i>Lena Schlipf</i>	17
Delaunay and other triangulations of moving point sets: What's going on? <i>Micha Sharir</i>	18
Towards Elastic Shape Matching <i>Fabian Stehn</i>	19
Extended Formulations for polytopes <i>Hans Raj Tiwary</i>	19
Reverse stabbing queries in galleries <i>Suresh Venkatasubramanian</i>	20
A Faster Algorithm for Computing Motorcycle Graphs <i>Antoine Vigneron</i>	20
Towards Understanding Gaussian Weighted Graph Laplacian <i>Yusu Wang</i>	20
Open Problems	21
Participants	23

3 Overview of Talks

3.1 Union of Random Minkowski Sums and Network Vulnerability Analysis

Pankaj Kumar Agarwal (Duke University – Durham, US)

License  Creative Commons BY 3.0 Unported license
© Pankaj Kumar Agarwal

Let $C = \{C_1, \dots, C_n\}$ be a set of n pairwise-disjoint convex s -gons, for some constant s , and let π be a probability density function (pdf) over the non-negative reals. For each i , let K_i be the Minkowski sum of C_i with a disk of radius r_i , where each r_i is a random non-negative number drawn independently from the distribution determined by π . We show that the expected complexity of the union of K_1, \dots, K_n is $O(n \log n)$, for any pdf π ; the constant of proportionality depends on s , but not on the pdf.

Next, we consider the following problem that arises in analyzing the vulnerability of a network under a physical attack. Let $G = (V, E)$ be a planar geometric graph where E is a set of n line segments with pairwise-disjoint relative interiors. Let $\phi : \mathbb{R}_{\geq 0} \rightarrow [0, 1]$ be an *edge failure probability function*, where a physical attack at a location $x \in \mathbb{R}^2$ causes an edge e of E at distance r from x to fail with probability $\phi(r)$; we assume that ϕ is of the form $1 - \Pi(x)$, where Π is a cumulative distribution function on the non-negative reals. The goal is to compute the most *vulnerable* location for G , i.e., the location of the attack that maximizes the expected number of failing edges of G . Using our bound on the complexity of the union of random Minkowski sums, we present a near-linear Monte-Carlo algorithm for computing a location that is an approximately most vulnerable location of attack for G .

3.2 Fast Point Location for Easy Points

Boris Aronov (Polytechnic Inst. of NYU, US)

License  Creative Commons BY 3.0 Unported license
© Boris Aronov

Joint work of Aronov, Boris; de Berg, Mark; Roeloffzen, Marcel; Speckmann, Bettina

Main reference B. Aronov, M. de Berg, M. Roeloffzen, B. Speckmann, “Distance-Sensitive Planar Point Location,” *WADS 2013*, to appear.

URL <http://www.wads.org>

Let \mathcal{S} be a connected planar polygonal subdivision with n edges and of total area 1. We present a data structure for point location in \mathcal{S} where queries with points far away from any region boundary are answered faster. More precisely, we show that point location queries can be answered in time $O(1 + \min(\log \frac{1}{\Delta_p}, \log n))$, where Δ_p is the distance of the query point p to the boundary of the region containing p . Our structure is based on the following result: any simple polygon P can be decomposed into a linear number of convex quadrilaterals with the following property: for any point $p \in P$, the quadrilateral containing p has area $\Omega(\Delta_p^2)$.

3.3 Geometry-driven collapses for simplifying Cech complexes

Dominique Attali (GIPSA Lab – Saint Martin d’Hères, FR)

License © Creative Commons BY 3.0 Unported license
© Dominique Attali

Joint work of Attali, Dominique; Lieutier, André

Main reference D. Attali, A. Lieutier, “Geometry driven collapses for converting a cech complex into a triangulation of a nicely triangulable shape,” arXiv:1304.3680v1 [cs.CG], 2013.

URL <http://arxiv.org/abs/1304.3680v1>

In many practical situations, the object of study is only known through a finite set of possibly noisy sample points. It is then desirable to try to recover the geometry and the topology of the object from this information.

In this talk, we will focus on an approach that approximates a shape from a set of sample points by returning the Rips complex of the points. Given a point set P and a scale parameter r , the Rips complex is the simplicial complex whose simplices are subsets of points in P with diameter at most $2r$. Rips complexes have generally a size and dimension much too large to allow an explicit representation. Nonetheless, Rips complexes enjoy the property to be completely determined by the graph of their vertices and edges which thus provide a compressed form of storage (quadratic in the number of data points and linear in the ambient dimension). This suggests to reconstruct a shape by first building the Rips complex of the data points at some scale (encoded with its vertices and edges) and second by simplifying the result through a sequence of elementary operations. In previous work, we formulated conditions under which the Rips complex of the point set at some scale reflects the homotopy type of the shape [1, 2]. In this talk, we give conditions under which the complex can be transformed by a sequence of collapses into a triangulation of the shape [3].

References

- 1 D. Attali and A. Lieutier. Reconstructing shapes with guarantees by unions of convex sets. In *Proc. SoCG 2010*, pp. 344–353, 2010.
- 2 D. Attali, A. Lieutier, and D. Salinas. Vietoris-Rips complexes also provide topologically correct reconstructions of sampled shapes. *Computational Geometry: Theory and Applications (CGTA)*, 46:448–465, 2012.
- 3 D. Attali and A. Lieutier. Geometry driven collapses for converting a cech complex into a triangulation of a nicely triangulable shape. *arXiv preprint arXiv:1304.3680*, 2013.

3.4 Exact Symbolic-Numeric Computation of Planar Algebraic Curves

Eric Berberich (MPI für Informatik – Saarbrücken, DE)

License © Creative Commons BY 3.0 Unported license
© Eric Berberich

Joint work of Berberich, Eric; Emelyanenko, Pavel; Kobel, Alexander; Sagraloff, Michael

Main reference Eric Berberich, Pavel Emelyanenko, Alexander Kobel, Michael Sagraloff, “Exact Symbolic-Numeric Computation of Planar Algebraic Curves,” arXiv:1201.1548v1 [cs.CG], 2012.

URL <http://arxiv.org/abs/1201.1548v1>

We present a certified and complete algorithm to compute arrangements of real planar algebraic curves. It computes the decomposition of the plane induced by a finite number of algebraic curves in terms of a cylindrical algebraic decomposition. From a high-level perspective, the overall method splits into two main subroutines, namely an algorithm denoted Bisolve to isolate the real solutions of a zero-dimensional bivariate system, and an algorithm denoted GeoTop to compute the topology of a single algebraic curve. Compared to existing

approaches based on elimination techniques, we considerably improve the corresponding lifting steps in both subroutines. As a result, generic position of the input system is never assumed, and thus our algorithm never demands for any change of coordinates. In addition, we significantly limit the types of symbolic operations involved, that is, we only use resultant and gcd computations as purely symbolic operations. The latter results are achieved by combining techniques from different fields such as (modular) symbolic computation, numerical analysis and algebraic geometry. We have implemented our algorithms as prototypical contributions to the C++-project Cgal. We exploit graphics hardware to expedite the remaining symbolic computations. We have also compared our implementation with the current reference implementations, that is, Lgp and Maple’s Isolate for polynomial system solving, and Cgal’s bivariate algebraic kernel for analyses and arrangement computations of algebraic curves. For various series of challenging instances, our exhaustive experiments show that the new implementations outperform the existing ones.

3.5 Four Soviets Walk the Dog – with an Application to Alt’s Conjecture

Kevin Buchin (TU Eindhoven, NL)

License  Creative Commons BY 3.0 Unported license
© Kevin Buchin

Joint work of Buchin, Kevin; Buchin, Maïke; Meulemans, Wouter; Mulzer, Wolfgang

Main reference K. Buchin, M. Buchin, W. Meulemans, W. Mulzer, “Four Soviets Walk the Dog-with an Application to Alt’s Conjecture,” arXiv:1209.4403v2 [cs.CG], 2012.

URL <http://arxiv.org/abs/1209.4403v2>

Given two polygonal curves in the plane, there are several ways to define a measure of similarity between them. One measure that has been extremely popular in the past is the Fréchet distance. Since it has been proposed by Alt and Godau in 1992, many variants and extensions have been described. However, even 20 years later, the original $O(n^2 \log n)$ algorithm by Alt and Godau for computing the Fréchet distance remains the state of the art (here n denotes the number of vertices on each curve). This has led Helmut Alt to conjecture that the associated decision problem is 3SUM-hard. In recent work, Agarwal et al. show how to break the quadratic barrier for the discrete version of the Fréchet distance, where we consider sequences of points instead of polygonal curves. Building on their work, we give an algorithm to compute the Fréchet distance between two polygonal curves in time $O(n^2(\log n)^{(1/2)}(\log \log n)^{(3/2)})$ on a pointer machine and in time $O(n^2(\log \log n)^2)$ on a word RAM. Furthermore, we show that there exists an algebraic decision tree for the Fréchet problem of depth $O(n^{(2-\epsilon)})$, for some $\epsilon > 0$. This provides evidence that computing the Fréchet distance may not be 3SUM-hard after all and reveals an intriguing new aspect of this well-studied problem.

3.6 Trajectory Grouping Structures

Maike Buchin (TU Eindhoven, NL)

License © Creative Commons BY 3.0 Unported license
© Maike Buchin

Joint work of Buchin, Kevin; Buchin, Maike; van Kreveld, Marc; Speckmann, Bettina; Staals, Frank
Main reference K. Buchin, M. Buchin, M. van Kreveld, B. Speckmann, F. Staals, “Trajectory Grouping Structures,” arXiv:1303.6127v1 [cs.CG], 2013.
URL <http://arxiv.org/abs/1303.6127v1>

The collective motion of a set of moving entities like people, birds, or other animals, is characterized by groups arising, merging, splitting, and ending. Given the trajectories of these entities, we define and model a structure that captures all of such changes using the Reeb graph, a concept from topology. The trajectory grouping structure has three natural parameters that allow more global views of the data in group size, group duration, and entity inter-distance. We prove complexity bounds on the maximum number of maximal groups that can be present, and give algorithms to compute the grouping structure efficiently. We also study how the trajectory grouping structure can be made robust, that is, how brief interruptions of groups can be disregarded in the global structure, adding a notion of persistence to the structure. Furthermore, we showcase the results of experiments using data generated by the NetLogo flocking model and from the Starkey project. The Starkey data describe the movement of elk, deer, and cattle. Although there is no ground truth for the grouping structure in this data, the experiments show that the trajectory grouping structure is plausible and has the desired effects when changing the essential parameters. Our research provides the first complete study of trajectory group evolution, including combinatorial, algorithmic, and experimental results.

3.7 Approximate Shortest Descending Paths

Siu-Wing Cheng (HKUST – Kowloon, HK)

License © Creative Commons BY 3.0 Unported license
© Siu-Wing Cheng

Joint work of Cheng, Siu-Wing; Jin, Jiongxin
Main reference S.-W. Cheng, J. Jin, “Approximate Shortest Descending Paths,” in Proc. of the 24th Annual ACM-SIAM Symp. on Discrete Algorithms (SODA’13), pp. 144–155, SIAM, 2013.
URL <http://knowledgecenter.siam.org/0236-000023/>

We present an approximate algorithm for the shortest descending path problem. Given a source s and a destination t on a terrain, a shortest descending path from s to t is a path of minimum Euclidean length on the terrain subject to the constraint that the height decreases monotonically as we traverse that path from s to t . Given any $\epsilon \in (0, 1)$, our algorithm returns in $O(n^4 \log(n/\epsilon))$ time a descending path of length at most $1 + \epsilon$ times the optimum. This is the first algorithm whose running time is polynomial in n and $\log(1/\epsilon)$ and independent of the terrain geometry.

3.8 Graph Induced Complex on Point Data

Tamal K. Dey (Ohio State University, US)

License © Creative Commons BY 3.0 Unported license
© Tamal K. Dey

Joint work of Dey, Tamal K.; Fan, Fengtao; Wang, Yusu

Main reference T.K. Dey, F. Fan, Y. Wang, “Graph Induced Complex on Point Data,” arXiv:1304.0662v1 [cs.CG]; to appear in Proc. of the 29th Annual Symp. on Computational Geometry 2013.

URL <http://arxiv.org/abs/1304.0662v1>

The efficiency of extracting topological information from point data depends largely on the complex that is built on top of the data points. From a computational viewpoint, the most favored complexes for this purpose have so far been Vietoris-Rips and witness complexes. While the Vietoris-Rips complex is simple to compute and is a good vehicle for extracting topology of sampled spaces, its size is huge—particularly in high dimensions. The witness complex on the other hand enjoys a smaller size because of a subsampling, but fails to capture the topology in high dimensions unless imposed with extra structures. We investigate a complex called the *graph induced complex* that, to some extent, enjoys the advantages of both. It works on a subsample but still retains the power of capturing the topology as the Vietoris-Rips complex. It only needs a graph connecting the original sample points from which it builds a complex on the subsample thus taming the size considerably. We show that, using the graph induced complex one can (i) infer the one dimensional homology of a manifold from a very lean subsample, (ii) reconstruct a surface in three dimension from a sparse subsample without computing Delaunay triangulations, (iii) infer the persistent homology groups of compact sets from a sufficiently dense sample. We provide experimental evidences in support of our theory.

3.9 Geometric Input Models

Anne Driemel (Utrecht University, NL)

License © Creative Commons BY 3.0 Unported license
© Anne Driemel

Joint work of Driemel, Anne; Har-Peled, Sariel; Wenk, Carola; Raichel, Benjamin

The worst-case analysis of the running time and space complexities as a function of the input size is a fundamental method in algorithm design. However, it fails to describe the actual behavior when the worst case is a contrived geometric configuration which would never occur in practice. There are different approaches to reasoning about algorithms and data structures for real data that allow a theoretical analysis with provable bounds. I will outline some techniques we used and the results that we achieved. In particular, I will talk about two results (i) an approximation algorithm for the Fréchet distance [1] and (ii) bounding the complexity of Voronoi diagrams on terrains [2]. This research was carried out as a part of my PhD studies and is joint work with Carola Wenk (Tulane University), Sariel Har-Peled (UIUC) and Benjamin Raichel (UIUC).

References

- 1 Anne Driemel, Sariel Har-Peled, and Carola Wenk. Approximating the Fréchet distance for realistic curves in near linear time. *Discrete & Computational Geometry*, 48(1):94–127, 2012.
- 2 Anne Driemel, Sariel Har-Peled, and Benjamin Raichel. On the expected complexity of Voronoi diagrams on terrains. In *Proc. SoCG 2012*, pages 101–110, 2012.

3.10 Efficiently hex-meshing things with topology

Jeff Erickson (*University of Illinois – Urbana, US*)

License © Creative Commons BY 3.0 Unported license
© Jeff Erickson

Main reference J. Erickson, “Efficiently hex-meshing things with topology,” to appear in Proc. of the 28th Annual Symp. on Computational Geometry 2013.

URL <http://www.cs.uiuc.edu/~jeffe/pubs/hexmesh.html>

A topological quadrilateral mesh Q of a connected surface in \mathbb{R}^3 can be extended to a topological hexahedral mesh of the interior domain Ω if and only if Q has an even number of quadrilaterals and no odd cycle in Q bounds a surface inside Ω . Moreover, if such a mesh exists, the required number of hexahedra is within a constant factor of the minimum number of tetrahedra in a triangulation of Ω that respects Q . Finally, if Q is given as a polyhedron in \mathbb{R}^3 with quadrilateral facets, a topological hexahedral mesh of the polyhedron can be constructed in polynomial time if such a mesh exists. All our results extend to domains with disconnected boundaries. Our results naturally generalize results of Thurston, Mitchell, and Eppstein for genus-zero and bipartite meshes, for which the odd-cycle criterion is trivial.

3.11 Theory Meets Practice: Two Videos

Sándor Fekete (*TU Braunschweig, DE*)

License © Creative Commons BY 3.0 Unported license
© Sándor Fekete

Joint work of Fekete, Sándor; Friedrichs, Stephan; Kröller, Alexander; Schmidt, Christiane; Borrmann, Dorit; de Rezende, Pedro J.; de Souza, Cid C.; Tozoni, Davi C.; Becker, Aaron; Lee, SengKyou; McLurkin, James

One of the driving engines of Computational Geometry is the interaction with practical problems; one of the application areas with strong ties to geometry is the field of robotics. In this talk, I present two videos that document ongoing collaborations with colleagues from robotics.

The first [1] considers exploration and triangulation with a swarm of small robots with relatively few individual capabilities; we develop ideas, provide theory and present a practical demonstration of how such a swarm can be used to explore an unknown territory, and guard it. This is joint work with colleagues from Rice University (USA).

The second [2] shows how building detailed three-dimensional maps with a robot platform that carries a powerful laserscanner is related to the classical Art Gallery Problem (AGP). We develop different methods for solving such problems to optimality, and demonstrate the resulting application. This is joint work with colleagues from the University of Campinas (Brazil) and Jacobs University Bremen (Germany).

References

- 1 A. Becker, S.P. Fekete, A. Kröller, L.S. Kyou, J. McLurkin and C. Schmidt. Triangulating Unknown Environments Using Robot Swarms, Video and abstract. To appear in: *Proc. SoCG 2013*.
- 2 D. Borrmann, P.J. de Rezende, C.C. de Souza, S.P. Fekete, S. Friedrichs, A. Kröller, A. Nüchter, C. Schmidt and D.C. Tozoni. Point Guards and Point Clouds: Solving General Art Gallery Problems, Video and abstract. To appear in: *Proc. SoCG 2013*

3.12 Random hypergraphs and small silhouettes

Marc Glisse (INRIA Saclay – Île-de-France – Orsay, FR)

License © Creative Commons BY 3.0 Unported license
© Marc Glisse

Joint work of Devillers, Olivier; Glisse, Marc; Goaoc, Xavier; Lazard, Sylvain; Michel, Julien; Pouget, Marc

We present a new simple scheme for the analysis of random geometric structures, which we illustrate on convex hulls and Delaunay triangulations. We then introduce some refinements of the analysis which tighten the bounds and give large-deviation-related results. Those refinements are finally used to deduce a worst-case bound on the size of the silhouettes of a random polytope.

3.13 Simplifying inclusion-exclusion formulas

Xavier Goaoc (INRIA Lorraine, FR)

License © Creative Commons BY 3.0 Unported license
© Xavier Goaoc

Joint work of Goaoc, Xavier; Matoušek, Jiří; Paták, Pavel; Safernová, Zuzana; Tancer, Martin

Main reference X. Goaoc, J. Matoušek, P. Paták, Z. Safernová, M. Tancer, “Simplifying inclusion-exclusion formulas,” arXiv:1207.2591v1 [math.CO] .

URL <http://arxiv.org/abs/1207.2591v1>

Let $F = \{F_1, F_2, \dots, F_n\}$ be a family of n sets on a ground set X , such as a family of balls in \mathbb{R}^d . For every finite measure μ on X , such that the sets of F are measurable, the classical *inclusion-exclusion formula* asserts that $\mu(F_1 \cup F_2 \cup \dots \cup F_n) = \sum_{I: \emptyset \neq I \subseteq [n]} (-1)^{|I|+1} \mu(\bigcap_{i \in I} F_i)$; that is, the measure of the union is expressed using measures of various intersections. The number of terms in this formula is exponential in n , and a significant amount of research, originating in applied areas, has been devoted to constructing simpler formulas for particular families F . We provide an upper bound valid for an arbitrary F : we show that every system F of n sets with m nonempty fields in the Venn diagram admits an inclusion-exclusion formula with $m^{O(\log^2 n)}$ terms and with ± 1 coefficients, and that such a formula can be computed in $m^{O(\log^2 n)}$ expected time. We also construct systems of n sets on n points for which every valid inclusion-exclusion formula has the sum of absolute values of the coefficients at least $\Omega(n^{3/2})$.

3.14 Fréchet Queries in Geometric Trees

Joachim Gudmundsson (University of Sydney, AU)

License © Creative Commons BY 3.0 Unported license
© Joachim Gudmundsson

Joint work of Gudmundsson, Joachim; Smid, Michiel

Let T be a tree that is embedded in the plane and let $\Delta > 0$ be a real number. The aim is to preprocess T into a data structure, such that, for any query polygonal path Q , we can decide if T contains a path P whose Fréchet distance $\delta_F(P, Q)$ to Q is less than Δ . We present an efficient data structure that solves an approximate version of this problem, for the case when T is c -packed and each of the edges of T and Q has length $\Omega(\Delta)$ (not required if T is a path): If the data structure returns NO, then there is no such path P . If it returns YES, then $\delta_F(P, Q) \leq \sqrt{2}(1 + \epsilon)\Delta$ if Q is a line segment, and $\delta_F(P, Q) \leq 3(1 + \epsilon)\Delta$ otherwise.

3.15 Geometric properties of space-filling curves: some results and open problems

Herman J. Haverkort (TU Eindhoven, NL)

License © Creative Commons BY 3.0 Unported license
© Herman J. Haverkort

Main reference H. Haverkort, “Recursive tilings and space-filling curves with little fragmentation,” *Journal of Computational Geometry*, 2(1):92–127, 2011.

URL <http://www.jocg.org/index.php/jocg/article/view/68>

A space-filling curve is a continuous surjective function f that maps the unit interval $[0,1]$ to a higher-dimensional region, such as the unit square. Such a curve is usually defined on the basis of a recursive tiling, such that the curve traverses the tiles of each level of the tiling one by one, and the curve can be parameterized such that the union of all points $f(t)$ over all t in $[a,b]$ is a region of measure exactly $b-a$. In this presentation I focus on two open problems about space-filling curves.

1. The Arrwvid number of a three-dimensional space-filling curve is the smallest number A , such that any ball with volume B can be covered by A pieces of the curve of total size $O(B)$. The three-dimensional curve with the lowest known Arrwvid number has Arrwvid number 4. We can prove that this is a lower bound for any three-dimensional curve that traverses the tiles of a recursive tiling with convex tiles one by one. Can we also prove this lower bound for curves that are not based on convex tiles? We would be able to prove this if we could prove a certain relation between the number of vertices, the number of tiles and the number of vertex-tile incidences that holds for any “reasonable” tiling in three dimensions.

2. The dilation of a two-dimensional space-filling curve is the maximum, over all a,b in $[0,1]$, of the squared distance between $f(a)$ and $f(b)$, divided by $(b-a)$. We can prove that each two-dimensional space-filling curve must have dilation at least $4/\pi$. The curve with the lowest known dilation has dilation 4. Can we improve the lower bound—or can we find a space-filling curve with dilation less than 4?

3.16 Qualitative symbolic perturbations

Menelaos Karavelas (University of Crete – Heraklion, GR)

License © Creative Commons BY 3.0 Unported license
© Menelaos Karavelas

Joint work of Devillers, Olivier; Karavelas, Menelaos; Teillaud, Monique

Main reference O. Devillers, M. Karavelas, M. Teillaud, “Qualitative Symbolic Perturbation: a new geometry-based perturbation framework,” HAL, RR-8153, 2012.

URL <http://hal.inria.fr/hal-00758631/>

In classical *Symbolic Perturbations*, degeneracies are resolved by using a sequence of predicates obtained by algebraic substitution of polynomials in ε to the input. Instead of a single perturbation, we propose to use a sequence of (simpler) perturbations and to look at their effect geometrically instead of algebraically. We obtain solutions for Apollonius predicates which were not solvable using the algebraic approach.

3.17 Bottleneck Non-Crossing Matching in the Plane

Matthew J. Katz (Ben Gurion University – Beer Sheva, IL)

License © Creative Commons BY 3.0 Unported license
© Matthew J. Katz

Joint work of Abu-Affash, A. Karim; Carmi, Paz; Katz, Matthew J.; Trabelsi Yohai
Main reference A. Karim Abu-Affash, P. Carmi, M.J. Katz, Y. Trabelsi, “Bottleneck Non-Crossing Matching in the Plane,” in Proc. of the 20th Annual European Symp. on Algorithms (ESA’12), LNCS, Vol. 7501, pp. 36–47, Springer, 2012.

URL http://dx.doi.org/10.1007/978-3-642-33090-2_5

Let P be a set of $2n$ points in the plane, and let M_C (resp., M_{NC}) denote a bottleneck matching (resp., a bottleneck non-crossing matching) of P . We study the problem of computing M_{NC} . We first prove that the problem is NP-hard and does not admit a PTAS. Then, we present an $O(n^{1.5} \log^{0.5} n)$ -time algorithm that computes a non-crossing matching M of P , such that $bn(M) \leq 2\sqrt{10} \cdot bn(M_{NC})$, where $bn(M)$ is the length of a longest edge in M . An interesting implication of our construction is that $bn(M_{NC})/bn(M_C) \leq 2\sqrt{10}$.

3.18 On the Complexity of Higher Order Abstract Voronoi Diagrams

Rolf Klein (Universität Bonn, DE)

License © Creative Commons BY 3.0 Unported license
© Rolf Klein

Joint work of Bohler, Cecilia; Cheilaris, Panagiotis; Klein, Rolf; Liu, Chih-Hung; Papadopoulou, Evanthia; Zavershynskiy, Maksymolf
Main reference C. Bohler, P. Cheilaris, R. Klein, C. Liu, E. Papadopoulou, M. Zavershynskiy, “On the Complexity of Higher Order Abstract Voronoi Diagrams,” to appear in the Proc. of the 40th Int’l Colloquium on Automata, Languages and Programming (ICALP ’13), Riga, 2013.

Abstract Voronoi diagrams are based on bisecting curves enjoying simple combinatorial properties, rather than on the geometric notions of sites and circles. They serve as a unifying concept. Once the bisector system of any concrete type of Voronoi diagram is shown to fulfill the AVD properties, structural results and efficient algorithms become available without further effort. For example, the first optimal algorithms for constructing nearest Voronoi diagrams of disjoint convex objects, or of line segments under the Hausdorff metric, have been obtained this way.

In a concrete order- k Voronoi diagram, all points are placed into the same region that have the same k nearest neighbors among the given sites. This paper is the first to study abstract Voronoi diagrams of arbitrary order k . We prove that their complexity is upper bounded by $2k(n - k)$. So far, an $O(k(n - k))$ bound has been shown only for point sites in the Euclidean and L_p plane, and, very recently, for line segments. These proofs made extensive use of the geometry of the sites.

Our result on AVDs implies a $2k(n - k)$ upper bound for a wide range of cases for which only trivial upper complexity bounds were previously known, and a slightly sharper bound for the known cases.

3.19 On numerical algorithms for the topology of curves with simple singularities

Guillaume Moroz (INRIA Grand Est – Nancy, FR)

License  Creative Commons BY 3.0 Unported license
© Guillaume Moroz

Joint work of Moroz, Guillaume; Pouget, Marc

Let C be the planar curve defined by a polynomial equation $f(x, y) = 0$. If C is smooth, its topology can be computed with adaptive numerical algorithms. Otherwise, computing the topology requires a different set of tools that induce a significant gap between the analysis of smooth curves and singular curves, even with simple multiplicity structure. Such tools include the computation of a resultant, the subdivision until a global separation. We present works in progress to fill this gap for singular curves with simple self-intersections.

3.20 Output-Sensitive Well-Separated Pair Decompositions for Dynamic Point Sets

David M. Mount (University of Maryland – College Park, US)

License  Creative Commons BY 3.0 Unported license
© David M. Mount

Joint work of Park Eunhui; Mount, David M.

Main reference (Unpublished manuscript, submitted for publication)

The well-separated pair decomposition (WSPD) is a fundamental structure in computational geometry. Given a set of n points in d -dimensional space and a positive parameter s , it is known that there exists an s -WSPD of size $O(s^d n)$. While this is linear in n , the factor of s^d is a significant consideration when the dimension d is even a moderately large constant. The actual number of pairs may be much smaller than this worst-case bound, for example, if the points are clustered near a lower dimensional subspace. Batch WSPD constructions are output sensitive, but existing algorithms for maintaining the WSPD of a dynamic point set are not. In this paper we present output-sensitive algorithms for maintaining the WSPD of a dynamic point set under insertion and deletion.

3.21 Improved Approximation for Geometric Unique Coverage Problems

Yoshio Okamoto (University of Electro-Communications, Tokyo, JP)

License © Creative Commons BY 3.0 Unported license

© Yoshio Okamoto

Joint work of Ito, Takehiro; Nakano, Shin-ichi; Okamoto, Yoshio; Otachi, Yota; Uehara, Ryuhei; Uno, Takeaki; Uno, Yushi

Main reference T. Ito, S.-I. Nakano, Y. Okamoto, Y. Otachi, R. Uehara, T. Uno, Y. Uno, “A 4.31-Approximation for the Geometric Unique Coverage Problem on Unit Disks,” in Proc. of the 23rd Int’l Symp. on Algorithms and Computation (ISAAC’12), LNCS, Vol. 7676, pp. 372–381, Springer, 2012.

URL http://dx.doi.org/10.1007/978-3-642-35261-4_40

Main reference T. Ito, S.-I. Nakano, Y. Okamoto, Y. Otachi, R. Uehara, T. Uno, Y. Uno, “A Polynomial-Time Approximation Scheme for the Geometric Unique Coverage Problem on Unit Squares,” in Proc. of the 13th Scandinavian Symp. and Workshops on Algorithm Theory (SWAT’12), LNCS, Vol. 7357, pp. 24–35, Springer, 2012.

URL http://dx.doi.org/10.1007/978-3-642-31155-0_3

Given a set of points and a set of objects, both in the plane, we wish to find a subset of the objects that maximizes the number of points contained in exactly one object in the subset. Erlebach and van Leeuwen [1] introduced this problem as the geometric version of the unique coverage problem, and gave polynomial-time approximation algorithms. Their approximation ratios were 18 when the objects were unit disks, and 4 when the objects were axis-parallel unit squares (which was later improved to 2 by van Leeuwen [2]). We improve the approximation ratios to 4.31 for unit disks and $1 + \varepsilon$ for axis-parallel unit squares.

References

- 1 Erlebach, T., van Leeuwen, E.J. Approximating geometric coverage problems. In *Proc. SODA 2008*, pp. 1267–1276 (2008)
- 2 van Leeuwen, E. J. Optimization and approximation on systems of geometric objects. *Ph.D. Thesis*, University of Amsterdam, 2009

3.22 Kinetic data structures in the black-box model

Marcel J. M. Roeloffzen (TU Eindhoven, NL)

License © Creative Commons BY 3.0 Unported license

© Marcel J. M. Roeloffzen

Joint work of Berg, Mark de; Roeloffzen, Marcel; Speckmann, Bettina

Over the past decade, the kinetic-data-structures framework has become the standard in computational geometry for dealing with moving objects. A fundamental assumption underlying the framework is that the motions of the objects are known in advance. This assumption severely limits the applicability of KDSs. We study KDSs in the black-box model, which is a hybrid of the KDS model and the traditional time-slicing approach. In this more practical model we receive the position of each object at regular time steps and we have an upper bound on d_{\max} , the maximum displacement of any point in one time step.

In this talk we describe the black-box model and give an overview of the results obtained for maintaining the convex hull, Delaunay triangulation and compressed quadtree of a set of points in the black-box model. We also go into some more detail on the latest result on maintaining the Euclidean 2-center.

3.23 α -Visibility

Jörg-Rüdiger Sack (Carleton University – Ottawa, CA)

License © Creative Commons BY 3.0 Unported license
© Jörg-Rüdiger Sack

Joint work of Ghodsi, Mohammad; Maheshwari, Anil; Nouri, Mostafa; Sack, Jörg-Rüdiger; Zarrabi-Zadeh, Hamid
Main reference M. Ghodsi, A. Maheshwari, M. Nouri, J.-R. Sack, H. Zarrabi-Zadeh, “ α -Visibility,” in Proc. of the 13th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT’2012), LNCS, Vol. 7357, pp. 1–12, Springer, 2012.

URL http://dx.doi.org/10.1007/978-3-642-31155-0_1

We study a new class of visibility problems based on the notion of α -visibility. Given an angle α and a collection of line segments S in the plane, a segment t is said to be α -visible from a point p , if there exists an empty triangle with one vertex at p and the side opposite to p on t such that the angle at p is α . In this model of visibility, we study the classical variants of point visibility, weak and complete segment visibility, and the construction of the visibility graph. We also investigate the natural query versions of these problems, when α is either fixed or specified at query time.

3.24 Parallel computation of the Hausdorff distance between shapes

Ludmila Scharf (FU Berlin, DE)

License © Creative Commons BY 3.0 Unported license
© Ludmila Scharf

Joint work of Alt, Helmut; Scharf, Ludmila

We show that the Hausdorff distance for two sets of n non-intersecting line segments can be computed in parallel in $O(\log^2 n)$ time using $O(n)$ processors in a CREW-PRAM computation model. We discuss how some parts of the sequential algorithm can be performed in parallel using previously known parallel algorithms; and identify the so-far least efficiently solved part of the problem for the parallel computation, which is the following: Given two sets of x -monotone curve segments, red and blue, for each red segment find its extremal intersection points with the blue set, i.e. points with the minimal and maximal x -coordinate. Each segment set is assumed to be intersection free. The best known parallel algorithm for this problem has total work of $O(n \log^3 n)$ and uses $O(n \log^2 n)$ space. The algorithm presented here improves the theoretical time and space performance while still being practically feasible.

3.25 Convex Transversals

Lena Schlipf (FU Berlin, DE)

License © Creative Commons BY 3.0 Unported license
© Lena Schlipf

The talk gives an overview of the work on convex transversals. The question about convex transversals was initially posed by Arik Tamir at the Fourth NYU Computational Geometry Day (1987): “Given a collection of compact sets, can one decide in polynomial time whether there exists a convex body whose boundary intersects every set in the collection?” So far, there have been very few results. One of these rare results is an $O(n \log n)$ algorithm by Goodrich and Snoeyink [1] that solves this problem when the sets are n parallel line

segments. We show that when the sets are segments in the plane, deciding existence of the convex stabber is NP-hard (this is joint work with Arkin, Dieckmann, Knauer, Mitchell, Polishchuk, Yang [2]). The problem remains NP-hard when the sets are simple regular polygons. Additionally, we prove the problem to be NP-hard when the sets are disjoint bends in the plane.

References

- 1 M. T. Goodrich and J.Snoeyink. Stabbing parallel segments with a convex polygon. *Computer Vision, Graphics, and Image Processing*, 49(2):152–170, 1990.
- 2 E.M. Arkin, C. Dieckmann, C. Knauer, J.S.B. Mitchell, V. Polishchuk, L.Schlipf, and S. Yang. Convex Transversals. In *Proc. of WADS'11*, pp. 49–60, 2011.

3.26 Delaunay and other triangulations of moving point sets: What's going on?

Micha Sharir (Tel Aviv University, IL)

License  Creative Commons BY 3.0 Unported license
© Micha Sharir

Joint work of Sharir, Micha; Agarwal, Pankaj; Kaplan, Haim; Rubin, Natan; and several others

In this talk we review several recent works addressing the following problem: Given a set P of n moving points in the plane, where the motion of each point is semialgebraic of constant description complexity, we want to maintain the Delaunay (or some other) triangulation of P kinetically, updating it after each discrete change that it experiences.

The Delaunay triangulation, $DT(P)$, besides its many useful properties, is ideal for such a maintenance, because it admits local certification, requiring the circumdisk of each triangle to be empty. The main problem is to show that the number of discrete changes in $DT(P)$ is “small”, meaning nearly quadratic in n . (A quadratic lower bound is known.) This is still open, and is considered one of the hardest open problems in combinatorial and computational geometry.

We review several recent attempts to address this issue:

(1) Developing other triangulation schemes, with a provably near-quadratic number of changes. This has been done by Agarwal, Wang and Yu, and later by Kaplan, Rubin and Sharir.

(2) Maintaining the Delaunay triangulation of P under a polygonal, non-Euclidean norm. Here too one can show that the diagram experiences only a near-quadratic number of changes, and can be maintained efficiently, but it has several drawbacks. This goes back to Chew, and has been treated in a more general and complete manner by Agarwal, Kaplan, Rubin, and Sharir (work in progress).

(3) Maintaining only a “stable” portion of the Delaunay diagram, roughly corresponding to edges whose dual Voronoi edges are seen from their sites at a sufficiently large angle. Again, a near-quadratic bound on the number of changes can be established, and the stable portion has several drawbacks. This is work in progress by Agarwal and many other authors, originally presented at SoCG many years ago.

(4) Most importantly, we review a recent work by Rubin, where he manages to establish a near-quadratic bound for the Euclidean Delaunay triangulation for points moving at unit speeds. The analysis is quite involved, and we review some of its main technical ingredients.

3.27 Towards Elastic Shape Matching

Fabian Stehn (Universität Bayreuth, DE)

License  Creative Commons BY 3.0 Unported license
© Fabian Stehn

Joint work of Knauer, Christian; Stehn, Fabian

Geometric shape matching problems are one of the core research topics in the field of computational geometry. The general question of a geometric matching problem is as follows: given are two geometric objects – a pattern and a model – and a transformation class as well as a similarity measure. One seeks a transformation t of the given class such that the similarity measure of the pattern transformed by t to the model is maximized.

We introduced the concept of *elastic* (non-uniform) geometric shape matching problems. In an elastic geometric shape matching problem the pattern is not transformed by a single transformation, but by a so-called *transformation ensemble*. Transformation ensembles allow non-uniform deformations of the pattern – different parts of the pattern can be transformed by different mappings. Another benefit of transformation ensembles is the possibility to incorporate temporal dependencies of the pattern and changes of its shape over time in this modeling. This allows to compute registrations that are valid within a certain time frame even if the reference objects change during this time period. This is achieved by linking together transformation ensembles at different points in time and by applying suitable temporal and spacial interpolation methods.

The modeling as an elastic geometric shape matching problem has various benefits from a theoretical as well as practical point of view. *Classical* geometric shape matching problems form a special case in this modeling and hence allow a direct comparison to results of elastic geometric shape matching problems. On the other hand, many practical applications (such as navigated surgeries for example) will benefit from algorithms and data structures that compute transformation ensembles: in these applications one often has to deal with local deformations of the input (e.g. due to distortion caused by magnetic fields), as well as entities that vary of time (for example soft tissue deformation in the aforementioned context of navigated surgeries). In this talk we introduced first results on elastic shape matching problems for point sequences under translations.

3.28 Extended Formulations for polytopes

Hans Raj Tiwary (University of Brussels, BE)

License  Creative Commons BY 3.0 Unported license
© Hans Raj Tiwary

Joint work of Tiwary, Hans Raj; Avis, David

Main reference D. Avis, H. Raj Tiwary, “On the extension complexity of combinatorial polytopes,” arXiv:1302.2340v2 [math.CO]; accepted in ICALP 2013.

URL <http://arxiv.org/abs/1302.2340v2>

A polytope Q is said to be an extended formulation (EF) for a polytope P , iff P is the projection of Q . The notion of extended formulations are not only important in many areas of applied sciences but also interesting from a theoretical perspective. In a certain sense a compact EF encodes “faithful” linear programs for solving optimization problems. In this talk, I will discuss some basics, some recent new results, and a purposely vague open problem related to the existence of compact EFs.

3.29 Reverse stabbing queries in galleries

Suresh Venkatasubramanian (University of Utah, US)

License © Creative Commons BY 3.0 Unported license
© Suresh Venkatasubramanian

Joint work of Daruki, Samira; Hillyard, Peter; Patwari, Neal; Venkatasubramanian, Suresh

Radio Tomographic Imaging (RTI) is an emerging technology that locates moving objects in areas surrounded by simple and inexpensive radios. RTI is useful in emergencies, rescue operations, and security breaches, since the objects being tracked need not carry an electronic device. Tracking humans moving through a building, for example, could help firefighters save lives by locating victims quickly. RTI works by placing small inexpensive radios in a region of interest. The radios can send and receive wireless signals, and form a network of links that cover the region of interest. When a person walks through the region, they interfere with the links, creating a “shadow” of broken links that can be used to infer presence and track individuals.

This yields the following problem: given a collection of radios and a set of “visible” links, infer the trajectory of a person moving through the region. This inference must be robust under link errors and occlusion, as well as be performed in real time. In addition, there is an associated planning problem of where to place the radios in order to make the tracking algorithm as effective as possible.

In this talk, I present geometric algorithms for these questions. The key technical developments include a generalization of stabbing line and transversal problems, as well as a novel generalization of traditional art gallery problems.

3.30 A Faster Algorithm for Computing Motorcycle Graphs

Antoine Vigneron (KAUST – Thuwal, SA)

License © Creative Commons BY 3.0 Unported license
© Antoine Vigneron

Joint work of Vigneron, Antoine; Lie, Yan

Main reference To appear in the Proceedings of the Symp. on Computational Geometry 2013.

We present a new algorithm for computing motorcycle graphs. Its running time is $O(n^{4/3})$, where n is the size of the input. When the motorcycles start from the side of a simple polygon, and input coordinates are $O(\log n)$ -bit rational numbers, the time bound improves to $O(n \log^3 n)$. It yields an $O(n \log^3 n)$ expected time algorithm for computing the straight skeleton of a simple polygon.

3.31 Towards Understanding Gaussian Weighted Graph Laplacian

Yusu Wang (Ohio State University, US)

License © Creative Commons BY 3.0 Unported license
© Yusu Wang

The Gaussian-weighted graph Laplacian, as a special form of graph Laplacians with general weights, has been a popular empirical operator for data analysis applications, including semi-supervised learning, clustering, and denoising. There have been various studies of the properties and behaviors of this empirical operator; most notably, its convergence behavior as the number of points sampled from a hidden manifold goes to infinity.

In this talk we present two new results on the theoretical properties of the Gaussian-weighted Graph Laplacian. The first one [1] is about its behavior as the input points where we construct the graph are sampled from a, what we call, *singular manifold*; while previous theoretical study of the Gaussian-weighted Graph Laplacian typically assumes that the hidden domain is a compact smooth manifold. A singular manifold can consist of a collection of potentially intersecting manifolds with boundaries, and represents one step towards modeling more complex hidden domains.

The second result we present is about the stability of the Gaussian-weighted Graph Laplacian as the hidden manifold where input points are sampled from have certain small perturbation [2]. The goal is to understand how the spectrum of Gaussian-weighted Graph Laplacian changes with respect to perturbations of the domain.

References

- 1 M. Belkin, Q. Que, Y. Wang and X. Zhou. Towards understanding complex spaces: graph Laplacians on manifolds with singularities and boundaries. In *Proc. of COLT 2012*, pp. 36.1–36.26, 2012.
- 2 T. K. Dey, P. Ranjan and Y. Wang. Weighted graph Laplace operator under Topological noise. In *Proc. of SODA 2013*, 2013.

4 Open Problems

► **PROBLEM 1 (OTFRIED CHEONG).** *3-dimensional Kakeya problem:* Find a smallest-volume three-dimensional convex body K such that, for any direction u , $\text{width}(K, u) \geq 1$.

In two dimensions, the optimal body is the equilateral triangle of height 1. In three dimensions, a regular tetrahedron where the distance between opposite edges is 1 is not optimal. One can “shave off” the corners in order to decrease the area while maintaining the width-condition.

► **PROBLEM 2 (JEFF ERICKSON).** *Almost simple polygons:* A polygon P with vertices $p_1, p_2, \dots, p_n \in \mathbb{R}^2$ is *almost simple* if, for any $\varepsilon > 0$, there is a simple polygon Q with vertices q_1, q_2, \dots, q_n such that $\|p_i - q_i\| < \varepsilon$ for each index i . Equivalently, an n -gon P is weakly simple if there are simple n -gons with arbitrarily small Fréchet distance to P .

IS THERE A POLYNOMIAL-TIME ALGORITHM TO DETERMINE WHETHER A GIVEN SEQUENCE OF POINTS IS THE VERTEX SEQUENCE OF AN ALMOST-SIMPLE POLYGON?

There is an algorithm to decide whether a *spur-free* polygon is almost simple in $O(n \log n)$ time, where a *spur* is a vertex with a zero-degree angle, or equivalently, a pair of consecutive edges that overlap. A spur-free polygon is weakly simple if and only if it contains no crossing subwalks and its winding number is ± 1 . (Two spur-free walks $a_0 b_1 \dots b_k a_{k+1}$ and $c_0 b_1 \dots b_k c_{k+1}$ *cross* if either the triples a_0, b_1, c_0 and a_{k+1}, b_k, c_{k+1} have the same orientation, or $k \leq 1$ and the walks intersect transversely.) However, this characterization does not extend to polygons with spurs, in part because the winding number is not well-defined.

► **PROBLEM 3 (FABIAN STEHN).** Let $S = \{s_1, \dots, s_n\}$ be a set of segments in the plane. Compute n translations t_1, \dots, t_n such that the set $S' = \{s'_i \mid s'_i = s_i + t_i, i = 1, \dots, n\}$ is disjoint and the convex hull of S' has minimum area.

Remark: Two segments of S' are allowed to have a common endpoint and the endpoint of a segment in S' is allowed to lie on another segment.

► **PROBLEM 4 (SURESH VENKATASUBRAMANIAN).** *MDS:* Given a distance matrix $[d_{ij}]$, where d_{ij} which is the distance between the i th and the j th object, find an embedding of points $x_1, \dots, x_n \in \mathbb{R}^k$ such that $\sum (d_{ij} - \|x_i - x_j\|)^2$ is minimal.

Other versions of the problem are, e.g., minimize $\sum(d_{ij}^2 - \|x_i - x_j\|^2)$. Even the case where $k = 1$ is open.

► **PROBLEM 5 (MAARTEN LÖFFLER)**. Given a unit square, find a set of lines such that there is a disk of radius ε centered on each line inside the unit square and such that no disk intersects another disk or another line. What is the maximum number of lines that can be placed?

Known bounds: $\Omega(1/\varepsilon)$ and $O(1/\varepsilon^2)$.

► **PROBLEM 6 (GÜNTER ROTE)**. This problem is due to Sergio Cabello and Maria Saumell.

Let P be a polygon of area 1. Let C_{\max} be the area of the largest convex polygon contained in P . It is easy to see that

$$C_{\max}^2 \leq \text{probability}(x \text{ sees } y | x, y \in P).$$

The question is whether the following reverse bound holds

$$\text{probability}(x \text{ sees } y | x, y \in P) \leq O(C_{\max}).$$

Other variants:

a) P is star-shaped.

b) P is any region.

c) P is a polygonal region with holes. *Remark:* Sándor Fekete has pointed out that a convex region with many small holes (punctures) will be a counterexample.

► **PROBLEM 7 (YOSHIO OKAMOTO)**. Given a polygonal domain P with a total number of n vertices, what is the maximum number of local maxima of the geodesic distance function $d(p, q)$, $p, q \in P$, on P ? Known results: $O(n^7)$ and $\Omega(n^2)$ (the lower bound is tight for polygons without holes) [1].

► **PROBLEM 8 (ROLF KLEIN)**. *Lion problem:* We are given a $n \times n$ grid, each cell is contaminated or clean. Additionally, we are given a fixed number of lions. We consider discrete time steps; in each step the contamination of a cell spreads to its four adjacent cells. A lion can clean one adjacent cell per step.

HOW MANY LIONS ARE NEEDED TO CLEAN THE GRID?

It is obvious that n lions are enough but it is an open question whether $n - 1$ lions are enough. Dumitrescu et al. [4] proved that \sqrt{n} lions are not enough. Later on, Brass et al. [3], and independently Berger et al. [2], showed that $\lfloor n/2 \rfloor$ lions are not enough.

In general, it is assumed that in the beginning all cells are contaminated but there are many variants for the problem:

- different number of contaminated cells at the beginning
- remove boundaries
- etc.

References

- 1 S.W. Bae, M. Korman, and Y. Okamoto. The geodesic diameter of polygonal domains. In *Proc. ESA 2010*, pp. 500–511, 2010.
- 2 F. Berger, A. Grüne, and R. Klein. How many lions can one man avoid? *Technical Report 006, Department of Computer Science I, University of Bonn*, 2007.
- 3 P. Braß, K.D. Kim, H.S. Na, and C.S. Shin. Escaping off-line searchers and a discrete isoperimetric theorem. In *Proc. ISAAC 2007*, pp 65–74, 2007.
- 4 A. Dumitrescu, I. Suzuki, and P. Zylinski. Offline variants of the “lion and man” problem. In *Proc. SoCG 2007*, pp 102–111, 2007.

Participants

- Pankaj Kumar Agarwal
Duke University, US
- Boris Aronov
Polytechnic Inst. of NYU, US
- Dominique Attali
GIPSA Lab – Saint Martin
d’Hères, FR
- Sang Won Bae
Kyonggi University, KR
- Mikhail Belkin
Ohio State University –
Columbus, US
- Eric Berberich
MPI für Informatik –
Saarbrücken, DE
- Kevin Buchin
TU Eindhoven, NL
- Maike Buchin
TU Eindhoven, NL
- Siu-Wing Cheng
HKUST – Kowloon, HK
- Otfried Cheong
KAIST – Daejeon, KR
- Kenneth L. Clarkson
IBM Almaden Center –
San José, US
- Tamal K. Dey
Ohio State University –
Columbus, US
- Anne Driemel
Utrecht University, NL
- Alon Efrat
Univ. of Arizona – Tucson, US
- Jeff Erickson
Univ. of Illinois – Urbana, US
- Sándor Fekete
TU Braunschweig, DE
- Joachim Giesen
Universität Jena, DE
- Marc Glisse
INRIA Saclay – Île-de-France
– Orsay, FR
- Xavier Goaoc
INRIA Lorraine, FR
- Joachim Gudmundsson
The University of Sydney, AU
- Herman J. Haverkort
TU Eindhoven, NL
- Menelaos Karavelas
Univ. of Crete – Heraklion, GR
- Matthew J. Katz
Ben Gurion University – Beer
Sheva, IL
- Michael Kerber
Stanford University, US
- Rolf Klein
Universität Bonn, DE
- Maarten Löffler
Utrecht University, NL
- Kurt Mehlhorn
MPI für Informatik –
Saarbrücken, DE
- Guillaume Moroz
INRIA Grand Est – Nancy, FR
- Dmitriy Morozov
Lawrence Berkeley National
Laboratory, US
- David M. Mount
University of Maryland – College
Park, US
- Yoshio Okamoto
Univ. of
Electro-Communications –
Tokyo, JP
- Marcel J. M. Roeloffzen
TU Eindhoven, NL
- Günter Rote
FU Berlin, DE
- Jörg-Rüdiger Sack
Carleton Univ. – Ottawa, CA
- Ludmila Scharf
FU Berlin, DE
- Lena Schlipf
FU Berlin, DE
- Raimund Seidel
Universität des Saarlandes, DE
- Micha Sharir
Tel Aviv University, IL
- Bettina Speckmann
TU Eindhoven, NL
- Fabian Stehn
Universität Bayreuth, DE
- Monique Teillaud
INRIA Sophia Antipolis –
Méditerranée, FR
- Hans Raj Tiwary
University of Brussels, BE
- Suresh Venkatasubramanian
University of Utah, US
- Antoine Vigneron
KAUST – Thuwal, SA
- Yusu Wang
Ohio State University –
Columbus, US
- Carola Wenk
Tulane University, US
- Nicola Wolpert
University of Applied Sciences –
Stuttgart, DE



Report from Dagstuhl Seminar 13111

Scheduling

Edited by

Susanne Albers¹, Onno J. Boxma², and Kirk Pruhs³

1 HU Berlin, DE, albers@informatik.hu-berlin.de

2 TU Eindhoven, NL, o.j.boxma@tue.nl

3 University of Pittsburgh, US, kirk@cs.pitt.edu

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 13111 “Scheduling”. The primary objective of the seminar is to facilitate dialog and collaboration between researchers in two different mathematically-oriented scheduling research communities, the stochastic scheduling and queuing community, and the worst-case approximation scheduling community.

Seminar 10.–15. March, 2013 – www.dagstuhl.de/13111

1998 ACM Subject Classification C.4 Performance of Systems

Keywords and phrases Scheduling, Algorithms, Stochastic, Optimization, Queuing

Digital Object Identifier 10.4230/DagRep.3.3.24

1 Executive Summary

Susanne Albers

Onno J. Boxma

Kirk Pruhs

License  Creative Commons BY 3.0 Unported license
© Susanne Albers, Onno J. Boxma, and Kirk Pruhs

The primary objective of the seminar is to facilitate dialog and collaboration between researchers in two different mathematically-oriented scheduling research communities, the stochastic scheduling and queuing community, and the worst-case approximation scheduling community. To a large extent, the applications considered by the two communities are the same. The stochastic community considers questions related to determining stochastic information (like the expectation or tail bounds) about the performance of algorithms and systems from stochastic information about the input. The worst-case community considers questions related to determining the worst-case performance of algorithms and systems assuming no stochastic information about the input. Each community has developed its own set of mathematical techniques that are best suited to answer these different sorts of questions. While addressing similar problems, these communities tend to attend different conferences (e.g. SIGMETRICS vs. SODA/IPCO), and publish in different journals. Thus the organizers believed that each community would benefit from greater interaction with the other community, and this seminar was an opportunity to further such interaction. The seminar was attended by about 15 researchers from the stochastic community and 40 researchers from the worst-case community.



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

Scheduling, *Dagstuhl Reports*, Vol. 3, Issue 3, pp. 24–50

Editors: Susanne Albers, Onno J. Boxma, and Kirk Pruhs



DAGSTUHL
REPORTS

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

2 Table of Contents

Executive Summary

<i>Susanne Albers, Onno J. Boxma, and Kirk Pruhs</i>	24
--	----

Overview of Talks

Optimal Scheduling Problem for Scalable Queues <i>Samuli Aalto</i>	27
Online Myopic Network Covering <i>Konstantin Avrachenkov</i>	28
Scheduling with time-varying capacities <i>Urtzi Ayesta</i>	29
Online Scheduling in the Cloud <i>Yossi Azar</i>	29
Worst Case and Stochastic Analysis in Scheduling: Similarities, Differences, and Bridges <i>Nikhil Bansal, Urtzi Ayesta, and Onno Boxma</i>	29
Clustered Scheduling of Real-Time Tasks <i>Vincenzo Bonifaci</i>	30
Wireless Random-Access Algorithms: Fluid Limits and Delay Issues <i>Sem Borst</i>	30
Dominance rules for $1 \sum w_j C_j^\beta$ <i>Christoph Dürr</i>	31
Stochastic Knapsacks and Matchings <i>Anupam Gupta</i>	31
When Does Stochasticity/Discreteness Matter? Fluid Models and Scheduling Queueing Networks <i>John Hasenbein</i>	32
Complexity of generalized min-sum scheduling <i>Wiebke Höhn</i>	32
The generalization of scheduling with machine cost <i>Csanád Imreh</i>	33
Online Scheduling with General Cost Functions <i>Sungjin Im</i>	34
Scheduling of Users with Time-Varying Service Rates <i>Peter Jacko</i>	35
Employee scheduling and rescheduling in call centers <i>Ger Koole</i>	36
Online scheduling algorithms analyzed by Dual Fitting <i>Amit Kumar</i>	36
Approximation in Dynamic Stochastic Scheduling <i>Nicole Megow</i>	36
MapReduce and Distributed Scheduling <i>Benjamin Moseley</i>	37

Weakly Coupled Stochastic Decision Systems	
<i>Kamesh Munagala</i>	37
Stochastic k-TSP	
<i>Viswanath Nagarajan</i>	37
Job Scheduling Mechanisms for Large Computing Clusters	
<i>Seffi Naor</i>	38
Scheduling in queues with customer grouping	
<i>Sindo Nunez Queija</i>	38
Stochastic Comparison of Multicast Pull and Push	
<i>Kirk Pruhs</i>	39
Scheduling with time-varying cost: Deterministic and stochastic models	
<i>Roman Rischke</i>	39
Online scheduling of jobs with fixed start times on related machines	
<i>Jiri Sgall</i>	40
Scheduling Advertising Campaigns	
<i>Hadas Shachnai</i>	41
Optimal queue-size scaling in switched networks	
<i>Devavrat Shah</i>	41
On Carry Over	
<i>Frits C. R. Spieksma</i>	41
Stochastic Optimal Control for a Class of Dynamic Resource Allocation Problems	
<i>Mark S. Squillante</i>	42
Truthful Scheduling	
<i>Rob van Stee</i>	43
Online Stochastic Matching	
<i>Clifford Stein</i>	44
Split scheduling with uniform setup times	
<i>Suzanne van der Ster</i>	44
An infinite server system with customer-to-server packing constraints	
<i>Alexander Stolyar</i>	45
Analysis of Smith’s Rule in Stochastic Scheduling	
<i>Marc Uetz</i>	46
Appointment Scheduling with Slot Blocking	
<i>Peter van de Ven</i>	46
Learning in Stochastic Scheduling	
<i>Tjark Vredeveld</i>	47
FCFS infinite matching, queues with skill based routing, and organ transplants	
<i>Gideon Weiss</i>	47
Working Groups	48
Open Problems	49
Participants	50

3 Overview of Talks

The opening talk on Monday morning was presented by Nikhil Bansal, Urtzi Ayesta, Onno Boxma (with materials contributed by Adam Wiermam). The purpose of this talk was to introduce the research approaches of each community to the other community, to a large extent using the two community's research on scheduling speed scalable processors as a running example. Each of the remaining seven half-days contained one longer, 30-40 minute, talk from each community from a speaker invited by the organizers. The talks from the worst-case community were given by Amit Kumar, Anupam Gupta, Cliff Stein, Kamesh Munagala, Yossi Azar, Nicole Megow, and Ben Moseley. The talks from the stochastic community were given by Mark Squillante, Devavrat Shah, Gideon Weiss, John Hasenbein, Alexander Stoylar, Sem Borst, and Ger Koole. Again the goal was to try to introduce each community to important tools and approaches from the other community. The remaining meeting times were filled in by short, 7 minute, talks by the remaining participants. Speakers were encouraged to be forward leaning, and present ongoing or future work, or open problems. Many talks provoked lively discussion, which was allowed to continue to its natural conclusion. The time between lunch and the afternoon coffee break was left open for individual discussions and collaborations. Abstracts of the talks can be found below.

3.1 Optimal Scheduling Problem for Scalable Queues

Samuli Aalto (Aalto University, FI)

License  Creative Commons BY 3.0 Unported license
© Samuli Aalto

Modern wireless cellular systems are able to utilize the opportunistic scheduling gain originating from the variability in users' channel conditions. By favoring the users with good instantaneous channel conditions, the service capacity of the system can be increased with the number of users. On the other hand, for service systems with fixed service capacity, the system performance can be optimized by utilizing the job size information. Combining the advantages of size-based scheduling with the opportunistic scheduling gain has proven a challenging task.

Inspired by the opportunistic scheduling gain, Sadiq and de Veciana (2010) defined a new queueing model, which they called an $M/GI/C$ queue. Briefly said, it is a service system where the service capacity scales with the number of jobs in the system. Thus, we called it a *scalable queue* in Aalto et. al (2011, 2012). More precisely said, the service capacity is defined via *capacity regions* \mathcal{C}_n , where n denotes to the number of jobs in the system.

Sadiq and de Veciana (2010) considered the optimal scheduling problem for a specific class of capacity regions, nested polymatroids, and found the optimal policy that minimizes the total sojourn time in the transient setting (without any arrivals). Aalto et al. (2011) generalized the result to compact and symmetric capacity regions, however, utilizing an additional implicit assumption. Aalto et al. (2012) demonstrated that the additional assumption, indeed, is satisfied by a class of wireless system models stemming from a time-scale separation assumption.

The optimal scheduling problem for scalable queues in the dynamic setting (with arrivals) is still completely open, as well as the problems with other types of objective functions.

References

- 1 Bilal Sadiq, Gustavo de Veciana: Balancing SRPT prioritization vs opportunistic gain in wireless systems with flow dynamics. Proceedings of ITC-22, 2010
- 2 Samuli Aalto, Aleksi Penttinen, Pasi Lassila, Prajwal Osti: On the optimal trade-off between SRPT and opportunistic scheduling. Proceedings of ACM SIGMETRICS, 2011
- 3 Samuli Aalto, Aleksi Penttinen, Pasi Lassila, Prajwal Osti: Optimal size-based opportunistic scheduler for wireless systems. Queueing Systems 72:5–30, 2012

3.2 Online Myopic Network Covering

Konstantin Avrachenkov (INRIA Sophia Antipolis – Méditerranée, FR)

License  Creative Commons BY 3.0 Unported license
© Konstantin Avrachenkov

Efficient marketing or awareness-raising campaigns seek to recruit n influential individuals – where n is the campaign budget – that are able to cover a large target audience through their social connections. So far most of the related literature on maximizing this network cover assumes that the social network topology is known. Even in such a case the optimal solution is NP-hard. In practice, however, the network topology is generally unknown and needs to be discovered on-the-fly. In this work we consider an unknown topology where recruited individuals disclose their social connections (a feature known as one-hop lookahead). The goal of this work is to provide efficient online algorithms that recruit individuals as to maximize the size of target audience covered by the campaign.

We analyze the performance of a variety of online algorithms currently used to sample and search large networks. We also propose a new greedy online algorithm, Maximum Expected d -Excess Degree (MEED), and provide, to the best of our knowledge, the first detailed theoretical analysis of the cover size of a variety of well known network sampling algorithms on finite networks. Our proposed algorithm greedily maximizes the expected size of the cover. For a class of random power law networks we show that MEED simplifies into a straightforward procedure, which we denote MOD (Maximum Observed Degree). Thus our problem gives an interesting linkage between deterministic and stochastic scheduling.

We note that performance may be further significantly improved if the node degree distribution is known or can be estimated online during the campaign. Thus, one open question is which algorithm would achieve an optimal trade off between exploration and exploitation of the network structure? Another open question: what is the approximation ratio of the greedy online algorithm?

Preliminary results on the topic can be found at:

References

- 1 K. Avrachenkov, P. Basu, G. Neglia, B. Ribeiro and D. Towsley, “Online Myopic Network Covering”, <http://arxiv.org/abs/1212.5035>

3.3 Scheduling with time-varying capacities

Urtzi Ayesta (LAAS – Toulouse, FR)

License © Creative Commons BY 3.0 Unported license
© Urtzi Ayesta

Classical results in size-based scheduling in a single-server queue show that giving preference to short flows is optimal in a wide variety of settings. However all these results typically assume that the speed of the server is constant over time and independent of the state of the queue. In this short talk we will show that when the capacity of the system is time-varying (either as a function of the state or as an exogenous process) giving preference to short flows is no longer necessarily optimal, which opens several interesting questions.

3.4 Online Scheduling in the Cloud

Yossi Azar (Tel Aviv University, IL)

License © Creative Commons BY 3.0 Unported license
© Yossi Azar

Online task scheduling of jobs on cloud computing infrastructures poses new challenges theoretically and practically. In particular the number of machines or virtual machines (VM) is not fixed any more and the goal is to minimize the cost of the computation as well as minimize the delay or the load. We will discuss various models and questions in this area concerned with identical vs heterogeneous machines, fixed setup time vs arbitrary setup cost and single vs multi dimension job requirements.

The talk is based on three papers:

1. one will appear in SODA – unrelated machine scheduling with startup cost (paper is called Online mixed packing and covering)
2. second submitted to STOC – Online Vector Bin Packing
3. third (not submitted) is Cloud Scheduling with Setup Cost.

3.5 Worst Case and Stochastic Analysis in Scheduling: Similarities, Differences, and Bridges

Nikhil Bansal (TU Eindhoven, NL), Urtzi Ayesta (LAAS – Toulouse, FR), Onno J. Boxma (TU Eindhoven, NL)

License © Creative Commons BY 3.0 Unported license
© Nikhil Bansal, Urtzi Ayesta, and Onno Boxma
Joint work of Nikhil Bansal, Adam Wierman, Urtzi Ayesta, and Onno Boxma

In order to provide background for the workshop participants, this talk will give a quick introduction to the two areas of worst case analysis and stochastic analysis of scheduling policies. We will introduce the basic notions in each of these areas, and then provide some examples of topics that have been studied by both of the communities. The goal of the talk will be to highlight the differences and similarities in the approaches for addressing these topics. We will also describe some examples of topics where the techniques from one community have proven useful in the other.

3.6 Clustered Scheduling of Real-Time Tasks

Vincenzo Bonifaci (National Research Council – Rome, IT)

License  Creative Commons BY 3.0 Unported license
© Vincenzo Bonifaci

We consider the scheduling of hard real-time tasks on multiple identical processors. Each such task is described by a worst-case execution time, a relative deadline, and a minimum interarrival time, and recurrently generates jobs with the same features. In the global approach, jobs can be arbitrarily preempted and migrated among the processors. In the partitioned approach, tasks are statically assigned to processors and then each processor is scheduled according to some fixed online policy, such as, say, Earliest Deadline First. Preemption within one processor is always allowed.

Deciding whether the system can be feasibly scheduled is NP-hard even in very simple cases. Therefore, we typically look for assignment and scheduling algorithms that guarantee at least the same service that can be provided by an optimal assignment and scheduling to slightly slower, or fewer, processors.

More recently, a “clustered” approach is being suggested. In the clustered approach, each task is assigned to a cluster of machines (e.g., a set of processor cores sharing a cache) and then globally scheduled within the cluster. This allows for more flexibility than in partitioned scheduling, while avoiding the disadvantages of a fully global approach where the overheads due to task migrations can be large. However, a quantitative study of the trade-offs achievable with the clustered approach is yet to be undertaken.

Pointers into the literature:

References

- 1 Sanjoy K. Baruah, Nathan Fisher: The Partitioned Multiprocessor Scheduling of Deadline-Constrained Sporadic Task Systems. *IEEE Trans. Computers* 55(7): 918–923 (2006)
- 2 Jian-Jia Chen, Samarjit Chakraborty: Resource Augmentation Bounds for Approximate Demand Bound Functions. *RTSS 2011*: 272–281
- 3 Andrea Bastoni, Björn B. Brandenburg, James H. Anderson: An Empirical Comparison of Global, Partitioned, and Clustered Multiprocessor EDF Schedulers. *RTSS 2010*:14–24

3.7 Wireless Random-Access Algorithms: Fluid Limits and Delay Issues

Sem C. Borst (TU Eindhoven, NL)

License  Creative Commons BY 3.0 Unported license
© Sem Borst

Queue-based wireless random-access algorithms are relatively simple and inherently distributed, yet provide a striking capability to match the optimal throughput performance of centralized scheduling mechanisms in a wide range of scenarios. Unfortunately, the specific type of activation rules for which throughput optimality has been established, may result in extremely long queues and delays. The use of more aggressive/persistent access schemes can improve the delay performance, but does not provide any universal maximum-stability guarantees.

In order to gain qualitative insights and examine stability properties, we investigate fluid limits where the system dynamics are scaled in space and time. Several distinct types

of fluid limits can arise, ranging from ones with smooth deterministic features, to others which exhibit random oscillatory characteristics, depending on the topology of the network, in conjunction with the form of the activation rules. As we will show, these qualitatively different regimes are strongly related to short-term fairness measures and mixing times for random-access mechanisms with fixed activation rates, and carry significant implications for stability properties.

Note: based on joint work with Niek Bouman (TU/e), Javad Ghaderi (UIUC), Johan van Leeuwen (TU/e), Alexandre Proutiere (KTH), Peter van de Ven (IBM), Phil Whiting (Alcatel-Lucent Bell Labs), Alessandro Zocca (TU/e)

3.8 Dominance rules for $1 || \sum w_j C_j^\beta$

Christoph Dürr (UPMC, Lab. LIP6 – Paris, FR)

License © Creative Commons BY 3.0 Unported license
© Christoph Dürr

We consider a simple scheduling problem on a single machine. Each job j has some given processing time p_j and a priority w_j . The goal is to find an ordering of the given jobs, that minimizes $\sum w_j C_j^\beta$, where β is a fixed positive constant.

Consider two jobs i, j . We say that we have the *local ordering* $i \prec_\ell j$, if for any instance containing i, j , every optimal schedule in which i, j are adjacent, the job i precedes j . Similarly we say that we have the *global ordering* $i \prec_g j$ if for any instance containing i, j every optimal schedule in which i, j are adjacent, the job i is scheduled before job j .

These ordering conditions are useful for any algorithm solving this problem. What are the conditions that imply $i \prec_\ell j$ or even $i \prec_g j$? This short talk gives an overview of what is known, and asks the question whether $i \prec_\ell j$ always implies $i \prec_g j$.

Pointers into the literature:

References

- 1 Höhn, W., Jacobs, T. An experimental and analytical study of order constraints for single machine scheduling with quadratic cost. In Proc. of the 14th Workshop on Algorithm Engineering and Experiments (ALENEX 2012).

3.9 Stochastic Knapsacks and Matchings

Anupam Gupta (Carnegie Mellon University – Pittsburgh, US)

License © Creative Commons BY 3.0 Unported license
© Anupam Gupta

I will survey some work on two stochastic packing problems: (a) stochastic knapsack where jobs of uncertain sizes and/or rewards are packed into a knapsack, and (b) stochastic matchings where uncertain edges are packed given a set of constraints. I'll focus on the techniques used for these problems (and how they give solutions for extensions and generalizations), and the many open questions.

3.10 When Does Stochasticity/Discreteness Matter? Fluid Models and Scheduling Queueing Networks

John Hasenbein (University of Texas – Austin, US)

License  Creative Commons BY 3.0 Unported license
© John Hasenbein

We discuss the connection between scheduling multiclass stochastic and fluid (deterministic) networks. First we overview the relationship via a few of classic results and examples. Next, we present more recent research on a stochastic combinatorial scheduling problem in which the “macro” stochasticity must be taken into account, but the “micro” stochasticity is less important.

3.11 Complexity of generalized min-sum scheduling

Wiebke Höhn (TU Berlin, DE)

License  Creative Commons BY 3.0 Unported license
© Wiebke Höhn

We consider a single-machine scheduling problem with generalized min-sum objective. Given a set of jobs $j = 1, \dots, n$ with individual weights $w_j \geq 0$ and processing times $p_j \geq 0$, the goal is to assign the jobs to non-overlapping time intervals on the machine to minimize $\sum_{j=1}^n w_j g(C_j)$, where C_j denotes the completion time of job j in the schedule and where g is some non-decreasing cost function. Note that the question of allowing preemption does not play a role here, because the jobs do not have release times and so the possibility of preemption never leads to a cheaper optimal schedule.

Alternatively, we can interpret this problem as the scenario of linear cost but non-uniform processor speed. Assume that the processor speed is given by a nonnegative function s . Then, the total workload processed until time t is $S(t) := \int_0^t s(x) dx$. Conversely, if the total workload of job j and all jobs processed before it is P , then the completion time of j in the schedule is $S^{-1}(P)$. Hence, according to the above model, S^{-1} can be seen as cost function. Note that S^{-1} is always monotone, and it is continuous even if s is not. Moreover if s is increasing or decreasing then S^{-1} is concave or convex, respectively.

Finally, we would like to point out that this general scheduling problem also covers the *Airplane Refueling Problem* (ARP) whose complexity was proposed as an open problem by Gerhard Woeginger on the last Dagstuhl Scheduling Workshop. The ARP translates one-to-one into the maximization variant of our problem with cost function $1/C_j$, and further into the minimization variant with concave cost function $C - 1/C_j$ for some sufficiently large constant C . Of course, this transformation does not preserve approximation guarantees. However, it does preserve classic complexity results.

Despite many recent approximation results, there are still many open problems concerning the complexity of the problem. For linear and exponential cost functions, the problem is in P, and it is known to be weakly NP-hard for tardiness cost $C_j - d$ with common due date d , i.e., for convex functions. Moreover, the problem is strongly NP-hard for piece-wise linear cost functions alternating between two speeds. On the positive side, Megow and Verschae designed a PTAS for general cost functions and an FPTAS for piece-wise linear functions with a constant number of linear segments. For general convex cost functions, it is open whether

the problem is strongly NP-hard, and for concave function—and in particular (concave and convex) monomials C_j^k , $k > 0$ —the complexity is completely open.

Pointers into the literature:

References

- 1 Wiebke Höhn, Tobias Jacobs: On the performance of Smith's rule in single-machine scheduling with nonlinear cost. LATIN 2012:482–493.
- 2 Nicole Megow, José Verschae: Dual techniques for scheduling on a machine with varying speed. ICALP 2013.
CoRR abs/1211.6216, 2012. <http://arxiv.org/abs/1211.6216>
- 3 Michael H. Rothkopf: Scheduling Independent Tasks on Parallel Processors. Management Science 12:437–447, 1966.
- 4 W. E. Smith: Various optimizers for single-stage production. Naval Research Logistics Quarterly 3:59–66, 1956.
- 5 Gerhard J. Woeginger: The airplane refueling problem. Open Problems – Scheduling, 2010. <http://drops.dagstuhl.de/opus/volltexte/2010/2536>
- 6 J. Yuan: The NP-hardness of the single machine common due date weighted tardiness problem. System Science and Mathematical Sciences 5:328–333, 1992.

3.12 The generalization of scheduling with machine cost

Csanád Imreh (University of Szeged, HU)

License © Creative Commons BY 3.0 Unported license
© Csanád Imreh

In resource allocation problems one has to perform some tasks during the shortest possible time or by a minimal amount of resource. The tasks are usually modeled by rectangles where the sides belong to the processing time and the required amount of resources. If we have fixed amount of resources or we have fixed time to perform the tasks then the problem can be modeled by strip packing where the items has to be packed without rotation and without overlapping into a unit width strip with minimal height. Here we investigate the problem where neither the amount of resources nor the time is fixed. Then the rectangles can be packed into an arbitrary container rectangle and the goal is to minimize the sum of γ times the used resources, denoted by H , and the time, denoted by W , thus our objective is $\gamma H + W$, where $\gamma > 0$ is a fixed parameter. If $\gamma = 1$ then the cost equals to the half of the perimeter of the container rectangle. We consider the online version of this resource allocation problem where the rectangles arrive one by one according to a list L , and we have to pack each of them into the container without any information about the further rectangles.

This model can be considered as the extension of online scheduling with machine cost. In scheduling with machine cost the number of the machines is not given as part of the input but the algorithm has to purchase each machine for a fixed unit cost. The objective function is to minimize the sum of the number of machines and the makespan which belong to the sides of an including rectangle of the schedule (thus $\gamma = 1$ holds in this model). The incoming jobs can be assigned to any existing (already bought) machine, it means that $W = m$, the number of the purchased machines, and the width of each incoming rectangle is just 1. An another closely related problem where one of the sides of the container is fixed is the online strip packing problem.

Suppose (as above) that the vertical side of the incoming rectangle corresponds to the resource needed to execute the task, and the horizontal side of the rectangle belongs to the time needed to execute the task. In the real life there exist such situations when one can modify the properties of the task (the sides of the rectangle) to use less resource, but using less resource means that it takes more time to execute the task. Or similarly, it is also possible to decrease the time needed to execute the job, but in this case the execution needs more resource. To handle this kind of relaxation it is supposed that one is allowed to change the rectangle keeping its area fixed.

There is a shelf-based 7.4803-competitive online algorithm in the standard model. For the special case when the rectangles arrive in a list ordered by decreasing height another shelf algorithm exists which is 2.5-competitive. In case of modifiable rectangles (keeping their area fixed) the best known algorithm is $(2 - \sqrt{2}) \sqrt{\frac{5}{4}\sqrt{2} + 2} \approx 1.1371$ -competitive and it is proved that no online algorithm can have smaller competitive ratio than $\frac{3}{2\sqrt{2}} \approx 1.061$.

There are several further interesting questions. First it would be good to decrease the gaps, mainly by giving a reasonable lower bound in the standard model. A second question is whether allowing rotation can yield better algorithms. And finally it would be interesting to study such models where we have to pay some penalty for changing the size of the items.

Pointers into the literature:

References

- 1 Gy. Dósa, Y. He, New upper and lower bounds for online scheduling with machine cost, *Discrete Optimization*, **7(3)**, (2010), 125–135.
- 2 Gy. Dósa, Cs. Imreh, The generalization of scheduling with machine cost, <http://www.inf.u-szeged.hu/~cimreh/machcost.pdf>
- 3 Cs. Imreh, Online strip packing with modifiable boxes, *Operations Research Letters*, **29**, (2001), 79–86.

3.13 Online Scheduling with General Cost Functions

Sungjin Im (Duke University – Durham, US)

License  Creative Commons BY 3.0 Unported license
© Sungjin Im

We consider a general online scheduling problem where the goal is to minimize $\sum_j w_j g(F_j)$, where w_j is the weight/importance of job J_j , F_j is the flow time of the job in the schedule, and g is an arbitrary non-decreasing cost function. Numerous natural scheduling objectives are special cases of this general framework:

- **Weighted Flow Time:** When $g(x) = x$, the objective becomes the total weighted flow time. The total stretch is a special case of the total weighted flow time where $w_j = 1/p_j$.
- **Weighted Flow Time Squared:** If $g(x) = x^2$ then the scheduling objective is the sum of weighted squares of the flows of the jobs.
- **Weighted Tardiness with Equal Spans:** Assume that there is a deadline d_j for each job J_j that is equal to the release time of j plus a fixed span d . If $g(t) = 0$ for t not greater than the deadline d_j , and $g(t) = w_j(t - d_j)$ for t greater than the deadline $r_j + d$, then the objective is weighted tardiness.
- **Weighted Exponential Flow:** If $g(x) = a^x$, for some real value $a > 1$, then the scheduling objective is the sum of the exponentials of the flow, which has been suggested as an

appropriate objective for scheduling problems related to air traffic control, and quality control in assembly lines.

We show that the scheduling algorithm Highest Density First (HDF) is $(2 + \epsilon)$ -speed $O(1)$ -competitive for all cost functions g *simultaneously*; see the pointer below. We also show that the HDF algorithm and this analysis are essentially optimal.

This raises a natural question if one can obtain analogous results in the stochastic setting. More concretely, what is the strongest statement one can make about *simultaneous* optimality with standard stochastic assumptions?

Pointers into the literature:

References

- 1 Sungjin Im, Benjamin Moseley, Kirk Pruhs: Online scheduling with general cost functions. SODA 2012:1254–1265.

3.14 Scheduling of Users with Time-Varying Service Rates

Peter Jacko (Lancaster University, GB)

License © Creative Commons BY 3.0 Unported license
© Peter Jacko

We discuss the problem of developing a well-performing and implementable scheduler of users with wireless connection to the base station. The main feature of such real-life systems is the time-varying quality of the channel conditions, which turn into the time-varying service rate.

If the service rate is constant, the Smith's $c\mu$ -rule is optimal for the single-server case. The variant with time-varying service rates is significantly more difficult, and it is unlikely that an optimal scheduler maintains such a simple structure. For instance, even for single-class users, threshold policies (of giving higher priority to users with higher transmission rate) are not necessarily optimal.

Practically important schedulers, however, are those as simple as possible, since the scheduling decisions are taken at milliseconds scale. Several schedulers have been proposed, but there are no (sub)optimality performance guarantees known at the moment, except for maximal stability or asymptotic (fluid) optimality, which only indicate what to do in the channel condition with highest service rate. Open problems further include adaptation to incomplete information and the multi-server case.

Pointers into the literature:

References

- 1 Sem Borst: User-level performance of channel-aware scheduling algorithms in wireless data networks. IEEE/ACM Transactions on Networking, 13(3):636–647, 2005.
- 2 Fabio Cecchi, Peter Jacko: Scheduling of Users with Markovian Time-Varying Service Rates. ACM Sigmetrics 2013.

3.15 Employee scheduling and rescheduling in call centers

Ger Koole (VU – Amsterdam, NL)

License  Creative Commons BY 3.0 Unported license
© Ger Koole

In call centers, many parameters are still uncertain the moment employees are scheduled. This leads to the necessity of real-time adjustments to the schedule. This requires different forms of flexibility in the initial schedule. Ideally, when making agent schedules the right amount of flexibility should be introduced. In this talk we discuss the different forms of parameter uncertainty, different ways to do rescheduling, and how this can be incorporated in the initial schedule.

3.16 Online scheduling algorithms analyzed by Dual Fitting

Amit Kumar (Indian Inst. of Technology – New Dehli, IN)

License  Creative Commons BY 3.0 Unported license
© Amit Kumar

I shall talk about a general dual-fitting technique for analyzing online scheduling algorithms in the unrelated machines setting where the objective involves weighted flow-time and we allow the machines of the online algorithm to have slightly extra resources than the offline optimum (the resource augmentation model). In this framework, one can often analyze simple greedy algorithms by considering the dual (or Lagrangian dual) of the linear (or convex) program for the corresponding scheduling problem, and finding a feasible dual solution as the online algorithm proceeds. I shall also mention some recent applications of this technique for deadline scheduling problems.

3.17 Approximation in Dynamic Stochastic Scheduling

Nicole Megow (TU Berlin, DE)

License  Creative Commons BY 3.0 Unported license
© Nicole Megow

Stochastic scheduling is concerned with scheduling problems in which job processing times are modeled as random variables with known probability distributions. The actual processing times are revealed only upon completion of the jobs. Such problems have been addressed since the 70s, but only more recently approximation results were derived. We give an overview of results and methods for obtaining provably good scheduling policies. This involves linear programming, lower bounding techniques borrowed from online scheduling, and index-based dynamic allocation rules known from multi-armed bandit problems. We discuss open problems, further research directions, and possible connections to other areas.

3.18 MapReduce and Distributed Scheduling

Benjamin J. Moseley (Toyota Technological Institute – Chicago, US)

License  Creative Commons BY 3.0 Unported license
© Benjamin Moseley

Recently, the MapReduce parallel computing framework has become the de facto standard for processing large data. The MapReduce distributed framework consists of an elegant combination of sequential computation and network communication that naturally lends itself to efficient distributed data processing. In a MapReduce implementation there is a centralized job tracker that coordinates job scheduling. Designing new scheduling policies has been one of the active research topics in MapReduce because of the need to balance often contradictory needs, e.g., system utilization, fairness, and response times. In this talk, we will first focus on introducing the fundamentals of MapReduce. Then we will discuss several scheduling issues that arise in MapReduce as well as recent developments in the theoretical scheduling community that have addressed these issues.

3.19 Weakly Coupled Stochastic Decision Systems

Kamesh Munagala (Duke University – Durham, US)

License  Creative Commons BY 3.0 Unported license
© Kamesh Munagala

Several problems in stochastic optimization and decision theory have the property that they are composed of many independent decision sub-problems coupled together by a few constraints. We present several examples of such problems from diverse application areas such as wireless communication, design of experiments, mechanism design, and budgeted allocations. We present unifying solution techniques based on linear programming and duality, and show connections to well-known heuristics used in practice. The talk is self-contained.

3.20 Stochastic k -TSP

Viswanath Nagarajan (IBM TJ Watson Research Center – Yorktown Heights, US)

License  Creative Commons BY 3.0 Unported license
© Viswanath Nagarajan

We will discuss a stochastic variant of the k -TSP problem: given a set of locations with random rewards, find a path originating from a depot, that minimizes the expected distance to obtain a total reward of k . We present an approximation algorithm, and upper/lower bounds on the “adaptivity gap”. The currently known results seem far from best-possible, and it would be interesting to obtain better results even for simple distributions such as Bernoulli. We will also discuss an open question regarding a submodular extension of this problem.

Pointers to some adaptive covering problems:

References

- 1 Michel X. Goemans and Jan Vondrak, Stochastic Covering and Adaptivity, LATIN 2006, 532–543.

- 2 Zhen Liu, Srinivasan Parthasarathy, Anand Ranganathan, Hao Yang, Near-optimal algorithms for shared filter evaluation in data stream systems, SIGMOD 2008, 133–146.
- 3 Sungjin Im, Viswanath Nagarajan and Ruben van der Zwaan, Minimum Latency Submodular Cover, ICALP (1) 2012, 485–497.

3.21 Job Scheduling Mechanisms for Large Computing Clusters

Seffi Naor (Technion – Haifa, IL)

License  Creative Commons BY 3.0 Unported license
© Seffi Naor

We study mechanisms for online deadline-aware scheduling in large computing clusters. Batch jobs that run on such clusters often require guarantees on their completion time (i.e. deadlines). However, most existing scheduling systems implement fair-share resource allocation between users, an approach that ignores heterogeneity in job requirements and may cause deadlines to be missed. In our framework, jobs arrive dynamically and are characterized by their value and total resource demand (or estimation thereof), along with their reported deadlines. The scheduler’s objective is to maximize the aggregate value of jobs completed by their deadlines. We circumvent known lower bounds for this problem by assuming that the input has slack, meaning that any job could be delayed and still finish by its deadline. Under the slackness assumption, we design a preemptive scheduler with a constant-factor worst-case performance guarantee.

3.22 Scheduling in queues with customer grouping

Sindo Nunez Queija (CWI – Amsterdam, NL)

License  Creative Commons BY 3.0 Unported license
© Sindo Nunez Queija

We consider a queueing system where customers of the same type may be grouped in a single service. Delaying service is therefore advantageous to reduce the server load and, as consequence, the load the system can support is unbounded if the the number of different service types is finite. The adversarial effect of delaying service is, naturally, long waiting times of customers. Our performance measure is the mean (or, alternatively, the tail probability) of customer waiting times. Optimizing on performance (avoiding long waiting times) calls for a balanced trade-off of positive and negative effects of customer grouping. We are particularly interested in an extended version of this model in which the waiting times are bounded by deadlines (patience) that either are exactly known to the service system, or their distribution is given beforehand. The motivation for this model comes from content delivery data systems, in which identical content may be requested by many users within a short time interval. Transmitting the content to several users at once corresponds to customer grouping in our model. We will discuss some initial exploration results of this problem.

3.23 Stochastic Comparison of Multicast Pull and Push

Kirk Pruhs (University of Pittsburgh, US)

License © Creative Commons BY 3.0 Unported license
© Kirk Pruhs

We consider a client server system where the server communicates to the client by broadcast. So if the server broadcasts a data item (page), then all clients waiting for that page will receive the page. We assume that clients have requests, where each request specifies a time of arrival and a page. Assume that our objective is the average time that a request has to wait to be satisfied. In a *pull* system, clients instantly forward requests to the server. So the server always knows what requests are outstanding. In a *push* system, requests are not forwarded to the server. So the server never knows what requests are outstanding. Let us assume for simplicity that all pages are of unit size (although varying sized pages is also of interest). So the server must decide at each point in time what page to broadcast.

Taking a worst-case view, the push problem is not that interesting, and the pull problem is resolved, there is a scalable algorithm and no better result is possible. Let us turn to a stochastic setting where requests for each data item are independent and Poisson. So the input distributions are completely specified by one arrival rate parameter for each page. In the push setting, it is known that computing the optimal expected waiting time is NP-hard, and a polynomial-time approximation scheme is known.

I propose that it would be an interesting question to consider the power of pulling in the stochastic setting. That is, how much better can the average waiting time be if the server knows the requests? One could compare optimal schedules, or show that there is a push algorithm that is competitive with any pull algorithm, or show that there must be a large gap between any push algorithm and a particular pull algorithm.

Pointers into the literature:

References

- 1 Nicolas Schabanel: The Data Broadcast Problem with Preemption. STACS 2000:181–192.
- 2 Nikhil Bansal, Ravishankar Krishnaswamy, Viswanath Nagarajan: Better Scalable Algorithms for Broadcast Scheduling. ICALP 2010:324–335

3.24 Scheduling with time-varying cost: Deterministic and stochastic models

Roman Rischke (TU Berlin, DE)

License © Creative Commons BY 3.0 Unported license
© Roman Rischke

We consider a natural generalization of classical scheduling problems in which using a time slots for processing a job causes some time-dependent cost in addition to the standard scheduling cost. Adding the cost consideration to classical scheduling increases the problem complexity significantly. Nevertheless, we also propose a related two-stage stochastic scheduling model with recourse. Suppose that the exact scheduling instance is not known in the first stage, we only have probabilistic information about scenarios, where a scenario represents a scheduling instance. In the first stage we may reserve time slots at some cheap price. In the second stage a particular scenario becomes known and we may buy additional time slots at a

(scenario-dependent) higher price so as to find a feasible schedule for the realized scenario minimizing the total cost of buying and scheduling. Notice that the deterministic problem described above appears as the recourse problem in the second stage in the stochastic model. Although two-stage (multi-stage) stochastic optimization with recourse has received a lot of attention in the theory of approximation algorithms in the past decade, corresponding scheduling problems have hardly been addressed.

In this short presentation we want to advertise the very natural deterministic and stochastic scheduling problems taking into account costs for using/reserving time slots. Both problems are much more complex than classical scheduling problems and require new techniques. We give some first insights and show some relation to scheduling on a machine that may change its speed. We hope to foster further research activity in this hitherto little explored area.

This is joint work in progress with S. Leonardi, N. Megow, L. Stougie, C. Swamy and J. Verschae.

Pointers into the literature:

References

- 1 L. Epstein, A. Levin, A. Marchetti-Spaccamela, N. Megow, J. Mestre, M. Skutella, and L. Stougie. Universal sequencing on an unreliable machine. *SIAM J. Comput.*, 41(3):565–586, 2012.
- 2 A. Gupta, M. Pal, R. Ravi, and A. Sinha. Sampling and cost-sharing: Approximation algorithms for stochastic optimization problems. *SIAM J. Comput.*, 40(5):1361–1401, 2011.
- 3 J. Kulkarni and K. Munagala. Algorithms for cost aware scheduling. In *Proc. of WAOA 2012*, to appear 2013.
- 4 N. Megow and J. Verschae. Scheduling on a machine with varying speed: Minimizing cost and energy via dual schedules. *CoRR*, abs/1211.6216, 2012.

3.25 Online scheduling of jobs with fixed start times on related machines

Jiri Sgall (Charles University – Prague, CZ)

License © Creative Commons BY 3.0 Unported license
© Jiri Sgall

We consider online scheduling of identical jobs with fixed starting times revealed at those times on m uniformly related machines, with the objective of maximizing the number of completed jobs. A newly released job must be either assigned to start running immediately on a machine or otherwise it is dropped. It is also possible to drop an already scheduled job, but only completed jobs contribute their weights to the profit of the algorithm.

In the paper we show that a natural greedy algorithm is $4/3$ -competitive and optimal on $m=2$ machines, while for a large m , its competitive ratio is between 1.56 and 2. Furthermore, no algorithm is better than 1.5-competitive. It is an open problem to improve these bounds.

3.26 Scheduling Advertising Campaigns

Hadas Shachnai (Technion – Haifa, IL)

License © Creative Commons BY 3.0 Unported license
© Hadas Shachnai

An advertising campaign is a series of advertisement messages that share a single idea and theme which make up an integrated marketing communication. Given a large set of campaigns that can be potentially delivered to a media audience, a service provider attempts to fully deliver a subset of campaigns that maximizes the total revenue, while satisfying constraints on the placement of ads that belong to the same campaign, as well as possible placement constraints among conflicting campaigns. In particular, to increase the number of viewers exposed to an ad campaign, one constraint is that each commercial break contains no more than a single ad from this campaign. Each ad has a given length (=size), which remains the same, regardless of the commercial break in which it is placed. This generic assignment problem defines a family of all-or-nothing variants of the generalized assignment problem (GAP). We design for these variants approximation algorithms with constant-factor worst-case performance guarantees.

3.27 Optimal queue-size scaling in switched networks

Devavrat Shah (MIT – Cambridge, US)

License © Creative Commons BY 3.0 Unported license
© Devavrat Shah

We consider a switched (queueing) network in which there are constraints on which queues may be served simultaneously; such networks have been used to effectively model input-queued switches, wireless networks and more recently data-centers. The scheduling policy for such a network specifies which queues to serve at any point in time, based on the current state or past history of the system. Designing a scheduling policy with optimal average queue-size for switched network has been a question of interest for a while now. As the main result, we shall discuss a new class of online scheduling policies that achieve optimal scaling for average queue-size for a class of switched networks including input-queued switches.

Talk is based on work with Neil Walton (U of Amsterdam)+ Yuan Zhong (UC Berkeley).

3.28 On Carry Over

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

License © Creative Commons BY 3.0 Unported license
© Frits C. R. Spijksma

Any schedule for a round robin tournament involves an order in which each team meets its opponents. We say that team i gives a *carry-over effect* to team j , if some other team t 's match against i is followed by a match against team j . This is particularly relevant in physical, body-contact sports. For instance, if team i is a very strong, tough-playing side, one can imagine that its opponent, team t , is weakened by injuries or fatigue, which could be an advantage for its next opponent, team j . Moreover, the carry-over effect can also be

relevant in a strictly psychological interpretation, when team t loses confidence and morale after a severe loss against the strong team i , again to the benefit of their next opponent, team j . The opposite may be true if team i is a weak team. Clearly, carry-over effects are unavoidable in any schedule, however, schedules can differ in the extent to which carry-over effects are balanced over the teams. We define c_{ij} as the number of times that team i gives a carry-over effect to team j in a schedule. The degree to which the carry-over effects are balanced is typically measured by the so-called carry-over effects value, which is defined as $\sum c_{i,j}^2$ (Russell 1980).

The table below gives the best known values for the carry-over effect, depending upon the number of teams. An asterisk denotes that the given value is minimum, i.e., no schedule exists achieving a value lower than the reported number.

Number of teams	Carry-over value
4	12* (Russell 1980)
6	60* (Russell 1980)
8	56* (Russell 1980)
10	108* (Anderson 1999, Eggermont 2011)
12	170 (Eggermont 2011)
14	234 (Anderson 1999)
16	240* (Russell 1980)
18	340 (Anderson 1999)
20	380* (Anderson 1999)
22	462* (Anderson 1999)
24	644 (Anderson 1999)

Pointers to the literature:

References

- 1 Anderson, I. (1999), Balancing carry-over effects in tournaments, in: *Combinatorial designs and their applications*, Chapman and Hall/CRC, page 1–16.
- 2 Eggermont, C. (2011), reachability Problems in Scheduling and Planning, PhD thesis Eindhoven University of Technology.
- 3 Russel, K. (1980), Balancing carry-over effects in round robin tournaments, *Biometrika* 67, 127–131.

3.29 Stochastic Optimal Control for a Class of Dynamic Resource Allocation Problems

Mark S. Squillante (IBM TJ Watson Research Center – Yorktown Heights, US)

License  Creative Commons BY 3.0 Unported license
© Mark S. Squillante

We consider a class of general dynamic resource allocation problems within a stochastic optimal control theoretic framework. This class of problems arises in a wide variety of applications, each of which intrinsically involves resources of different types and demand with uncertainty and/or variability. The goal is to determine the allocation capacity for every resource type in order to serve the uncertain/variable demand and maximize the expected profit (utility) over a time horizon of interest based on the rewards, costs and flexibility

associated with the different resources. We derive the optimal (online) control policy within a singular stochastic optimal control setting, which includes simple expressions for governing the dynamic adjustments to resource allocation capacities over time. Numerical experiments investigate various issues of both theoretical and practical interest, quantifying the benefits of our approach over alternative optimization approaches.

This talk is based on joint work with Xuefeng Gao, Yingdong Lu, Mayank Sharma and Joost Bosman.

3.30 Truthful Scheduling

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

License  Creative Commons BY 3.0 Unported license
© Rob van Stee

A major question in algorithmic game theory is whether the presence of selfish agents affects the approximability of various classic optimization problems. Specifically, the following research agenda was suggested: “*to what extent is incentive compatible efficient computation fundamentally less powerful than ‘classic’ efficient computation?*”. Of particular interest are scheduling problems, where jobs are assigned for processing to agents, each controlling one machine, and who have some private information regarding their machines. In this paper, we consider the case of single-parameter agents with scheduling problems on uniformly related machines, which was among the first problems considered in the area of algorithmic mechanism design. The private information of an agent is the cost of processing one unit of work, which is also the inverse of the speed of the machine. We show that selfish agents do not affect the approximability of scheduling problems on uniformly related machines, by designing $(1 + \varepsilon)$ -approximation mechanisms for these problems for any $\varepsilon > 0$. Given that these problems are (strongly) NP-hard, this is the best possible result.

Non-preemptive scheduling problems on m uniformly related machines are defined as follows. We let the set of machines be denoted by $M = \{1, 2, \dots, m\}$. We are given a set of jobs $J = \{1, 2, \dots, n\}$, where each job j has a positive size p_j . The jobs need to be partitioned into m subsets S_1, \dots, S_m , with S_i being the subset of jobs assigned to machine i . We let s_i denote the (actual) speed of machine i , meaning that the processing of job j takes $\frac{p_j}{s_i}$ time units if j is assigned to machine i . For such a solution (also known as a schedule), we let $L_i = (\sum_{j \in S_i} p_j) / s_i$ be the *completion time* or *load* of machine i . The *work* of machine i is $W_i = \sum_{j \in S_i} p_j = L_i \cdot s_i$, that is, the total size of the jobs which are assigned to i . We consider objective functions which are functions of the machine loads, L_1, L_2, \dots, L_m . We consider a variety of objective functions (social goals), like minimizing the makespan (maximum load) and maximizing the minimum load (cover).

The setup of mechanism design for single-parameter agents operating uniformly related machines is as follows. Agents present bids to a mechanism, where the bid b_i of an agent i is the claimed cost per unit of work of its machine (the inverse of its claimed speed). Based on these bids, the mechanism allocates the jobs to the machines and also assigns payments to the agents. We assume that each agent is only interested in maximizing its own profit, which is its payment minus its (actual) cost of processing the jobs allocated to it. A mechanism is called *truthful* if reporting their true costs per unit of work is a dominant strategy for the agents. That is, this strategy maximizes the profit for each agent, regardless of the strategies of the other agents. In the case of single-parameter agents, a well-known necessary

and sufficient condition for truthfulness is that the allocation algorithm is *monotone*, that is, the allocation algorithm must have the property that if an agent i increases its claimed speed (i.e., decreases its bid) while all other bids are unchanged, the work allocated to i does not decrease. More precisely, in such a case there exist simple payment functions that can be coupled with the (monotone) allocation algorithm to give a truthful mechanism. If the allocation algorithm runs in polynomial time, and the payments can be computed in polynomial time as well, then the resulting truthful mechanism can be implemented in polynomial time.

Pointers into the literature:

References

- 1 Aaron Archer. *Mechanisms for Discrete Optimization with Rational Agents*. PhD thesis, Cornell University, 2004.
- 2 L. Epstein and J. Sgall. Approximation schemes for scheduling on uniformly related and identical parallel machines. *Algorithmica*, 39(1):43–57, 2004.

3.31 Online Stochastic Matching

Clifford Stein (Columbia University, US)

License  Creative Commons BY 3.0 Unported license
© Clifford Stein

Online matching is an important and well-studied problem and has received particular attention recently because of its connections to allocation problems in internet advertising. In the traditional worst-case model, on-line algorithms are well understood, with tight competitive ratios of $1/2$ for deterministic algorithms and $(e-1)/e$ for randomized algorithms.

Several recent works have considered relaxing the worst-case model and considering various probabilistic models for on-line matching. We will survey some of the models and results, and also describe an online algorithm in a model in which an adversary chooses the graph, but does not control the order in which the online nodes are revealed (This work is joint with Feldman, Korula, Henzinger, and Mirrokni). We will also show how these ideas generalize to online packing problems

3.32 Split scheduling with uniform setup times

Suzanne van der Ster (VU – Amsterdam, NL)

License  Creative Commons BY 3.0 Unported license
© Suzanne van der Ster

We consider the problem of scheduling a set of n jobs on m parallel machines with *job splitting* and *setup times*. This means that each job may be split into parts and multiple parts of the same job may be processed simultaneously. However, before a machine can start processing (a part of) a job, a fixed setup time s is required. (Note that the setup time is job-, machine-, and sequence-independent.) Job j has processing time p_j , for $j = 1, \dots, n$ and the objective is to minimise the sum of completion times of the jobs.

Given a schedule for the jobs, we define by M_j the set of all machines on which job j is scheduled. A job is called a d -job when it is scheduled on d machines. We call a job balanced

in a schedule if it completes at the same time on all machines on which it is processed. We assume that jobs are numbered in non-decreasing order of their processing times. This order is also called the Shortest Processing Time first (SPT) order.

There is a polynomial-time algorithm solving the problem for the case of 2 machines. The properties below give us a handle to find an optimal schedule.

- (a) There exists an optimal schedule such that on each machine the job parts are processed (started and completed) in SPT order of the corresponding jobs.
- (b) There exists an optimal schedule in which all jobs are balanced.
- (c) There are no 1-jobs after a 2-job in an optimal schedule satisfying the two properties above.

From these properties we derive that all 2-jobs are scheduled in SPT order at the end and the first 2-job is not shorter than the preceding 1-jobs. This implies that the 1-jobs can be scheduled in SPT order without increasing the completion time of the 2-jobs. We can simply consider each of the n jobs as the first 2-job, or perform a more careful update of consecutive solutions to improve on the running time.

In future research we want to study the problem for more than 2 machines. Already for 3 machines this is anything but straightforward. Properties (a) and (b) hold for any number of machines, but property (c) (or a generalization of it) does not. We have examples where in none of the optimal solutions $|M_j|$ is monotone in j . A few of the questions that could be considered:

- Does there exist an optimal schedule such that the jobs are started (or finished) in SPT order? (This is an option that is not ruled out by the counterexamples showing that a generalization of (c) does not hold.)
- What is the complexity of the problem for more than 2 machines? Or even, what is the complexity for only 3 machines?

The results given above, as well as a more extensive discussion of the problems we encounter already for three machines, can be found here.

References

- 1 F. Schalekamp, R. Sitters, S. van der Ster, L. Stougie, V. Verdugo, and A. van Zuylen. Split scheduling with uniform setup times. <http://arxiv.org/abs/1212.1754>

3.33 An infinite server system with customer-to-server packing constraints

Alexander Stolyar (Bell Labs – Murray Hill, US)

License © Creative Commons BY 3.0 Unported license
© Alexander Stolyar

The model is motivated by the problem of efficient “packing” of virtual machines into physical host machines in a network cloud data center. There is an infinite number of servers and multiple flows of arriving customers of different types. Each server can simultaneously serve several customers, subject to some “packing” constraints. Service times of different customers are independent – even if customers share a server. Customers leave after their service is complete. The underlying objective is to minimize the number of occupied servers. We show that some versions of a greedy packing strategy are asymptotically optimal as the system scale (the average total number of customers in service) goes to infinity.

3.34 Analysis of Smith’s Rule in Stochastic Scheduling

Marc Uetz (*University of Twente, NL*)

License  Creative Commons BY 3.0 Unported license
© Marc Uetz

In a landmark paper from 1986, Kawaguchi and Kyan show that scheduling jobs according to ratios weight over processing time – a.k.a. Smith’s rule – has a tight performance guarantee of 1.207 for minimizing the total weighted completion time in parallel machine scheduling. The talk addresses the performance of Smith’s rule for the variant where processing times are exponentially distributed random variables. The current status of the problem is as follows. The expected performance of Smith’s rule is no worse than $(2 - 1/m)$ times optimum, m being the number of machines Möhring, Schulz and Uetz, JACM 46, 1999, 924-942. This analysis is based on a LP relaxation, and hold for any random variable with coefficient of variation at most 1. On the other hand, the analysis of a (slightly modified) stochastic version of the Kawaguchi and Kyan instance yields an expected performance at least 1.243 times optimum Jagtenberg, Schwiegelshohn and Uetz, WAOA 2010. The open problem is to find the true performance bound for Smith’s rule in stochastic parallel machine scheduling.

3.35 Appointment Scheduling with Slot Blocking

Peter van de Ven (*IBM TJ Watson Research Center – Hawthorne, US*)

License  Creative Commons BY 3.0 Unported license
© Peter van de Ven

Consider an electronic appointment booking system that allows patients to book their doctor appointment online; the system displays the available time slots and the patient selects an acceptable slot or leaves if none of the available slots are acceptable to him. Patients typically have some set of acceptable slots, and surveys shows that certain slots are popular (deemed acceptable by more patients) than others. The popular slots tend to be booked early in the reservation process, leaving only less desirable slots open. As a result, patients that arrive at a later stage may not find any acceptable slot upon arrival, and at the end of the reservation proces several unpopular slots may be left unoccupied, reducing the utilization of this doctor.

We propose to block the popular slots during initial stages of the reservation process, forcing early patients to select less popular slots. Making available the popular slots to later arrivals then allows for more flexibility and may result in higher utilization. We study the tradeoff between sending away early patients (blocking too long) and having little flexibility later on (blocking too little). This model has an interesting connection with the online bipartite matching problem.

3.36 Learning in Stochastic Scheduling

Tjark Vredeveld (Maastricht University, NL)

License  Creative Commons BY 3.0 Unported license
© Tjark Vredeveld

We consider a scheduling problem in which K classes of independent jobs have to be processed non-preemptively. The processing times of the jobs are assumed to be exponentially distributed with parameters depending on the class of each job. However, the precise values of the parameters are unknown to the scheduler. The objective is to minimize the sum of expected completion times.

The scheduler has certain beliefs on the values of the parameters for the different classes and the belief on a parameter of a certain class is updated after the completion of a job of this class.

As the Shortest Expected Processing Time (SEPT) rule is optimal for the single machine problem without parameter uncertainty and it is a simple rule, we are interested in how variants of this rule perform in the presence of parameter uncertainty. It has been shown that on a single machine, two variants of SEPT have a performance guarantee of 2, i.e., SEPT obtains a schedule with expected value at most two times the optimal value, and that this bound is tight whenever there are only two different classes of jobs.

An interesting question is what the performance guarantees of several variants of SEPT are when there is more than one machine and more than 2 job classes.

Based on joint work with S. Marbán and C. Rutten.

Pointers into the literature:

References

- 1 S. Marbán, C. Rutten, T. Vredeveld. Learning in stochastic machine scheduling. WAOA 2011:21-34
- 2 S. Marbán (2012). Pricing and Scheduling under Uncertainty. PhD Thesis, Maastricht University. (Chapters 4, 5)

3.37 FCFS infinite matching, queues with skill based routing, and organ transplants

Gideon Weiss (Haifa University, IL)

License  Creative Commons BY 3.0 Unported license
© Gideon Weiss

This talk is a survey of recent work with Ivo Adan and several other collaborators, including Cor Hurkens, Marko Boon, Ana Busic, and Jean Mairesse. It is based on earlier work of Rishi Talreja and Ward Whitt, and of Rene Caldentey and Ed Kaplan. In recent years it has become very important to investigate service systems which serve customers of several types, and which employ servers of various skills. Such service systems are referred to in current literature as queues with skill based routing. These have several types of customers, and several types of servers, and a bipartite compatibility graph to indicate which types of servers can serve which types of customers. Applications include such varied fields as call centers, outsourcing, manufacturing process, cloud computing and health systems. A somewhat different model is the matching of applicants and positions, of organ donors and patients, of adoptive parents and children, the so called marriage model. These two types of

service models motivate our research, we note that these applications are also studied by schedulers and combinatorial optimizers. There is a significant difference between these two models: While in a queueing model the sequence of interarrival times and the sequence of service times get hopelessly entangled through the busy/idle cycles of the servers, this is absent from the matching applications: Here customers and servers play a symmetric role and each customer server encounter has no after effects on the rest of the system. The taxi rank is a simple example of that. A further simplification is to consider just the sequence of customers and of servers, ordered by arrivals and classified by types: We think of those as i.i.d infinite sequences of customer types (all the customers that will appear in the future) and of server types (all the services that will be given). For these sequences it is natural to define a FCFS matching: The first server is matched to the first compatible customer, and the n th server is matched to the first compatible customer that none of the previous $n - 1$ servers picked up. This infinite FCFS model has a much more combinatorial flavor than the standard queueing models. It turns out that it is quite tractable. The key property which we discovered is that this infinite matching model is in some sense time reversible. Reversibility plays a key role in the theory of Markov chains and of queueing models. It is the property which underlies product form results. It is a sad fact that most queueing models are complicated and often intractable. Research on queueing networks would have gone nowhere if it weren't for Jackson's discovery that Jackson networks have steady state distribution. From that time onwards, product form results have been keenly sought after, as they seem the best way of getting explicit solutions and useful insight for more general models. Product form is always related to some form of reversibility. The reversibility of the infinite matching model underlies all our further results. Coming back to queues with skilled based routing, we focus on FCFS-ALIS, first come first served, assign longest idle server policy. This means that whenever a server becomes available he will go to the longest waiting customer in the system which he can serve, and whenever a customer arrives to find several idle servers he will be assigned to the longest idle compatible server. This policy has several attractive features: First and foremost it is fair to both customers and servers, in many systems e.g. organ donations, public housing assignment, FCFS is dictated by law. ALIS is the best way to equalize the efforts of the servers, and thus it encourages diligent service. The policy is also very natural and easy to implement, and it requires minimal information about the parameters of the system and its current state. As a result it is useful in systems in which load and staffing keep changing over the operating horizon. Our exact results are that under a FCFS-ALIS policy when arrivals are Poisson and services exponential we have product form queue length distributions, and closed form expressions for waiting times. Furthermore, these results can be used to obtain excellent approximations for much more general queueing models, under many server scaling. These include general link dependent service times, abandonments, and efficiency driven systems. We will present some examples of how to use these results in the design of call centers, and in the planning of an organ transplants policy.

4 Working Groups

Our experience is that it takes at least a year before one can start to assess the impact of a seminar with any confidence. But some initial collaborations arising from the seminar are:

- Based on collaborations in preparing for the seminar, Bert Zwart (from the stochastic community) and Nikhil Bansal (from the worst-case community) wrote a successful grant proposal for funding of one joint Ph.D. student.

- Nicole Megow and Leen Stougie (from the worst-case community) began a research collaboration on 2-stage stochastic optimization problems.
- Frits Spijksma (from the worst-case community) and Onno Boxma (from the stochastic community) have started a research collaboration on scheduling shipping locks.
- Martin Skutella, Maxim Sviridenko, and Marc Uetz (from the worst-case community) began a research collaboration on stochastic scheduling of unrelated machines, resulting in a paper, <http://eprints.eemcs.utwente.nl/23246/>
- Ger Koole and Rhonda Righter started a research collaboration on assessing various performance measures for call centers with abandonment.
- Urtzi Ayesta and Rhonda Righter started a research collaboration on scheduling of bandit processes in a random environment.
- Samir Khuller, and Christoph Dürr each began research with one of their Ph.D. students on one of the open problems presented in the short talks.

5 Open Problems

Open problems are included in the talk abstracts above.

Participants

- Samuli Aalto
Aalto University, FI
- Susanne Albers
HU Berlin, DE
- Konstantin Avrachenkov
INRIA Sophia Antipolis –
Méditerranée, FR
- Urtzi Ayesta
LAAS – Toulouse, FR
- Yossi Azar
Tel Aviv University, IL
- Nikhil Bansal
TU Eindhoven, NL
- Vincenzo Bonifaci
National Research Council –
Rome, IT
- Sem C. Borst
TU Eindhoven, NL
- Onno J. Boxma
TU Eindhoven, NL
- Marek Chrobak
University of California –
Riverside, US
- Christoph Dürr
UPMC, Lab. LIP6 – Paris, FR
- Fritz Eisenbrand
EPFL – Lausanne, CH
- Thomas Erlebach
University of Leicester, GB
- Naveen Garg
Indian Inst. of Technology –
New Dehli, IN
- Anupam Gupta
Carnegie Mellon University –
Pittsburgh, US
- John Hasenbein
University of Texas – Austin, US
- Wiebke Höhn
TU Berlin, DE
- Sungjin Im
Duke University – Durham, US
- Csanad Imreh
University of Szeged, HU
- Peter Jacko
Lancaster University, GB
- Samir Khuller
University of Maryland, US
- Amit Kumar
Indian Inst. of Technology –
New Dehli, IN
- Fei Li
George Mason University –
Fairfax, US
- Alberto Marchetti-Spaccamela
University of Rome “La
Sapienza”, IT
- Monaldo Mastrolilli
IDSIA – Lugano, CH
- Nicole Megow
TU Berlin, DE
- Rolf H. Möhring
TU Berlin, DE
- Benjamin J. Moseley
Toyota Technological Institute –
Chicago, US
- Kamesh Munagala
Duke University – Durham, US
- Viswanath Nagarajan
IBM TJ Watson Res. Center –
Yorktown Heights, US
- Seffi Naor
Technion – Haifa, IL
- Sindo Nunez Queija
CWI – Amsterdam, NL
- Kirk Pruhs
University of Pittsburgh, US
- Rhonda Righter
University of California –
Berkeley, US
- Roman Rischke
TU Berlin, DE
- Jiri Sgall
Charles University – Prague, CZ
- Hadas Shachnai
Technion – Haifa, IL
- Devavrat Shah
MIT – Cambridge, US
- René Sitters
VU – Amsterdam, NL
- Martin Skutella
TU Berlin, DE
- Frits C. R. Spieksma
K.U. Leuven, BE
- Mark S. Squillante
IBM TJ Watson Res. Center –
Yorktown Heights, US
- Clifford Stein
Columbia University, US
- Sebastian Stiller
TU Berlin, DE
- Alexander Stolyar
Bell Labs – Murray Hill, US
- Leen Stougie
CWI – Amsterdam, NL
- Maxim I. Sviridenko
University of Warwick, GB
- Marc Uetz
University of Twente, NL
- Peter van de Ven
IBM TJ Watson Res. Center –
Hawthorne, US
- Suzanne van der Ster
VU – Amsterdam, NL
- Rob van Stee
MPI für Informatik –
Saarbrücken, DE
- Tjark Vredeveld
Maastricht University, NL
- Gideon Weiss
Haifa University, IL
- Gerhard J. Woeginger
TU Eindhoven, NL
- Prudence W. H. Wong
University of Liverpool, GB



Report from Dagstuhl Seminar 13121

Bidimensional Structures: Algorithms, Combinatorics and Logic

Edited by

Erik D. Demaine¹, Fedor V. Fomin², MohammadTaghi Hajiaghayi³, and Dimitrios M. Thilikos⁴

1 MIT – Cambridge, US, edemaine@mit.edu

2 University of Bergen, NO, fomin@ii.uib.no

3 University of Maryland – College Park, US, hajiagha@cs.umd.edu

4 National and Kapodistrian University of Athens, GR, sedthilk@thilikos.info

Abstract

We provide a report on the Dagstuhl Seminar 13121: *Bidimensional Structures: Algorithms, Combinatorics and Logic* held at Schloss Dagstuhl in Wadern, Germany between Monday 18 and Friday 22 of March 2013. The report contains the motivation of the seminar, the abstracts of the talks given during the seminar, and the list of open problems.

Seminar 18.–22. March, 2013 – www.dagstuhl.de/13121

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems, G.2.2 Graph Theory

Keywords and phrases Graph Minors, Treewidth, Graph algorithms, Parameterized Algorithms

Digital Object Identifier 10.4230/DagRep.3.3.51

1 Executive Summary

Erik D. Demaine

Fedor V. Fomin

MohammadTaghi Hajiaghayi

Dimitrios M. Thilikos

License © Creative Commons BY 3.0 Unported license

© Erik D. Demaine, Fedor V. Fomin, MohammadTaghi Hajiaghayi, and Dimitrios M. Thilikos

The monumental Graph Minors project developed by Robertson and Seymour in the 1980s is one of the most fundamental achievements of Combinatorics. The project had several groundbreaking consequences for Theoretical Computer Science. However, the wide spread opinion in the algorithmic community, expressed by David S. Johnson in his NP-Completeness Column (J. Algorithms 1987), was that it is mainly of theoretical importance. It took some time to realize that the techniques developed in Graph Minors can be used in the design of efficient and generic algorithms. One of the main techniques extracted from Graph Minors is based on the structural results explaining the existence (or the absence) of certain grid-like or bidimensional structures in graphs. The usage of bidimensional structures and the related width parameters in many areas of Computer Science and Combinatorics makes such techniques ubiquitous.

Historically, the first applications of bidimensional structures are originated in Graph Minors of Robertson and Seymour, because of the structure of the graphs excluding some fixed some graph as a minor. There is still an on-going work in Combinatorics on obtaining new



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

Bidimensional Structures: Algorithms, Combinatorics and Logic, *Dagstuhl Reports*, Vol. 3, Issue 3, pp. 51–74

Editors: Erik D. Demaine, Fedor V. Fomin, MohammadTaghi Hajiaghayi, and Dimitrios M. Thilikos



DAGSTUHL
REPORTS

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

structural theorems. There are much more examples in Combinatorics, where bidimensional structures and width parameters play a crucial role like in obtaining Erdős-Pósa type of results. Reed used bidimensional structures to settle Erdős-Hajnal conjecture on near-bipartite graphs. Kawarabayashi and Reed used bidimensional structures to bound the size of a minimal counterexample to Hadwiger's conjecture. Demaine and Hajiaghayi optimized the original grid-exclusion theorem on H -minor free graphs.

The usage of bidimensional structures and width parameters in Algorithms goes back to the parameter of treewidth, introduced in the Graph Minors series. Treewidth is now ubiquitous in algorithm design and expresses the degree of topological resemblance of a graph to the structure of a tree. Its algorithmic importance dates back in the early 90's to the powerful meta-algorithmic result of Courcelle asserting that all graph problems expressible in Monadic Second Order Logic can be solved in linear time on graphs of bounded treewidth. Bounded treewidth can be guaranteed by the exclusion of certain bidimensional structures. Intuitively, this exclusion is what enables the application of a series of classic algorithmic techniques (divide-and-conquer, dynamic programming, finite automata) for problems of certain descriptive complexity. This phenomenon was perhaps the first strong indication of the deep interleave between graph structure and logic in graph algorithms. However, a deeper understanding of it became more evident during the last decade and produced powerful meta-algorithmic techniques.

Apparently, graph-theoretic fundamentals emerging from the Graph Minors project developed by Robertson and Seymour, are used currently in several areas of Computer Science and Discrete Mathematics. Algorithmic fertilization of these ideas occurred mostly in the context of parameterized complexity and its foundational links to logic. The course of developing a structural algorithmic graph theory revealed strong connections between Graph Theory, Algorithms, Logic, and Computational Complexity and joined a rapidly developing community of researchers from Theoretical Computer Science and Discrete Mathematics.

Dagstuhl seminar 13121 brought together some of the most active researchers on this growing field.

2 Table of Contents

Executive Summary

<i>Erik D. Demaine, Fedor V. Fomin, MohammadTaghi Hajiaghayi, and Dimitrios M. Thilikos</i>	51
---	----

Overview of Talks

A welcome to treewidth <i>Dimitrios M. Thilikos</i>	55
Bidimensionality and its applications I: Bidimensionality: Yesterday, Today and Tomorrow <i>Saket Saurabh</i>	55
Bidimensionality and its applications II: Graph Surgery and Kernelization <i>Daniel Lokshtanov</i>	55
Treewidth and dimension <i>Gwenaël Joret</i>	56
Canonical tree-decompositions, k -blocks and tangles <i>Reinhard Diestel</i>	56
Subexponential-time parameterized algorithm for Steiner Tree on planar graphs <i>Erik Jan van Leeuwen</i>	56
Rank based algorithms for bounded treewidth graphs <i>Marek Cygan</i>	57
Decomposing quantified conjunctive (or disjunctive) formulas <i>Victor Dalmau</i>	57
Branch Decompositions and Linear Matroids <i>Illya V. Hicks</i>	57
A 5-approximation for treewidth using linear time, single exponential in the treewidth <i>Hans Bodlaender</i>	58
Contraction Decomposition: a new technique for H-minor-free graphs <i>Erik Demaine</i>	58
Topological problems in tournaments <i>Michał Pilipczuk</i>	58
Kernelization using structural parameters on sparse graph classes <i>Felix Reidl</i>	59
A new proof for the weak-structure theorem with explicit constants <i>Paul Wollan</i>	59
Approximability and fixed parameter algorithms: a new look <i>Rajesh Chitnis</i>	59
Excluded vertex-minors for graphs of linear rank-width at most k <i>Sang-il Oum</i>	60
Definability of numerical graph parameters and various notions of width <i>Johann A. Makowsky</i>	60

The k -disjoint paths problem in directed planar graphs <i>Dániel Marx</i>	61
Surface split decompositions: fast dynamic programming over branch decompositions for graphs of bounded genus <i>Paul Bonsma</i>	61
Tools for multicoloring with applications to planar graphs and bounded treewidth graphs <i>Guy Korstarz</i>	61
What makes normalized weighted satisfiability tractable? <i>Iyad A. Kanj</i>	62
An excluded grid theorem for digraphs with forbidden minors <i>Stephan Kreutzer</i>	63
Exclusion theorems for immersions on surface embedded graphs <i>Archontia Giannopoulou</i>	63
Packing edge-disjoint odd S-cycles in 4-edge-connected graphs <i>Yusuke Kobayashi</i>	63
Tree-width of hypergraphs and surface duality <i>Frédéric Mazoit</i>	64
Open problems	64
Program	72
Participants	74

3 Overview of Talks

3.1 A welcome to treewidth

Dimitrios M. Thilikos (University of Athens, GR)

License  Creative Commons BY 3.0 Unported license
© Dimitrios M. Thilikos

This talk is a short introduction to the various fields where the graph invariant of treewidth has been important. This includes: (i) the deep combinatorial results of the Graph Minors Series of Robertson and Seymour, (ii) the meta-algorithmic framework initiated by Courcelle's theorem, (iii) the derivation of dynamic programming algorithms and its multiple applications in parameterized algorithms and complexity, and (iv) the design of EPTAS for NP-hard problems. It is stressed that, in all above fields, combinatorial results concerning *bidimensional structures* in graphs, such as the *Grid Exclusion Theorem*, have played an important role.

3.2 Bidimensionality and its applications I: Bidimensionality: Yesterday, Today and Tomorrow

Saket Saurabh (University of Bergen, NO)

License  Creative Commons BY 3.0 Unported license
© Saket Saurabh

In this talk we give the chronological survey of the development of Bidimensionality as a field. This talk will set up all the necessary definitions for talks to come and explain the key developments in the area in details.

3.3 Bidimensionality and its applications II: Graph Surgery and Kernelization

Daniel Lokshantov (University of Bergen, NO)

License  Creative Commons BY 3.0 Unported license
© Daniel Lokshantov

Over the last few years there has been quite a few results in parameterized algorithms and kernelization that are based on cutting away a piece of the input instance, analyzing it, and replacing it by a smaller, equivalent piece. In this talk we outline a language in which it is natural to talk about such operations when the considered instances are graphs. In particular we define bounded graphs together with some operations on them, and show how these operations have some nice algebraic properties. We then proceed to use this language to show that a large class of parameterized problems admit linear kernels on any class of graphs excluding a fixed H as a minor.

3.4 Treewidth and dimension

Gwenaël Joret (Université Libre de Bruxelles, BE)

License  Creative Commons BY 3.0 Unported license
© Gwenaël Joret

Over the last 30 years, researchers have investigated connections between dimension for posets and planarity for graphs. Here we extend this line of research to the structural graph theory parameter tree-width by proving that the dimension of a finite poset is bounded in terms of its height and the tree-width of its cover graph.

3.5 Canonical tree-decompositions, k -blocks and tangles

Reinhard Diestel (University of Hamburg, DE)

License  Creative Commons BY 3.0 Unported license
© Reinhard Diestel

A k -block in a graph G is a set X of at least k vertices no two of which are separated in G by fewer than k other vertices (which may or may not lie in X). In joint work with Carmesin, Hundertmark and Stein I recently proved that, for each k , every graph has a ‘canonical’ tree-decomposition whose adhesion sets have order $< k$ and separate any two k -blocks, which thus come to lie in different parts of the decomposition. The decompositions are *canonical* in that the automorphisms of the graph act on their sets of parts inducing automorphisms of the decomposition trees. The following algorithmic problems may be of interest, given an integer k and a graph G : to find all the k -blocks in G ; to construct a canonical tree-decomposition as above; to use these decompositions to solve Graph Isomorphism in polynomial time for suitable classes of graphs.

3.6 Subexponential-time parameterized algorithm for Steiner Tree on planar graphs

Erik Jan van Leeuwen (Max-Planck Institut für Informatik, Saarbrücken, DE)

License  Creative Commons BY 3.0 Unported license
© Erik Jan van Leeuwen

Bidimensionality theory provides a method for designing fast, subexponential-time parameterized algorithms for a vast number of NP-hard problems on sparse graph classes such as planar graphs, bounded genus graphs, or, more generally, graphs with a fixed excluded minor. However, in order to apply the bidimensionality framework the considered problem needs to fulfill a special density property. Some well-known problems do not have this property, unfortunately, with probably the most prominent and important example being the Steiner Tree problem. Hence the question whether a subexponential-time parameterized algorithm for Steiner Tree on planar graphs exists has remained open. In this talk, we answer this question positively and develop an algorithm running in time subexponential in k and polynomial space, where k is the size of the Steiner tree. Our algorithm does not rely on tools from bidimensionality theory or graph minors theory, apart from Baker’s classical approach.

Instead, we introduce new tools and concepts to the study of the parameterized complexity of problems on sparse graphs.

3.7 Rank based algorithms for bounded treewidth graphs

Marek Cygan (University of Warsaw, PL)

License  Creative Commons BY 3.0 Unported license
© Marek Cygan

Joint work of Marek Cygan, Hans Bodlaender, Jesper Nederlof, Stefan Kratsch

During the talk we will present a new approach to algorithms in bounded treewidth graph, which relates a DP computation with a rank computation of an appropriately defined matrix. Using this tool we obtain deterministic $c^{\text{tw}(G)}n^{O(1)}$ time algorithms in graphs of treewidth tw for problems such as Steiner Tree, Hamiltonicity or Feedback Vertex Set. Moreover we show how to solve weighted variants of those problems in the same running time and describe a different approach which allowed us to compute the number of Hamiltonian cycles in $c^{\text{tw}(G)}n^{O(1)}$ time.

3.8 Decomposing quantified conjunctive (or disjunctive) formulas

Victor Dalmau (Universitat Pompeu Fabra, ES)

License  Creative Commons BY 3.0 Unported license
© Victor Dalmau

Model checking-deciding if a logical sentence holds on a structure-is a basic computational task that is well-known to be intractable in general. For first-order logic on finite structures, it is PSPACE-complete, and the natural evaluation algorithm exhibits exponential dependence on the formula. We study model checking on the quantified conjunctive fragment of first-order logic, namely, prenex sentences having a purely conjunctive quantifier-free part. Following a number of works, we associate a graph to the quantifier-free part; each sentence then induces a prefixed graph, a quantifier prefix paired with a graph on its variables. We give a comprehensive classification of the sets of prefixed graphs on which model checking is tractable, based on a novel generalization of treewidth, that generalizes and places into a unified framework a number of existing results.

3.9 Branch Decompositions and Linear Matroids

Illya V. Hicks (Rice University, US)

License  Creative Commons BY 3.0 Unported license
© Illya V. Hicks

Joint work of Illya V. Hicks, Edray Goins, Jing Ma, Susan Margulies

This talk gives a general overview of practical computational methods for computing branch decompositions for linear matroids and their usage for solving integer programs. The concept of branch decompositions and its related invariant branch width were first introduced by Robertson and Seymour in their proof of the Graph Minors Theorem and can be easily generalized for any symmetric submodular set function.

3.10 A 5-approximation for treewidth using linear time, single exponential in the treewidth

Hans L. Bodlaender (Utrecht University, NL)

License  Creative Commons BY 3.0 Unported license
© Hans Bodlaender

Joint work of Pål Grønås Drange, Markus S. Dregi, Fedor Fomin, Daniel Lokshtanov and Michał Pilipczuk

We give an algorithm that, given a graph $G = (V, E)$ and an integer k , either finds a tree decomposition of G of width at most $5k + 4$, or decides that the treewidth of G is more than k . The algorithm uses $O(c^k \cdot n)$ time on graphs with n vertices. This is the first algorithm of the type that is both single exponential in the treewidth and linear in the number of vertices. Earlier algorithms either use quadratic time, or use $2^{\Omega(k \log k)}$ steps when time is measured as a function of k . The algorithm uses various techniques, including a data structure that allows several queries to be executed in $O(\log n)$ time, a table lookup technique for large values of n . As a consequence of our result, many problems allow algorithms whose time is linear in the number of vertices and single exponential in the treewidth; the result removes the need of a tree decomposition given as part of the input.

3.11 Contraction Decomposition: a new technique for H-minor-free graphs

Erik Demaine (MIT, US)

License  Creative Commons BY 3.0 Unported license
© Erik Demaine

Many problems are closed under contractions but not deletions, suggesting that we develop a Graph Contraction Theory to parallel Graph Minor Theory. Alas, some theorems like well-quasi-ordering do not hold in this setting. Nonetheless, we have been able to develop many contraction analogs to minor results. One powerful such result is contraction decomposition, which splits the edges of any graph into k pieces such that contracting any piece results in a graph of bounded treewidth. This approach has led to the best approximation algorithms for Traveling Salesman Problem on graphs.

3.12 Topological problems in tournaments

Michał Pilipczuk (University of Bergen, NO)

License  Creative Commons BY 3.0 Unported license
© Michał Pilipczuk

The containment theory for tournaments was developed recently by Chudnovsky, Fradkin, Kim, Scott, and Seymour. It appears that the natural containment notions in this setting form well-quasi-orderings, and correspond to two natural width measures, namely pathwidth and cutwidth. This creates possibilities for many algorithmic applications, including XP and FPT algorithms. During the talk, we will survey the status of algorithmic results on topological problems in tournaments, with particular focus on fixed-parameter tractability and similarities with bidimensionality

3.13 Kernelization using structural parameters on sparse graph classes

Felix Reidl (RWTH Aachen, DE)

License  Creative Commons BY 3.0 Unported license
© Felix Reidl

Meta-theorems for polynomial (linear) kernels have been the subject of intensive research in parameterized complexity. Heretofore, there were meta-theorems for linear kernels on graphs of bounded genus, H -minor-free graphs, and H -topological-minor-free graphs. To the best of our knowledge, there are no known meta-theorems for kernels for any of the larger sparse graph classes: graphs of bounded expansion, locally bounded expansion, and nowhere dense graphs. In this paper we prove meta-theorems for these three graph classes. More specifically, we show that graph problems that have finite integer index (FI) have linear kernels on graphs of bounded expansion when parameterized by the size of a modulator to constant-treewidth graphs. For graphs of locally bounded expansion, our result yields a quadratic kernel and for nowhere dense graphs, a polynomial kernel. While our parameter may seem rather strong, we show that a linear kernel result on graphs of bounded expansion with a weaker parameter will necessarily fail to include some of the problems included in our framework. Moreover, we only require problems to have FI on graphs of constant treewidth. This allows us to prove linear kernels for problems such as LONGEST PATH/CYCLE, EXACT (s, t) -PATH, TREewidth, and PATHwidth which do not have FI in general graphs.

3.14 A new proof for the weak-structure theorem with explicit constants

Paul Wollan (University of Rome “La Sapienza”, IT)

License  Creative Commons BY 3.0 Unported license
© Paul Wollan

The Weak Structure Theorem of Robertson and Seymour is the cornerstone of many of the algorithmic applications of graph minors techniques. The theorem states that any graph which has both large tree-width and excludes a fixed size clique minor contains a large, nearly planar subgraph. In this talk, we will discuss a new proof of this result which is significantly simpler than the original proof of Robertson and Seymour. As a testament to the simplicity of the proof, one can extract explicit constants to the bounds given in the theorem ensuring a linear relationship between the size of the grid minor and the size of the planar subgraph guaranteed by the theorem.

3.15 Approximability and fixed parameter algorithms: a new look

Rajesh Chitnis (University of Maryland, US)

License  Creative Commons BY 3.0 Unported license
© Rajesh Chitnis

Traditionally fixed-parameter algorithms (FPT) and approximation algorithms have been considered as different approaches for dealing with NP-hard problem. The area of fixed-parameter approximation algorithms tries to tackle problems which are intractable to both

these techniques. In this talk we will start with the formal definitions of fixed-parameter approximation algorithms and give a brief survey of known positive and negative results. Then (under standard conjectures in computational complexity) we show the first fixed-parameter inapproximability results for Clique and Set Cover, which are two of the most famous fixed-parameter intractable problems. On the positive side we obtain polynomial time $f(OPT)$ -approximation algorithms for a number of $W[1]$ -hard problems such as Minimum Edge Cover, Directed Steiner Forest, Directed Steiner Network, etc. Finally we give a natural problem which is $W[1]$ -hard, does not have a constant factor approximation in polynomial time, but admits a constant factor FPT-approximation.

3.16 Excluded vertex-minors for graphs of linear rank-width at most k

Sang-il Oum (KAIST – Daejeon, KR)

License  Creative Commons BY 3.0 Unported license
© Sang-il Oum

Linear rank-width is a graph width parameter, which is a variation of rank-width by restricting its tree to a caterpillar. As a corollary of known theorems, for each k , there is a finite set \mathcal{O}_k of graphs such that a graph G has linear rank-width at most k if and only if no vertex-minor of G is isomorphic to a graph in \mathcal{O}_k . However, no attempts have been made to bound the number of graphs in \mathcal{O}_k for $k \geq 2$. We construct, for each k , $2^{\Omega(3^k)}$ pairwise locally non-equivalent graphs that are excluded vertex-minors for graphs of linear rank-width at most k . Therefore the number of graphs in \mathcal{O}_k is at least double exponential.

3.17 Definability of numerical graph parameters and various notions of width

Johann A. Makowsky (Technion – Haifa, IL)

License  Creative Commons BY 3.0 Unported license
© Johann A. Makowsky

Graph parameters (and properties) definable in Monadic Second Order Logic (possibly with modular counting) are FPT computable for graphs of bounded width (where the width notion may vary depending on the logical presentation of the (hyper)-graph. It is therefore desirable to be able to show not only definability, but also non-definability. We present a method to show this for graph parameters which take values in a field which is even new for graph properties. For a graph parameter f and a binary operation on graphs \square the Hankel matrix $H(f, \square)$ is the infinite matrix where rows and columns are labeled by graphs G_i and the entry $H_{i,j}$ is given by $f(G_i, G_j)$. The methods are based in Hankel matrices (aka connection matrices) and the finite Rank Theorem (B. Godlin, T. Kotek and J.A. Makowsky, 2008) which states the MSOL-definable graph parameters have Hankel matrices of finite rank, provided \square behaves nicely with respect to the logic. We show that many examples of well studied graph parameters are not MSOL-definable even on ordered structures and with modular counting. A striking example is the chromatic number, for which non-definability was known before only using the complexity assumption ETH.

(Joint work with T. Kotek, published in the Proceedings of CSL'2012, LIPICS)

3.18 The k -disjoint paths problem in directed planar graphs

Dániel Marx (MTA – Budapest, HU)

License © Creative Commons BY 3.0 Unported license
© Dániel Marx

Joint work of Marek Cygan, Marcin Pilipczuk, and Michał Pilipczuk

Given a graph G and k pairs of vertices $(s_1, t_1), \dots, (s_k, t_k)$, the k -vertex-disjoint paths problem asks for pairwise vertex disjoint paths P_1, \dots, P_k such that P_i goes from s_i to t_i . Schrijver proved that the k -vertex-disjoint paths problem on planar directed graphs can be solved in time $O(n^k)$. We give an algorithm with running time $2^{2^{O(k^2)}} \cdot n^{O(1)}$ for the problem, that is, we show the fixed-parameter tractability of the problem.

3.19 Surface split decompositions: fast dynamic programming over branch decompositions for graphs of bounded genus

Paul Bonsma (University of Twente, NL)

License © Creative Commons BY 3.0 Unported license
© Paul Bonsma

Surface split decompositions are a special kind of branch decomposition, for graphs embedded on a surface of bounded genus. They are a direct generalization of sphere cut decompositions for planar graphs. Surface split decompositions have been introduced in [P. Bonsma, Surface split decompositions and subgraph isomorphism in graphs on surfaces, STACS 2012], where also a surprisingly simple method is introduced to obtain improved complexity bounds for various dynamic programming algorithms, provided that the given branch decomposition is a surface split decomposition. (In the aforementioned paper, this is only applied to the Subgraph Isomorphism problem.) In this talk, I will first give an introduction to surface split decompositions, and discuss their algorithmic applications in general. Next, I will introduce the following conjecture: Given a branch decomposition of a graph embedded in a surface of bounded genus, in polynomial time it is possible to construct a surface split decomposition where the width is increased by at most an (additive) constant. A proof of this conjecture, combined with known bidimensionality techniques, would make it possible to easily prove strong (e.g. subexponential) complexity bounds for many problems on graphs of bounded genus. This would give simplified proofs for various best known complexity bounds of this kind, and also enable new results in this direction.

3.20 Tools for multicoloring with applications to planar graphs and bounded treewidth graphs

Guy Kortsarz (Rutgers University-Camden, US)

License © Creative Commons BY 3.0 Unported license
© Guy Kortsarz

In this paper, we study scheduling jobs with conflicts on two fundamental classes of graphs: planar graphs and bounded tree width graph. The problem is represented by a graph in which a vertex is a job and has a work demand $p(v)$, meaning it takes $p(v)$ time units to finish

the job v . A solution for v is, $p(v)$ integers, the least of which is 1 or larger, that represent the $p(v)$ rounds in which v is active. And a general solution is an assignment for every v . The jobs (vertices) compete on resources and thus at every round only an independent set can be processed. The independent set has no cardinality bound meaning no bound on the number of processors (but the case of bounded number of processors can be handled easily). The colors of v can be non-preemptive $i, i + 1, \dots, j - 1, j$ so that $j - i + 1 = p(v)$, or arbitrary (preemptive) Perhaps our main contribution is designing very *general tool* for multicoloring graphs that are of independent interest. These results *should be* by other papers (at least if they know our paper). The max measure calls for minimizing the maximum number of every vertex, hence the makespan. The sum version required to minimize the sum of largest numbers assigned to a vertex. Hence the sum of completing times of jobs. For the preemptive makespan multicoloring of bounded tree width graphs we give a PTAS with quite unique properties. The coloring depends on the independent sets chosen and the $p(v)$ but never on the edges of G . Also there are always at most $O(\log n)$ preemptions. For the non-preemptive minimum makespan of bounded tree width graphs we get a FPTAS that applies to a large collection of functions (not only max and sum). For sum multicoloring (both preemptive and non preemptive) of Planar graphs and bounded tree width graphs we provide a PTAS. Our algorithms are quite non-trivial.

3.21 What makes normalized weighted satisfiability tractable?

Iyad A. Kanj (DePaul University – Chicago, US)

License © Creative Commons BY 3.0 Unported license
© Iyad A. Kanj

We consider the weighted monotone and antimonotone satisfiability problems on normalized circuits of depth at most $t \geq 2$, abbreviated $\text{WSAT}^+[t]$ and $\text{WSAT}^-[t]$, respectively. These problems model the weighted satisfiability of monotone and antimonotone propositional formulas (including weighted monotone/antimonotone CNF-SAT in a natural way, and serve as the canonical problems in the definition of the parameterized complexity hierarchy. In particular, $\text{WSAT}^+[t]$ ($t \geq 2$) is $W[t]$ -complete for even t and $W[t - 1]$ -complete for odd t , and $\text{WSAT}^-[t]$ ($t \geq 2$) is $W[t]$ -complete for odd t and $W[t - 1]$ -complete for even t . Moreover, several well-studied problems, including important graph problems, can be modeled as $\text{WSAT}^+[t]$ and $\text{WSAT}^-[t]$ problems in a straightforward manner. We characterize the parameterized complexity of $\text{WSAT}^+[t]$ and $\text{WSAT}^-[t]$ with respect to the genus of the circuit. For $\text{WSAT}^-[t]$, we give a precise characterization: $\text{WSAT}^-[t]$ is fixed-parameter tractable (FPT) on circuits whose genus is $n^{o(1)}$, where n is the number of the variables in the circuit, and it has the same W -hardness as the general $\text{WSAT}^-[t]$ problem (i.e., with no restriction on the genus) on circuits whose genus is $n^{\Omega(1)}$. For $\text{WSAT}^+[2]$ (i.e., weighted monotone cnf-sat and $\text{WSAT}^+[2]$, which are both $W[2]$ -complete, the characterization is also precise: $\text{WSAT}^+[2]$, and $\text{WSAT}^+[3]$, are FPT if the genus is $n^{o(1)}$ and $W[2]$ -complete if the genus is $n^{\Omega(1)}$. For $\text{WSAT}^+[t]$ where $t > 3$, we show that it is FPT if the genus is $O(\sqrt{\log n})$, and that it has the same W -hardness as the general $\text{WSAT}^+[t]$ problem if the genus is $n^{\Omega(1)}$. The above results give, via standard parameterized reductions, precise characterizations of the parameterized complexity of several problems with respect to the genus of the underlying graph, as shown in the current paper.

3.22 An excluded grid theorem for digraphs with forbidden minors

Stephan Kreutzer (TU Berlin, DE)

License  Creative Commons BY 3.0 Unported license
© Stephan Kreutzer

One of the fundamental results in graph structure theory is the excluded grid theorem, proved by Robertson and Seymour, which states that every graph of sufficiently high tree-width contains a large grid as a minor. This theorem provides the structural foundation for algorithmic techniques such as bidimensionality and is also the basis for more advanced structure theorems in the graph minor series. In 1997, Reed and later, in 2001, Johnson, Robertson, Seymour and Thomas conjectured an excluded grid theorem for directed graphs, i.e. the existence of a function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that every digraph of directed tree-width at least $f(k)$ contains a directed grid of order k . Johnson et al. proved the conjecture for planar digraphs in 2001 but for all other cases the conjecture remained open. In this talk we present a proof of the conjecture for the case of digraphs excluding a fixed undirected graph as a minor. We present the main proof ideas and give examples of possible algorithmic applications.

This is joint work with Ken-ichi Kawarabayashi.

3.23 Exclusion theorems for immersions on surface embedded graphs

Archontia Giannopoulou (University of Bergen, NO)

License  Creative Commons BY 3.0 Unported license
© Archontia Giannopoulou

In this talk, we provide a structural characterization of graphs that forbid a grid as an immersion and can be embedded in a surface of Euler genus g . In particular, we prove that a graph G that excludes some grid as an immersion and is embedded in a surface of Euler genus g has either “small” treewidth (bounded by a function of H and g) or “small” edge connectivity (bounded by 3). By generalizing these techniques we also provide a structural characterization for the case where the excluded graph is any graph H .

3.24 Packing edge-disjoint odd S-cycles in 4-edge-connected graphs

Yusuke Kobayashi (University of Tokyo, JP)

License  Creative Commons BY 3.0 Unported license
© Yusuke Kobayashi

We develop theory for 4-edge-connected graphs. It has been known that for a 4-edge-connected graph G ,

1. there is a much simpler polynomial-time algorithm for the k edge-disjoint paths problem for fixed k (compared to a graph minor algorithm), and
2. the Erdős-Pósa property holds for edge-disjoint odd cycles, and there is a simple polynomial-time algorithm to test whether or not G has k edge-disjoint odd cycles for fixed k .

Note that the Erdős-Pósa property does not hold for edge-disjoint odd cycles in general. In this paper, we generalize the above results as follows:

1. there is a simple polynomial-time algorithm for the PARITY k EDGE-DISJOINT PATHS problem for fixed k (i.e., the length of each path is of a specified parity), and
2. the Erdős-Pósa property holds for edge-disjoint odd S -cycles (i.e., odd cycles through a vertex in a specified vertex set S), and there is a simple polynomial-time algorithm to test whether or not G has k edge-disjoint odd S -cycles for fixed k .

This is joint work with Naonori Kakimura and Ken-ichi Kawarabayashi.

3.25 Tree-width of hypergraphs and surface duality

Frédéric Mazoit (Université Bordeaux, FR)

License  Creative Commons BY 3.0 Unported license
© Frédéric Mazoit

In Graph minors III, Robertson and Seymour write: “It seems that the tree-width of a planar graph and the tree-width of its geometric dual are approximately equal? indeed, we have convinced ourselves that they differ by at most one”. They never gave a proof of this. We recently published a generalisation of this statement to embedding of hypergraphs on general surfaces, and we prove that our bound is tight. Although the result is purely graph theoretical, the proof uses a “surface-cut” decomposition which may find algorithmic applications for graphs on surfaces.

4 Open problems

We give a list of the problems presented on Monday, March 18, 2013 and Thursday, March 21, 2013 at the open-problem session of the Seminar on *Bidimensional Structures: Algorithms, Combinatorics and Logic*, held at Schloss Dagstuhl in Wadern, Germany.

Canonical tree-decompositions, k -blocks and tangles

Reinhard Diestel, University of Hamburg, Hamburg, R.Diestel@math.uni-hamburg.de

A k -block in a graph G is a set X of at least k vertices no two of which are separated in G by fewer than k other vertices (which may or may not be in X). Such k -blocks might be of interest in applications; in a computer network, for example, they might specify the nodes at which one would place those servers that should be able to communicate with each other even when some $\ell < k$ other nodes have failed. We thus pose the following algorithmic problem:

Algorithmic Problem 1. Given an integer k and a graph G , find all the k -blocks in G . A simple ad-hoc algorithm can do this in $O(kn^4)$ time [1], but maybe this can be improved.

An important feature of k -blocks is that their connectivity need not lie in the subgraph they induce, but will more typically be provided by the ambient graph. Thus, while k -connected subgraphs are obvious candidates for k -blocks, other k -blocks might induce no edge at all. See [1] for examples of quite different types.

Although k -blocks were already considered by Mader [5], they came to the fore more recently as a key ingredient of a solution offered in [3] to an old problem in graph connectivity:

the problem of how to decompose a $(k - 1)$ -connected graph into its ‘ k -connected pieces’ in some structured way. The main results of [3] can be summarized as follows.

Theorem 1. For every integer k , every graph has a canonical tree-decomposition that distinguishes all its k -blocks efficiently.

(A tree-decomposition is *canonical* if it is invariant under the automorphisms of the graph. It *distinguishes* two k -blocks X_1, X_2 if it has an adhesion set S of order $< k$ that separates them, and it does so *efficiently* if S is no larger than the smallest X_1 – X_2 separator in the graph (which will always have order $< k$). Thus, every k -block lies in some part of the tree-decomposition, and distinct k -blocks lie in different parts.)

Algorithmic Problem 2. Given k and a graph G , find in G a tree-decomposition such as in Theorem 1.

We believe we can do this in $O(k^3 n^4)$ time, but have not pursued the matter.

One may wonder whether every graph even has one unified tree-decomposition that distinguishes all its k -blocks, for all k simultaneously. This is not true in this generality, but it is almost true:

Theorem 2. Every graph has a canonical tree-decomposition that distinguishes all its robust blocks efficiently.

(*Robust k -blocks* are defined in [3], and the definition is technical. But most k -blocks are robust; for example, all k -blocks of size at least $\frac{3}{2}k$ are. Loosely speaking, non-robust blocks are a rare technical phenomenon that can occur but usually does not. A *block* is a k -block for some k .)

Algorithmic Problem 3. Find a tree-decomposition such as in Theorem 2 for a given input graph.

An analysis of the proofs of Theorems 1 and 2 due to Hundertmark [4] shows that the only information about k -blocks we really use is how they relate to the $(< k)$ -separations of the given graph G . Every k -block *orients* every $(< k)$ -separation of G towards the side that contains it. For each k -block, these orientations are consistent in two ways. Given two nested separations (A, B) and (C, D) such that $B \supseteq D$ and (C, D) is oriented towards D , then (A, B) is oriented towards B . Given two crossing separations (A, B) and (C, D) such that (A, B) is oriented towards B and (C, D) is oriented towards D , the ‘corner separation’ $(A \cup C, B \cap D)$ will be oriented towards $B \cap D$ if it is oriented at all, i.e., if it has order $< k$. (If G is $(k - 1)$ -connected, which is an important special case when we consider k -blocks, that will in fact be so.)

Call a set of orientations, one of each $(< k)$ -separation of G , a *k -profile* if it satisfies these two consistency requirements. Thus, every k -block defines a k -profile. Hundertmark [1] showed that we can adapt the proofs of Theorems 1 and 2 so as to establish the existence of canonical tree-decompositions that distinguish all the k -profiles for any given k (Theorem 1), or all the ‘robust’ k -profiles for all k simultaneously. (A tree-decomposition *distinguishes* two k -profiles if it induces a separation of order $< k$ – one corresponding to an edge of the decomposition tree – that is oriented differently by the two profiles.)

The advantage of this more abstract approach is that there are k -profiles that do not come from k -blocks; so Theorems 1 and 2 become more comprehensive. For example, it is easy to see that every tangle of order k is a k -profile, and so we have canonical tree-decompositions that also distinguish all the maximal tangles in a graph (as well as their robust blocks). This strengthens a theorem of Robertson and Seymour [6], who proved the existence of such a

tree-decomposition which, however, is not canonical. (In order to select the required nested subset of ($< k$)-separations from the set of all these separations, they need a ‘tie-breaker’ that depends on a chosen vertex enumeration.)

All we need in order to define a k -profile on a set is to have a notion of separations of this set, and a notion of the order of such separations. These things are also given for matroids, and hence all our results generalize to matroids too [4].

Profiles not only generalize tangles, they are also special cases of preferences [4], which in turn are the prime examples of brambles. They are the weakest known way of consistently orienting all the ($< k$)-separations of a graph that still allows some tree-decomposition (canonical or not) to distinguish them all: there is in general no tree-decomposition that does this for preferences or brambles, since a graph of order n may have more than n preferences and brambles (but any minimal tree-decomposition separating them would have at most n parts).

References

- 1 J. Carmesin, R. Diestel, M. Hamann and F. Hundertmark, Forcing k -blocks in graphs by minimum and average degree conditions, preprint 2013.
- 2 J. Carmesin, R. Diestel, M. Hamann and F. Hundertmark, Canonical tree-decompositions of finite graphs, preprint 2013.
- 3 J. Carmesin, R. Diestel, F. Hundertmark and M. Stein, Connectivity and tree structure in finite graphs, *Combinatorica* (to appear) and arXiv:1105.1611.
- 4 F. Hundertmark, Profiles, preprint 2013.
- 5 W. Mader. Über n -fach zusammenhängende Eckenmengen in Graphen. *J. Combin. Theory (Series B)* **25** (1978), 74–93.
- 6 N. Robertson and P.D. Seymour, Graph minors. X. Obstructions to tree-decomposition. *J. Combin. Theory (Series B)* **52** (1991), 153–190.

Hyperbolicity

Blair D. Sullivan, Oak Ridge National Lab, sullivanb@ornl.gov

An unweighted undirected graph G is δ -hyperbolic if, for any four vertices u, v, w, x ordered so that $d(u, v) + d(w, x) \geq d(u, w) + d(v, x) \geq d(u, x) + d(v, w)$, we have $[d(u, v) + d(w, x)] - [d(u, w) + d(v, x)] \leq \delta$. Define the *hyperbolicity* $\delta(G)$ to be the minimum such δ . Define $\nu(G)$ to be the maximum hyperbolicity over all metric cycles of G . (A cycle in G is *metric* if the distance between any two of its vertices is the same in G and in the cycle. So $\nu(G) \approx |V(G)|/4$.)

1. Is $\delta(G) < \text{tw}(G) + \nu(G) - 1$? Or more generally, is $\delta = O(\text{tw}(G) + \nu(G) + 1)$?
2. Can treewidth and treelength be simultaneously approximated when the hyperbolicity δ is small? That is, is there a tree decomposition whose width is within a constant factor of the treewidth, and whose length is within a constant factor of the treelength? The *length* of a tree decomposition is the largest diameter of a bag (whereas width measures the maximum size of a bag). With large hyperbolicity, no simultaneous approximation is possible [1].

References

- 1 Yon Dourisboure and Cyril Gavoille. Tree-decompositions with bags of small diameter. *Discrete Mathematics* 307:2008–2029, 2007. <http://dept-info.labri.fr/~gavoille/article/DG07.pdf>

Spanning Trees Respecting Faces

Illya V. Hicks, Rice University, ivhicks@rice.edu

Given an embedded planar graph, what is the complexity of finding a spanning tree T that minimizes $\max\{d_T(u, v) : u, v \text{ on a common face}\}$? If it is easier, you can assume that the graph is bipartite.

Blair Sullivan asks whether there is a polynomial-time $O(1)$ -approximation.

The problem is motivated by a possible relation to computing planar branchwidth.

Odd Immersion Hadwiger Conjecture

Bojan Mohar, Simon Fraser University, mohar@sfu.ca

A graph H *immerses* in G if there are distinct vertices v_1, v_2, \dots, v_h in G corresponding to the h vertices $\{1, 2, \dots, h\}$ in H , and there are edge-disjoint paths p_{ij} joining v_i to v_j in G for every edge (i, j) in H . Such an immersion is *odd* if all the paths p_{ij} have odd length.

Conjecture: If G has chromatic number at least k , then K_k oddly immerses into G .

This conjecture is a stronger (odd) form of the immersion Hadwiger conjecture [1, 2], which in turn is a variation of the classic Hadwiger's Conjecture: if G has chromatic number at least k , then K_k is a minor of G . The immersion Hadwiger conjecture, for example, is proved for $k \leq 7$ [3].

References

- 1 Faisal N. Abu-Khzam and Michael A. Langston. Graph coloring and the immersion order. In *Proceedings of the International Conference on Computing and Combinatorics*, 2003, pages 394–403.
- 2 Faisal N. Abu-Khzam and Michael A. Langston. Immersion containment and connectivity in color-critical graphs. *Discrete Mathematics and Theoretical Computer Science*, 14(2):155–164, 2012.
- 3 Matt Devos, Ken-ichi Kawarabayashi, Bojan Mohar, and Haruko Okamura. Immersing small complete graphs. *ARS Mathematica Contemporena*, 3:139–146, 2010.

Grid Minors

MohammadTaghi Hajiaghayi, University of Maryland, College Park, hajiaghayi@gmail.com

What is the smallest integer function f such that, if G is a graph of treewidth $f(r)$, then it has an $r \times r$ grid minor? In particular, is $f(r)$ polynomial?

Conjecture: $f(r) = O(r^3)$ [2].

Robertson, Seymour, and Thomas [4] proved that $f(r) \leq 2^{O(r^5)}$ and $f(r) = \Omega(r^2 \log r)$, and conjectured that the latter bound may be closer to the truth. Kawarabayashi and Kobayashi [3] improved the upper bound to $f(r) \leq 2^{O(r^2 \log r)}$. For H -minor-free graphs, $f(r) = O(r)$ [1]. For map graphs, $f(r) = O(r^3)$ [2]. Constant powers of such graphs also have polynomial bounds on $f(r)$ [2].

References

- 1 Erik D. Demaine and MohammadTaghi Hajiaghayi. Linearity of grid minors in treewidth with applications through bidimensionality. *Combinatorica* 28(1):19–36, January 2008.
- 2 Erik D. Demaine, MohammadTaghi Hajiaghayi, and Ken-ichi Kawarabayashi. Algorithmic graph minor theory: improved grid minor bounds and Wagner's contraction. *Algorithmica* 54(2):142–180, June 2009.

- 3 Ken-ichi Kawarabayashi and Yusuke Kobayashi. Linear min-max relation between the treewidth of H -minor-free graphs and its largest grid. In *Proceedings of the 29th International Symposium on Theoretical Aspects of Computer Science*, 2012, pages 278–289. <http://dx.doi.org/10.4230/LIPIcs.STACS.2012.278>
- 4 Neil Robertson, Paul Seymour, Robin Thomas. Quickly excluding a planar graph. *Journal of Combinatorial Theory, Series B* 62(2):323–348, November 1994.

Subexponential FPT Algorithms on Planar Graphs

Dániel Marx, MTA, Budapest, dmarx@cs.bme.hu

Is there a fixed-parameter algorithm with running time $2^{\tilde{O}(\sqrt{k})}n^{O(1)}$ for the following problems?

1. k disjoint paths in directed planar graphs.
2. Steiner tree in planar graphs, parameterized by the size of the tree, or even by the number of terminals.
3. Subgraph isomorphism in planar graphs, parameterized by the size of the small graph.
4. Weighted versions of e.g. independent set or vertex cover.
5. Tree spanner: find a spanning tree of distortion at most k , parameterized by k . [Posed by Fedor Fomin.]

Disjoint Paths in Planar Graphs

Daniel Lokshtanov, University of Bergen, daniello@ii.uib.no

For the k disjoint paths problem in planar graphs, is there a fixed-parameter algorithm with running time $2^{k^{O(1)}}n^{O(1)}$?

We have a $k^{2^{O(k)}}n^{O(1)}$ -time algorithm [1].

References

- 1 Isolde Adler, Stavros G. Kolliopoulos, Philipp Klaus Krause, Daniel Lokshtanov, Saket Saurabh, and Dimitrios M. Thilikos. Tight bounds for linkages in planar graphs. In *Proceedings of the 38th International Colloquium on Automata, Languages and Programming*, LNCS 6755, Zürich, Switzerland, July 2011, pages 110–121.

Spaghetti Treewidth

Hans Bodlaender, Utrecht University, n.l.bodlaender@uu.nl

A tree decomposition is *spaghetti* if every vertex appears in a path of bags (instead of a general subtree). *Spaghetti treewidth* is the minimum width of a spaghetti tree decomposition.

Is the class of graphs of spaghetti treewidth 2 closed under minors? If so, can we find the excluded minors?

A related, previously studied notion is “special treewidth” [1]. A rooted tree decomposition is *special* if every vertex appears in a rooted path of bags. The resulting parameter *special treewidth* falls somewhere between treewidth and pathwidth.

References

- 1 Bruno Courcelle. On the model-checking of monadic second-order formulas with edge set quantifications. *Discrete Applied Mathematics* 160(6):866–887, 2012.

Bounded Rank

Hans Bodlaender, Utrecht University, n.l.bodlaender@uu.nl

The *bounded rank* approach leads to several singly exponential dynamic-programming algorithms [1].

Extend the applicability of this approach to the widest possible variety of problems.

References

- 1 Hans L. Bodlaender, Marek Cygan, Stefan Kratsch, Jesper Nederlof. Solving weighted and counting variants of connectivity problems parameterized by treewidth deterministically in single exponential time. arXiv:1211.1505.

Fast Treewidth

Hans Bodlaender, Utrecht University, n.l.bodlaender@uu.nl

Is there a fixed-parameter algorithm with running time $2^{o(k^3)}n$ for computing the treewidth k in a general graph? Or is there some sort of lower bound? (There is an algorithm with running time $2^{O(k^3)}n$ [1].)

Also recall the famous open problem: is treewidth NP-hard for planar graphs? Only a 1.5-approximation is known [2].

References

- 1 Hans L. Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM Journal on Computing* 25(6):1305–1317.
- 2 P. D. Seymour and R. Thomas. Call routing and the ratcatcher. *Combinatorica* 14(2):217–241, 1994.

Disjoint Paths in Tournaments

Michal Pilipczuk, University of Bergen, michal.pilipczuk@ii.uib.no

What is the parameterized complexity of the k disjoint paths problem in tournaments? In particular, is it fixed-parameter tractable? Chudnovsky, Scott, and Seymour [1] proved that it is in XP (there is an algorithm of running time $n^{f(k)}$).

The k disjoint paths problem in general directed graphs is known to be NP-hard even for $k = 2$ [2]. For directed acyclic graphs, it is known to be in XP but $W[1]$ -hard [3].

References

- 1 Maria Chudnovsky, Alex Scott, and Paul Seymour. Disjoint paths in tournaments. Submitted manuscript, 2012. <http://www.columbia.edu/~mc2775/tournpaths.pdf>
- 2 S. Fortune, J. Hopcroft, and J. Wyllie. The directed subgraphs homeomorphism problem. *Theoretical Computer Science* 10:111–121, 1980.
- 3 Aleksandrs Slivkins. Parameterized tractability of edge-disjoint paths on directed acyclic graphs. *SIAM Journal on Discrete Mathematics* 24(1):146–157, 2010.

Grid Theorem for Immersions

Paul Wollan, University of Rome “La Sapienza”, paul.wollan@gmail.com

Wollan [1] proved that, if G is 3-connected and does not contain a “large” wall as an immersion, then it should have large degree and large treewidth.

Is the class of graphs with small treewidth and small degree algorithmically interesting?

Are there parameterized problems that, while hard in general, are fixed-parameter tractable when restricted to this class?

Are there problems that, when parameterized with respect to treewidth or maximum degree are $W[1]$ -hard, while they admit a fixed-parameter algorithm when both maximum degree and treewidth are bounded?

Is it possible to develop a bidimensionality theory based on immersions instead of minors?

References

- 1 Paul Wollan. The structure of graphs not admitting a fixed immersion. arXiv:1302.3867.

Local Search Variant of Vertex Cover

Fedor Fomin, University of Bergen, fomin@ii.uib.no

In the k -LOCAL SEARCH FOR VERTEX COVER, we are given a graph G and a vertex cover S of G , the question is if k -exchange neighborhood of S contains a smaller vertex cover. In other words, is it possible to remove k vertices from S and add at most $k - 1$ such that the resulting set is still a vertex cover? On general graphs the problem is known to be $W[1]$ -hard parameterized by k . However, for planar graphs, it admits a singly exponential parameterized algorithm, i.e., of running time $2^{O(k)}n^{O(1)}$ [1]. Is there a subexponential parameterized algorithm for this problem, i.e., of running time $2^{o(k)}n^{O(1)}$? Similar questions can be asked for local search variants of Dominating Set and Odd Cycle Transversal. The other question is if the “local search” variants of other “connectivity” problems such as PLANAR TSP, LONGEST PATH, and PLANAR FEEDBACK VERTEX SET are FPT on planar graphs? See [2, 3] for $W[1]$ -hardness of TSP on non-planar graphs.

References

- 1 Michael R. Fellows, Fedor V. Fomin, Daniel Lokshtanov, Frances A. Rosamond, Saket Saurabh, and Yngve Villanger. Local search: Is brute-force avoidable? *Journal of Computer and System Sciences* 78(3):707–719, 2012.
- 2 Jiong Guo, Sepp Hartung, Rolf Niedermeier, Ondrej Suchy, The Parameterized Complexity of Local Search for TSP, More Refined. ISAAC 2011: 614-623
- 3 Daniel Marx: Searching the k -change neighborhood for TSP is $W[1]$ -hard. *Oper. Res. Lett.* 36(1): 31-36 (2008)

Upper and Lower Bounds for Algorithms on Parameterized Problems

Marek Cygan, Uniwersytetu Warszawskiego, cygan@mimuw.edu.pl

Problem 1: Is there a $3^k n^{O(1)}$ algorithm for FEEDBACK VERTEX SET?

Problem 2: It is known that there is a $c^{\text{tw}(G)} n^{O(1)}$ algorithm for HAMILTONIAN CYCLE. Let c be the smallest (if exists) constant for which such an algorithm exists. Current lower and upper bounds establish that $2 + \sqrt{2} \leq c \leq 4$ (the lower bound is up to the Strong Exponential Time Hypothesis). Find better estimations of the constant c .

Problem 3: Can we count the number of vertex feedback sets of size at most k of a graph in $c^{\text{tw}(G)} n^{O(1)}$ time?

Problem 4: Recent advances [1, 2] on dynamic programming on tree width reveal that problems like HAMILTONIAN CYCLE, STEINER TREE, and FEEDBACK VERTEX SET can be solved by singly exponential randomized parameterized algorithms when parameterized by the treewidth of the input graph. Specifically, for each of these problems Π , there is a constant c_Π such that Π can be solved by a randomized algorithm in $c_\Pi^{\text{tw}(G)} \cdot n$ time. For the same parameterized problems, there are also deterministic singly exponential algorithms, but the constants are worse. Is it possible to improve the deterministic constant to match the randomized constant?

References

- 1 Hans L. Bodlaender, Marek Cygan, Stefan Kratsch, and Jesper Nederlof. Solving weighted and counting variants of connectivity problems parameterized by treewidth deterministically in single exponential time. arXiv:1211.1505.
- 2 Marek Cygan, Jesper Nederlof, Marcin Pilipczuk, Michał Pilipczuk, Johan van Rooij, and Jakub Onufry Wojtaszczyk. Solving connectivity problems parameterized by treewidth in single exponential time. arXiv:1103.0534.

Planar Completion to Bounded Diameter

Dimitrios M. Thilikos, National & Kapodistrian U. Athens, sedthilk@thilikos.info

Given a plane graph G and a nonnegative integer k , is it possible to add at most k edges such that the resulting graph remains plane and has diameter at most k ?

If (G, k) is a YES instance of the above problem, and G' is a minor of G , then (G', k) is also a YES instance of the same problem. From the meta-algorithmic consequence of the Graph Minors series, it follows that the above problem is fixed-parameter tractable, i.e., there exists an algorithm with running time $f(k)n^{O(1)}$ (see e.g. [1]). However, so far no such algorithm has been constructed.

A possible approach is to use the fact that, if (G, k) is a YES instance, then G has treewidth bounded by some function of k . Then a dynamic programming algorithm for PLANAR COMPLETION TO BOUNDED DIAMETER on graphs of bound treewidth would result in the construction of the desired algorithm.

More generally, how can we design bounded-treewidth dynamic-programming algorithms for planar completion problems?

References

- 1 Michael R. Fellows and Michael A. Langston. Nonconstructive tools for proving polynomial-time decidability. *Journal of the ACM* 35(3):727–739, 1988.

5 Program

Monday, March 18

- [09:00–09:15] A welcome to Dagstuhl
 [09:15–09:30] **Dimitrios M. Thilikos:** *A welcome to treewidth*
 [09:30–09:45] **Saket Saurabh:** *Bidimensionality and its applications (part I)*
 [09:45–11:00] **Daniel Lokshtanov:** *Bidimensionality and its applications (part II)*
 [11:00–11:30] Break
 [11:30–12:00] **Gwenaël Joret:** *Treewidth and dimension*
 [12:15] Lunch
 [15:00–15:40] **Reinhard Diestel:** *Canonical tree-decompositions, k-blocks and tangles*
 [15:45–16:30] Cake
 [16:30–17:00] **Erik Jan van Leeuwen:** *Subexponential-time parameterized algorithm for Steiner Tree on planar graphs*
 [17:00–18:00] Open Problem Session

Tuesday, March 19

- [09:00–10:00] **Marek Cygan:** *Rank based algorithms for bounded treewidth graphs*
 [10:00–10:30] Break
 [10:30–11:00] **Victor Dalmau** *Decomposing quantified conjunctive (or disjunctive) formulas*
 [11:00–11:30] **Illya V. Hicks** *Branch Decompositions and Linear Matroids*
 [11:30–12:00] **Hans L. Bodlaender** *A 5-approximation for treewidth using linear time, single exponential in the treewidth*
 [12:00] **Photo** *in front of the Chapel*
 [12:15] Lunch
 [15:00–15:45] **Erik Demaine:** *Contraction Decomposition: a new technique for H-minor-free graphs*
 [15:45–16:30] Cake
 [16:30–17:00] **Michał Pilipczuk:** *Topological problems in tournaments*
 [17:00–17:30] **Felix Reidl** *Kernelization using structural parameters on sparse graph classes*

Wednesday, March 20

- [09:00–10:00] **Paul Wollan:** *A new proof for the weak-structure theorem with explicit constants*
 [10:00–10:30] Break
 [10:30–11:00] **Rajesh Chitnis** *Approximability and fixed parameter algorithms: a new look*
 [11:00–11:30] **Sang-Il Oum:** *Excluded vertex-minors for graphs of linear rank-width at most k*
 [11:30–12:00] **Janos Makowski:** *Definability of numerical graph parameters and various notions of width*
 [12:15] Lunch

Thursday, March 21

- [09:00–10:00] **Dániel Marx:** *The k -disjoint paths problem in directed planar graphs*
[10:00–10:30] Break
[10:30–11:00] **Paul Bonsma:** *Surface split decompositions: fast dynamic programming over branch decompositions for graphs of bounded genus*
[11:00–11:30] **Guy Kortsarz:** *Tools for multicoloring with applications to planar graphs and bounded treewidth graphs*
[11:30–12:00] **Iyad A. Kanj:** *What makes normalized weighted satisfiability tractable?*
[12:15] Lunch
[15:30–17:00] Cake
[17:00–18:00] Open Problem Session

Friday, March 22

- [09:00–10:00] **Stephan Kreutzer:** *An excluded grid theorem for digraphs with forbidden minors*
[10:00–10:30] Break
[10:30–11:00] **Archontia C. Giannopoulou:** *Exclusion theorems for immersions on surface embedded graphs*
[11:00–11:30] **Yusuke Kobayashi:** *Packing edge-disjoint odd S -cycles in 4-edge-connected graphs*
[11:30–12:00] **Frédéric Mazoit:** *Tree-width of hypergraphs and surface duality*
[12:15] Lunch

Participants

- Isolde Adler
Goethe-Universität Frankfurt am Main, DE
- Hans L. Bodlaender
Utrecht University, NL
- Paul Bonsma
University of Twente, NL
- Jianer Chen
CSU-Changsha, CN
- Rajesh Chitnis
University of Maryland, US
- Marek Cygan
University of Warsaw, PL
- Victor Dalmau
Univ. Pompeu Fabra – Barcelona, ES
- Anuj Dawar
University of Cambridge, GB
- Erik D. Demaine
MIT, US
- Reinhard Diestel
Universität Hamburg, DE
- Frederic Dorn
SINTEF – Trondheim, NO
- Fedor V. Fomin
University of Bergen, NO
- Archontia Giannopoulou
University of Bergen, NO
- Petr A. Golovach
University of Bergen, NO
- Alexander Grigoriev
Maastricht University, NL
- MohammadTaghi Hajiaghayi
University of Maryland – College Park, US
- Illya V. Hicks
Rice University – Houston, US
- Gwenaël Joret
Université Libre de Bruxelles, BE
- Marcin Kaminski
Université Libre de Bruxelles, BE
- Iyad A. Kanj
DePaul University – Chicago, US
- Mamadou Moustapha Kanté
Blaise Pascal Univ. – Aubiere, FR
- Yusuke Kobayashi
University of Tokyo, JP
- Sudeshna Kolay
The Institute of Mathematical Sciences – Chennai, IN
- Guy Kortsarz
Rutgers University-Camden, US
- Stephan Kreutzer
TU Berlin, DE
- O-joung Kwon
KAIST – Daejeon, KR
- Daniel Lokshtanov
University of Bergen, NO
- Johann A. Makowsky
Technion – Haifa, IL
- Dániel Marx
MTA – Budapest, HU
- Frédéric Mazoit
Université Bordeaux, FR
- Bojan Mohar
Simon Fraser University – Burnaby, CA
- Sang-il Oum
KAIST – Daejeon, KR
- Geevarghese Philip
MPI für Informatik – Saarbrücken, DE
- Marcin Pilipczuk
University of Warsaw, PL
- Michał Pilipczuk
University of Bergen, NO
- Felix Reidl
RWTH Aachen, DE
- Peter Rossmanith
RWTH Aachen, DE
- Ignasi Sau Valls
CNRS, Univ. Montpellier II, FR
- Saket Saurabh
The Institute of Mathematical Sciences – Chennai, IN
- Konstantinos Stavropoulos
HU Berlin, DE
- Blair D. Sullivan
Oak Ridge National Lab., US
- Hisao Tamaki
Meiji University – Kawasaki, JP
- Jan Arne Telle
University of Bergen, NO
- Dimitrios M. Thilikos
National and Kapodistrian University of Athens, GR
- Ioan Todinca
Université d'Orleans, FR
- Erik Jan van Leeuwen
MPI für Informatik – Saarbrücken, DE
- Yngve Villanger
University of Bergen, NO
- Paul Wollan
University of Rome “La Sapienza”, IT
- Christian Wulff-Nilsen
University of Copenhagen, DK



Future Internet

Edited by

Jon Crowcroft¹, Markus Fidler², Klara Nahrstedt³, and
Ralf Steinmetz⁴

1 University of Cambridge, UK, Jon.Crowcroft@cl.cam.ac.uk

2 Leibniz Universität Hannover, DE, markus.fidler@ikt.uni-hannover.de

3 University of Illinois – Urbana, US, klara@illinois.edu

4 TU Darmstadt, DE, ralf.steinmetz@kom.tu-darmstadt.de

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 13131 “Future Internet”. At the seminar, about 40 invited researchers from academia and industry discussed the promises, approaches, and open challenges of the Future Internet. This report gives a general overview of the presentations and outcomes of discussions of the seminar.

Seminar 24.–27. March, 2013 – www.dagstuhl.de/13131

1998 ACM Subject Classification C.2.1 Network Architecture and Design – Packet switching networks, C.2.6 Internetworking – Standards

Keywords and phrases Future Internet, Network Design, Testbed, Software Defined Networking, Virtualization, OpenFlow

Digital Object Identifier 10.4230/DagRep.3.3.75

1 Executive Summary

Jon Crowcroft

Markus Fidler

Klara Nahrstedt

Ralf Steinmetz

License © Creative Commons BY 3.0 Unported license
© Jon Crowcroft, Markus Fidler, Klara Nahrstedt, and Ralf Steinmetz

The recent vision of the “Future Internet” attracts significant networking research and causes controversial debates on the actions to be taken. Clean-slate initiatives envision a fresh start that put fundamental principles of networking into question. Avoiding any constraints of the current Internet implementation, the ambition of the clean-slate approach is to understand and design the ‘right’ network architecture. Evolutionary approaches, on the other hand, seek incremental improvements, assuming that the Internet can –as in the past– be fixed to accommodate the changing needs of users and applications.

Numerous initiatives on the Future Internet, like FIND, GENI funded by the NSF, FIRE, 4WARD by the EU, and G-LAB by the BMBF, reflect the importance of the topic. Characteristic for numerous Future Internet initiatives is an experimental approach using testbed facilities such as the GENI or the G-Lab platform.

Challenges that are of central importance for the Future Internet fall into the following categories:



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

Future Internet, *Dagstuhl Reports*, Vol. 2, Issue 3, pp. 75–99

Editors: Jon Crowcroft, Markus Fidler, Klara Nahrstedt, and Ralf Steinmetz



Dagstuhl Reports

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

Network design: computer networks and the Internet obey certain architectural guidelines that reflect experience gained in the art of network design, such as layered reference models or the Internet end-to-end argument. While these principles are backed up by the success of the Internet, it has to be noted that the network exhibits major architectural restrictions, e.g., regarding mobility, security, and quality of service. Computer networking as a relatively young field of research can benefit significantly from architectural reconsiderations that are initiated by clean-slate initiatives. While today, network theory is largely descriptive, this Dagstuhl seminar investigated the potentialities of a prescriptive network theory, which could justify a methodical rule/equation-based approach for the design of future networks.

Virtualization: the virtualization paradigm revolutionized the use of computers and data centers, where the flexibility and mobility of virtual machines offers tremendous potential, posing, however, significant new challenges for networking. On the other hand, the virtualization paradigm has already many applications in networking, e.g., in virtual private networks or overlay networks. Currently, virtualization finds its way into network components, e.g., routers, itself, where the decoupling of logical entities from the physical substrate enables major innovations, e.g., concurrent (possibly post-IP) networks, infrastructure as a service, redundant shadow configurations, in network management, and in energy efficiency. Furthermore, the provisioning of service-oriented virtual networks across multiple infrastructure providers creates the need for separation between the network operations and the physical infrastructure. This is expected to change the way that virtual networks are managed, debugged, and operated. The Dagstuhl seminar contrasted different approaches to network virtualization and investigated their applications.

Experimental research: the Internet standardization process relies on running code and real world verification. An essential prerequisite for the transfer of research results is building of large scale testbed networks. These are frequently implemented as virtual, Software Defined Networks that run concurrently to a production network using the same hardware. The Dagstuhl seminar revisited the experimental approach and gathered lessons learned and best practices.

During the seminar, we discussed and (partly) answered the following questions:

Is a prescriptive network theory feasible? Today, network research is largely descriptive, e.g., there exist methods and tools to model communication networks and protocols, to analyze their performance, or to verify their correctness. The design of new networks, however, lacks a prescriptive network theory that provides necessary rules and equations that specify how a network for a given purpose has to be built. Instead, network design relies heavily on previous experience and best practices frequently resulting in incremental works. In contrast, the clean-slate Future Internet approach seeks to build a new Internet architecture from scratch. In this case the design space is entirely open requiring decisions regarding functional and non-functional aspects, e.g.,

- Where to implement reliable/unreliable and connectionless/connection-oriented?
- Where (end systems or network) and in which layer to keep state information?
- Where and how to achieve security, quality of service, and dependability?
- How to split locators and identifiers?

Given the examples above, we discussed:

- How can a prescriptive approach to network theory be formulated?
- What are the perspectives and the fundamental limits of the candidate approaches?
- What are the prospects of the approach if successful?

Which insights can the experimental, testbed-based approach reveal? Many approaches to the Future Internet are experimentally driven and centered around a testbed that ideally if successful becomes the first running prototype of the Future Internet. Clearly, testbeds are indispensable to implement running code as a proof-of-concept, whereas their use for understanding networking and for establishing new principles and paradigms can be debated. In the seminar we elaborated on this question to provide answers to:

- Which insights can be expected?
- Which exemplary fundamental insights did emerge from testbeds?
- For which use cases are testbeds meaningful, e.g., to engineer details, to approach concepts weakly understood, to understand the impact of users, etc.?
- How should a testbed platform look like, which properties must be provided?
- How can testbeds be benchmarked to achieve comparability and validity?

What are the challenges for wide-area service-oriented virtual networks? The virtualization paradigm gained a lot of attention in networking as it provides numerous useful applications and promises to solve a number of important issues, such as the gradual deployment of new networking solutions in parallel to a running production network. Considering existing networking technologies, it becomes apparent that virtual networks and virtual network components are already being used in a multitude of different ways and in different layers, e.g., Virtual LANs (VLANs), Virtual Private Networks (VPNs), the Virtual Router Redundancy Protocol (VRRP), or in form of overlay networks to name a few. Furthermore, the abundance of resources offered by commodity hardware can turn it into a powerful and highly programmable platform for packet processing and forwarding. The virtualization of such programmable network elements can provide network slices which are highly customized for particular network services and applications. The topics that were discussed at the seminar include:

- Resource discovery and provisioning of virtual networks across multiple domains, given that infrastructure providers will not be willing to expose their topology, resource information and peering relationships to third-parties;
- Virtualization of network components (e.g., resource allocation, isolation issues);
- Scaling of virtual resources to meet variable service demand;
- Use cases of network virtualization.

This report provides an overview of the talks that were given during the seminar. Also, the seminar comprised a one minute madness session for introduction and for statements on the Future Internet, a breakout session for group work on the topic of prescriptive network theory, as well as podium discussions on experimentally driven research and on the use cases of SDN. The discussions, viewpoints, and results that were obtained are also summarized in the sequel.

We would like to thank all presenters, scribes, and participants for their contributions and lively discussions. Particular thanks go to the team of Schloss Dagstuhl for their excellent organization and support. We also would like to thank Anil Madhavapeddy for his feedback and comments on SDN.

2 Table of Contents

Executive Summary

<i>Jon Crowcroft, Markus Fidler, Klara Nahrstedt, and Ralf Steinmetz</i>	75
--	----

Overview of Talks

Prescriptive Network Theories <i>Jon Crowcroft</i>	80
Introduction to Prescriptive Network Theory <i>Markus Fidler</i>	80
Software-Defined Networking: Challenges and Opportunities <i>David Hausheer</i>	80
<provoke>Research Testbeds Considered Harmful</provoke> <i>Markus Hofmann</i>	81
Wide-Area Distributed System Deployments Yield Fundamental Insights <i>Brad Karp</i>	81
A Formal Model for Network Communication Mechanisms <i>Martin Karsten</i>	82
Opportunities and Challenges for Software Defined Network Systems <i>Wolfgang Kellerer</i>	82
Automatic Energy Efficiency Management of Data Center Servers Operated in Hot and Cold Standby with DVDS <i>Paul Kühn</i>	83
Experimentally driven Networking Research <i>Jörg Liebeherr</i>	84
SDN++: Beyond Programmable Plumbing <i>Laurent Mathy</i>	84
Congestion Exposure (ConEx) – An Experimental Protocol for the Future Internet <i>Michael Menth</i>	84
A Virtual Environment for Distributed Systems Research <i>Paul Müller</i>	85
Software-defined Networks for Distributed Interactive Multimedia Environments <i>Klara Nahrstedt</i>	85
Towards Wide-Area Network Virtualization <i>Panagiotis Papadimitriou</i>	86
The NEBULA Future Internet Architecture <i>Jonathan M. Smith</i>	86
Multi-Mechanism Adaptation for the Future Internet – MAKI <i>Ralf Steinmetz</i>	87
Some Reflections on Experimentally driven Research and Testbeds in Future Network Research <i>Phuoc Tran-Gia</i>	87

Overlay Networks as Innovation Engines <i>Oliver P. Waldhorst</i>	87
Working Groups on Prescriptive Network Theory	
Group A <i>Ruben Cuevas-Rumin</i>	88
Group B <i>Markus Fidler</i>	89
Group C <i>Florin Ciucu</i>	90
Group D <i>Tobias Hoßfeld</i>	91
Group E <i>Oliver P. Waldhorst</i>	92
Podium Discussions	
Experimentally driven research <i>Tobias Hoßfeld</i>	93
Use cases of SDN <i>Jon Crowcroft</i>	96
Seminar Programme	98
Participants	99

3 Overview of Talks

3.1 Prescriptive Network Theories

Jon Crowcroft (University of Cambridge, UK)

License  Creative Commons BY 3.0 Unported license
© Jon Crowcroft

We have many descriptive theories which help us with the design process for future networks, but mostly they work as descriptive or analytic tools rather than generative systems.

What we want is some way to feed in some constraints and out comes a prescription for a network architecture – we can then fire up the traditional tools (graph theory, queueing theory, control theory, etc) on the various components, to get a more detailed understanding of how the system will operate.

Key ideas that might go into such a theory come from the work by John Doyle on Highly Optimised Tolerant systems, and in particular, a recent idea on “Constraints that de-constrain”, such as the choice of the lowest common denominator in the layered architecture of the Internet as the waist of the hour glass. And the use of unstructured addresses in the original choice of how IP destinations (and sources) might be interpreted.

How could we embed such ideas in a generative theory?

3.2 Introduction to Prescriptive Network Theory

Markus Fidler (Leibniz Universität Hannover, DE)

License  Creative Commons BY 3.0 Unported license
© Markus Fidler

The term ‘prescriptive network theory’ is not commonly used and it may be the lack of such a theory that explains fundamental uncertainties in network design. A prescriptive theory, as opposed to a descriptive theory, is concerned with what should be, rather than with what is. Prescriptive theories are often viewed as high-level guidelines or rules, e.g., for the design of systems. Formally, a prescriptive theory is axiomatic. As such, it is a framework to derive certain conclusions, yet, it cannot be proven in itself. This presentation gives an introduction to the topic and raises a number of fundamental questions for group work discussion.

3.3 Software-Defined Networking: Challenges and Opportunities

David Hausheer (TU Darmstadt, DE)

License  Creative Commons BY 3.0 Unported license
© David Hausheer

Main reference D. Hausheer, A. Parekh, J.C. Walrand, G. Schwartz, “Towards a compelling new Internet platform,” *Integrated Network Management 2011:1224–1227*.

URL <http://dx.doi.org/10.1109/INM.2011.5990569>

One of the biggest challenges related to the Future Internet in general and Software-Defined Networking (SDN) more specifically, is that the current Internet architecture is ossified, i.e. it can hardly be changed. New protocols to improve security, QoS, or mobility rarely get deployed across administrative boundaries. Moreover, fundamental problems of coordination

among providers lead to implementation within a single domain only. Such issues include interoperability of network elements, resolve conflicts of interest, and revenue sharing. SDN as a new technology will face the same challenges, i.e. a major problem will be to support deployment of SDN across multiple domains. This requires amongst other issues a “East/West”-SDN-Controller Interface to address coordination problems among multiple providers.

3.4 <provoke>Research Testbeds Considered Harmful</provoke>

Markus Hofmann (Alcatel-Lucent Bell Labs – Holmdel, US)

License  Creative Commons BY 3.0 Unported license
© Markus Hofmann

The best research solves the best problems. To find the best problems, look within the known wide-spread or large-scale problems. Almost everyone will know about the high-level problem; it is the successful researcher who looks at the details to find a game-changing specific problem whose solution leads to a breakthrough.

In our profession, many of the best problems become apparent only when a system has been deployed for large-scale use. For example, it is hard to predict the creativity of hackers in a lab behind closed doors. Scalability issues often become apparent only after millions or billions of users adopt a solution and start relying on it for their daily life. Cost effectiveness – often one of the hardest problems, often leading to exciting inventions – raises in importance when commercial deployments are being considered. Resiliency of a network is non-trivial to determine on a whiteboard or even in a lab. How do the various components that make up a network react to dirt, dust, extreme temperatures? Will they hold up? Will we have to design the system differently? All these questions have to be addressed when designing networks for the real-world. They are some of the hardest problems to solve.

As such, experimental research, running code, and large-scale tests are key ingredients to making progress, having real impact, being relevant. While lab experiments provide initial insights into the feasibility of an approach, they rarely can address questions around scalability and resiliency. As a result, we often set-up larger-scale overlay testbeds that run as kind of a virtual network on top of production networks.

Using examples from earlier work with experimental overlay testbeds and real-life production networks, this talk will address the need for experimental research but also discuss the challenges and shortcomings of overlay-based approaches.

3.5 Wide-Area Distributed System Deployments Yield Fundamental Insights

Brad Karp (University College London, UK)

License  Creative Commons BY 3.0 Unported license
© Brad Karp

The first decade of the new millennium saw a burst of concentrated research activity on wide-area distributed systems: the handful of core Distributed Hash Table (DHT) designs; a rich set of applications built atop DHTs; peer-to-peer file sharing schemes such as BitTorrent; distributed network measurement services such as iPlane and Hubble; and content distribution

and caching systems like Coral, CoDeeN, and CoBlitz. While it was clearly “fashionable” to do research on these topics during this period (especially in the DHT and P2P area), I believe another factor contributed to this explosion in activity: the widespread availability of wide-area testbeds like PlanetLab and RON.

In this talk, I will briefly review my experience (particularly with the OpenDHT DHT service, which ran on PlanetLab and was publicly accessible between 2004 and 2009) and that of others deploying systems on wide-area testbeds. It is rare for publications in the systems area to describe which aspects of a problem definition or design were motivated by deployment experience in the wide area. That paucity of examples in the literature shouldn't be confused with testbeds' not leading researchers to such fundamental insights, however. The thesis I will present in this talk is that experience deploying distributed systems on wide-area testbeds routinely leads to fundamental insights of research significance, as measured in new problem definitions and algorithmic solutions to those problems.

3.6 A Formal Model for Network Communication Mechanisms

Martin Karsten (University of Waterloo, CA)

License © Creative Commons BY 3.0 Unported license
© Martin Karsten

Joint work of Karsten, Martin; Keshav, S.; Prasad, Sanjiva; Beg, Mirza

Main reference M. Karsten, S. Keshav, S. Prasad, M. Beg, “An axiomatic basis for communication,” in Proc. of the 2007 Conf. on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM'07), pp. 217–228, ACM, 2007.

URL <http://dx.doi.org/10.1145/1282380.1282405>

In 2007, a group of collaborators and I presented a paper at the Sigcomm conference that introduced an axiomatic description of network communication. While we have continually received very positive feedback about this work, the ACM digital library shows only 5 proper citations almost 6 years later and in general, there seems to be very little directly related work. Other formal models are typically concerned with more specific problem areas and protocols, such as BGP-style routing. At the same time, it has been surprisingly difficult (for me) to evolve the original communication model into a complete yet simple language that offers tangible benefits. In this talk, I will present recent progress in designing such a formal model/language for basic network communication mechanisms. More importantly, I want to use the opportunity to have a discussion about the usefulness and feasibility of this approach.

3.7 Opportunities and Challenges for Software Defined Network Systems

Wolfgang Kellerer (TU München, DE)

License © Creative Commons BY 3.0 Unported license
© Wolfgang Kellerer

Whereas the Internet has emerged to an economic factor for industries across all disciplines of our information society, its current system architecture fails to support such emerging application requirements in a flexible and dynamic way. In this respect, Software Defined Networking (SDN) marks a fundamental paradigm shift in information and communication

networking technology. It introduces an open interface between network hardware realizing data forwarding and the corresponding control software overcoming several limitations of current network architectures. The SDN concept allows for the first time to implement a completely dynamic control of communication networks. Forwarding rules are pushed in runtime from a logically-centralized external control entity to the distributed network hardware. From a network programming point of view, this mechanism allows to adapt the communication infrastructure flexibly and rapidly with respect to changing service demand created by the users of the network. Moreover, the concept of SDN is not limited to the basic switches and routers, but can be viewed as a general concept to increase flexibility and dynamic adaptation in communication networks spanning all network infrastructure including servers and storage.

This presentation reviews the opportunities provided by SDN addressing not only the SDN controller southbound interface (e.g. OpenFlow), but, in particular, the northbound interface towards support for network system applications of heterogeneous application domains. Standardization has already started and many network enterprises expect SDN to solve all problems they have with the Internet. In fact, however, in many respects SDN research is still at its initial research phase with many basic challenges to be resolved.

3.8 Automatic Energy Efficiency Management of Data Center Servers Operated in Hot and Cold Standby with DVDS

Paul Kühn (Univ. Stuttgart, DE)

License  Creative Commons BY 3.0 Unported license
© Paul Kühn

The operation of large Data Centers with thousands of servers is very costly in terms of energy consumption and cooling requirements. Distributed Data Centers will play a major role in the future Internet for Content Distribution and as Service Delivery Platforms, also known as Cloud Networks. Energy efficiency and cloud network performance are largely reciprocal to each other and have to be optimally managed to save energy costs at one hand and to meet Service Level Agreements (SLA) in terms of throughput and delay at the other hand. Moreover, volatile service requests need a careful load balancing by process/data migration between the distributed Data Centers causing additional costs for data replication and communications. The participating components span a wide field for the management of Distributed Data Centers.

In this contribution a generalized model for resource management of Data Center servers is presented which allows an automatic server consolidation by a load-dependent control of server activations/deactivations using multi-parallel hystereses thresholds, cold and hot server standby, and Dynamic Voltage and Frequency Scaling (DVFS). For the analysis of energy efficiency and performance a multi-server queuing model is defined which is controlled by a Finite State Machine (FSM) and a two-dimensional system state variable which allows for an automatic adaptation to the current system load. The FSM is constructed such that all requirements of the above stated functionalities can be considered and influenced by proper parameterization.

The model can be analysed exactly under Markovian traffic assumptions from which all relevant performance and energy efficiency metrics can be derived. The focus of the presentation is drawn mainly to numerical results for various parametric studies which

demonstrate the usefulness of the modeling approach with respect to Data Center Management and shows the tradeoff between the conflicting aims of Energy Efficiency and Performance. The mathematical methodology will be reported in more detail in a forthcoming publication.

An outlook is given to current work on the management of Distributed Data Centers within a Cloud Network operating on the current load or on system-state based process migration between the Data Centers which is controlled by a mapping of virtualized servers on the physical server resources for load balancing maintaining SLA agreements.

3.9 Experimentally driven Networking Research

Jörg Liebeherr (University of Toronto, CA)

License  Creative Commons BY 3.0 Unported license
© Jörg Liebeherr

Experiments in experimental networking research often follow an engineering design process rather than the classical scientific experiment approach. The different discovery process which involves (in-)validating the realization of a design, as opposed the rejecting/accepting of a hypothesis puts different requirements on experiments. The talk reviews challenges and risks of experiments in the engineering design discovery process.

3.10 SDN++: Beyond Programmable Plumbing

Laurent Mathy (University of Liège, BE)

License  Creative Commons BY 3.0 Unported license
© Laurent Mathy

SDN, as exemplified by OpenFlow, has focused on control plane programmability. While this is a tremendous step forward towards greater networking application flexibility as well as network virtualization, much can also be gained from data plane programmability. This talk explores some whys and hows in this context.

3.11 Congestion Exposure (ConEx) – An Experimental Protocol for the Future Internet

Michael Menth (Univ. Tübingen, DE)

License  Creative Commons BY 3.0 Unported license
© Michael Menth

The congestion and delay experienced in the Internet is often caused by a small fraction of heavy users that overuse limited shared resources. Internet Service Providers (ISPs) counteract with measures that are costly or not very effective, impose unnecessary restrictions, or violate network neutrality. This calls for novel congestion management solutions that need input about the congestion state of the network. The IETF is currently developing the Congestion Exposure (ConEx) protocol that reveals the congestion caused by a flow to all nodes on its path. Intermediate boxes may leverage that information for novel congestion management

approaches. The talk explains how congestion exposure works, gives example for its use, and illustrates ConEx-based congestion policing in more detail.

3.12 A Virtual Environment for Distributed Systems Research

Paul Müller (TU Kaiserslautern, DE)

License © Creative Commons BY 3.0 Unported license
© Paul Müller

Joint work of Mueller, Paul; Cappos, Justin; Schwerdel, Dennis

Networks and distributed systems are an important field of research. To enable experimental research in this field we propose to build a virtual research environment for distributed systems. It should support researchers from various branches of science who investigate distributed systems by providing a virtual environment for their research. With the topology management tool (ToMaTo), a flexible and efficient networking testbed has been built that can be used as a core for the proposed research environment. Using various virtualization technologies, ToMaTo is able to provide realistic components that can run real-world software as well as lightweight components that can be used to analyze algorithms at large scale. This paper describes how an additional virtualization technology from the Seattle testbed has been added to ToMaTo to allow even larger experiments with distributed algorithms.

3.13 Software-defined Networks for Distributed Interactive Multimedia Environments

Klara Nahrstedt (University of Illinois – Urbana, US)

License © Creative Commons BY 3.0 Unported license
© Klara Nahrstedt

Joint work of Nahrstedt, Klara; Ahsan Arefin; Raoul Rivas

Main reference A. Arefin, R. Rivas, R. Tabassum, K. Nahrstedt, OpenSession, “Cross-layer Multi-stream Management Protocol for Interactive 3D Teleimmersion,” Technical Report, University of Illinois at Urbana-Champaign, 2013.

URL <https://www.ideals.illinois.edu/handle/2142/43348>

New Distributed Interactive Multimedia Environments (DIMEs) such as 3D teleimmersive environments are emerging, consisting of multiple 3D video streams, audio, graphics and other sensory information to connect geographically distributed users into a joint virtual space for virtual interaction and remote collaboration. These DIMEs have a very large demand on bandwidth/CPU resources and tight delay requirements for enabling interactivity. Multi-stream, multi-modality and multi-view requirements add to the complexity of the DIME’s application model. In the past couple of years, these types of environments have been using extensively the underlying Best Effort Internet infrastructure with bandwidth and delay-aware overlay session-based protocols since underlying IP-based protocols and frameworks such as RSVP/IntServ, DiffServ, or MPLS technologies have not been easily accessible for distributed interactive multimedia application developers (e.g., multiplayer gaming, video-conferencing developers). With Software-Defined Networks (SDN) and their Open Flow implementation, an opportunity emerges for DIME developers to have a much stronger access to lower level protocol stack and QoS control over Internet networks, and hence coordinate the bandwidth and delay controls between session-based and IP-based protocols and layers. The question is ‘How’ one achieves this cross-layer SDN-DIME coordination and

control. In this talk, we will discuss challenges, problems and directions that we see as we explore cross-layer SDN-DIME coordination and control issues in our TEEVE (Teleimmersion for Everybody) testbed.

3.14 Towards Wide-Area Network Virtualization

Panagiotis Papadimitriou (Leibniz Universität Hannover, DE)

License © Creative Commons BY 3.0 Unported license
© Panagiotis Papadimitriou

Joint work of Papadimitriou, Panagiotis; Dietrich, David; Rizk, Amr; Bozakov, Zdravko; Mathy, Laurent; Werle, Christoph; Bless, Roland

Main reference D. Dietrich, A. Rizk, P. Papadimitriou, “Multi-Domain Virtual Network Embedding with Limited Information Disclosure,” IFIP Networking 2013.

The ever-increasing need to diversify the Internet has recently revived the interest in network virtualization. Most virtual network (VN) provisioning techniques assume a single administrative domain and, as such, the deployed VNs are limited to the geographic footprint of the substrate provider. To enable VN provisioning across multiple substrate providers, network virtualization architectures need to address the intricacies of inter-domain aspects, i.e., how to provision VNs with limited control and knowledge of any aspect of the physical infrastructure.

In this talk, we discuss the challenges of multi-provider VN provisioning and present a framework that circumvents the difficulty of VN embedding with limited information disclosure (LID) from substrate providers. We further present a technique for inter-domain virtual link setup based on NSIS.

Given the increasing interest in network programmability and control, we discuss the benefits and challenges of SDN virtualization. In this context, we outline the design of a distributed SDN hypervisor that facilitates the embedding, deployment, configuration and operation of virtual SDNs.

3.15 The NEBULA Future Internet Architecture

Jonathan M. Smith (University of Pennsylvania, US)

License © Creative Commons BY 3.0 Unported license
© Jonathan M. Smith

URL <http://nebula.cis.upenn.edu>

The NEBULA Future Internet Architecture is a clean-slate design focused on secure and resilient networking to support cloud computing. It is designed to support applications that would be difficult or impossible today, such as cloud-based real-time control of medical devices. It is comprised of three parts: (1) N CORE, a core network interconnecting data centers; (2) NDP, a new data plane; and (3) NVENT, a new control plane. Two flavors of NDP have been implemented, Icing and TorIP. Icing uses cryptographic markers on each packet for policy enforcement. The proof of consent (PoC) is generated iff every network element agrees with the policy. This has been implemented and integrated with the RapidNet declarative networking system. TorIP (“Tor instead of IP”) likewise requires receiver interest, but uses onion routing to defend against malicious ISPs. NEBULA is moving forward rapidly and is looking to deploy over the next two years.

3.16 Multi-Mechanism Adaptation for the Future Internet – MAKI

Ralf Steinmetz (TU Darmstadt, DE)

License © Creative Commons BY 3.0 Unported license
 © Ralf Steinmetz
URL <http://www.maki.tu-darmstadt.de>

The Collaborative Research Centre (CRC) MAKI (Multi-Mechanism-Adaption for the Future Internet) addresses this challenge. In particular, it investigates all kinds of mechanisms in communication systems, the adaptation, interaction, constant optimization, and evolution thereof. The term mechanism describes both, communication protocols and parts thereof – defining the functionality of communication systems – and the functional aspects of the distributed systems realized on top. We witness a constant development of novel mechanisms. Yet, mechanisms providing equivalent functionality under different conditions coexist, since an adaptation of legacy mechanisms to traffic conditions, bandwidth, etc. is limited. Particularly mobile usage induces highly fluctuating conditions, which would require the online adaptation of the communication system by means of transitions between functionally equivalent mechanisms – which is mostly impossible as of today. Interactions between mechanisms that jointly depend on each other are even more complex and require coordinated transitions in groups of equivalent mechanisms, so-called multi-mechanism adaptation.

3.17 Some Reflections on Experimentally driven Research and Testbeds in Future Network Research

Phuoc Tran-Gia (Univ. Würzburg, DE)

License © Creative Commons BY 3.0 Unported license
 © Phuoc Tran-Gia

During the past decade we witness numerous activities towards next generation network. A huge portion of them deals with building generic testbeds and experimental platforms, with a number of experimental facilities initiated in several countries, e.g. in the US, EU, Japan etc. The talk gives some reflections on the pros and cons of such ventures and concludes with a view on how the current development of test methodologies emerges.

3.18 Overlay Networks as Innovation Engines

Oliver P. Waldhorst (KIT – Karlsruhe Institute of Technology, DE)

License © Creative Commons BY 3.0 Unported license
 © Oliver P. Waldhorst
Joint work of Waldhorst, Oliver P.; Zitterbart, Martina

The current Internet architecture has proven to be quite stable due to the invariants imposed by IP. Innovations are mainly seen in new networking technologies on the lower layers and in applications, i.e., on the very bottom and top of the architecture, respectively. Nevertheless, novel applications on the one hand impose new requirements on the Internet, e.g., robustness, security, and mobility support, and call for sophisticated network services, e.g., group communication or in network data processing. On the other hand, applications have to be ready to be deployed on top of evolving networking technologies leading to potentially

very heterogeneous networks. In this talk, we argue that application layer overlay networks can support an application developer in coping with these challenges. We introduce the Spontaneous Virtual Networks (SpoVNet) architecture as an illustrating example. Beyond what has been achieved with overlays, we outline steps towards our vision of generating and deploying a networked application automatically from a high level specification.

4 Working Groups on Prescriptive Network Theory

4.1 Group A

Ruben Cuevas-Rumin (Universidad Carlos III de Madrid, ES)

License  Creative Commons BY 3.0 Unported license

© Ruben Cuevas-Rumin

Joint work of Bozakov, Zdravko; Crowcroft, Jon; Cuevas-Rumin, Ruben; de Meer, Hermann; Müller, Paul; Stingl, Dominik

The discussion focused on two main aspects:

First, we discussed what is the most appropriate approach to design a future Internet architecture. The research effort to define the future internet architecture should combine short and long term approaches. In particular, the short term approach is an evolutionary one that improves or extends the current architecture to address new requirements in the short term. It must take into account current technical as well as external (e.g., business model or policies) constrains. Instead, the long term approach should present a clean state design that should not consider current external constrains (e.g. the current business model of Content Providers may very well change in few years from now). Furthermore, the definition of a new architecture should come from a “new point of view”. A possible “point of view” may be the Software Engineering discipline that provides more flexibility.

Second, we discussed few important requirements that need to be consider in the design of the future Internet Architecture:

- Privacy is a new requirement that has gained relevance in the last years associated to the ‘boom’ of OSNs. A future Internet architecture should be able to provide privacy as a service.
- A power-aware architecture is needed. At the current growth of bandwidth, processing capacities, etc. the current infrastructure will not be able to dissipate the generated heat. Then, we need to take in consideration power consumption as a design requirement of the future Internet architecture.
- High availability: the current Internet guarantees a decent level of availability for a non critical service, however the availability guarantees offered for critical services is deficient. The Future Internet architecture has to deal with this issue and offer a very high availability.

As summary of the discussion, we concluded that any new architecture should provide evolvability/flexibility/adaptiveness and therefore self-organization/autonomic operation is a key necessity.

4.2 Group B

Markus Fidler (*Leibniz Universität Hannover, DE*)

License © Creative Commons BY 3.0 Unported license
© Markus Fidler

Joint work of Fidler, Markus; Gross, Christian; Hofmann, Markus; Pries, Rastin; Sarrar, Nadi; Tran-Gia, Phuoc; Mühlhäuser, Max

Group B discussed what we would refer to as the ‘meta level’ of prescriptive network theories, where we focused on:

1. the different means for construction of (Future Internet) networks, not only theories;
2. the spectrum of interesting issues (foci) that such means could concentrate on;
3. the possible degree of generality and of breadth.

With the goal of a formal theory in mind, we questioned whether there may be any fundamental laws in networking that lay out the foundation of such a theory, like the laws of nature in physics. During the discussion a number of examples of technical systems came up where basic constraints, rules, and/or theories exist that have greatly advanced the respective fields and in some cases have even been used in a ‘generative way.’ Examples include Shannon’s capacity, Erlang’s formulas, e.g., for dimensioning of telephony switches, and the AIMD-law in TCP congestion control.

Regarding networking, we first clustered the problems/fields that we want to address and based on this we formulated the kinds of means that we want to have. The fields of interest are (without claim to be complete):

1. high level
 - scalability, evolvability
 - divide & conquer vs. (inevitable) interdependencies
2. functional
 - multicast
 - multihoming
 - mobility
3. non-functional
 - energy efficiency
 - security, privacy, quality of service (bandwidth, latency, fairness, etc.)
 - economy

In the following discussion, it came out that there are numerous formal theories and models that can provide a deep understanding how and in which way certain mechanisms work, e.g., queueing theory, graph theory, game theory, and many others. These methods are frequently used to analyze non-functional aspects. The question how to engineer a certain function resorts much less to theory and rather relies on cookbooks of known mechanisms that exist, e.g., for medium access control, ARQ, congestion control, and so on. While many non-functional aspects of these methods are well-understood, e.g., the throughput of ALOHA vs. CSMA, we seem to lack a generative theory that enables finding new mechanisms. Finally, considering network architecture on a high level, we face problems that are hard to formalize and hard to anticipate so that we are confined to guidelines and best practices. In conclusion, we formulated how formal we expect a prescriptive network theory to become with respect to the complexity of the problem:

Computing a solution: optimizing a protocol, naming scheme, etc.;

Engineering a solution: designing a custom network, e.g., WSN, enterprise, etc.;

Crafting a solution: designing *the* Future Internet.

4.3 Group C

Florin Ciucu (TU Berlin, DE)

License © Creative Commons BY 3.0 Unported license
© Florin Ciucu

Joint work of Ciucu, Florin; Hausheer, David; Hollick, Matthias; Karp, Brad; Karsten, Martin; Mathy, Laurent; Santini, Silvia

The group questioned existential aspects of a generic ‘Formal Theory’ such as the necessity of explaining vs. designing the ‘reality’, or its temporal positioning. One can in particular question the usefulness of a theory either retrospectively (as being meant to ‘explain’ what has been done so far) or prospectively (as being meant to help what has not been done yet). The group also addressed questions related to its limits and scope, and also attempted to define benchmarks for its usefulness. For instance, in the context of a hypothetical network theory, one should consider temporal aspects such as design time or the frequency of break downs.

One interesting analogy that emerged during the discussion is that while physicists take for granted inherent laws that nature obeys, network scientists are confronted with a lack of uniform laws since artificial entities are by default subject to many unknown and uncontrollable inputs. Moreover, while nature has inevitably led to unfortunate ‘designs’, those have been corrected by evolution; in turn, the role of ‘evolution’ in the context of artificial entities remains elusive.

The group then debated about the appropriateness of defining a ‘prescriptive theory’ in either a strict sense (i.e., axiomatically driven) or more generally by considering in particular conducive facets for system design. The follow-up discussion assumed the existence of such a theory, whose axioms could be conceivably applied to building a network, but which eventually would turn out as bogus. This apparent contradiction was resolved by arguing that an ‘accurate’ theory can only produce a single network, in which case ‘restrictive’ would be semantically more appropriate than ‘prescriptive’, as a qualifier for ‘theory’. The rather contrived line of thought that emerged was followed by an enlightening contrast between architects who can both follow and break design principles, and network scientists who are confronted at the very least with the lack of specifications for network input.

In attempting to exemplify success stories of prescriptive theories, the group first mentioned control theory and its key role to transport protocols, meta-routing, or XCP. The initial consensus was that control theory is only partly prescriptive and that it is restricted by the need of making many simplifications and assuming complete knowledge. In an effort to underline the challenges faced by network scientists, it was humorously speculated that control theorists believe that ‘everything’ (including networks) can be controllable or at the very least explained. The second example was game theory which was thought to be as being more prescriptive than other theories by arguing that it starts out with the axiom that players act in their own interest. However, by reflecting on applications to selfish routing, the prescriptive facet of game theory started being questioned by considering the position of ISPs. Queuing theory was the last example, and that was argued to be characterized by descriptive rather than prescriptive aspects since it is based on axioms of mathematics.

Another separate line of discussion concerned broken aspects and possible fixes in the current Internet. Immediate facts being mentioned included the existence of too many protocols for the same purpose, the network overloading with functionality, the lack of bounded latency, the apparent impossibility of building proper firewalls, the lack of extreme availability and reliability, or the unclear role of admission control. A possible cause for the current frustrating state of the Internet was identified to be that functionality itself is

indistinguishable from the infrastructure. More concretely, the whole network appears to look like an embedded system, conceivably implying that who knows the infrastructure also knows the functionality. In plastic terms, the Internet was compared to a washing machine, and it was further recommended to adopt similar principles as the airline industry, i.e., immediate redesign after failure.

The discussion overall spurred some meta-thoughts which apart from being plausible certainly deserve being reflected upon:

- All theories are limited in scope and also in time.
- Internet research is mostly empirical.
- Technology changes much faster than theorists can adapt.
- We need evolution theory to model the Internet.
- Lets just accept there will not be a clean-slate Internet.

4.4 Group D

Tobias Hoßfeld (Universität Würzburg, DE)

License  Creative Commons BY 3.0 Unported license
© Tobias Hoßfeld

Joint work of Dietrich, David; Hoßfeld, Tobias; Klug, Karl; Nahrstedt, Klara; Papadimitriou, Panagiotis

This report documents the discussions and outcomes of the working group D with the participants above during the Dagstuhl seminar on “Future Internet”. Thereby, the group agreed to discuss on a methodology towards prescriptive network design and had a good mixture between researchers from industry and academia.

The term “prescriptive network design” reflects a theoretical approach to design a full-service, comprehensive network close from scratch based on a set of constraints. As a result of the group discussions, five basic steps are to be fulfilled for prescriptive network design (PND). In particular, the constraints for network design have to be de-constrained in order to derive a minimum required functionality and an appropriate function split. As a pre-requisite, the elements of PND have to be identified, while the actual constraints and requirements are derived for a set of networks, use-cases, and models which are stored in a common database. Then, an interface between prescriptive and descriptive network design has to be defined in order to allow a flexible and adaptive solution (that may be seen as an iterative process where ‘top-down’ and ‘bottom-up’ are converging to a pragmatic, near-optimal solution). Thus, the recommended methodology for PND is as follows.

- Step 1. What are the elements of prescriptive network design?
- Step 2. Database of networks, use-cases, models, etc.
- Step 3. Minimum function split
- Step 4. Interface between pre-/descriptive
- Step 5. Recommendations from theory

Step 1 addresses the elements of PND. It has to be derived which are the basic elements, what are the properties of the elements, the capabilities of the architectural elements, and which are the required fundamental services. Thereby, it can be distinguished between a) functional and b) non-functional elements, constraints, or requirements. An example of a functional element is the negotiation of resources, quality guarantees, security, etc. Thereby, the location of functions has to be considered in PND, too. Non-functional elements address

for example scalability. However, it has to be clarified which are the basic elements of PND, e.g. protocols, topology, switches, etc.

Step 2 recommends a database of networks, use-cases, models, etc. In particular, the requirements for those use-cases are to be derived and collected. This includes non-functional and functional constraints or requirements, as well as business model constraints or requirements. In general, models for different use-cases, models for network dynamics, etc. have to be taken into account in this database. A concrete use-case is for example an educational campus with changing traffic demands according to the students habits and living on the campus.

Step 3 aims at de-constraining the constraints and tries to derive the minimum function split. This minimum function split reflects the smallest denominator in order to compose higher functions and services, i.e., it is the basis for functional composition. In particular, a classification of the constraints and requirements from the use-cases is to be developed. Then, the function split of the PND is based on this classification. Thereby, the minimum number of functions is to be derived, but the complexity of functions has to be considered too.

Step 4 of the recommend methodology is the definition of an interface between prescriptive and descriptive network design. Since the different services know best their requirements, it seems to be straightforward that the services describe their requirements themselves. The prescription is implemented via an interface between different services in order to design the network from a holistic point of view for all services. This interface allows an adaptive and flexible network design.

Finally, Step 5 of the PND are recommendations by the theory. In particular, the theory of PND should answer questions like the following. “How to design the network protocol stack?” “What is an appropriate network topology for the basic elements?” “How shall concrete mechanisms, e.g. schedulers, look like?” As an outcome of the PND theory, tools for network planning but also for network operation may be derived. This is in particular relevant in the context of network virtualization and the different stakeholders like physical infrastructure providers, virtual network provider, and virtual network operator. From a business point of view, use-cases may be defined with prescriptive description of networks, which lead to valid business cases for network operators.

4.5 Group E

Oliver P. Waldhorst (KIT – Karlsruhe Institute of Technology, DE)

License © Creative Commons BY 3.0 Unported license

© Oliver P. Waldhorst

Joint work of Godfrey, Philip Brighten; Kühn, Paul; Liebeherr, Jörg; Menth, Michael; Smith, Jonathan; Waldhorst, Oliver P.

Motivated by the statements of the seminar participants in the opening session the group selected ‘latency’ as an example requirement for future Internet architectures and discussed reasons for latency as well as ways of reducing it using architectural and methodical approaches. The group identified transmission / propagation delay, switching delay, queuing delays, and security handshakes as main sources of latency. While some of these sources cannot be changed, e.g., propagation delay is bounded by the speed of light or handshakes cannot be avoided without sacrificing security, some other can be tackled by changes architecture or paradigms. This holds in particular for queuing delays introduced by router buffers.

Architectural changes for reducing queuing delays include a circuit switching architecture without any buffers, which means trading delay for loss, or losing the statistical multiplexing gain. Obviously, controlling latency is expensive, since it requires sophisticated control mechanisms. Another way to reduce latency in general is multiplexing resource usage across multiple resources, for example by requesting a web site from multiple servers and using the first response. From a methodical point of view, a formalism is required to compare one architecture for reducing latency to another. One example is the enhancement of the Axiomatic Basis for Communication proposed by Martin Karsten et al. with means to compute latency.

5 Podium Discussions

5.1 Experimentally driven research

Tobias Hofffeld (Universität Würzburg, DE)

License  Creative Commons BY 3.0 Unported license
© Tobias Hofffeld

On the second day of seminar, experimentally driven research was one of the major discussion topics. To this end, five different talks were given to stimulate the discussions among the participants:

- Martina Zitterbart and Oliver Waldhorst: “overlay networks as innovation engines”
- Jörg Liebeherr: “experimentally driven networking research”
- Brad Karp: “wide-area distributed system deployments yield fundamental insights”
- Markus Hofmann: “<provoke>research testbeds considered harmful</provoke>”
- Phuoc Tran-Gia: “some reflections on experimentally-driven research and testbeds in future network research”.

Finally, a podium discussion with the speakers above took place in order to raise questions to the speakers and to express own opinions about experimentally driven research.

Just before the podium discussion, Phuoc Tran-Gia gave a talk and shared his reflections on experimentally-driven research and testbeds in future network research. Thereby, the following five key observations and statements were made by Phuoc Tran-Gia.

1. Analytical studies are necessary for future Internet research. A recommended structure of experimental-driven research projects includes two different facets that are (A) research (sub-)projects and (B) the experimental facility. On one hand, the research (sub-)projects (A) are needed to investigate particular questions in the context of future network research. This does not necessarily mean that those research projects need to be tested within an experimental facility, but the answer to those research questions may also be derived from theoretical studies, i.e. analytical approaches or simulations. One particular future internet research question addresses scalability which requires in particular simulations and analytical models to overcome limitations of testbeds in size. On the other hand, the experimental facility (B) is to be developed. Sometimes this facility is then used by the research projects from (A). This structure of experimental-driven research was successfully applied in the G-Lab project combining future internet research and experimentation, whereby the experimentation consumed 10-15% of the overall efforts in the G-Lab project. As a prerequisite for experimental-driven research projects, test cases and test interfaces in the experimental facility need to be defined. Only parts of the developed solutions in the research projects (A), e.g. some virtualization mechanisms, may then be tested

- later within the experimental facility (B). Hence, successful experimental-driven research needs a combination of measurements, simulations, analysis.
2. Generic experimental platforms have the following advantages. They enable large-scale experimental platforms which can be reused for several test-cases. This reduces the setup time for the testbed initiation phase. Further, the generic experimental platform provides a clean environment which allows to generate reproducible results. The reproducibility of the experiments is however questionable. Another advantage of generic experimental platforms is adaptability which allows to enlarge a number of nodes in the test phase easier, to quickly change test images in the platform, to enable federation with testbeds, etc. Finally, generic experimental platforms can in principle be designed for sustainability.
 3. Major problems with generic experimental platforms however outweigh the advantages. Often, generic experimental platforms cannot be used in most emerging test cases, as they need e.g. special requirements not offered by the platform. The lifetime of generic experimental platforms is limited due to the hardware and software lifecycles. Further, the question is raised how generic is a generic testbed, how to do scalability analysis, whether experiments are really reproducible. The question arises how “real” are experimental platforms? Conducting measurements e.g. in Planetlab and via crowdsourcing show that the obtained results are not representative. Another major drawback of generic experimental platforms is the fact that the focus is often only on the testbeds itself while classic methods are often neglected. For example, testbeds can only partly support overload-preventive design process and hardly support scalability analysis. Another example is network-aware application design which makes testbeds more complex. Thus, it is an open question whether application-aware network design & network-aware application design benefit from testbeds.
 4. Some observations on generic testbeds were stated. First, most testbeds are underutilized, many testbeds are not used by industries and there is a lack of users and experimenters of the testbed. Another observation is that generic testbeds start to break down in several disjoint testbeds. Further, maintenance costs for the testbed are significant and large efforts have to be spent for sustainability and life-time of the testbed.
 5. Possible solutions to overcome the limitations of generic experimental platforms are the following. First it is important to design research projects together with experimental facility research (see statement 1. above). This includes to design and finance test projects together with the testbed setup. Further, software-defined experimental platforms are easy to extend and to change and may be combined with other techniques. For example, the human cloud and real users may be integrated in experimental facilities, e.g. by means of crowdsourcing. The integration of the human cloud is sometimes required, e.g. when looking at Quality of Experience (QoE), and cannot be offered by a machine cloud. Another recommendation is the federation and integration of special testbeds in generic experimental platforms in order to extend their capabilities and possibilities.

Directly after the presentation by Phuoc Tran-Gia, the podium discussion took place with the speakers on experimental-driven research in front of the audience: Martina Zitterbart, Oliver Waldhorst, Jörg Liebeherr, Brad Karp, Markus Hofmann, Phuoc Tran-Gia (in chronological order of presentations).

The first question to the panel was raised by Tobias Hofffeld: “Is real user behavior important in testbeds? Do we have to include real users in testbeds? Can crowdsourcing help to identify problems like signaling storm?” There was, however, no agreement in the panel and in the audience about the integration of real users in testbeds. Jörg Liebeherr mentioned that the purpose of a testbed is to find limits of technologies, e.g., security

features of cars are not tested using real users but dummies. In contrast, Phuoc Tran-Gia mentioned that the car example does not hold as other features are tested in fact by real users. Further, Markus Hofmann explained that real-world problems are overseen without real users. To this end, Brad Karp mentioned that there is a continuum between basic research and experiments and that crowd testing and machine testing are important. Another viewpoint was given by Oliver Waldhorst, as the application developer (as a real user) has to be considered in experiment-driven research, too.

The second question was asked by Matthias Hollick and focused on the presentation by Markus Hofmann: “Do you provide a testbed for application developers?” Markus Hofmann explained that this is not done by their company, as they are only considering the requirements, expectations, and so on. Nevertheless, there is a change in research, as customers have to be asked about scenarios. Further, it is hard to predict the consequences of application development. In contrast, e.g., to car industry, scalability is an important issue in future networking research.

Then, the discussion focused on scalability. Phuoc Tran-Gia mentioned that scalability analysis requires test theory, analytical models, etc.; it cannot be tested in testbeds. Scalability is a big challenge, because 1 Mio testing devices is not possible in a testbed, although it can be extended to a certain extent by crowdsourcing. Jon Crowcroft asked what will change when you test a system with 10 Mio users. Will unknown effects happen in large-scale, e.g., synchronization effects? Klara Nahrstedt mentioned the example of group dynamics from social science. In particular, group dynamics change depending on size of group and a change is observed between 1, 5, 10, 50, 100, 150 users, while above 150 users no additional changes happen. Jörg Liebeherr criticized that any scalability study leaves the reader unsatisfied, e.g., due to model assumptions. However, Paul Müller argued as mathematician that we can believe in ‘small numbers’ and draw conclusions for scalability analysis, since there are statistical methods to investigate scalability and to provide valid results. Markus Hofmann finally mitigates the discussion on scalability analysis, as it is not a binary decision to have large-scale experiments or not, since we need all different methods in research. It will not help to use just one or the other.

Paul Kühn stimulated the discussions then in two different directions. First, the effects of social networks have to be considered in future Internet research which was commented by Jon Crowcroft that social networks and communication networks can be combined. Second, errors in software deployment and error propagation in software-defined networking need to be considered in (SDN) experiments. When considering Software Defined Networking, we have to deal with a lot of software errors. This will be problematic, because we first have to identify software problems. Brad Karp answered that when looking at Cisco, we have millions of lines of code, which we cannot take a look at. This will not be the case with SDN. However, Paul Kühn argued that the frequency of changes in software is a lot higher in SDNs. Paul Müller mentioned that a formal verification of software and tools for software verification are required, which is a general software problem, but tools are partly available as pointed by Brad Karp.

Next, Phuoc Tran-Gia asked the following: “Is reverse-engineering of (human-designed and implemented) mechanisms science or engineering? Example: YouTube re-engineering of mechanism . . . and then the mechanism is changed with the next version.” This was commented by an additional question by Paul Müller: “Is academia running behind industry, e.g., Skype?” However, Klara Nahrstedt noted that there are a lot of success stories coming out from university and academia. Markus Hofmann mentioned that industry does not care about optimal solutions, but pragmatic solutions. Industry often

gets along with a non-optimal version, while university often tries to find the optimum. The next question was raised by Klara Nahrstedt: “Heterogeneous devices are not supported by testbeds. How to plugin more realistic traffic patterns? How to bring applications closer to network testbeds?” Paul Müller mentioned the Seattle tool for this, while Phuoc Tran-Gia also stated that many testbeds are ‘too low’ and neglect applications and users. The last question in the podium discussion was queried by Johnathan M. Smith: “How to sample and privatize data?” There was a lively discussion on sampling and privacy. Markus Hofmann answered that a large amount of data is available for research, but often not accessible due to privacy concerns, regulations, and so on. He additionally asked: “Can we make providers (like YouTube) share data?” and claimed that sampling and anonymization tools are required that are broadly acceptable. Even within companies, it is difficult to get data from other groups. Brad Karp agreed that anonymization is a hard problem and that it is unclear how to ensure anonymization of data.

5.2 Use cases of SDN

Jon Crowcroft ((University of Cambridge, UK))

License  Creative Commons BY 3.0 Unported license
© Jon Crowcroft

On the last day of the seminar, we had a session where we tried to enumerate some of the use-cases for SDN, and this is a partial list of what was captured:

1. multi tenant data centers have a need for specialised networking for different tenants who may have different network service requirements – just picking a few examples, one tenant might need multicast for serving IPTV, while another might need delay bounds for traders or gamers, whilst still another might need in-network processing to support better MapReduce task deterministic throughput.
2. Multi User VR and Massive Multiplayer Online Role Playing Games have large scale dynamic network requirements whose parameters keep changing – one could certainly imagine supporting this partly through SDN.
3. Live Migration of Virtual Machines is increasingly used both for load balancing and for high-availability/persistence during scheduled maintenance. VM Migration puts a spike load onto the net, and could interfere with routine steady state traffic unless serviced specially, and requires itself, specialised support for continued access (re-routing existing flows, and redirecting new clients, for example). This could be supported as an SDN function.
4. The Internet of Things (IoT) is touted as a big driver for networking (e.g. IPv6 deployment to support the massive increase in globally reachable devices). IoT will also require specialized network functionality (securing access, delay bounding access for some sensor/actuator cyber-physical applications, and many as yet unforeseen applications that will no doubt emerge).
5. Content Distribution Networks (CDN) need management. Some CDNs place difficult load on an ISP – especially CDNs that are all or partly P2P, perhaps interfering with peering arrangements. Managing this could easily be seen as a good match to SDN.
6. Middlebox management is sorely needed. Middleboxes increasingly are the main cause of ossification of the network – the inability to deploy new transport layer protocols or extensions is mainly traceable to the ad-hoc nature of the large number of different

middle box functions deployed throughout the Internet. Bringing these into a coherent framework would allow some semblance of progress to be made. SDN again would be the tool to replace the ad-hoc functionality with a clean slate programmable system with a public, open standard API.

7. This is the metaclass of the SDN use case – management of multiple SDNs will itself be a challenge.
8. Hybrid clouds where the cloud supports a set of user applications and the set of SDN apps is an interesting case of meta-management. Co-existence of the applications and the customised, specialised support for these applications in the software defined network is a key requirement, for example, in today's' data centers.
9. Resilience in networks currently relies on ad-hoc approaches to providing replication. SDN can help unify a set of mechanisms under one control plane.
10. Cross-layer design of distributed applications requires potentially more open, possibly reflective APIs in SDN, so that the Application and the SDN can co-evolve efficiently.
11. Enterprise infrastructure setup is a key need in large scale private intranets. There are many such systems in the world, and often their owners incur high costs to provide customised network services. SDN offers a way to build more flexible networks that could be matched to an enterprise's needs more as a matter of configuration than bespoke engineering.
12. Improved security may be on the cards if SDN takes on board improved software practises, using safer programming languages, and trusted computing bases, and techniques for information flow analysis, software verification, and so forth.
13. One simple hope for SDN might be to take research results in policy routing (e.g. meta-routing) and do a one-time replacement of BGP.

We also considered SDN in data plane.

1. Data plane middleboxes already interfere with TCP/IP packets in an unstructured way (frequently, to improve operation of protocols in wireless networks such as 3G and 4G nets, but also interfering with the ability to deploy new versions of TCP. SDN could include data plane packet processing, at least near the edges of the network where performance requirements can be met affordably.
2. Fine grain media control could be another SDN data plane activity – e.g. video and audio re-coding for different receivers with different rendering capabilities.
3. Network as a Service (NaaS) for data center (e.g. mapreduce) in-net application code has also been suggested as an SDN data plane task. The TCP incast problem can be solved by processing a fixed number of shuffle phase data packets in switches, with relatively simple tasks.

In addition to numerous use-cases or applications on top of SDN, SDN was also found to have the potential to connect technologies below the SDN implementation. But there are broader questions that we leave unanswered here:

1. Can we use SDN to Connect IP and non-IP networks?
2. Could we do layer 2 and layer 4 SDN via OpenFlow?
3. What could make SDN harmful?
4. What are the key SDN business cases?

6 Seminar Programme

Monday	Prescriptive network theory
09:00-09:30	Welcome and general introduction
09:30-10:30	One minute madness: introduction of participants & Future Internet statement Scribes: Zdravko Bozakov and David Dietrich
11:00-12:00	Markus Fidler: introduction to prescriptive network theory Martin Karsten: a formal model of network communication mechanisms Scribes: Christian Groß and Dominik Stingl
14:00-15:15	Ralf Steinmetz: multi-mechanism adaptation for the future Internet Jon Crowcroft: prescriptive network theories Scribes: Christian Groß and Dominik Stingl
15:15-15:30	Definition of topics for group work and opinions Scribe: Zdravko Bozakov
16:00-17:45	Breakout sessions in group work
19:00-20:00	Wrap-up of group work results in plenum

Tuesday	Experimentally driven research (in overlays)
09:00-10:15	Martina Zitterbart/Oliver Waldhorst: overlay networks as innovation engines Jörg Liebeherr: experimentally driven networking research Scribe: Nadi Sarrer
10:45-12:00	Brad Karp: wide-area distributed system deployments yield fundamental insights Markus Hofmann: <provoke>research testbeds considered harmful</provoke> Scribe: Florin Ciucu
13:45-14:20	Phuoc Tran-Gia: some reflections on experimentally driven research and testbeds in future network research Scribes: Rastin Pries and Tobias Hofffeld
14:20-15:10	Podium discussion Scribes: Rastin Pries and Tobias Hofffeld
16:00-17:45	Paul Müller: a virtual environment for distributed systems research Jonathan M. Smith: NEBULA future Internet Michael Menth: Conextion Exposure (ConEx) – an experimental protocol for the future Internet Paul Kühn: automatic energy efficiency management of data center servers operated in hot and cold standby with DVDS Scribes: Panagiotis Papadimitriou and Ruben Cuevas-Rumin
17:45-18:15	Discussion and opinions

Wednesday	SDN, virtualization, and OpenFlow
08:30-10:00	Wolfgang Kellerer: opportunities and challenges for Software Defined Network systems Klara Nahrstedt: Software Defined Networks for distributed interactive multimedia environments Laurent Mathy: SDN++: beyond programmable plumbing Scribe: David Dietrich
10:30-11:00	Panagiotis Papadimitriou: towards wide area network virtualization Scribe: David Hausheer
11:00-12:00	Discussion, opinions, and use cases of SDN Scribe: David Hausheer
12:00-12:15	Seminar resume and farewell

Participants

- Zdravko Bozakov
Leibniz Univ. Hannover, DE
- Florin Ciucu
TU Berlin, DE
- Jon Crowcroft
University of Cambridge, UK
- Ruben Cuevas-Rumin
Universidad Carlos III de Madrid, ES
- Hermann de Meer
Universität Passau, DE
- David Dietrich
Leibniz Univ.t Hannover, DE
- Markus Fidler
Leibniz Univ. Hannover, DE
- Philip Brighten Godfrey
Univ. of Illinois – Urbana, US
- Christian Gross
TU Darmstadt, DE
- David Hausheer
TU Darmstadt, DE
- Markus Hofmann
Alcatel-Lucent Bell Labs – Holmdel, US
- Matthias Hollick
TU Darmstadt, DE
- Tobias Hoffeld
Universität Würzburg, DE
- Brad Karp
University College London, UK
- Martin Karsten
University of Waterloo, CA
- Wolfgang Kellerer
TU München, DE
- Karl Klug
Siemens – München, DE
- Paul J. Kühn
Universität Stuttgart, DE
- Jörg Liebeherr
University of Toronto, CA
- Laurent Mathy
University of Liège, BE
- Martin Mauve
Heinrich-Heine-Universität Düsseldorf, DE
- Michael Menth
Univ. Tübingen, DE
- Max Mühlhäuser
TU Darmstadt, DE
- Paul Müller
TU Kaiserslautern, DE
- Klara Nahrstedt
Univ. of Illinois – Urbana, US
- Panagiotis Papadimitriou
Leibniz Univ. Hannover, DE
- Rastin Pries
VDI/VDE Innovation + Technik GmbH – München, DE
- Ivica Rimac
Bell-Labs Alcatel-Lucent, DE
- Silvia Santini
TU Darmstadt, DE
- Nadi Sarrar
TU Berlin, DE
- Jonathan M. Smith
University of Pennsylvania, US
- Ralf Steinmetz
TU Darmstadt, DE
- Dominik Stingl
TU Darmstadt, DE
- Phuoc Tran-Gia
Universität Würzburg, DE
- Oliver P. Waldhorst
KIT – Karlsruhe Institute of Technology, DE
- Klaus Wehrle
RWTH Aachen, DE
- Michael Zink
University of Massachusetts – Amherst, US
- Martina Zitterbart
KIT – Karlsruhe Institute of Technology, DE

