

Report from Dagstuhl Seminar 17072

Applications of Topology to the Analysis of 1-Dimensional Objects

Edited by

Benjamin Burton¹, Maarten Löffler², Carola Wenk³, and Erin Moriarty Wolf Chambers⁴

1 The University of Queensland, AU, bab@maths.uq.edu.au

2 Utrecht University, NL, m.loffler@uu.nl

3 Tulane University, US, cwenk@tulane.edu

4 St. Louis University, US, echambe5@slu.edu

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 17072 “Applications of Topology to the Analysis of 1-Dimensional Objects”.

Seminar February 12–17, 2017 – <http://www.dagstuhl.de/17072>

1998 ACM Subject Classification I.3.5 Computational Geometry and Object Modeling

Keywords and phrases curves, graph drawing, homotopy, knot theory

Digital Object Identifier 10.4230/DagRep.7.2.64

Edited in cooperation with Hsien-Chih Chang

1 Executive Summary

Benjamin Burton

Maarten Löffler

Carola Wenk

Erin Moriarty Wolf Chambers

License  Creative Commons BY 3.0 Unported license

© Benjamin Burton, Maarten Löffler, Carola Wenk, and Erin Moriarty Wolf Chambers

Description of the seminar

One-dimensional objects embedded in higher-dimensional spaces are one of the most natural phenomena we encounter: ranging from DNA strands to roads to planetary orbits, they occur at all granularities throughout the sciences. Computer-assisted analysis of one-dimensional data is now standard procedure in many sciences; yet the underlying mathematics are not always well understood, preventing the most powerful analytical tools from being used.

Adding to the confusion, one-dimensional objects are studied under different names in different areas of mathematics and computer science (knots, curves, paths, traces, trajectories). In mathematics, 1-dimensional objects are well-understood, and research endeavors have moved on to higher dimensions. On the other hand, many fundamental applications demand solutions that deal with 1-dimensional objects, and these computational problems have largely been studied in separate communities by those unaware of all of the mathematical foundations.

The main goal of the proposed seminar was to identify connections and seed new research collaborations along the spectrum from knot theory and topology, to computational topology



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

Applications of Topology to the Analysis of 1-Dimensional Objects, *Dagstuhl Reports*, Vol. 7, Issue 2, pp. 64–88
Editors: Benjamin Burton, Maarten Löffler, Carola Wenk, and Erin Moriarty Wolf Chambers



DAGSTUHL
REPORTS

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

and computational geometry, all the way to graph drawing. Each of the invited speakers explored synergies in algorithms concerning 1-dimensional objects embedded in 2- and 3-dimensional spaces, as this is both the most fundamental setting in many applications, as well as the setting where the discrepancy in computational complexity between generic mathematical theory and potential algorithmic solutions is most apparent. In addition, each talk proposed a set of open questions from their research area that could benefit from attention from the other communities, and participants of the seminar were invited to propose their own research questions.

Below, we (the organizers) briefly describe the three main areas bridged; the abstracts of talks in the seminar and preliminary results from the working groups are also outlined later in this report.

Curves in Trajectory Analysis

Applications of computational topology are on the rise; examples include the analysis of GIS data, medical image analysis, graphics and image modeling, and many others. Despite how fundamental the question of topological equivalence is in mathematics, many of the relatively simple settings needed in computational settings (such as the plane or a 2-manifold) have been less examined in mathematics, where computability is known but optimizing algorithms in such “easy” settings has not been of interest until relatively recently.

Homotopy is one of the most fundamental problems to consider in a topological space, as this measure captures continuous deformation between objects. However, homotopy is notoriously difficult, as even deciding if two curves are homotopic is undecidable in a generic 2-complex. Nonetheless, many application settings provide restrictions that make computation more accessible. For example, most GIS applications return trajectories in a planar setting, at which point finding optimal homotopies (for some definition of optimal) becomes tractable.

Homology has been more recently pursued, as finding good homologies reduces to a linear algebra problem which can be solved efficiently. An example of this in the 1-dimensional setting is the recent work by Pokorný on clustering trajectories based on relative persistent homology. However, it is not always clear that optimal homologies provide as intuitive a notion for similarity measures compared with homotopy, and further investigations into applications settings is necessary.

Curves in Knot Theory

A fundamental question in 3-manifold topology is the problem of isotopy. Testing if two curves are ambiently isotopic is a foundational problem of *knot theory*: essentially, this asks whether two knots in 3-space are topologically equivalent. Problems in knot theory are tightly related to problems in 3-manifold topology, a field that has seen major breakthroughs in recent years, including Perelman’s 2002 solution to the geometrisation and Poincaré conjectures, and Agol’s recent 2012 proof of the virtual Haken conjecture. Algorithms and computation in these fields are now receiving significant attention from both mathematicians and computer scientists.

Complexity results are surprisingly difficult to come by. For example, one of the most fundamental and best-known problems is detecting whether a curve is knotted. This is known to be in both NP and co-NP; the former result was shown by Hass, Lagarias and Pippenger in 1999, but the latter was proven unconditionally by Lackenby just this year. Finding a polynomial time algorithm remains a major open problem. Hardness results are known for

some knot invariants, but (despite being widely expected) no hardness result is known for the general problem of testing two knots for equivalence. Techniques such as randomisation and parameterised complexity are now emerging as fruitful methods for understanding the inherent difficulty of these problems at a deeper level.

Algorithmically, many fundamental problems in knot theory are solved by translating to 3-manifold topology. Here there have been great strides in practical software in recent years: software packages such as *SnapPy* and *Regina* are now extremely effective in practice for moderate-sized problems, and have become core tools in the mathematical research process. Nevertheless, their underlying algorithms have significant limitations: *SnapPy* is based on numerical methods that can lead to numerical instability, and *Regina* is based on polytope algorithms that can suffer from combinatorial explosions. It is now a major question as to how to design algorithms for knots and 3-manifolds that are exact, implementable, and have provably viable worst-case analyses.

Curves in Graph Drawing

On the computer science end of the spectrum, the study of one-dimensional objects is closely related to Graph Drawing.

Graph Drawing studies the embedding of zero- and one-dimensional features (vertices and edges of graphs) into higher-dimensional spaces; both from an analytic (given an embedding, what can we say about it) and synthetic (come up with a good embedding) point of view. Computational questions (how can we embed a given graph such that it satisfies certain properties / optimises certain criteria) and fundamental questions (which classes of graphs admit which styles of embeddings) have been studied extensively, and a large body of algorithmic results is readily available.

Planarity (non-crossing edges) is a central theme in graph drawing. There is a rich literature discussing which graphs can be drawn planarly, when, and how, as well as how to avoid crossings or other undesirable features in a drawing, such as non-rational vertices. Traditionally, edges have always been embedded as straight line segments; however, there is a recent trend to consider different shapes and curves, drastically increasing the space of possible drawings of a graph. The potential benefits of this broader spectrum are obvious, but the effects (both computational and fundamental) are still ill understood.

Connections between graph drawing and knot theory have long been recognised, yet are still being actively explored. Already in 1983, Conway and Gordon showed that every spatial representation of K_7 contains at least one knotted Hamiltonian cycle. Based on this, in 2013, Politano and Rowland characterised which knots appear as Hamiltonian cycles in canonical book embeddings of complete graphs (as defined by Otsuki in 1996).

Goals and Results of this Seminar

Now is an exciting time for computational and algorithmic knot theory: practical algorithms are showing their potential through experimentation and computer-assisted proofs, and we are now seeing key breakthroughs in our understanding of the complex relationships between knot theory and computability and complexity theory. Early interactions between mathematicians and computer scientists in these areas have proven extremely fruitful, and as these interactions deepen it is hoped that major unsolved problems in the field will come within reach.

Similarly, applications for graph drawing and trajectory analysis are in great demand, especially given the rise of massive amounts of data through GIS systems, map analysis, and

many other application areas. However, despite the fact that many problems on curves are seen as mathematically trivial, there are few CS researchers who are truly familiar with the deeper topological results from mathematics. It is likely that many algorithmically interesting questions can benefit from an understanding of this rich history and toolset.

This seminar brought together a group of researchers from computer science and mathematics that study algorithms and mathematical properties of curves in various settings, as the interplay between these two groups is recent. In addition, we invited researchers in applications domains, who often do heuristic analysis of 1-dimensional objects in a variety of settings. Working groups were formed organically, but often allowed participants from various subfields to swap both open problems and favorite tools, and the overview talks discussed favorite tools and techniques from subdomains that may be useful to those in other areas. Concretely, we hope that in addition to the work begun in the working groups, many of these new collaborations will have positive long-term effects on all areas.

2 Table of Contents

Executive Summary

Benjamin Burton, Maarten Löffler, Carola Wenk, and Erin Moriarty Wolf Chambers 64

Overview of Talks

Geometric Realizations and Reconfigurations <i>Anna Lubiw</i>	69
Untangling Graphs and Curves on Surfaces via Local Moves <i>Hsien-Chih Chang</i>	69
Embeddings in 3-Space <i>Eric Sedgwick</i>	70
Telling 3-manifolds apart: new algorithms to compute Turaev-Viro invariants <i>Jonathan Spreer</i>	70
Persistent Cohomology and Circle-valued Coordinates <i>Mikael Vejdemo-Johansson</i>	71
Similarity Measures for Curves on Surfaces <i>Erin Moriarty Wolf Chambers</i>	71

Working groups

Trajectory Clustering <i>Anne Driemel, Maike Buchin, Brittany Terese Fasy, Florian T. Pokorny, Mikael Vejdemo-Johansson, and Carola Wenk</i>	71
Simplifying Curves on Surface via Local Moves <i>David Letscher, Gregory R. Chambers, Hsien-Chih Chang, Arnaud de Mesmay, Anne Driemel, Brittany Terese Fasy, Jessica S. Purcell, Saul Schleimer, Eric Sedgwick, Dylan Thurston, Stephan Tillmann, and Birgit Vogtenhuber</i>	72
Homotopy Height <i>Wouter Meulemans, Benjamin Burton, Tim Ophelders, Bettina Speckmann, Marc van Kreveld, and Erin Moriarty Wolf Chambers</i>	74
Convexifying Planar Drawings with Few Convexity-Increasing Linear Morphs <i>Lena Schlipf, Linda Kleist, Boris Klemz, Irina Kostitsyna, Anna Lubiw, Regina Rotman, Frank Staals, Darren Strash, and Carola Wenk</i>	76
The Pachner Graph of the Three-Sphere and Related Questions <i>Jonathan Spreer, Benjamin Burton, Arnaud de Mesmay, Parker Evans, Jessica S. Purcell, and Saul Schleimer</i>	79
Quadratic Genus with Linear Boundary <i>Dylan Thurston, Benjamin Burton, Maarten Löffler, Saul Schleimer, and Stephan Tillmann</i>	83
Lombardi Drawings of Knots and Links <i>Birgit Vogtenhuber, Philipp Kindermann, Stephen G. Kobourov, Maarten Löffler, Martin Nöllenburg, André Schulz, and Dylan Thurston</i>	84

Participants	88
------------------------	----

3 Overview of Talks

3.1 Geometric Realizations and Reconfigurations

Anna Lubiw (University of Waterloo, CA)

License © Creative Commons BY 3.0 Unported license
© Anna Lubiw

Main reference Soroush Alamdari, Patrizio Angelini, Fidel Barrera-Cruz, Timothy M. Chan, Giordano Da Lozzo, Giuseppe Di Battista, Fabrizio Frati, Penny Haxell, Anna Lubiw, Maurizio Patrignani, Vincenzo Roselli, Sahil Singla, Bryan T. Wilkinson, “How to Morph Planar Graph Drawings”, arXiv:1606.00425v1 [cs.CG], 2016.

URL <https://arxiv.org/abs/1606.00425>

Main results on drawing planar graphs deal with drawing edges as straight-line segments and restricting vertices to a small grid. I will discuss these issues for the problem of morphing (or “reconfiguring”) one drawing of a planar graph to another. This can be done – while preserving a straight-line planar drawing – by means of a sequence of $O(n)$ linear morphs, where a linear morph moves each vertex at uniform speed along a straight line. [“How to Morph Planar Graph Drawings”, to appear, SIAM J. Computing]. Restricting vertex positions (between the linear morphs) to a small grid is an open problem. Going beyond planarity to simultaneous planarity or intersection graphs of segments, we arrive at problems where finding a realization with straight line segments is complete for existential theory of the reals, and the reconfiguration space becomes disconnected.

3.2 Untangling Graphs and Curves on Surfaces via Local Moves

Hsien-Chih Chang (University of Illinois – Urbana-Champaign, US)

License © Creative Commons BY 3.0 Unported license
© Hsien-Chih Chang

Joint work of Hsien-Chih Chang, Jeff Erickson

Main reference Hsien-Chih Chang, Jeff Erickson, “Untangling planar curves”, in Proc. of the 32nd Int. Symp. on Computational Geometry (SoCG’16), LIPIcs, Vol. 51, pp. 29:1–29:15, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2016.

URL <http://dx.doi.org/10.4230/LIPIcs.SocG.2016.29>

Any continuous deformation of one closed curve to another on the same surface can be decomposed into a finite sequence of local transformations called homotopy moves. We are interested in the number of homotopy moves required to simplify a generic closed curve with n self-crossings as much as possible on an arbitrary surface. In the plane, an $O(n^2)$ upper bound is implicit in the classical work of Steinitz on polyhedra; a later result of Hass and Scott extended this upper bound to contractible curves on arbitrary surfaces.

Electrical transformations – the collection of degree-1 reductions, series-parallel reductions, and ΔY transformations – was studied intensively due to its use in optimization problems on planar graphs. Again we are interested in the number of electrical transformations required to reduce a plane graph with n vertices as much as possible. Using arguments of Noble and Welsh, we can relate the number of electrical transformations required to reduce a plane graph to the number of homotopy moves required to simplify its medial graph, viewed as curves embedded in the plane. A major open problem due to Feo and Provan is whether $O(n^{3/2})$ electrical transformations are always sufficient.

In this talk we will survey the results on these two closely related problems, including the three classical approaches in the plane, the $\Theta(n^{3/2})$ bound on the number of homotopy moves

required to simplify a plane curve, and a new result that simplifying a contractible curve in the annulus requires $\Omega(n^2)$ homotopy moves and its connection to the Feo and Provan conjecture.

This is a joint work with Jeff Erickson. Some of the results are published in our previous SoCG paper and its earlier preprint; the newer results can be found in our upcoming paper.

3.3 Embeddings in 3-Space

Eric Sedgwick (DePaul University – Chicago, US)

License  Creative Commons BY 3.0 Unported license
© Eric Sedgwick

The author, along with Matušek, Tancer and Wagner, showed that $\text{EMBED}_{2 \rightarrow 3}$, the problem of determining whether 2-complexes embed in 3-space is decidable. Here we discuss the obstacles, the intuition behind the solution, and the connection with Kirby diagrams, framed graphs embedded in 3-space that describe 3-manifolds. Finally, some open problems about embeddings and Kirby diagrams are stated.

3.4 Telling 3-manifolds apart: new algorithms to compute Turaev-Viro invariants

Jonathan Spreer (FU Berlin, DE)

License  Creative Commons BY 3.0 Unported license
© Jonathan Spreer

Joint work of Benjamin Burton, Clément Maria, Jonathan Spreer

Main reference Clément Maria, Jonathan Spreer, “A polynomial time algorithm to compute quantum invariants of 3-manifolds with bounded first Betti number”, in Proc. of the 28th Annual ACM-SIAM Symp. on Discrete Algorithms (SODA’17), pp. 2721–2732, SIAM, 2017.

URL <http://dx.doi.org/10.1137/1.9781611974782.180>

In low-dimensional topology, distinguishing between manifolds is a fundamental problem, which is remarkably difficult to solve in dimensions beyond two. As a result, topologists rely on simpler invariants to solve this task. In dimension three, the Turaev-Viro invariants are amongst the most powerful invariants, but traditional algorithms to compute them have prohibitive running times for numerous instances occurring in practice.

In this talk I present a fixed parameter tractable algorithm to compute one of these invariants in polynomial time for manifolds with bounded Betti number. Moreover, I discuss further ideas and approaches for new algorithms and applications.

References

- 1 Clément Maria, Jonathan Spreer, *A polynomial time algorithm to compute quantum invariants of 3-manifolds with bounded first Betti number*. Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA 2017), pages 2721–2732, 2017.
- 2 Clément Maria, Jonathan Spreer, *Admissible colourings of 3-manifold triangulations for Turaev-Viro type invariants*. Proceedings of the 24th Annual European Symposium on Algorithms (ESA 2016), 64:1–64:16.
- 3 Benjamin A. Burton, Clément Maria, Jonathan Spreer, *Algorithms and complexity for Turaev-Viro invariants*. Automata, Languages, and Programming: 42nd International Colloquium (ICALP 2015), Proceedings, Part 1, pp. 281–293.

3.5 Persistent Cohomology and Circle-valued Coordinates

Mikael Vejdemo-Johansson (CUNY College of Staten Island, US)

License  Creative Commons BY 3.0 Unported license
© Mikael Vejdemo-Johansson

We review the classical definition of simplicial homology and cohomology, demonstrate how the circle being the classifying space of $H^1(-; \mathbb{Z})$ produces equivalence classes of coordinate maps $[- \rightarrow S^1]$, and give examples from geometry, dynamics and motion capture time series.

3.6 Similarity Measures for Curves on Surfaces

Erin Moriarty Wolf Chambers (St. Louis University, US)

License  Creative Commons BY 3.0 Unported license
© Erin Moriarty Wolf Chambers

The question of how to measure similarity between curves in various settings has received much attention recently, motivated by applications in GIS data analysis, medical imaging, and computer graphics. While geometric measures such as the Hausdorff and Frechet distance have efficient algorithms, measures that take the underlying topology of the ambient space into account are less well understood. Several candidates have been proposed in recent years that are based on homotopy or homology, but many of these are only tractable in restricted settings, and surprisingly little is known about their practicality. In this talk, we will survey known results (both geometric and topological), and then focus on some of the recent algorithmic results and remaining open questions for the topological measures.

4 Working groups


4.1 Trajectory Clustering

Anne Driemel (TU Eindhoven, NL), Maike Buchin (Ruhr-Universität Bochum, DE), Brittany Terese Fasy (Montana State University – Bozeman, US), Florian T. Pokorny (KTH Royal Institute of Technology – Stockholm, SE), Mikael Vejdemo-Johansson (CUNY College of Staten Island, US), and Carola Wenk (Tulane University, US)

License  Creative Commons BY 3.0 Unported license
© Anne Driemel, Maike Buchin, Brittany Terese Fasy, Florian T. Pokorny, Mikael Vejdemo-Johansson, and Carola Wenk

The starting point of our discussion was the intention to find a mathematically founded definition of clustering for curves which can be computed efficiently. Many clustering formulations are based on the definition of a centroid or median. This lead us to re-thinking the definition of a median for a set of curves. We observed that one can extend Tukey's definition of a median to our setting as follows. Recall that the Tukey depth of a point p in a finite set of points $P \subset \mathbb{R}^d$ is defined as

$$\min_{\substack{h \in H \\ p \in h}} |P \cap h|,$$



■ **Figure 1** 4 line segments (black) contained in a bisector represented by 2 line segments (blue).

where H is the set of half-spaces in R^d . The Tukey median is then defined as the point with largest depth. We observe that a half-space can be represented using the bisector of two points, which is the set of points that are equidistant to two fixed points a and b . This notion of half-space partition of P naturally extends to any distance metric defined on curves. To solve the above optimization problem we are thus interested in all 2-cell Voronoi partitions of P , where the Voronoi partitions are formed under a certain distance measure. We can represent a bisector by two curves a and b which we call *bisector representatives*. Given a bisector in this implicit way, we can compute the Voronoi partition by simply computing the distances of all points in P to the representatives a and b . In fact the count of points on both sides of the bisector can be estimated very efficiently by using random sampling on P . Most importantly, these computations can be done without computing the bisector explicitly. To solve the optimization problem we initially focused on the special case of line segments in R^2 and the Fréchet distance to measure distances between line segments. We think that one can compute all 2-cell Voronoi partitions by finding bisectors that contain subsets of points from P . In particular we believe that it is sufficient to consider subsets of either 4 or 5 line segments, and for each configuration it is sufficient to compute a constant number of bisector representatives. This directly implies a polynomial time algorithm for determining the median of a set of line segments. Furthermore we investigated geodesics in this space and how to project a curve onto its closest point on a given bisector.

4.2 Simplifying Curves on Surface via Local Moves

David Letscher (St. Louis University, US), Gregory R. Chambers (University of Chicago, US), Hsien-Chih Chang (University of Illinois – Urbana-Champaign, US), Arnaud de Mesmay (University of Grenoble, FR), Anne Driemel (TU Eindhoven, NL), Brittany Terese Fasy (Montana State University – Bozeman, US), Jessica S. Purcell (Monash University – Clayton, AU), Saul Schleimer (University of Warwick – Coventry, GB), Eric Sedgwick (DePaul University – Chicago, US), Dylan Thurston (Indiana University – Bloomington, US), Stephan Tillmann (University of Sydney, AU), and Birgit Vogtenhuber (TU Graz, AT)

License © Creative Commons BY 3.0 Unported license

© David Letscher, Gregory R. Chambers, Hsien-Chih Chang, Arnaud de Mesmay, Anne Driemel, Brittany Terese Fasy, Jessica S. Purcell, Saul Schleimer, Eric Sedgwick, Dylan Thurston, Stephan Tillmann, and Birgit Vogtenhuber

A total of 12 seminar participants participated in discussions about the following question presented at the open problem session: “How many homotopy moves are needed to reduce a

generic closed curve on a surface to have minimal number of self-intersections?”



Homotopy moves $1 \rightarrow 0$, $2 \rightarrow 0$, and $3 \rightarrow 3$.


Any closed curve on a surface can be *reduced* to another one that has minimum number of self-intersections within its homotopy class. Hass and Scott [5] showed that such curves can be reduced through a finite sequence of local transformations called *homotopy moves*. For generic curves in the plane with n self-intersections, a proof that $O(n^2)$ moves are always sufficient to reduce the curve is implicit in Steinitz’s proof that every 3-connected planar graph is the 1-skeleton of a convex polyhedron [6, 7]. Specifically, Steinitz proved that any non-simple closed curve with no empty loops contains a minimal *bigon* which can be reduced and removed by a sequence of homotopy moves. This upper bound was later improved to $O(n^{3/2})$ by Chang and Erickson, which is the best possible in the worst case [1]. For curves on the annulus, de Graaf and Schrijver [3] showed that $O(n^2)$ moves are always sufficient. Chang and Erickson [2] found a matching lower bound (which extends to curves in any higher genus surfaces).

Our group focused on finding upper bounds for the problem on various surfaces. We started with the torus and were able to show that $O(n^2)$ moves are again sufficient. Using similar techniques we were then able to prove the same quadratic upper bounds on the Möbius strip, the Klein bottle, and in the projective plane. In the final few meetings we considered the case of a curve with n self-intersections in a orientable surface of genus g . Using techniques from combinatorial group theory, we showed that there is an $O(g^2n^3)$ upper bound on number of homotopy moves required to reduce the given curve. Our main technical contribution is to extend Steinitz’s bigon reduction technique to *singular bigons*, whose existence is guaranteed by an earlier result of Hass and Scott’s [4].


A subset of participates of this group plans to write up the results and submit them to an appropriate conference or journal.

References

- 1 Hsien-Chih Chang and Jeff Erickson. Untangling planar curves. *Proc. 32nd Int. Symp. Comput. Geom.*, 29:1–29:15, 2016. Leibniz International Proceedings in Informatics 51.
- 2 Hsien-Chih Chang and Jeff Erickson. Unwinding annular curves and electrically reducing planar networks. Preprint, 2017.
- 3 Maurits de Graaf and Alexander Schrijver. Making curves minimally crossing by Reidemeister moves. *J. Comb. Theory Ser. B* 70(1):134–156, 1997.
- 4 Joel Hass and Peter Scott. Intersections of curves on surfaces. *Israel J. Math.* 51:90–120, 1985.
- 5 Joel Hass and Peter Scott. Shortening curves on surfaces. *Topology* 33(1):25–43, 1994.
- 6 Ernst Steinitz. Polyeder und Raumeinteilungen. *Enzyklopädie der mathematischen Wissenschaften mit Einschluss ihrer Anwendungen* III.AB(12):1–139, 1916.
- 7 Ernst Steinitz and Hans Rademacher. *Vorlesungen über die Theorie der Polyeder: unter Einschluss der Elemente der Topologie*. Grundlehren der mathematischen Wissenschaften 41. Springer-Verlag, 1934. Reprinted 1976.



■ **Figure 2** The Euclidean setting: two boundary paths (black), a start and end leash (s and t), and point obstacles (squares) with cost indicated by size. Two intermediate leashes are indicated.



■ **Figure 3** Special case with collapsed boundary paths.

4.3 Homotopy Height


Wouter Meulemans (TU Eindhoven, NL), Benjamin Burton (The University of Queensland, AU), Tim Ophelders (Eindhoven Univ. of Technology, NL), Bettina Speckmann (TU Eindhoven, NL), Marc van Kreveld (Utrecht University, NL), and Erin Moriarty Wolf Chambers (St. Louis University, US)

License © Creative Commons BY 3.0 Unported license
 © Wouter Meulemans, Benjamin Burton, Tim Ophelders, Bettina Speckmann, Marc van Kreveld, and Erin Moriarty Wolf Chambers


We initially considered the following problem, as posed in the open problem session: Given a triangulation of a disk, we wish to find the best way to move the left side of the boundary of the disk to the right side of the boundary, via flips across the faces or spikes along the edges. This problem is known as the *homotopy height problem*, as the sweep encodes a homotopy across the disk. One can also consider shrinking the boundary of the disk to a point, or the variant where one sweeps across an annulus, moving the outer boundary to the inner one step at a time.

In an effort to focus on a situation which allows for more geometric intuition, we studied the Euclidean case, in which we have a polygonal boundary enclosing a planar domain containing point obstacles (Figure 2). In this setting the boundary is divided into four parts: two *leashes* interleaved with two *boundary paths*. Our goal is now to continuously transform the starting leash s into the target leash t via a *homotopy*, in which the endpoints of the intermediate leashes travel along the boundary paths. The *cost* of a leash is its length plus an additional cost per obstacle it encounters; we refer to this additional cost as the *weight* of the obstacle. Our goal is to find a homotopy that minimizes the cost of the most expensive intermediate leash.

It follows from previous work that there is actually an optimal homotopy that is an *isotopy*. Furthermore, there is such an optimal isotopy in which the leashes move monotonically. We consider two different scenarios: all obstacles have the same (unit) weight, or obstacles can have different weights.



■ **Figure 4** (a) A leash may need many inflection points in the variable-weight case. (b) The same principle in the convex case, with one boundary path and the two initial leashes collapsed to a point. Closest point and corresponding distance circles indicated for each obstacle.



■ **Figure 5** Special case with collapsed s and t , and unit-weight obstacles.

Variable-weight obstacles

If the two boundary paths collapse to points (Figure 3), then we can compute an optimal homotopy in polynomial time using a simple greedy strategy. Our result builds on the observation that in this setting the leashes do not have inflection points which are not induced by the boundary.

In the more general case, the leashes have no simple characterization anymore: specifically, any optimal homotopy might require a leash with linearly many inflection points. For this, consider the example in Figure 4(a). Here the weights are decreasing in the upward direction, to ensure that the best position of going over an obstacle is when the leash is as short as possible, only wrapping around the lower points. These best positions alternate between the first (left) half and the second (right) half of the boundary path, to create inflection points. This same principle can even be applied if the boundary is convex, only one boundary path is not a point, and both s and t are points (see Figure 4(b)). This leads us to conjecture that the problem is NP-Hard in this most general setting.

Unit-weight obstacles

For unit-weight obstacles we can compute an optimal homotopy in polynomial time, for the general case. Our results build on the same greedy strategy as in the variable-weight case.

For unit-weight obstacles, the leash does not need more than one inflection point that is not caused by the boundary, and this only in particular situations. Combining these observations with the monotonicity of an optimal homotopy allows us to solve this problem via dynamic programming.

We also found that some cases can be solved more efficiently. For example, if s and t collapse to a point (Figure 5), and the boundary paths together form a convex polygon, we need to find only an antipodal pair that splits the problem into two simpler ones: how do we shrink the leash at this antipodal pair to s (and analogously, to t)? This again follows the greedy strategy, and afterwards we only need to glue the two homotopies together, which may require at most one inflection point on the leash.

4.4 Convexifying Planar Drawings with Few Convexity-Increasing Linear Morphs

Lena Schlipf (FernUniversität in Hagen, DE), Linda Kleist (TU Berlin, DE), Boris Klemz (FU Berlin, DE), Irina Kostitsyna (Free University of Brussels, BE), Anna Lubiw (University of Waterloo, CA), Regina Rotman (University of Toronto, CA), Frank Staals (Aarhus University, DK), Darren Strash (Colgate University – Hamilton, US), and Carola Wenk (Tulane University, US)

License © Creative Commons BY 3.0 Unported license
© Lena Schlipf, Linda Kleist, Boris Klemz, Irina Kostitsyna, Anna Lubiw, Regina Rotman, Frank Staals, Darren Strash, and Carola Wenk

Abstract

We study the problem of *convexifying* drawings of planar graphs. Given a planar straight-line drawing of a graph G , we wish to *morph* the drawing to a planar straight-line drawing of G in which all faces are convex, while maintaining planarity at all times. Furthermore, we want the morph to be *convexity-increasing*, meaning that the set of convex angles in the drawing never decreases. We give a polynomial time algorithm to construct such a morph.

Problem Definition and Background

A *morph* between two planar straight-line drawings Γ_0 and Γ_1 of a graph G is a continuous movement of the vertices from one to the other, with the straight-line edges determined by the vertex positions. If each vertex moves along a straight line at uniform speed, the morph is called *linear*. (Note that different vertices may move at different speeds, and some may remain stationary.) If, in addition, all the lines along which vertices move are parallel, then the morph is called *unidirectional*. A morph is *planar* if it preserves planarity of the drawing. Alamdari et al. [2] gave a polynomial time algorithm to find a planar morph of Γ_0 to Γ_1 using a sequence of $O(n)$ unidirectional morphs, where n is the number of vertices of the graph. One disadvantage of this algorithm is that vertices become almost coincident and there is no bound on the number of bits required for the vertex coordinates in the $O(n)$ intermediate drawings between the unidirectional morphs. Ideally, we would hope for the intermediate drawings to lie on a polynomial-sized grid (i.e. with a logarithmic number of bits for vertex coordinates). A weaker but still open question is to find planar morphs that can be specified with polynomially-bounded space (measuring bit complexity).

We say that a planar morph *convexifies* a given straight-line graph drawing if the end result is a convex graph drawing, i.e. a straight-line graph drawing in which the angles of

all internal faces are convex and the angles of the external face are reflex. Note that we do not require strict convexity – we allow angles of 180° . We say that a planar morph is *convexity-increasing* if the set of face angles that are internal and convex, or external and reflex, never decreases, i.e. the progress towards a convex drawing, as measured by the number of face angles violating the convexity condition, is non-decreasing.

Throughout, we assume that our input graph has a convex drawing with the same faces and the same outer face as the input drawing. Necessary and sufficient conditions for the existence of such a convex drawing were given by Tutte [10], Thomassen [9], and Hong and Nagamochi [7]. These conditions can be tested in linear time by the algorithm of Chiba et al. [4]. Such conditions are usually stated for a fixed convex drawing of the outer face, but the conditions become simpler when, as in our case, the drawing of the outer face may be chosen to have no 3 consecutive collinear vertices. The conditions simplify further when no internal vertex has degree 2, and this can be assumed without loss of generality since an internal vertex of degree 2 must be drawn as a point in the interior of the straight line segment formed by its two incident edges. With these simplifications, the result stated by Hong and Nagamochi [7] is that a plane graph G with outer face C and with no internal degree-2 vertex has a convex drawing with outer face C if and only if the graph is *internally 3-connected*, i.e., the graph is 2-connected and any pair of cut vertices u, v has the properties that u and v lie on the outer face and every connected component of $G - \{u, v\}$ has a vertex of the outer face.

The algorithm of Alamdari et al. (or any other algorithm to morph graph drawings) can be used to convexify a given planar graph drawing Γ_0 , since we can construct a convex drawing Γ_1 of the graph and morph Γ_0 to Γ_1 . However, all known morphing algorithms triangulate the drawing, and hence will fail to be convexity-increasing in general¹. Furthermore, morphing to *some* convex drawing is a weaker condition than morphing to a *particular* convex drawing, and may give us more freedom to keep to a small grid.

It is an open question to find convexity-increasing morphs. In the special case when the graph consists of a single cycle the problem is solved by the result of Aichholzer et al. [1] that morphs a polygon to a convex polygon without losing any visibilities between pairs of vertices.

Progress and Preliminary Results

At the seminar we outlined an algorithm to convexify a given planar graph drawing via a convexity-increasing morph that consists of $O(n)$ unidirectional morphs. Furthermore, each unidirectional morph moves vertices in either the horizontal or vertical direction, which means that the trajectory of each vertex during the complete morph is a path consisting of horizontal and vertical segments.

We will first discuss the reason why we concentrate on unidirectional morphs. After that we discuss the main idea of our algorithm.

Restricting to linear morphs seems like a sensible way to discretize morphs – essentially, it asks for the vertex trajectories to be piece-wise linear. At first glance, the restriction to unidirectional morphs, seems arbitrary and restrictive. However, it turns out to be easier to prove the existence of unidirectional morphs, for the following reason. Suppose we do a unidirectional morph in the direction parallel to the x -axis. Then every vertex must keep its

¹ We note that there is an algorithm to morph one convex drawing to another [3], which does not triangulate the graph, but this does not solve the problem of convexifying a non-convex drawing.

y -coordinate. This simplifies the planarity requirements because Lemma 13 of the paper by Alamdari et al. [2] states that the linear morph between two planar straight-line drawings Γ_1 and Γ_2 is planar if every line parallel to the x -axis crosses the same set of edges and vertices in the same order in both drawings. Note that this condition requires in particular that every vertex is at the same y -coordinate in both drawings, and the condition implies that the morph is unidirectional.

This means that, after committing to a direction for a unidirectional morph, we are free to choose a new drawing that satisfies the above conditions – and a planar unidirectional morph is guaranteed.

We use an existing algorithm to create new drawings. Planar straight-line drawings with vertices at fixed y -coordinates are called *layered drawings of hierarchical graphs*. Hong and Nagamochi [7] gave an algorithm to construct a convex layered drawing of any *hierarchical st-graph* – a hierarchical graph in which the boundary of every face consists of two upward chains.

Our algorithm proceeds by a sequence of steps where each step is as follows: choose the horizontal or vertical direction; augment the graph to a hierarchical *st-graph*; appeal to Hong and Nagamochi to produce a new convex layered drawing of this augmented graph; and perform a linear morph to the new drawing. A horizontal step will convexify any reflex angle formed by three vertices whose y -coordinates are increasing. Similarly, a vertical step will convexify any reflex angle formed by three vertices whose x -coordinates are increasing. No step will make a convex angle reflex. We may need to apply a shear transformation (which is a unidirectional morph) before each step in order to guarantee that there is at least one reflex angle whose vertices have increasing x - or y -coordinates.

Our main contribution is the following theorem.

► **Theorem 1.** *Any planar straight-line drawing of an internally 3-connected graph can be convexified via a convexity-increasing morph that consists of $O(b)$ unidirectional morphs, where b is the number of face angles that are internal and reflex or external and convex, and each unidirectional morph moves vertices in the horizontal or vertical direction. Furthermore, there is a polynomial-time algorithm to find such a morph.*

We have a family of examples, based on those of Alamdari et al. [2], to show that $\Omega(b)$ is a lower bound on the number of linear morphs that may be required to convexify a straight-line planar graph drawing.

Open Questions and Future Work

Although our algorithm improves on the general morphing algorithm [2] in that we do not use the technique of “almost” contracting vertices, still, we do not seem to have a polynomial bound on the bit-complexity of the intermediate drawings of our morph. We can design a family of examples, based on those of Lin and Eades [8] to show that one of our unidirectional morphs may unavoidably blow up the the width of the drawing from $O(n)$ to $\Omega(n!)$. This is still polynomial bit complexity, but the danger is that successive steps might cause exponential growth. In an attempt to get better bounds on the grid size/bit complexity of the intermediate drawings of the morph, we tried replacing Hong and Nagamochi’s algorithm to find convex layered drawings by using an extension of Tutte’s planar graph drawing algorithm [11] to more general edge weights ([5] or see [6]), and choosing the edge weights to place vertices at the desired layers. We plan to explore this issue further, but so far it remains an open question to achieve polynomially-bounded bit complexity.

Acknowledgments. We wish to thank André Schulz for helpful discussions on generalizations of Tutte’s algorithm.

References

- 1 O. Aichholzer, G. Aloupis, E. D. Demaine, M. L. Demaine, V. Dujmovic, F. Hurtado, A. Lubiw, G. Rote, A. Schulz, D. L. Souvaine, and A. Winslow. Convexifying polygons without losing visibilities. In *CCCG*, 2011.
- 2 S. Alamdari, P. Angelini, F. Barrera-Cruz, T. M. Chan, G. Da Lozzo, G. Di Battista, F. Frati, P. Haxell, A. Lubiw, M. Patrignani, V. Roselli, S. Singla, and B. T. Wilkinson. How to morph planar graph drawings. *arXiv:1606.00425*, 2016. to appear in *SIAM J. Computing*.
- 3 P. Angelini, G. Da Lozzo, F. Frati, A. Lubiw, M. Patrignani, and V. Roselli. Optimal morphs of convex drawings. In L. Arge and J. Pach, editors, *Proceedings of the 31st International Symposium on Computational Geometry (SoCG 2015)*, volume 34 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 126–140, Dagstuhl, Germany, 2015. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- 4 N. Chiba, K. Onoguchi, and T. Nishizeki. Drawing plane graphs nicely. *Acta Informatica*, 22(2):187–201, 1985.
- 5 M. S. Floater. Parametric tilings and scattered data approximation. *International Journal of Shape Modeling*, 4(03n04):165–182, 1998.
- 6 S. J. Gortler, C. Gotsman, and D. Thurston. Discrete one-forms on meshes and applications to 3D mesh parameterization. *Computer Aided Geometric Design*, 23(2):83–112, 2006.
- 7 S.-H. Hong and H. Nagamochi. Convex drawings of hierarchical planar graphs and clustered planar graphs. *Journal of Discrete Algorithms*, 8(3):282–295, 2010.
- 8 X. Lin and P. Eades. Towards area requirements for drawing hierarchically planar graphs. *Theoretical computer science*, 292(3):679–695, 2003.
- 9 C. Thomassen. Plane representations of graphs. In J. A. Bondy and U. S. R. Murty, editors, *Progress in Graph Theory*, pages 43–69. Academic Press, 1984.
- 10 W. T. Tutte. Convex representations of graphs. *Proceedings of the London Mathematical Society*, 3(1):304–320, 1960.
- 11 W. T. Tutte. How to draw a graph. *Proceedings of the London Mathematical Society*, 3(1):743–767, 1963.

4.5 The Pachner Graph of the Three-Sphere and Related Questions


Jonathan Spreer (FU Berlin, DE), Benjamin Burton (The University of Queensland, AU), Arnaud de Mesmay (University of Grenoble, FR), Parker Evans (Tulane University – New Orleans, US), Jessica S. Purcell (Monash University – Clayton, AU), and Saul Schleimer (University of Warwick – Coventry, GB)

License © Creative Commons BY 3.0 Unported license
 © Jonathan Spreer, Benjamin Burton, Arnaud de Mesmay, Parker Evans, Jessica S. Purcell, and Saul Schleimer

Background

It is known that any two triangulations of a surface with the same number of vertices can be connected by a sequence of *diagonal flips*, or *2-2 moves*, given by exchanging the diagonal of a quadrilateral, as shown on the left of Figure 6. For example, see [1].

Similarly, any closed orientable 3-manifold can be decomposed into tetrahedra, and any two such decompositions with the same number of vertices can be related by a sequence of 3-dimensional flips, namely a *Pachner moves*, or 2-3 (3-2) moves, as shown on the right of



■ **Figure 6** Left: A 2-dimensional diagonal flip. Right: A 3-dimensional flip, called a 2-3 move, or Pachner move. Two tetrahedra become three.

Figure 6. This is shown by Matveev in [6]. [**Note:** There are additional moves that are also called Pachner moves, such as the 4-1 and 1-4 moves, but these create and destroy vertices, so we ignore these for the purposes of this report.]

For a given 3-manifold and a given number of vertices, we build a graph related to these moves called the *Pachner graph*. Each vertex corresponds to a triangulation of the 3-manifold, and two vertices are connected by an edge if there is a single Pachner move changing one to the other. Matveev’s result implies that the graph is connected.

Questions

The motivating major question is: what is the “shape” of the Pachner graph? Given two triangulations of a manifold with the same number of vertices, what is the shortest path between them? What is the shortest path to a “canonical” triangulation of the manifold?

These questions are hard, and wide open. For this project, we restrict to the 3-sphere S^3 , and (typically) restrict to triangulations with a single vertex. Let T denote such a triangulation: a one-vertex triangulations of S^3 , with n tetrahedra ($n > 1$).

Upper bounds

- There is a sequence of papers by Mijatovic (starting in 2003) which uses normal and almost normal surface theory to show that any triangulation of S^3 is connected to a constant size standard triangulation by a sequence of at most $A \cdot e^{Bn^2}$ moves [7, 8]. The bound comes from the upper bounds on the complexity of normal and almost normal spheres in T . Thus one might be able to find better upper bounds by finding smaller normal 2-spheres. However, there are examples of triangulations in which the smallest 2-sphere has exponentially large complexity (cf the examples of [5] and [4]).
- A natural tactic is to restrict to triangulations with nice properties. For example, the manifold could have a small Cheeger constant, or the dual graph to T could have low tree width. In this case, we can find a separator of small size. It is still unclear how to use such small separators to simplify triangulations.
- A final observation – much of the discussion here is similar to the story of Reidemeister simplification of diagrams of the unknot U . In that area there has been recent progress, due to Lackenby [3]. For any n crossing diagram of U there is a sequence of Reidemeister moves of length at most $O(n^{11})$ taking it to the trivial diagram. Can we use ideas similar to Lackenby’s proof (using the combinatorics of a foliated spanning disk) to generalise to the 3-sphere? Or perhaps even the solid torus, where there is a simpler foliation?

So far we have the following possible correspondences:

Lackenby	Proposed 1	Proposed 2
unknot diagram	trig. of S^3	trig. of solid torus
unknot	1-skeleton	1-skeleton or μ
spanning disk	2-skeleton	2-skeleton or merid. disk
core curve	\emptyset	core curve
foliation by pages / book	foliation by spheres	foliation by annuli

Idea 1: Use Lackenby’s short core curves [2] and annulus foliations to simplify triangulations of solid tori. [Aside: What is the algorithmic complexity of finding the short core curve?]

Idea 2: Take Lackenby’s setup from [3] to prove Lackenby’s theorem about core curves in [2], which Lackenby proves using other techniques.

Further questions:

1. **Unknotted edges:**

- Does T having an unknotted edge help? Let T be a one-vertex triangulation of a 3-sphere with an unknotted edge. Is there a polynomial time algorithm to simplify T to a smaller triangulation? Is there a polynomial time algorithm to simplify T to a smaller triangulation T' , where T' has an unknotted edge?
- When does T have an unknotted edge? Ben Burton reports that every one-vertex triangulation of S^3 , with at most nine tetrahedra, has an unknotted edge. The unknotted edge is not necessarily the highest valence edge (as we had expected).
- Related: how do $2 - 3$ and $3 - 2$ moves change the complexity of the edges of T as knots in S^3 (for various notions of complexity)? Can we build a triangulation where edges are all knotted? Where they have arbitrarily high knot complexity (for various notions of complexity)?

2. **Tree width one:** Can we simplify 3-sphere triangulations with tree-width one in polynomial (quadratic) time?

In the tree-width one case, Burton’s thesis tells us that we either have a layered triangulation (which trivially can be simplified) or a “hat” (two triangles identified along two faces) with tetrahedra inside and outside. The latter complex can be simplified by a thickening-pulling-flattening move. Can we express this move in a sequence of Pachner moves?

3. **Other:**

- Can we simplify locally constructible (collapsible) 3-sphere triangulations in polynomial time?
- Let T be a 3-sphere with optimal Morse function with $\leq k$ critical faces. Can we simplify T in $O(f(k)n^{O(1)})$ ($O(f(k)e^{O(n)})$) time?

Lower bounds

Given a one vertex triangulation of S^3 , can we find a lower bound on the number of moves required to simplify the triangulation? So far, the largest known examples take $n + 2$ moves to simplify. Can we find something that requires more moves? We considered several families of triangulations of S^3 .

1. Gluing a pair of Fibonacci layered tori, following ideas of Letscher [5]. These triangulations have several interesting properties, e.g. low tree-width, lots of very small separators, but *no* small two-spheres (even almost normal ones). However, these examples simplify directly (n moves). Interestingly there seems to be only one location for the simplification

to take place, namely between the layered tori. That is, there is some concentrated positive curvature.


2. Lackenby and Souto constructed “expansive spheres” [4]. These do not seem to appear in the literature, but we worked through some details of the construction, involving expansive Cayley graphs of the groups $\mathrm{PSL}(2, p^n)$. These examples do not have small tree-width and so do not have small separators of any kind. However, again these still simplify directly – there is concentrated positive curvature along the doubling sphere. This again seems to be the only place to do Pachner moves.
3. If a triangulation has only edges of degree five and higher, then there is no 3 – 2-move possible, so we must increase the number of tetrahedra (at least twice!) before decreasing. For example the 600 cell has this property. [And leads to the minimal triangulation of the Poincaré homology sphere Σ^3 .] Here the curvature seems to be well distributed. Of course, the triangulation has 120 vertices, not one. Note that Regina and Snappy both reduce this triangulation to bounded size (two tetrahedra) immediately.
We did a search for one-vertex triangulations with only high valence edges, and found several with all edges of valence four and higher among the census with seven, eight, and nine tetrahedra. However, we found no triangulations of S^3 with edges of valence five and higher. We don’t know of any obstruction to the existence of such triangulations.
4. Can we count the number of one-vertex, n -tetrahedra triangulations of S^3 ? Can we prove that the number of such triangulations grows super-exponentially? Does this imply lower bounds?

Other

1. Instead of studying the Pachner graph of 1-vertex 3-sphere triangulations, look at their vertex links and how they change. This will give us a subset of the vertex set of the Pachner graph of 2-sphere triangulations. How sparse is this subset? What can we say for the edges in this complex?
2. The flip graph of n -gon triangulations is isomorphic to the 1-skeleton of the associahedron. Is there a similar polytope (or related object) for the Pachner graph of S^3 (M^3 , etc.)?

References

- 1 Allen Hatcher, *On triangulations of surfaces*, Topology Appl. **40** (1991), no. 2, 189–194. MR 1123262
- 2 Marc Lackenby, *Core curves of triangulated solid tori*, Trans. Amer. Math. Soc. **366** (2014), no. 11, 6027–6050. MR 3256192
- 3 Marc Lackenby, *A polynomial upper bound on Reidemeister moves*, Ann. of Math. (2) **182** (2015), no. 2, 491–564. MR 3418524
- 4 Marc Lackenby and Juan Souto, *Expander knots and metrics*, Unpublished.
- 5 David Letscher, *Fibonacci triangulations of s^3* , personal communication.
- 6 Sergei Matveev, *Algorithmic topology and classification of 3-manifolds*, second ed., Algorithms and Computation in Mathematics, vol. 9, Springer, Berlin, 2007. MR 2341532
- 7 Aleksandar Mijatović, *Simplifying triangulations of S^3* , Pacific J. Math. **208** (2003), no. 2, 291–324. MR 1971667
- 8 Aleksandar Mijatović, *Triangulations of Seifert fibred manifolds*, Math. Ann. **330** (2004), no. 2, 235–273. MR 2089425



■ **Figure 7** A parking garage, shown for $k = 3$ and $l = 3$. Edges with corresponding marking are glued; this can be achieved by stacking the sheets in 3 dimensions and attaching connecting ramps.

4.6 Quadratic Genus with Linear Boundary

Dylan Thurston (Indiana University – Bloomington, US), Benjamin Burton (The University of Queensland, AU), Maarten Löffler (Utrecht University, NL), Saul Schleimer (University of Warwick – Coventry, GB), and Stephan Tillmann (University of Sydney, AU)

License © Creative Commons BY 3.0 Unported license
 © Dylan Thurston, Benjamin Burton, Maarten Löffler, Saul Schleimer, and Stephan Tillmann

In the course of analyzing multi-level motion planning [3, 1, 2], the following problem naturally arises.

► **Question 1.** *Suppose you are given a surface Σ smoothly embedded in \mathbb{R}^3 so that the vertical projection of Σ to \mathbb{R}^2 is an immersion with polygonal boundary made of m line segments. (That is, the vertical direction is transverse to Σ .) Is $\text{genus}(\Sigma)$ bounded by a linear function of m ?*

We answered Question 1 in the negative; the best bound is quadratic. It turns out that the restriction to surfaces that are embedded in \mathbb{R}^3 is inessential.

► **Theorem 2.** *Let Σ be a surface with boundary and let $f: \Sigma \rightarrow \mathbb{R}^2$ be an immersion on the interior of Σ so that $f(\partial\Sigma)$ is a polygonal path with m line segments. Then $\text{genus}(\Sigma) \leq m(m+1)/4$. Furthermore, there are examples coming from embeddings in \mathbb{R}^3 with $\text{genus}(\Sigma) = (m/8 - 1)^2$.*

The examples achieving quadratic genus growth are “parking garages” $P_{k,l}$, as shown in Figure 7:

- take k parallel rectangular sheets;
- cut out l slits from each sheet (stacked on top of each other); and
- rejoin across the slits, shifting down one level as you go.

We can apply the Gauss–Bonnet theorem to the metric on Σ coming from the map to \mathbb{R}^2 . We use the special case when the curvature vanishes on the interior and the curvature of the boundary is zero except at the polygonal corners.

► **Theorem 3 (Gauss–Bonnet, flat version).** *Let Σ be a surface with a locally Euclidean metric and polygonal boundary, with corners at c_i with interior angle θ_i . Then the Euler characteristic of Σ is*

$$\chi(\Sigma) = 2 - 2\text{genus}(\Sigma) - \#\partial\Sigma = \frac{1}{2\pi} \sum_i (\pi - \theta_i),$$

where $\#\partial\Sigma$ is the number of components in the boundary of Σ .

Here, $\pi - \theta_i$ should be thought of as the bending angle at c_i : zero if there is no actual corner, positive if the corner is convex as on the boundary of a convex polygon in the plane, and negative if the corner is concave. Some of the corners in $P_{k,l}$ are very concave, with a total internal angle of approximately $2l\pi$. The result of this computation is that $\text{genus}(P_{k,l}) = (k-1)(l-1)$. Furthermore, $P_{k,l}$ can be realized with a polygonal boundary with $4k + 4l$ corners.

For the upper bounds on genus, we again apply Theorem 3 and give an upper bound on the interior angles θ_i . To do this, we first bound the total multiplicity in any region, the degree by which it is covered by Σ . The multiplicity at a point $x \in \mathbb{R}^2$ can be computed by sending a ray out to infinity in either direction from x , and so is at most $m/2$. The angle θ_i at a corner c_i is bounded by 2π times the multiplicity in any adjoining region. This yields the stated upper bound on genus.

References

- 1 A. Hillebrand, M. van den Akker, R. Geraerts, and H. Hoogeveen, *Performing multicut on walkable environments*, Combinatorial Optimization and Applications. COCOA 2016 (Cham) (T-H. Chan, M. Wang, and L. Wang, eds.), Lecture Notes in Computer Science, vol. 10043, Springer, 2016, pp. 311–325.
- 2 A. Hillebrand, M. van den Akker, R. Geraerts, and H. Hoogeveen, *Separating a walkable environment into layers*, Proceedings of the 9th International Conference on Motion in Games (New York, NY, USA), MIG'16, ACM, 2016, pp. 101–106.
- 3 W. van Toll, A.F. Cook, and R. Geraerts, *Navigation meshes for realistic multi-layered environments*, 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems, Sept 2011, pp. 3526–3532.

4.7 Lombardi Drawings of Knots and Links

Birgit Vogtenhuber (TU Graz, AT), Philipp Kindermann (FernUniversität in Hagen, DE), Stephen G. Kobourov (University of Arizona – Tucson, US), Maarten Löffler (Utrecht University, NL), Martin Nöllenburg (TU Wien, AT), André Schulz (FernUniversität in Hagen, DE), and Dylan Thurston (Indiana University – Bloomington, US)


License  Creative Commons BY 3.0 Unported license
 © Birgit Vogtenhuber, Philipp Kindermann, Stephen G. Kobourov, Maarten Löffler, Martin Nöllenburg, André Schulz, and Dylan Thurston

Introduction

This work is motivated by the following question posed by Benjamin Burton: Given a drawing of a knot, how can it be redrawn *nicely* without changing the given topology of the drawing?

A *knot* is an embedding of a simple closed curve in 3-dimensional Euclidean space R^3 , considered up to continuous transformations, which cannot be untangled to the simple loop, also known as the *unknot*. Similarly, a *link* is a collection of simple closed curves in R^3 that cannot be untangled. A *drawing of a knot (link)* is a mapping of the knot (link) to the Euclidean plane R^2 such that for any point of R^2 , at most two points of the curve(s) are mapped to it [6, 5, 1].

It is easy to see that drawings of links and knots are 4-regular plane multigraphs that contain neither loops nor split vertices. Likewise, every 4-regular planar multigraph without loops and split vertices can be interpreted as a link.



■ **Figure 8** Two different 2-Lombardi drawings of knot 4_1 , which by Theorem 2 does not admit a Lombardi-drawing.

A *Lombardi drawing* of a (multi-)graph $G = (V, E)$ is a drawing of G in the Euclidean plane with the following properties:

1. The vertices are represented as distinct points in the plane
2. The edges are represented as circular arcs connecting the representations of their end vertices (and not containing the representation of any other vertex); note that a straight-line segment is a valid circular arc with radius infinity.
3. Every vertex has *perfect angular resolution*, that is, its incident edges are equiangularly spaced. For links and knots this means that the angle between any two consecutive edges is $\pi/2$.

Lombardi drawings have been introduced by Duncan et al. [3] who showed a number of positive results (e.g., all d -regular graphs with $d \not\equiv 2 \pmod{4}$ have circular Lombardi drawings and all 2-degenerate graphs have Lombardi drawings) and negative results (e.g., there are planar graphs that do not have planar Lombardi drawings). Eppstein [4] showed that every (simple) planar graph with maximum degree three has a plane Lombardi drawing. Further, he showed that a certain class of 4-regular planar graphs (the medial graphs of polyhedral graphs) also admit plane Lombardi drawings and presented an example of a 4-regular planar graph that does not have a plane Lombardi drawing.

k-Lombardi drawings are a generalization of Lombardi drawings in which every edge is a sequence of at most k circular arcs that meet at a common tangent. Duncan et al. [2] showed that every planar graph has a 3-Lombardi drawing.


Results

The main question we are considering in this work is motivated by applying the Lombardi drawing style to knot and link drawings: Given a 4-regular planar multigraph G with a fixed combinatorial embedding (without loops and split-vertices), does G admit a plane Lombardi drawing with the same embedding? And what can still be done if this is not the case?

By the results of Duncan et al. [2], every link trivially admits a 3-Lombardi drawing. As our first result, we showed that every link also admits a 2-Lombardi drawing.

► **Theorem 1.** *Every 4-regular planar multigraph G (without loops and split-vertices) with a fixed combinatorial embedding admits a plane 2-Lombardi drawing with the same embedding.*

Concerning the original Lombardi drawings, we know from Eppstein [4] that there exist 4-regular planar graphs that do not admit a Lombardi drawing. However, the example by Eppstein represents a link, not a knot. When searching for whether or not all knots admit a Lombardi drawing, we obtained a surprising negative result: The 4-knot is not Lombardi. Moreover, the following stronger statement holds:



■ **Figure 9** Two examples of Lombardi drawings of knots and the according 4-regular graphs whose existence follows from Theorem 3.

► **Theorem 2.** *Every 4-regular planar multigraph G that contains K_4 as a subdrawing does not admit a plane Lombardi drawing.*

On the positive side, we were able to extend the result from Eppstein [4] to a larger class of graphs: Every plane drawing D of a 4-regular planar multigraph can be interpreted as the medial graph of a multigraph and its dual. If one of those graphs is simple, then D admits a Lombardi-drawing.

► **Theorem 3.** *Let D be a 4-regular planar multigraph G (without loops and split-vertices) with a fixed combinatorial embedding and let M and M' be the primal-dual multigraph pair for which D is the medial graph. If one of M and M' is simple, then D admits a plane Lombardi drawing with the same embedding.*

We remark that neither the result from Theorem 2 nor the one from Theorem 3 is tight: We found a 4-regular planar multigraph G that does not contain the 4-knot as a subdrawing and still does not admit a Lombardi drawing, and we found 4-regular planar multigraphs admitting a Lombardi drawing whose primal-dual pair M and M' both contain parallel edges.

Open problems and ongoing work

There are many open questions remaining which we plan to consider in this context. As main questions concerning Lombardi drawings we have the following.

► **Question 4.** *Can we give a complete characterization of 4-regular planar multigraphs that admit a Lombardi drawing?*

► **Question 5.** *What is the complexity of deciding whether a given 4-regular planar multigraph admits a Lombardi drawing?*

The next question is about the transition between Lombardi drawings and 2-Lombardi drawings.

► **Question 6.** *Given a 4-regular planar multigraph, what is the minimum number of edges consisting of two circular arcs in any 2-Lombardi drawing?*

We conclude with a question about a different relaxation of Lombardi drawings for drawings of 4-regular planar multigraphs.

► **Question 7.** *Can we redraw every drawing of a 4-regular planar multigraph using circular arcs as edges such that at every vertex, every pair of non-adjacent edges emanates in opposite directions? And if yes, what is the maximum smallest angle between consecutive edges that we can guarantee?*

References

- 1 P. R. Cromwell. *Knots and links*. Cambridge University Press, 2004.
- 2 C. A. Duncan, D. Eppstein, M. T. Goodrich, S. G. Kobourov, and M. Löffler. Planar and poly-arc Lombardi drawings. In M. van Kreveld and B. Speckmann, editors, *Graph Drawing (GD'11)*, volume 7034 of *LNCS*, pages 308–319. Springer, 2012.
- 3 C. A. Duncan, D. Eppstein, M. T. Goodrich, S. G. Kobourov, and M. Nöllenburg. Lombardi drawings of graphs. *Journal of Graph Algorithms and Applications*, 16(1):37–83, 2012.
- 4 D. Eppstein. Planar Lombardi drawings for subcubic graphs. In W. Didimo and M. Patrignani, editors, *GD 2012*, number 7704 in *LNCS*, pages 126–137. Springer-Verlag Berlin Heidelberg, 2013.
- 5 D. Rolfsen. *Knots and links*. American Mathematical Soc., 1976.
- 6 R. G. Scharein. *Interactive Topological Drawing*. PhD thesis, Department of Computer Science, The University of British Columbia, 1998.

Participants

- Maïke Buchin
Ruhr-Universität Bochum, DE
- Benjamin Burton
The University of
Queensland, AU
- Gregory R. Chambers
University of Chicago, US
- Hsien-Chih Chang
University of Illinois –
Urbana-Champaign, US
- Arnaud de Mesmay
University of Grenoble, FR
- Anne Driemel
TU Eindhoven, NL
- Parker Evans
Tulane University –
New Orleans, US
- Brittany Terese Fasy
Montana State University –
Bozeman, US
- Philipp Kindermann
FernUniversität in Hagen, DE
- Linda Kleist
TU Berlin, DE
- Boris Klemz
FU Berlin, DE
- Stephen G. Kobourov
University of Arizona –
Tucson, US
- Irina Kostitsyna
Free University of Brussels, BE
- David Letscher
St. Louis University, US
- Maarten Löffler
Utrecht University, NL
- Anna Lubiw
University of Waterloo, CA
- Wouter Meulemans
TU Eindhoven, NL
- Martin Nöllenburg
TU Wien, AT
- Tim Ophelders
Eindhoven University of
Technology, NL
- Florian T. Pokorny
KTH Royal Institute of
Technology – Stockholm, SE
- Jessica S. Purcell
Monash University –
Clayton, AU
- Regina Rotman
University of Toronto, CA
- Saul Schleimer
University of Warwick –
Coventry, GB
- Lena Schlipf
FernUniversität in Hagen, DE
- André Schulz
FernUniversität in Hagen, DE
- Eric Sedgwick
DePaul University – Chicago, US
- Bettina Speckmann
TU Eindhoven, NL
- Jonathan Spreer
FU Berlin, DE
- Frank Staals
Aarhus University, DK
- Darren Strash
Colgate University –
Hamilton, US
- Dylan Thurston
Indiana University –
Bloomington, US
- Stephan Tillmann
University of Sydney, AU
- Marc van Kreveld
Utrecht University, NL
- Mikael Vejdemo-Johansson
CUNY College of
Staten Island, US
- Birgit Vogtenhuber
TU Graz, AT
- Carola Wenk
Tulane University, US
- Erin Moriarty Wolf Chambers
St. Louis University, US

