Facets of Planar Graph Drawing

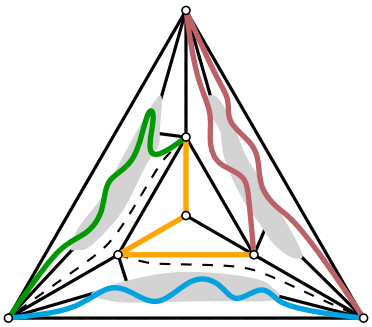

Dissertation zur Erlangung des akademischen Grades
eines Doktors der Naturwissenschaften

vorgelegt am

Fachbereich Mathematik und Informatik
der Freien Universität Berlin

2020

von

Boris Klemz

Institut für Informatik
Freie Universität Berlin
Takustraße 9
14195 Berlin
klemz@inf.fu-berlin.de

Betreuer: Prof. Dr. Günther Rothe
Institut für Informatik
Freie Universität Berlin
Takustraße 9
14195 Berlin
Germany
rote@inf.fu-berlin.de
phone: +49 30 838 75150

Gutachter: Prof. Dr. Günther Rothe
Institut für Informatik
Freie Universität Berlin
Takustraße 9
14195 Berlin
Germany
rote@inf.fu-berlin.de
phone: +49 30 838 75150

Prof. Dr. Martin Nöllenburg
Institute of Logic and Computation
Technische Universität Wien
Favoritenstraße 9–11, E192-01
1040 Wien
Austria
noellenburg@ac.tuwien.ac.at
phone: +43(1)58801–192120

Tag der Disputation: 30. April 2020

Erklärung

Ich erkläre hiermit, dass ich alle Hilfsmittel und Hilfen angegeben habe und versichere auf dieser Grundlage die Arbeit selbstständig verfasst zu haben. Die Arbeit habe ich nicht in einem früheren Promotionsverfahren eingereicht.

Berlin, 17. Januar 2020

Boris Klemz

Contents

Abstract

This thesis makes a contribution to the field of Graph Drawing, with a focus on the planarity drawing convention. The following three problems are considered.

Ordered Level Planarity. We introduce and study the problem Ordered Level Planarity which asks for a planar drawing of a graph such that vertices are placed at prescribed positions in the plane and such that every edge is realized as a y-monotone curve. This can be interpreted as a variant of Level Planarity in which the vertices on each level appear in a prescribed total order. We establish a complexity dichotomy with respect to both the maximum degree and the level-width, that is, the maximum number of vertices that share a level. Our study of Ordered Level Planarity is motivated by connections to several other graph drawing problems. With reductions from Ordered Level Planarity, we show $\mathcal{NP}$-hardness of multiple problems whose complexity was previously open, and strengthen several previous hardness results. In particular, our reduction to Clustered Level Planarity generates instances with only two nontrivial clusters. This answers a question posed by Angelini, Da Lozzo, Di Battista, Frati, and Roselli [2015]. We settle the complexity of the Bi-Monotonicity problem, which was proposed by Fulek, Pelsmajer, Schaefer, and Štefankovič [2013]. We also present a reduction to Manhattan Geodesic Planarity, showing that a previously [2009] claimed polynomial time algorithm is incorrect unless $\mathcal{P} = \mathcal{NP}$.

Two-page book embeddings of triconnected planar graphs. We show that every triconnected planar graph of maximum degree five is a subgraph of a Hamiltonian planar graph or, equivalently, it admits a two-page book embedding. In fact, our result is more general: we only require vertices of separating 3-cycles to have degree at most five, all other vertices may have arbitrary degree. This degree bound is tight: we describe a family of triconnected planar graphs that cannot be realized on two pages and where every vertex of a separating 3-cycle has degree at most six. Our results strengthen earlier work by Heath [1995] and by Bauernöppel [1987] and, independently, Bekos, Gronemann, and Raftopoulou [2016], who showed that planar graphs of maximum degree three and four, respectively, can always be realized on two pages. The proof is constructive and yields a quadratic time algorithm to realize the given graph on two pages.

Convexity-increasing morphs. We study the problem of convexifying drawings of planar graphs. Given any planar straight-line drawing of an internally 3-connected graph, we show how to morph the drawing to one with strictly convex faces while maintaining planarity at all times. Our morph is convexity-increasing, meaning that once an angle is convex, it remains convex. We give an efficient algorithm that constructs such a morph as a composition of a linear number of steps where each step either moves vertices along horizontal lines or moves vertices along vertical lines. Moreover, we show that a linear number of steps is worst-case optimal.

Acknowledgements

First of all, I would like to thank my advisor Günter Rote for offering me the chance to do a PhD with him, and for providing me with a lot of freedom and numerous great opportunities. In particular, I appreciated his "yes"-saying attitude, and I thank him for proofreading this thesis.

My academic journey started with my studies of computer science at the Karlsruhe Institute of Technology. It were Dorothea Wagner's lecture on Algorithms for Planar Graphs, and Martin Nöllenburg's lecture on Computational Geometry in particular that led to my interest in the topics that later became my research focus. I thank Martin, who also co-advised my Studienarbeit for which I conducted my first original research, for an overall great introduction to the TCS world back in my days in Karlsruhe, and, now, for his continued interest in my research and for agreeing to review this thesis.

During my PhD studies, I had the opportunity to attend several research workshops in places all over Europe. I would like to thank the organizers for making these events possible and for inviting me to participate—in particular the Dagstuhl seminars 17072 and 19352, the two Fixed-Parameter Computational Geometry workshops at Lorentz center, and the four Geometric Graph Weeks that I attended were very fruitful for me. These meetings presented a great way to meet and work with new people, many of which later ended up being my coauthors. In particular, I had the pleasure to get to know Michael Hoffmann. I thank Michael for our exciting joint projects, for his collaborative attitude, and for inviting me to Zürich for a research stay. Our research sessions in various places all over the world and via Skype were among the most inspiring and enjoyable moments of the recent years. Speaking of which, another activity that falls into this category was teaching. In this context, I would like to express my dearest gratitude to both Frank Hoffmann and Klaus Kriegel for the very exciting and rewarding opportunity to co-lecture our first semester course Logic and Discrete Mathematics along their sides.

I thank my numerous coauthors, in particular those that were involved in the projects that are discussed in this thesis: Michael Hoffmann, Linda Kleist, Anna Lubiw, Günter Rote, Lena Schlipf, Frank Staals, and Darren Strash.

Special thanks go to my colleagues of the theoretical computer science group at Freie Universität Berlin for providing an excellent research environment, and an overall friendly and productive atmosphere. In particular, I thank Wolfgang Mulzer, who was always eager to answer my research, teaching, or career related questions, and Nadja Seiferth for always being open for discussions on all sorts of topics, and for joining me in the attempt to solve the research question that refused to be solved.

Last but not least I thank Sara—for everything.

Chapter 1

Introduction

This thesis makes a contribution to Graph Drawing, a research area devoted to algorithmic and combinatorial questions related to the visualization of graphs. The exact origins of Graph Drawing do not appear to be well known. In 1994, Di Battista, Eades, Tamassia, and Tollis [44] published an annotated bibliography that already cited more than 300 publications related to Graph Drawing, most of them from the 80s and early 90s. The same group of authors state in [45], that "Knuth's 1963 paper [94] on drawing flowcharts was perhaps the first paper to present an algorithm for drawing a graph for visualization purposes". Graph Drawing has grown to be an active research field with its own techniques and terminology. Its community has brought forth multiple dedicated text books [45, 87, 102, 119] and an annual international symposium [1].

Certainly the main motivation for finding graph drawings with specific properties is automated visualization. According to Di Battista et al. [45], application areas "include software engineering (data flow diagrams, subroutine-call graphs, program nesting tress, object-oriented class hierarchies), databases (entity-relationship diagrams), information systems (organization charts), real-time systems (Petri nets, state-transition diagrams), decision support systems (PERT networks, activity trees), VLSI (circuit schematics), artificial intelligence (knowledge-representation diagrams), and logic programming (SLD-trees)". Graph Drawing also has applications in layouting. Here, the goal is not necessarily to find a visually pleasing drawing—instead, the drawing (layout) is supposed to satisfy certain constraints that make it actually usable in practice. As an example, Nishizeki and Rahman [102] mention the design of printed circuit boards and state that planarity is required for realizing a layout on a single layer.

According to Di Battista et al. [45], a typical Graph Drawing problem can be described by three types of parameters: *conventions*, *constraints*, and *aesthetics*.

Drawing *conventions* are described by Di Battista et al. [45] as basic rules that describe how the vertices and edges are represented. In the following we list some examples of drawing conventions.

- Each edge is realized as a line segment / polygonal chain / curve composed of vertical and horizontal line segments / y-monotone curve / Bézier spline.
- Vertices are placed on integer coordinates.
- Between each pair of vertices there exists a path that is monotone in some direction [6].
- No pair of edges intersects except at a common endpoint (i.e., the drawing is planar).

In case of visualization problems, these rules aim to make the drawing readable, that is, the reader should be able to easily understand the relationship between the represented objects and quickly perform simple tasks such as determining a path between two given vertices. In layouting problems, conventions might also be required due to technical limitations as in the above example about printed circuit boards.

Constraints, according to Di Battista et al. [45], refer to additional information about the graph that is encoded as part of the input and supposed to be reflected by the drawing. For instance, in a tool for generating metro maps, the user should be able to specify the desired vertex positions (at least in a relative sense), e.g. the city center should appear close to the center of the drawing. When visualizing hierarchies, we may desire a drawing in which objects higher up in the hierarchy are placed at larger y-coordinates in the drawing.

Finally, Di Battista et al. [45] describe *aesthetics* as drawing criteria that one might want to optimize to improve the readability of the drawing. For instance, when dealing with a non-planar graph, we might want to minimize the number of crossings, and when using the convention that the vertices have to be placed on an integer grid, it might be desirable to minimize the area of the grid.

1.1 Thesis outline

In this thesis, we focus on the fundamental drawing convention of planarity. Planar drawings have empirically proven visual qualities [105, 106, 125]. We also saw above that the desire for drawings without crossings may arise naturally in the context of layouting. Each of the three main chapters of the thesis is devoted to a particular Graph Drawing problem with the goal to generate drawings that are planar while also satisfying additional constraints and conventions.

We proceed with a very brief overview of the three main chapters. A slightly more detailed overview is given in the Sections 1.2–1.4. The chapters are partially based on published results. The references to these publications can be found in Section 1.5.

In Chapter 3, we introduce and study a problem called Ordered Level Planarity which asks for a planar drawing of a graph such that vertices are placed at prescribed positions in the plane and such that every edge is realized as a y-monotone

curve. This can be interpreted as a variant of Level Planarity in which the the vertices of each level appear in a prescribed total order. Our main result is a $\mathcal{NP}$-hardness statement for a very constrained special case. We motivate this result by establishing connections to several other graph drawing problems that have been studied in the literature. For a slightly more detailed overview, see Section 1.2.

Chapter 4 is concerned with two-page book embeddings. In such a drawing, all vertices are placed on a vertical line and each edge is realized either in the left or the right half-plane. Improving a series of previous results, we give a new sufficient condition for the two-page book embeddability of triconnected planar graphs. We show that the condition is, in some sense, best possible. A slightly more detailed overview is given in Section 1.3.

Finally, Chapter 5 is concerned with a morphing problem: we are given a straight-line planar drawing and want to find a continuous deformation that preserves straight-line planarity at all times and turns the initial drawing into another one in which each face is a convex polygon. Motivated by visualization goals, we want to ensure that an inner angle that is already convex at some time during the motion never becomes nonconvex later. We give an efficient algorithm that finds such a deformation. The algorithm is optimal in the worst-case with respect to the number of "piece-wise linear" steps.

Despite being considered an independent research field, Graph Drawing is a very multifaceted discipline that borrows techniques from many other areas such as classical graph theory, algorithms and complexity, computational geometry, information visualization, the theory of topological and geometric graphs, knot theory, topology, order theory, physics, and algebra. While many of these connections might seem very natural, some of them could also be considered surprising. In spite of the unifying theme of planarity, each of the chapters in this thesis has a very distinct flavor as it relies on techniques of a particular set of the above facets. In particular, the chapter about book embeddings makes heavy use of many concepts of classic graph theory: connectivity, separators, edge contractions, decomposition trees, graph augmentation, Hamiltonian cycles, and more. In contrast, the instances generated by our reduction to show that Ordered Level Planarity is $\mathcal{NP}$-hard do not exhibit any meaningful graph theoretic properties. Hence, we have to rely on direct planarity arguments to prove the correctness of the reduction, which is reminiscent of problems related to topological graphs. Finally, the objects of desire in the chapter about morphing are straight-line drawings in which each face is described by a convex polygon. Naturally, here we rely on techniques from computational geometry among others. We discuss these aspects more in the final, concluding Chapter 6 of this thesis, where we also summarize our main results and discuss the most intriguing open problems.

Most definitions are given in the beginning of the chapter in which they are needed. In Chapter 2, we define only the most basic concepts that are relevant for all chapters.

1.2 Ordered Level Planarity (Chapter 3)

The well known Upward Planarity problem asks for a planar drawing of a directed graph with the drawing convention that each edge (u, v) is realized as a y-monotone curve that goes upward from u to v [62]. Such upward planar drawings provide a natural way of visualizing hierarchies, workflow diagrams, and partial orders on sets of items (posets). As shown by Garg and Tamassia, Upward Planarity is $\mathcal{NP}$-hard [62]. However, when we add the constraint that the y-coordinate, called level, of each vertex is prescribed as part of the input, the problem can be solved efficiently. This variant of Upward Planarity is called Level Planarity, and the study of its complexity has a long history [46, 59, 76, 82, 83]. Ultimately, a linear-time algorithm was presented by Jünger, Leipert, and Mutzel [83].

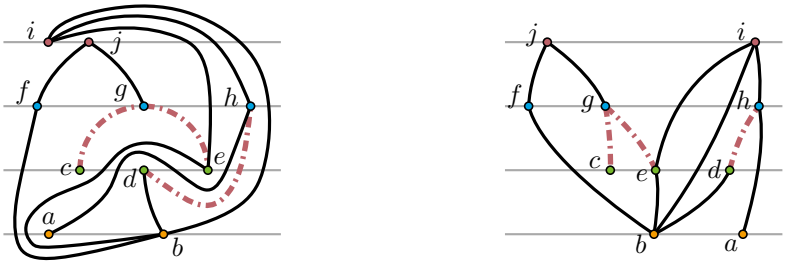

Figure 1: In Level Planarity the order of the vertices of a common level is not fixed. Finding a good ordering is an essential part of finding a solution. The ordering suggested in (a) is not realizable since the edge (d, h) cannot be drawn without crossing (c, g) or (e, g). (b) A valid realization of the instance in (a).

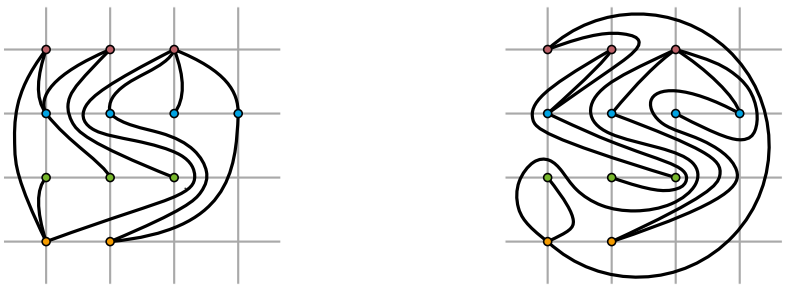

Figure 2: (a) A realization of the Ordered Level Planarity instance in (b).

As part of the solution, each algorithm for Level Planarity needs to determine an appropriate left-to-right order for the set of vertices of each level, see Figure 1. Hence, it might seem that prescribing these orderings (which can be thought of as prescribing both the x-coordinate and the y-coordinate of each vertex, see Figure 2)

should make the problem easier. However, we show that this variant of Level Planarity, which we call Ordered Level Planarity, is $\mathcal{NP}$-hard, even in very constrained special cases.

We motivate our result by establishing connections to several other Graph Drawing problems. Ordered Level Planarity, especially if restricted to the instances for which we show it to be $\mathcal{NP}$-hard, is an elementary problem that readily reduces to several variants of Level Planarity and other problems with constraints that (partially) fix vertex coordinates in advance. With such reductions, we show $\mathcal{NP}$-hardness of multiple problems whose complexity was previously open. We also strengthen several previous hardness results. Some of our results answer open problems posed by members of the Graph Drawing community. In more detail, we establish the following:

- The drawing convention in the problem Manhattan Geodesic Planarity is that edges are realized as rectilinear paths, that is, curves composed of vertical and horizontal segments, that are geodesics with respect to the L_1-norm. Additionally, as a constraint, each vertex has to be placed at a prescribed position in the plane. Previously, it was claimed that Manhattan Geodesic Planarity is polynomial-time solvable if the input graph is a matching [86]. We provide a reduction from Ordered Level Planarity to Manhattan Geodesic Planarity that shows that this claim is incorrect unless $\mathcal{P} = \mathcal{NP}$.

- The problem Bi-Monotonicity is very similar to Manhattan Geodesic Planarity. Its complexity status was posed as an open question by Fulek, Pelsmajer, Schaefer, and Štefankovič [59], who introduced the problem. Our reduction to Manhattan Geodesic Planarity extends to Bi-Monotonicity and answers the question by Fulek et al.

- We provide reductions to T-Level Planarity, Constrained Level Planarity, and Clustered Level Planarity. The first two of these variants of Level Planarity impose (partial) ordering constraints on the vertices of each level, which immediately suggests a connection to Ordered Level Planarity. On the other hand, Clustered Level Planarity, at first glance, does not seem related to Ordered Level Planarity: it combines Level Planarity with Clustered Planarity. The goal in the latter problem is to find a planar drawing of a graph while simultaneously visualizing a clustering hierarchy on its set of vertices. Notably, the clustering hierarchy is not necessarily related to the hierarchy suggested by the directed edges.

  Our reductions strengthen previous hardness results. In particular, the reduction to Clustered Level Planarity produces instances with only two non-nested clusters. This answers a question posed by Angelini, Da Lozzo, Di Battista, and Roselli [9].

Very recently, Da Lozzo, Di Battista, and Frati [39] studied another generalization of Ordered Level Planarity and applied our result to show that their problem is $\mathcal{NP}$-hard. We expect, that the $\mathcal{NP}$-hardness of Ordered Level Planarity will serve as a useful tool for further reductions.

1.3 Two-page book embeddings of triconnected planar graphs (Chapter 4)

Intuitively, a book embedding [22, 75, 131] is a drawing of a graph in a book. Formally, the corresponding drawing convention requires that all vertices are embedded on a line in $\mathbb{R}^3$ called spine, and every edge is embedded in a half-plane, called page, bounded by the spine. No two edges (on the same page) are allowed to cross. If k pages are used, then the corresponding embedding is a k-page book embedding. Applications for book embeddings include VLSI design [35], bioinformatics [74], and other Graph Drawing problems [11, 128].

Obviously, every graph has a book embedding: we can simply realize each edge on a separate page. The book thickness of a graph is the smallest number of pages that suffices to realize the graph in form of a book embedding. Yannakakis showed, improving a series of earlier results, that the book thickness of planar graphs is at most four [131]. A corresponding lower bound is still elusive, in spite of initial claims [130].

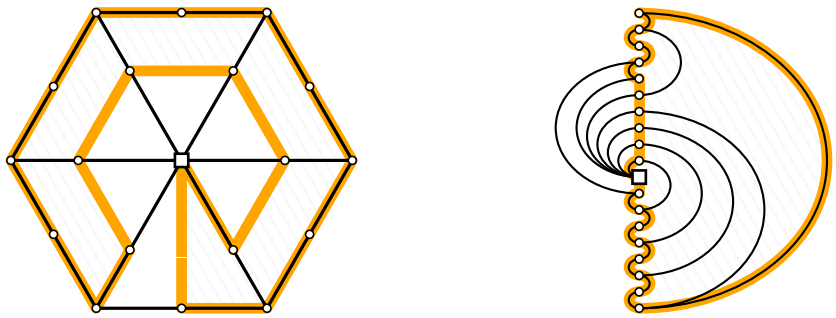

(a) (b)

Figure 3: (a) A nonhamiltonian graph with a subhamiltonian cycle (a Hamiltonian cycle in a plane augmentation) in orange, and (b) a corresponding two-page book embedding. In particular, the order of the vertices along the cycle corresponds to the ordering of the vertices along the spine. The interior of the cycle corresponds to the right page and its exterior corresponds to the left page.

Each graph that admits a two-page book embedding is planar. Bernhart and Kainen [22] characterized those graphs that can be embedded on two pages as the subhamiltonian planar graphs. A graph G is subhamiltonian planar if it is a subgraph of a Hamiltonian planar graph H on the same vertex set, see Figure 3. This turns

the problem of embedding a graph on two pages into a graph augmentation problem. A Hamiltonian cycle in an augmentation H is called a subhamiltonian cycle for G. Observe that if G is maximal planar, then it is subhamiltonian planar if and only if it is Hamiltonian. This implies that the recognition of graphs with book thickness at most two is $\mathcal{NP}$-hard [127].

In this thesis, we are interested in sufficient conditions for two-page book embeddability. Sufficient conditions for the existence of a (classic) Hamiltonian cycle typically require fairly large vertex degrees. For instance, Dirac's Theorem [49] states that graphs with minimum degree $\geq n/2$ are always Hamiltonian. Similarly, previously established sufficient conditions for being subhamiltonian planar also involve assumptions about vertex degrees [20, 21, 75]. However, these assumptions involve the maximum degree of the graph, rather than the minimum degree. Intuitively, the reason is that when looking for a subhamiltonian cycle in a given planar graph G, we may create any desired edges that are not part of G. In contrast, unwanted edges that belong to G may act as obstructions.

Heath [75] showed that planar graphs of maximum degree three are always subhamiltonian. Later, Bauernöppel [20] and, independently, Bekos, Gronemann, and Raftopoulou [21] showed that maximum degree four is also a sufficient condition for a planar graph to be subhamiltonian. The ultimate goal is to determine the largest k such that every planar graph with maximum degree at most k is subhamiltonian.

We take a considerable step towards solving this goal by showing that all triconnected planar graphs of maximum degree five are subhamiltonian. In fact, we prove a more general statement: we only require the vertices of separating 3-cycles to have degree at most five, all other vertices may have arbitrary degrees. This degree bound is tight: we describe a family of triconnected planar graphs that are not subhamiltonian planar and where every vertex of a separating triangle has degree at most six. Our result is constructive: we describe how to compute a subhamiltonian cycle in the given degree bounded graph in quadratic time.

1.4 Convexity-increasing morphs (Chapter 5)

A morph is a continuous deformation of a graph drawing that preserves planarity and straight-line edges at all times [3, 28, 57, 67, 120]. Graph morphing is motivated by applications in animation and computer graphics [65]. Hence, previous literature focuses on the problem of finding morphs between two given drawings of a common plane graph. The existence of such morphs has been established long ago [28, 120]. More recent results aim to improve the visual quality [57, 67] and decrease the complexity of the deformation [3].

Instead of morphing between two given drawings, we are interested in transforming an initial drawing into one that satisfies the drawing convention that each face is described by a strictly convex polygon. Such a drawing is called strictly convex. This

task by itself is easy to achieve (given the previous literature), however, it has been established by empirical investigations [107] that, for the purposes of visualization, it is important to maintain the viewer's "mental map", which means changing as little as possible while making observable progress towards a goal. Hence, we would like our morphs to be convexity-increasing, meaning that once an angle has its desired convexity status, this status is maintained. We illustrate such a morph in Figure 4.

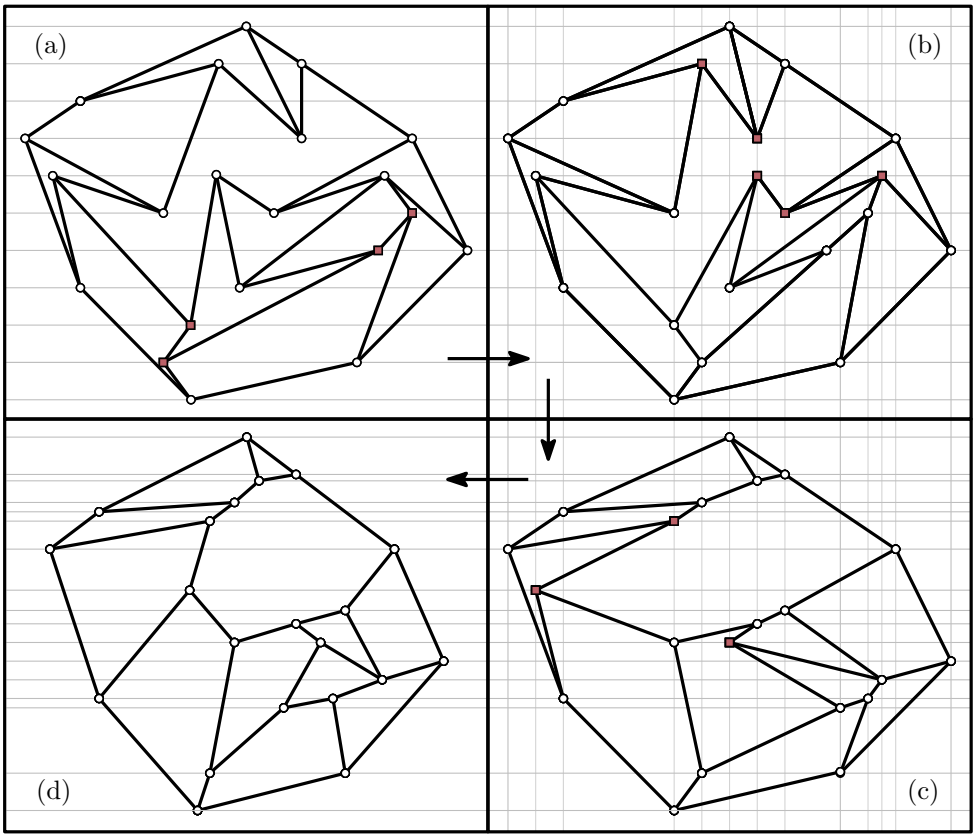

Figure 4: A sequence of convexity-increasing morphs (horizontal, vertical, horizontal) that morph a straight-line drawing of a graph G (a) into a strictly convex drawing of G (d). The marked vertices have reflex angles that are eliminated in the next step. All convex angles are preserved.

This idea falls in line with research done by Connelly, Demaine, and Rote [37] and Streinu [117, 118], who studied expansive deformations that preserve the length of each edge, and by Aichholzer et al. [2], who studied visibility-increasing morphs. Both these types of morphs are also convexity-increasing. However, the algorithms in all of the above references are restricted to the case that the graph is a cycle, i.e., the given drawing is a polygon. Hence, our goal is much more general, as we need to maintain and increase the convexity of multiple interconnected polygons at the same time.

A necessary conditions for the existence of a convexity-increasing morph to some strictly convex drawing is that the given graph actually admits a strictly convex drawing. We show that this condition is also sufficient: we present an efficient algorithm that finds a convexity-increasing morph to some strictly convex redrawing of the initial drawing, whenever such a drawing exists. Morphs are usually discretized into "piece-wise linear" steps. Our algorithm produces morphs composed of at most a linear number of such steps, which we show to be worst-case optimal.

1.5 Publications

Chapter 3 is based on work done with Günter Rote. A full version appeared [93] in *ACM Transaction on Algorithms*, which differs content-wise only slightly from the material in Chapter 3. A preliminary version, which only sketches the main ideas, appeared [92] in the proceedings of the symposium on *Graph Drawing and Network Visualization* (*GD*) 2017. An extended abstract appeared [91] in the proceedings of the *European Workshop on Computational Geometry* (*EuroCG*) 2017.

- [93] B. Klemz and G. Rote. Ordered level planarity and its relationship to geodesic planarity, bi-monotonicity, and variations of level planarity. *ACM Trans. Algorithms*, 15(4):53:1–53:25, 2019. doi:10.1145/3359587.

- [92] B. Klemz and G. Rote. Ordered level planarity, geodesic planarity and bi-monotonicity. In F. Frati and K. Ma, editors, *Graph Drawing and Network Visualization - 25th International Symposium, GD 2017, Boston, MA, USA, September 25–27, 2017, Revised Selected Papers*, volume 10692 of *Lecture Notes in Computer Science*, pages 440–453. Springer, 2017. doi:10.1007/978-3-319-73915-1_34.

- [91] B. Klemz and G. Rote. Ordered level planarity and geodesic planarity. In *Proceedings of the 33rd European Workshop on Computational Geometry (EuroCG 2017)*, pages 269–272, Apr. 2017.

Chapter 4 is based on work done with Michael Hoffmann. A full version has not been published yet. A preliminary version, which only sketches the main ideas, appeared [77] in the proceedings of the *European Symposium on Algorithms* (*ESA*) 2019.

- [77] M. Hoffmann and B. Klemz. Triconnected planar graphs of maximum degree five are subhamiltonian. In M. A. Bender, O. Svensson, and G. Herman, editors, *27th Annual European Symposium on Algorithms, ESA 2019, September 9–11, 2019, Munich/Garching, Germany*, volume 144 of *LIPIcs*, pages 58:1–58:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.ESA.2019.58.

Chapter 5 is based on work done with Linda Kleist, Anna Lubiw, Lena Schlipf, Frank Staals, and Darren Strash. A journal version has been published [90] in *Computational Geometry - Theory and Applications* (*CGTA*). Chapter 5 contains several minor corrections compared to the version in [90]. The main algorithm was also extended to be convexity-increasing, rather than just weakly convexity-increasing. This extension is described in Section 5.5. A preliminary version of [90], which only sketches the main ideas, appeared [88] in the proceedings of the *Workshop on Graph-Theoretic Concepts in Computer Science* (*WG*) 2018. An extended abstract appeared [89] in the proceedings of the *European Workshop on Computational Geometry* (*EuroCG*) 2018.

- [90] L. Kleist, B. Klemz, A. Lubiw, L. Schlipf, F. Staals, and D. Strash. Convexity-increasing morphs of planar graphs. *Comput. Geom.*, 84:69–88, 2019. doi:10.1016/j.comgeo.2019.07.007.

- [88] L. Kleist, B. Klemz, A. Lubiw, L. Schlipf, F. Staals, and D. Strash. Convexity-increasing morphs of planar graphs. In A. Brandstädt, E. Köhler, and K. Meer, editors, *Graph-Theoretic Concepts in Computer Science - 44th International Workshop, WG 2018, Cottbus, Germany, June 27–29, 2018, Proceedings*, volume 11159 of *Lecture Notes in Computer Science*, pages 318–330. Springer, 2018. doi:10.1007/978-3-030-00256-5_26.

- [89] L. Kleist, B. Klemz, A. Lubiw, L. Schlipf, F. Staals, and D. Strash. Convexity-increasing morphs of planar graphs. In *Proceedings of the 34th European Workshop on Computational Geometry (EuroCG 2018)*, pages 65:1–65:6, Apr. 2018.

Chapter 2

Terminology and notation

We assume familiarity with basic concepts of mathematics and computer science such as algorithms and complexity, naive set theory, geometry, and basic graph theory. For a comprehensive introduction to these topics, we refer to dedicated text books: Cormen, Leiserson, Rivest, and Stein [38] provide an extensive introduction to algorithms and their analysis. A comprehensive guide to the theory of $\mathcal{NP}$-completeness is given by Garey and Johnson [60]. For an introduction to discrete mathematics, we refer to the book by Rosen [113]. One of the standard textbooks on graph theory was written by Diestel [48]. Finally, an overview of many important data structures and algorithms related to computational geometry can be found in the book by de Berg, Cheong, van Kreveld, and Overmars [42].

Most definitions are introduced in the beginning of the chapters in which they are needed. In this section, we introduce some notation and recall basic terminology that is used throughout the entire thesis.

Graphs. We denote by $\mathrm{V}(G)$ the vertex set and by $\mathrm{E}(G)$ the edge set of a graph G. For a set of edges $E \subseteq \mathrm{E}(G)$ we use $\mathrm{V}(E)$ to denote the set of vertices that are incident to at least one edge in E. For a vertex v of G let $\mathrm{N}_G(v) \subset \mathrm{V}(G)$ denote the set of neighbors of v in G, and similarly the neighborhood of a vertex set $U \subseteq \mathrm{V}(G)$ by $\mathrm{N}_G(U) = \bigcup_{u \in U} \mathrm{N}_G(u)$.

Let $G = (V, E)$, and $V' \subseteq V$ and $E' \subseteq E$. Moreover, let G' and G'' be subgraphs of G. We use $G \setminus V'$ to denote the subgraph of G that is induced by $V \setminus V'$. We use $G \setminus E'$ to denote the graph $(V, E \setminus E')$. Finally, we use $G \setminus G'$ to denote the subgraph of G that is induced by $V \setminus \mathrm{V}(G')$. Moreover, we use $G' \cup G''$ to denote the subgraph $(\mathrm{V}(G') \cup \mathrm{V}(G''), \mathrm{E}(G') \cup \mathrm{E}(G''))$ of G. We also use $G' \cup V'$ to denote the subgraph $(\mathrm{V}(G') \cup V', \mathrm{E}(G'))$ of G. Finally, we use $G' \cup E'$ to denote the subgraph $(\mathrm{V}(G') \cup \mathrm{V}(E'), \mathrm{E}(G') \cup E')$ of G.

In Chapters 3 and 5 we use the usual notations $\{u, v\}$ and (s, t) to denote undirected and directed edges, respectively. We deviate from this convention in Chapter 4, where we use the shorthand uv to denote an undirected edge between u

and v. We also use shorthands such as *stuvw* to denote a path (or cycle) (s, t, u, v, w) (it will be clear from the context whether a path or cycle is meant). The reason for this inconsistency in notation is that in Chapter 4, we often have to argue about small constant sized configurations, where the shorthands help to remove clutter from the notation. In contrast, in Chapters 3 and 5, we encounter large graphs with heavily indexed labeling, where the classic notation improves readability.

Drawings and embeddings. In this paragraph, we recall basic terminology from the field of Graph Drawing. More information can be found in the textbooks [45, 87, 102, 119].

A *drawing* Γ of a graph G maps the vertices $\mathrm{V}(G)$ to pairwise distinct points in the plane and each edge $\{u, v\}$ to a simple curve joining the points representing u and v. In a *straight-line* drawing, these curves are line segments. We identify vertices and edges with their geometric representations. A drawing is *planar* if no pair of edges intersects, except at a common endpoint (which is necessary if the edges are incident to a common vertex). A graph that admits a planar drawing is called *planar*.

A planar drawing partitions the plane into regions, which are called *faces*. We also identify faces with their geometric representations. In each planar drawing there is exactly one unbounded face, called the *outer* face. The remaining faces are called *internal* or *inner* faces. The *boundary* of each face f can be uniquely described by a counterclockwise sequence of edges, or by multiple such sequences of edges in case the graph is not connected. In the connected case, we use ∂f to denote the boundary of f. If the graph is 2-connected, then ∂f is a *simple* cycle. Otherwise, edges may be visited twice when traversing the boundary of a face.

A planar drawing Γ determines a circular ordering of the neighbors of each vertex. The set of these orderings together with the set of face boundaries is called the *combinatorial embedding* of Γ. A 3-connected planar graph has exactly two distinct combinatorial embeddings, which are mirror images of each other [54, 126]. Two drawings may have the same combinatorial embedding, but different outer faces. A *plane* graph is a planar graph equipped with a combinatorial embedding and a distinguished outer face.

The DCEL. A combinatorial embedding may be efficiently stored and traversed by means of a *doubly connected edge list* (DCEL) [42]. We describe this data structure for connected graphs, which is sufficient in the context of this thesis. The full description, which also handles the disconnected case, is given in [42]. Every edge $\{u, v\}$ is stored as two directed *half-edges* (u, v) and (v, u). There is also a record for every vertex and face. Each half-edge (u, v) is associated with the face to its left. Each vertex has a pointer to one of its outgoing half-edges. Each face f has a pointer to one of the half-edges that are associated with f. Each half-edge (u, v) has a pointer to its *twin* (v, u), to its associated face f, to its origin u, and to its successor (v, w) and predecessor (t, u) along ∂f.

The DCEL allows for efficient traversal of the stored embedding. For instance, a face boundary ∂f may be traversed in $O(|\partial f|)$ time by starting with its associated half-edge (u, v) and then following successor pointers until encountering (u, v) again, and one may iterate over all neighbors of u in $O(|\mathrm{N}_G(u)|)$ time by starting with its associated half-edge (u, v) and then repeatedly accessing the predecessor's twin.

Throughout the thesis, whenever an algorithm is given a plane graph as part of the input, we assume it to be encoded by means of a DCEL.

Chapter 3

Ordered Level Planarity

3.1 Introduction

In this chapter, we introduce the problem Ordered Level Planarity and study its complexity. The problem can be interpreted as a variant of Level Planarity. We establish connections to several other graph drawing problems (for an overview, see Figure 5), which we survey in the following Sections 3.1.1–3.1.4. We proceed from general problems to more and more constrained ones. In Section 3.1.1, we recall the definition of the original Level Planarity problem. Section 3.1.2 discusses several constrained variations of Level Planarity. Ordered Level Planarity is defined in Section 3.1.3. The problems Geodesic Planarity and Bi-Monotonicity, which are closely related to one another, are discussed in Section 3.1.4. Section 3.1.5 summarizes the main results of this chapter and gives an overview of its remaining sections.

3.1.1 Upward Planarity and Level Planarity

Upward Planarity. A planar drawing of a directed graph (all graphs in this chapter are simple) is called *upward* if each edge $e = (u, v)$ is realized as a y-monotone curve that goes upward from u to v [62]. Such a drawing provides a natural way of visualizing a partial order on a set of items. The problem Upward Planarity of testing whether a directed graph has an upward planar drawing is $\mathcal{NP}$-complete [62]. However, if the y-coordinate of each vertex is prescribed, the problem can be solved in polynomial time [83]. This is captured by the notion of level graphs, which are discussed in the next paragraph.

Level Planarity. A *level graph* $\mathcal{G} = (G, \gamma)$ is a directed graph $G = (V, E)$ together with a *level assignment*, i.e. a surjective map $\gamma\colon V \to \{0, \dots, h\}$ with $\gamma(u) < \gamma(v)$ for every edge $(u, v) \in E$. The value h is the *height* of $\mathcal{G}$. The vertex set $V_i = \{v \mid \gamma(v) = i\}$ is called the i-th *level* of $\mathcal{G}$ and $\lambda_i = |V_i|$ is its *width*. The *level-width* λ of $\mathcal{G}$ is the

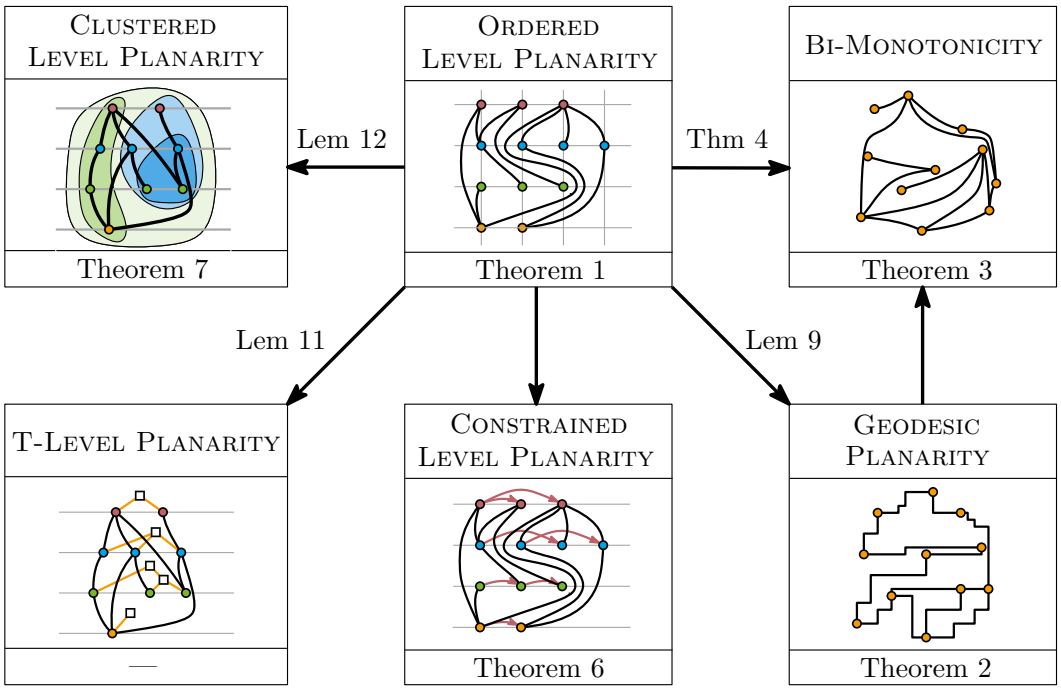

Figure 5: Ordered Level Planarity is a special case of several other graph drawing problems.

maximum width among all levels of $\mathcal{G}$. We say that $\mathcal{G}$ is *acyclic* if the underlying undirected graph of G is acyclic. A *level* planar drawing of $\mathcal{G}$ is an upward planar drawing of G where the y-coordinate of each vertex v is $\gamma(v)$, for an illustration see Figure 6(b). The horizontal line with y-coordinate i is denoted by L_i. The problem Level Planarity asks whether a given level graph has a level planar drawing [46, 59, 76, 82, 83], cf. Figures 6(a) and 6(b).

The study of the complexity of Level Planarity has a long history [46, 59, 76, 82, 83], culminating in a linear-time algorithm by Jünger, Leipert, and Mutzel [83]. Their algorithm is based on work for the special case of single-source level graphs by Di Battista and Nardelli [46]. There was an earlier attempt by Heath and Pemmaraju [76] to extend the work by Di Battista and Nardelli [46] to general level graphs. However, Jünger et al. [82] pointed out gaps in this construction. All these approaches utilize PQ-trees. Various simpler but asymptotically slower approaches to solve Level Planarity have been considered, see the work of Fulek, Pelsmajer, Schaefer, and Štefankovič [59] for one of these approaches (cf. Section 3.1.4) and a more comprehensive summary. Level Planarity has been extended to drawings of level graphs on surfaces different from the plane [8, 16, 17]. In particular, Radial Level Planarity [16], Cyclic Level Planarity [8, 17], and Torus Level Planarity [8] arrange levels on a standing cylinder, a rolling cylinder, and a torus, respectively.

Proper instances. An important special case are *proper* level graphs, that is, level graphs in which $\gamma(v) = \gamma(u) + 1$ for every edge $(u, v) \in E$. Instances of Level Planarity can be assumed to be proper without loss of generality by subdividing long edges [83]. However, in variations of Level Planarity where we impose additional constraints, the assumption that instances are proper can have a strong impact on the complexity of the respective problems [9]. The definition of proper instances naturally extends to the following variations of level graphs.

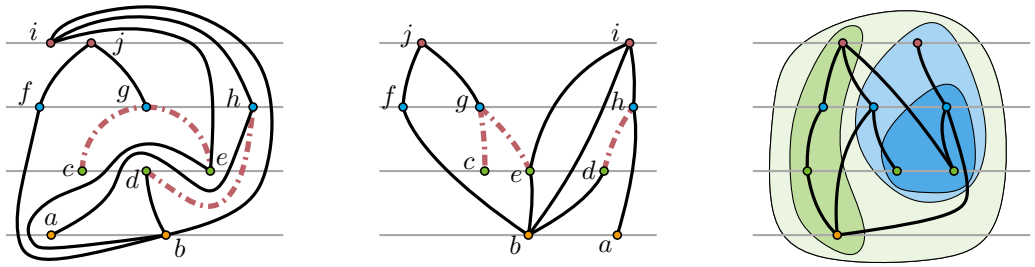

Figure 6: In Level Planarity the order of the vertices of a common level V_i is not fixed. Finding a good ordering is an essential part of finding a solution. The ordering suggested in (a) is not realizable since the edge (d, h) cannot be drawn without crossing (c, g) or (e, g). (b) A level planar drawing of (a). (c) A clustered level drawing.

3.1.2 Level Planarity with various constraints

Clustered Level Planarity. Forster and Bachmaier [58] introduced a version of Level Planarity that allows the visualization of vertex clusterings. A *clustered* level graph $\mathcal{G}$ is a triple $(G = (V, E), \gamma, T)$ where (G, γ) is a level graph and T is a *cluster hierarchy*, i.e. a rooted tree whose leaves are the vertices in V. Each internal node of T is called a *cluster*. We call the cluster of the root *trivial* as it contains all vertices. All other clusters are called *non-trivial*. The *vertices* of a cluster c are the leaves of the subtree of T rooted at c. A cluster hierarchy is *flat* if all leaves have distance at most two from the root, i.e. if non-trivial clusters are not nested. A *clustered* level planar drawing of a clustered level graph $\mathcal{G}$ is a level planar drawing of (G, γ) together with a closed simple curve for each cluster that encloses precisely the vertices of the cluster such that the following three conditions hold: (i) no two cluster boundaries intersect, (ii) every edge crosses each cluster boundary at most once, and (iii) the intersection of any cluster with the horizontal line L_i through level V_i is either a line segment or empty for any level V_i, see Figure 6(c).

The problem Clustered Level Planarity asks whether a given clustered level graph has a clustered level planar drawing. Forster and Bachmaier [58] presented an $O(h|V|)$-time algorithm for a special case of proper clustered level graphs, where h is

the height of $\mathcal{G}$. Angelini, Da Lozzo, Di Battista, Frati, and Roselli [9] provided a quartic-time algorithm for all proper instances. The general version of Clustered Level Planarity is $\mathcal{NP}$-complete even for clustered level graphs with maximum degree $\Delta = 2$ and level-width $\lambda = 3$, and for 2-connected series-parallel clustered level graphs [9]. In this thesis, we further strengthen these previous results (Theorem 7).

T-Level Planarity. This variation of Level Planarity considers consecutivity constraints for the vertices on each level. A T-*level* graph $\mathcal{G}$ is a triple $(G = (V, E), \gamma, \mathcal{T})$ where (G, γ) is a level graph and $\mathcal{T} = (T_0, \ldots, T_h)$ is a set of trees where the leaves of T_i are V_i. A T-*level planar* drawing of a T-level graph $\mathcal{G}$ is a level planar drawing of (G, γ) such that, for every level V_i and for each node u of T_i, the leaves of the subtree of T_i rooted at u appear consecutively along L_i.

The problem T-Level Planarity asks whether a given T-level graph has a T-level planar drawing. Wotzlaw, Speckenmeyer, and Porschen [129] introduced the problem and provided a quadratic-time algorithm for proper instances with constant level-width. Angelini et al. [9] give a quartic-time algorithm for proper instances with unbounded level-width. For general T-level graphs the problem is $\mathcal{NP}$-complete [9] even for T-level graphs with maximum degree $\Delta = 2$ and level-width $\lambda = 3$, and for 2-connected series-parallel T-level graphs.

Constrained Level Planarity. Recently, Brückner and Rutter [27] explored a variant of Level Planarity in which the left-to-right order of the vertices on each level has to be a linear extension of a given partial order. They refer to this problem as Constrained Level Planarity and they provide an efficient algorithm for single-source level graphs and show $\mathcal{NP}$-completeness for connected proper level graphs.

3.1.3 A common special case: Ordered Level Planarity

We introduce a natural variant of Level Planarity that specifies a total order for the vertices on each level. An *ordered* level graph $\mathcal{G}$ is a triple $(G = (V, E), \gamma, \chi)$ where (G, γ) is a level graph and $\chi \colon V \to \{0, \ldots, \lambda - 1\}$ is a *level ordering* for G. We require that χ maps each level V_i $(= \gamma^{-1}(i))$ bijectively to $\{0, \ldots, \lambda_i - 1\}$. An *ordered* level planar drawing of an ordered level graph $\mathcal{G}$ is a level planar drawing of (G, γ) where for every $v \in V$ the x-coordinate of v is $\chi(v)$. Thus, the position of every vertex is fixed.

The problem Ordered Level Planarity asks whether a given ordered level graph has an ordered level planar drawing, see Figures 7(a–b). In this thesis, we show that Ordered Level Planarity is a common special case of all the Level Planarity variants defined in Section 3.1.2 (Theorem 5); and we provide a complexity dichotomy with respect to both the level-width and the maximum degree (Theorem 1).

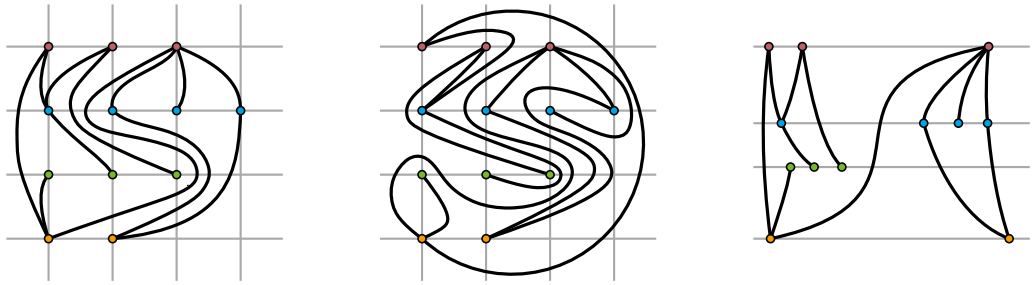

(a) (b) (c)

Figure 7: (a) An ordered level drawing of the instance given in (b). (c) An equivalent drawing for the relaxed version of the problem.

Order and realizability. In the above definition, the x-coordinates assigned via χ merely act as a convenient way to encode a total order for the vertices of each level V_i. Similarly, the y-coordinates assigned via γ encode a total preorder (i.e. a total ordering that allows ties) for the set of all vertices. In terms of realizability, the problem is equivalent to a generalized version where χ and γ range over arbitrary real numbers. In other words, the fixed vertex positions can be any points in the plane. All reductions and algorithms in this chapter carry over to these generalized versions, if we pay the cost for presorting the vertices according to their coordinates. There is another equivalent version that is even more relaxed: we only require that the vertices appear according to the prescribed orderings without insisting on specific coordinates, see Figures 7(a–c). For the sake of visual clarity, many of the figures in this chapter make use of this last equivalence, i.e. the vertices are arranged according to the orderings, but do not necessarily appear at the corresponding exact coordinates.

3.1.4 Geodesic Planarity and Bi-Monotonicity

In this section, we discuss the problems Geodesic Planarity and Bi-Monotonicity. These problems are closely related, as explained at the end of the section.

Geodesic Planarity. Let $S \subset \mathbb{Q}^2$ be a finite set of directions that is symmetric with respect to the origin, i.e. for each direction $s \in S$, the reverse direction $(-s)$ is also contained in S. A planar drawing of a graph is *geodesic* with respect to S if every edge is realized as a polygonal path p composed of line segments with two adjacent directions from S. Two directions of S are *adjacent* if they appear consecutively in the projection of S to the unit circle. The name *geodesic* comes from the fact that such a path p is a shortest path with respect to some polygonal norm (a norm whose unit ball is a centrally symmetric polygon), which depends on S.

An instance of the decision problem Geodesic Planarity is a 4-tuple $\mathcal{G} = (G = (V, E), x, y, S)$ where G is a graph, x and y map from V to the reals, and S is a set of directions as stated above. The task is to decide whether $\mathcal{G}$ has a *geodesic drawing*,

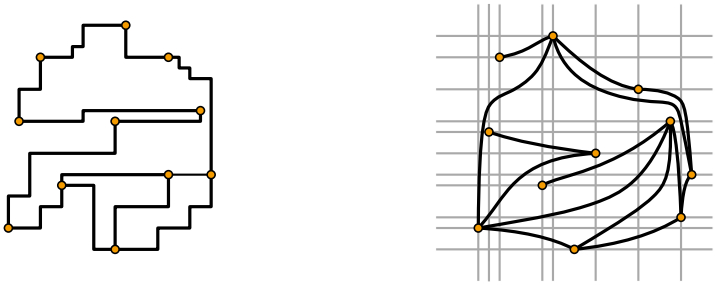

(a) (b)

Figure 8: (a) A manhattan geodesic drawing, and (b) a bi-monotone drawing.

that is, G has a geodesic drawing with respect to S in which every vertex $v \in V$ is placed at $(x(v), y(v))$.

Katz, Krug, Rutter, and Wolff [86] study Manhattan Geodesic Planarity, which is the special case of Geodesic Planarity where the set S consists of the two horizontal and the two vertical directions, see Figure 8(a). Geodesic drawings with respect to this set of directions are also referred to as *orthogeodesic* drawings [63, 64]. Katz et al. [86] show that a variant of Manhattan Geodesic Planarity in which the drawings are restricted to the integer grid is $\mathcal{NP}$-hard even if G is a perfect matching. The proof is by reduction from 3-Partition and makes use of the fact that the number of edges that can pass between two vertices on a grid line is bounded. In contrast, it is claimed in [86] that the standard version of Manhattan Geodesic Planarity is polynomial-time solvable for perfect matchings [86, Theorem 5]. To this end, the authors sketch a plane sweep algorithm that maintains a linear order among the edges that cross the sweep line. When a new edge is encountered it is inserted as low as possible subject to the constraints implied by the prescribed vertex positions. When we asked the authors for more details, they informed us that they are no longer convinced of the correctness of their approach. Theorem 2 of this thesis implies that the approach is indeed incorrect unless $\mathcal{P} = \mathcal{NP}$.

Bi-Monotonicity. Fulek, Pelsmajer, Schaefer, and Štefankovič [59] study so-called *y-monotone* drawings, that is, upward drawings in which all vertices have distinct y-coordinates. They present a Hanani-Tutte theorem for these types of drawings and accompany their result with a simple and efficient algorithm for Y-Monotonicity, which can be defined as (Ordered) Level Planarity restricted to instances with level-width $\lambda = 1$. Moreover, they show that, even without the restriction on λ, Level Planarity is equivalent to Y-Monotonicity by providing an efficient reduction from Level Planarity. Altogether, this results in a simple quadratic time algorithm for Level Planarity.

Fulek et al. [59] propose the problem Bi-Monotonicity and leave its complexity as an open problem. Bi-Monotonicity combines Y-Monotonicity and

X-Monotonicity, which is defined analogously to Y-Monotonicity. More precisely, the input of Bi-Monotonicity is a triple $\mathcal{G} = (G = (V, E), x, y)$ where G is a graph, and x and y are *injective* maps from V to the reals. The task is to decide whether $\mathcal{G}$ has a planar *bi-monotone* drawing, that is, a planar drawing in which edges are realized as curves that are both x-monotone and y-monotone, and in which every vertex $v \in V$ is placed at $(x(v), y(v))$, see Figure 8(b).

Comparison. Bi-Monotonicity is very similar to Manhattan Geodesic Planarity. One difference is that Manhattan Geodesic Planarity imposes an implicit bound on the number of adjacent edges leading in similar directions, i.e. a vertex can have at most two neighbors in a single quadrant. The overall degree of each vertex is at most four. Another difference is that Bi-Monotonicity requires the coordinate mappings x and y to be injective. When both these additional constraints are satisfied, the problems are equivalent. In this thesis, we exploit this relationship between the two problems to settle the question by Fulek et al. [59] regarding the complexity of Bi-Monotonicity (Theorem 3).

3.1.5 Main results

In Section 3.4 we study the complexity of Ordered Level Planarity. While Upward Planarity is $\mathcal{NP}$-complete [62] in general but becomes polynomial-time solvable [83] for prescribed y-coordinates, we show that prescribing both x-coordinates and y-coordinates renders the problem $\mathcal{NP}$-complete. We complement our result with efficient approaches for some special cases of ordered level graphs and, thereby, establish a complexity dichotomy with respect to the level-width and the maximum degree.

Theorem 1. Ordered Level Planarity *is $\mathcal{NP}$-complete, even for acyclic ordered level graphs with maximum degree $\Delta = 2$ and level-width $\lambda = 2$. The problem can be solved in linear time if the given level graph is proper, or if the level-width is $\lambda = 1$, or if $\Delta^+ = \Delta^- = 1$, where Δ^+ and Δ^- are the maximum in-degree and out-degree respectively.*

Ordered Level Planarity, especially if restricted to instances with $\lambda = 2$ and $\Delta = 2$, is an elementary problem that readily reduces to several other graph drawing problems. The remainder of this chapter is dedicated to demonstrating the centrality of Ordered Level Planarity by providing reductions to all the problems listed in Sections 3.1.2 and 3.1.4. All these reductions heavily rely on either a small value of Δ or λ and they produce very constrained instances of the targeted problems. Thereby, we are able to solve multiple open questions that were posed by the graph drawing community. We expect that Theorem 1 may serve as a suitable basis for more reductions in the future.

In Section 3.2 we study Geodesic Planarity and obtain:

Theorem 2. Geodesic Planarity *is $\mathcal{NP}$-hard for any set of directions S with $|S| \geq 4$, even for perfect matchings in general position.*

Observe the aforementioned discrepancy between Theorem 2 and the claim in [86] that Manhattan Geodesic Planarity for perfect matchings is in $\mathcal{P}$.

Bi-Monotonicity is closely related to a special case of Manhattan Geodesic Planarity. With a simple corollary we settle the complexity of Bi-Monotonicity and, thus, answer the open question by Fulek et al. [59].

Theorem 3. Bi-Monotonicity *is $\mathcal{NP}$-hard, even for perfect matchings.*

Theorem 4. Ordered Level Planarity *reduces to* Bi-Monotonicity *in linear time. The reduction can be carried out such that the input graph is identical to the output graph, that is, only the coordinates are modified.*

In Section 3.3 we establish Ordered Level Planarity as a special case of all the variations of Level Planarity described in Section 3.1.2.

Theorem 5. Ordered Level Planarity *reduces in linear time to* Constrained Level Planarity *and* T-Level Planarity*, and in quadratic time to* Clustered Level Planarity*.*

The reduction to Constrained Level Planarity is immediate, which also yields:

Theorem 6. Constrained Level Planarity *is $\mathcal{NP}$-hard even for acyclic level graphs with maximum degree $\Delta = 2$ and level-width $\lambda = 2$ and prescribed* total *orderings.*

Angelini, Da Lozzo, Di Battista, Frati, and Roselli [9] propose the complexity of Clustered Level Planarity for clustered level graphs with a flat cluster hierarchy as an open question. Our reduction to Clustered Level Planarity provides the following answer.

Theorem 7. Clustered Level Planarity *is $\mathcal{NP}$-hard even for acyclic clustered level graphs with maximum degree $\Delta = 2$, level-width $\lambda = 2$ and a flat cluster hierarchy that partitions the vertices into two non-trivial clusters.*

In general, we can consider two different versions of all of the above problems: we may prescribe a combinatorial embedding or allow an arbitrary embedding. Our results apply to both of these versions, as in most cases the instances are just disjoint unions of paths and, thus, the embedding is unique. The only exception is the linear time algorithm for proper instances of Ordered Level Planarity. However, in this case, yes-instances have a (combinatorially) unique drawing and we only need to check if it respects the given embedding.

To reduce from Ordered Level Planarity to Geodesic Planarity, our main reduction (to Ordered Level Planarity) is tailored to achieve a small maximum degree of $\Delta = 2$. As a consequence, the resulting graphs are not connected. At the cost of an increased maximum degree, it is possible to make our instances connected by inserting additional edges. We discuss these adaptations in Section 3.5.

Theorem 8. *The following problems are $\mathcal{NP}$-hard even for connected instances with maximum degree $\Delta = 4$:*

- Ordered Level Planarity*, even for level-width $\lambda = 2$;*
- Constrained Level Planarity*, even for level-width $\lambda = 2$ and prescribed* total *orderings;*
- Clustered Level Planarity*, even for level-width $\lambda = 2$ and flat cluster hierarchies that partition the vertices into two non-trivial clusters; and*
- Bi-Monotonicity*.*

3.2 Reducing to Geodesic Planarity & Bi-Monotonicity

In this section, we establish that deciding whether an instance $\mathcal{G} = (G, x, y, S)$ of Geodesic Planarity has a geodesic drawing is $\mathcal{NP}$-hard, even if G is a perfect matching and the coordinates assigned via x and y are in *general position*, that is, no two vertices lie on a common line with a direction from S. The $\mathcal{NP}$-hardness of Bi-Monotonicity for perfect matchings follows as a simple corollary. Our results are obtained via a reduction from Ordered Level Planarity.

3.2.1 Reducing to Geodesic Planarity

We start with our reduction to Geodesic Planarity.

Lemma 9. *Let $S \subset \mathbb{Q}^2$ with $|S| \geq 4$ be a finite set of directions which is symmetric with respect to the origin.* Ordered Level Planarity *with maximum degree $\Delta = 2$ and level-width $\lambda = 2$ reduces to* Geodesic Planarity *such that the resulting instances are in general position and consist of a perfect matching and direction set S. The reduction can be carried out using a linear number of arithmetic operations.*

Proof. We first prove our claim for the classical case that S contains exactly the four horizontal and vertical directions. Afterwards, we discuss the necessary adaptations for the general case.

Our reduction is carried out in two steps. Let $\mathcal{G}_o = (G_o = (V, E), \gamma, \chi)$ be an Ordered Level Planarity instance with maximum degree $\Delta = 2$ and level-width $\lambda = 2$. In Step (i) we turn $\mathcal{G}_o$ into an equivalent Geodesic Planarity instance $\mathcal{G}'_g = (G_o, x', \gamma, S)$. In Step (ii) we transform $\mathcal{G}'_g$ into an equivalent Geodesic

Planarity instance $\mathcal{G}_g = (G_g, x, y, S)$ where G_g is a perfect matching and the vertex positions assigned via x and y are in general position.

Step (i): To transform $\mathcal{G}_o$ into $\mathcal{G}'_g$, we apply a horizontal shearing transformation to the vertex positions specified by χ and γ. More precisely, for every $v \in V$ we define $x'(v) = \chi(v) + 2\gamma(v)$, see Figures 9(a) and 9(b).

Clearly, every geodesic drawing of $\mathcal{G}'_g$ can be turned into an ordered level planar drawing of $\mathcal{G}_o$. For the other direction, consider an ordered level planar drawing Γ_o of $\mathcal{G}_o$. Without loss of generality, we can assume that in Γ_o all edges are realized as polygonal paths in which bend points occur only on the horizontal lines L_i through the levels V_i where $0 \leq i \leq h$. Further, since $\chi(V) \subseteq \{0, 1\}$ we may assume that all bend points have x-coordinates in the open interval $(-\frac{1}{2}, \frac{3}{2})$. We shear Γ_o by translating the bend points and vertices of level V_i by $2i$ units to the right for $0 \leq i \leq h$, see Figure 9(b). In the resulting drawing Γ'_o, the vertex positions match those of $\mathcal{G}'_g$. Furthermore, all edge-segments have a positive slope. Thus, since the maximum degree is $\Delta = 2$, we can replace all edge-segments with L_1-geodesic rectilinear paths that closely trace the segments and we obtain a geodesic drawing Γ'_g of $\mathcal{G}'_g$, see Figure 9(c).

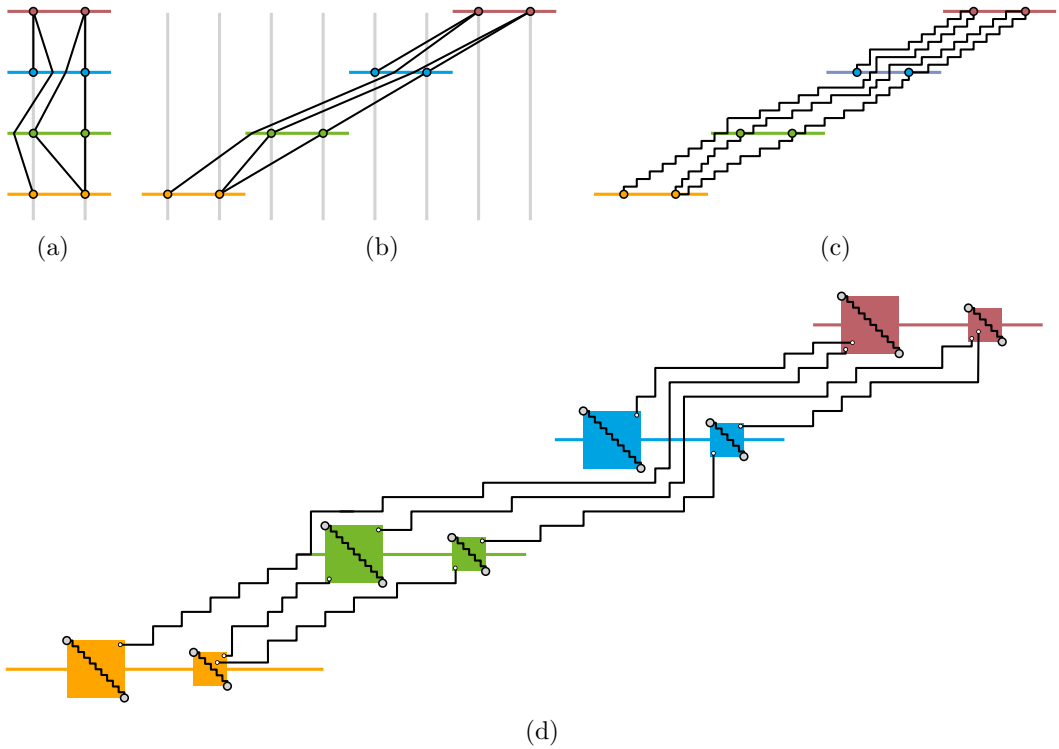

Figure 9: The reduction from Ordered Level Planarity to Manhattan Geodesic Planarity in Lemma 9 is carried out in two steps. Step (i) is illustrated in (a)–(c), and Step (ii) is illustrated in (d).

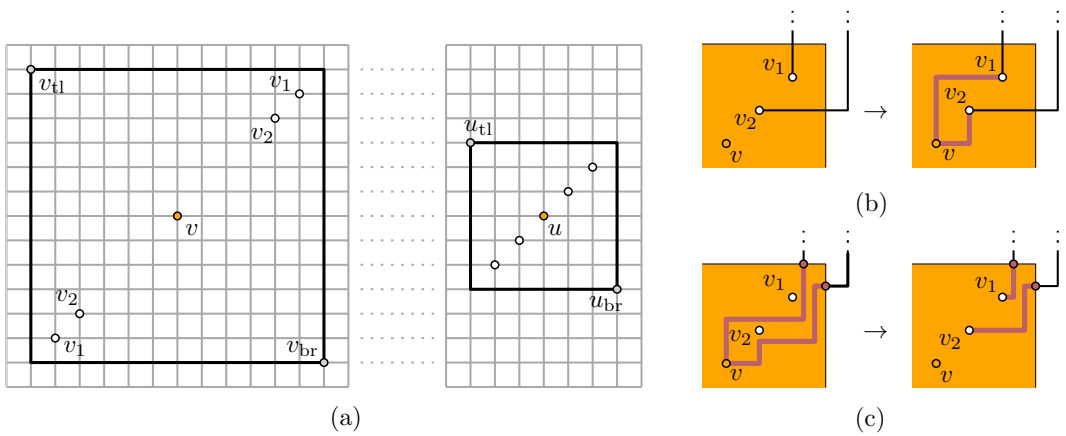

Figure 10: (a) The two gadget squares of each level. Grid cells have size $\frac{1}{48} \times \frac{1}{48}$. (b) Turning a drawing of $\mathcal{G}_g$ into a drawing of $\mathcal{G}'_g$ and (c) vice versa.

Step (ii): To turn $\mathcal{G}'_g = (G_o = (V, E), x', \gamma, S)$ into the equivalent instance $\mathcal{G}_g = (G_g, x, y, S)$, we transform G_o into a perfect matching. To this end, we split each vertex $v \in V$ by replacing it with a small gadget that fits inside a square r_v centered on the position $p_v = (x'(v), \gamma(v))$ of v, see Figure 9(d). We call r_v the *square* of v and use p_v^{tr}, p_v^{tl}, p_v^{br}, and p_v^{bl} to denote the top-right, top-left, bottom-right, and bottom-left corner of r_v, respectively. We use two different sizes to ensure general position. The size of the gadget square is $\frac{1}{4} \times \frac{1}{4}$ if $\chi(v) = 0$, and it is $\frac{1}{8} \times \frac{1}{8}$ if $\chi(v) = 1$. The gadget contains a degree-1 vertex for every edge incident to v.

In the following we explain the gadget construction in detail, for an illustration see Figure 10(a). Let $\{v, u\}$ be an edge incident to v. We create an edge $\{v_1, u\}$ where v_1 is a new vertex that is placed at $(p_v^{\text{tr}} - (\frac{1}{48}, \frac{1}{48}))$ if u is located to the top-right of v and it is placed at $(p_v^{\text{bl}} + (\frac{1}{48}, \frac{1}{48}))$ if u is located to the bottom-left of v. Similarly, if v is incident to a second edge $\{v, u'\}$, we create an edge $\{v_2, u'\}$ where v_2 is placed at $(p_v^{\text{tr}} - (\frac{1}{24}, \frac{1}{24}))$ or $(p_v^{\text{bl}} + (\frac{1}{24}, \frac{1}{24}))$ depending on the position of u'. We refer to v_1 and v_2 as the *gadget vertices* of v and its square r_v. Finally, we create two new vertices v_{tl} and v_{br} and a *blocking* edge $\{v_{\text{tl}}, v_{\text{br}}\}$ where v_{tl} is placed at p_v^{tl} and v_{br} is placed at p_v^{br}. All the assigned coordinates are distinct in both components, and hence the points are in general position. The construction can be carried out in linear time.

Assume that $\mathcal{G}_g$ has a geodesic drawing Γ_g. By construction, for each blocking edge, one of its vertices is located to the top-*left* of the other. In contrast, for each non-blocking edge, one of its vertices is located to the top-*right* of the other. As a result, a non-blocking edge $e = \{v, u\}$ cannot pass through any gadget square r_w where $w \notin \{v, u\}$, since otherwise e would have to cross the blocking edge of r_w. Accordingly, it is straightforward to obtain a geodesic drawing of Γ'_g: we remove the blocking edges, reinsert the vertices of V according to the mappings x' and γ and connect them to the gadget vertices of their respective squares in a geodesic fashion.

This can always be done without crossings. Figure 10(b) shows one possibility. If the edge from v_2 passes to the left of v_1, we may have to choose a reflected version. Finally, we remove the vertices v_1 and v_2, which now act as subdivision vertices.

For the other direction, let Γ'_g be a geodesic planar drawing of $\mathcal{G}'_g$. Without loss of generality, we can assume that each edge $\{u, v\}$ intersects only the squares of u and v. Furthermore, for each $v \in V$ we can assume that its incident edges intersect the boundary of r_v only to the top-right of $(p_v^{\text{tr}} - (\frac{1}{48}, \frac{1}{48}))$ or to the bottom-left of $(p_v^{\text{bl}} + (\frac{1}{48}, \frac{1}{48}))$, see Figure 10(c). Thus, we can simply replace the parts of the edges *inside* the gadget squares by connections to the gadget vertices v_1 and v_2 in a geodesic fashion, see Figure 10(c).

This concludes the proof for the case that S contains exactly the four horizontal and vertical directions.

The general case. It remains to discuss the adaptations for the case that S is an arbitrary set of directions that is symmetric with respect to the origin. By applying a linear transformation we can assume without loss of generality that $(1, 0)$ and $(0, 1)$ are adjacent directions in S. Accordingly, all the remaining directions point into the top-left or the bottom-right quadrant. Further, by vertical scaling we can assume that no direction is parallel to $(1, -1)$. Observe that if we do not insist on a coordinate assignment in general position, the reduction for the restricted case discussed above is already sufficient.

To guarantee general position, we have to avoid *conflicting* vertices, i.e. distinct vertices whose positions lie on a common line with a direction from S. This requires some simple but somewhat technical modifications of our construction.

Let s_1 be the flattest slope of any direction in $S \setminus \{(1, 0), (0, 1)\}$, i.e. the slope with the smallest absolute value (note that all the slopes are negative). Further, let s_2 be the steepest slope of any direction in $S \setminus \{(1, 0), (0, 1)\}$, i.e. the slope with the largest absolute value.

Assume that c', d' are conflicting vertices such that c' belongs to the gadget square r_c of $c \in V$, and d' belongs to the gadget square r_d of $d \in V$. Consider Figure 11(a). Since no direction of S points to the top-right or bottom-left quadrant, $\gamma(c) = \gamma(d)$. It is possible that $c = d$.

To guarantee general position, we apply the following two modifications.

Modification (a): We first cover the case $c \neq d$, that is, we show how to avoid conflicts between two vertices c', d' that belong to *distinct* squares of the same level. To this end, we simply scale the positions of the gadget squares in the x-direction. More precisely, instead of using the coordinates $(2i, i)$ and $(2i + 1, i)$ for the centers of the two squares r_v and r_u of level i, we use the positions $(2ki, i)$ and $(2ki + k, i)$ where $k \geq 1$ is chosen large enough that p_u^{bl} is above the line ℓ with slope s_1 through p_v^{tr}, see Figure 11(a).

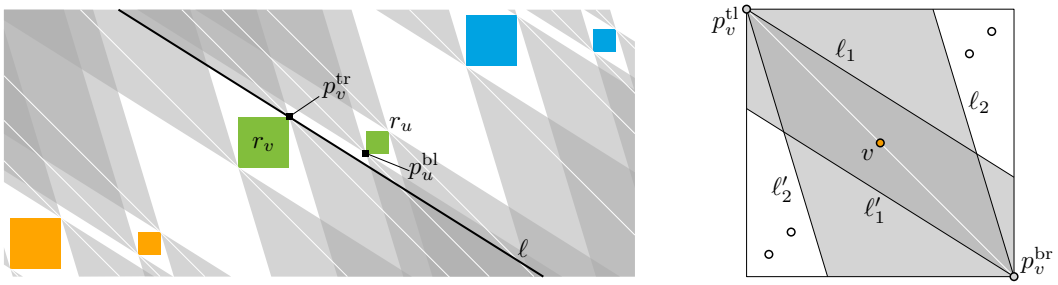

(a) (b)

Figure 11: Modifications (a) and (b) for the general case. The shaded regions indicate the union of all lines with directions from S that pass through (a) the top-right or bottom-left corner or (b) the top-left or bottom-right corner of a gadget square.

Modification (b): It remains to cover the case $c = d$, i.e. to avoid conflicts between vertices c', d' that belong to the same gadget square r_v. To this end, we modify the placement of the gadget vertices inside the gadget squares as follows. We change the offset to the gadget square corners from $\pm(\frac{1}{48})$ and $\pm(\frac{1}{24})$ to $\pm(\frac{z}{48})$ and $\pm(\frac{z}{24})$ where $0 < z < 1$ is chosen small enough such that each gadget vertex is placed either

- above the line ℓ_1 with slope s_1 through p_v^{tl} and above the line ℓ_2 with slope s_2 through p_v^{br}; or
- below the line ℓ_1' with slope s_1 through p_v^{br} and below the line ℓ_2' with slope s_2 through p_v^{tl};

see the white regions in Figure 11(b). □

The bit size of the numbers involved in the calculations of our reduction is linearly bounded in the bit size of the directions of S. Together with Theorem 1 we obtain the proof of Theorem 2.

Theorem 2. Geodesic Planarity *is $\mathcal{NP}$-hard for any set of directions S with $|S| \geq 4$, even for perfect matchings in general position.* □

3.2.2 Reducing to Bi-Monotonicity

The instances generated by Lemma 9 are in general position. In particular, this means that the mappings x and y are injective. We obtain an immediate reduction to Bi-Monotonicity. The correctness follows from the fact that every L_1-geodesic rectilinear path can be transformed into a bi-monotone curve and vice versa. Thus, we obtain Theorem 3.

Theorem 3. Bi-Monotonicity *is $\mathcal{NP}$-hard, even for perfect matchings.* □

By combining Lemma 9 and the remarks in the previous paragraph, we obtain a reduction from Ordered Level Planarity to Bi-Monotonicity. However, the intermediate reduction via Manhattan Geodesic Planarity requires the original Ordered Level Planarity instance to have a maximum out-degree of $\Delta^+ \leq 2$ and a maximum in-degree of $\Delta^- \leq 2$ (otherwise, our reduction would produce Manhattan Geodesic Planarity instances with vertices that have more than two neighbors in the same quadrant and such instances are never realizable, see Section 3.1.4).

In Section 3.5, we require a reduction that accepts more general instances of Ordered Level Planarity. For this reason, we state the following direct (and, in fact, much simpler) reduction from Ordered Level Planarity to Bi-Monotonicity.

Theorem 4. Ordered Level Planarity *reduces to* Bi-Monotonicity *in linear time. The reduction can be carried out such that the input graph is identical to the output graph, that is, only the coordinates are modified.*

Proof. Let $\mathcal{G} = (G = (V, E), \gamma, \chi)$ be an ordered level graph with level-width λ and height h. We create an instance of Bi-Monotonicty as follows.

The graph G remains unchanged. The new vertex-coordinates are obtained by applying the following linear function f to the assignment given by χ and γ. The function f is a linear deformation of the plane that scales the original coordinates and rotates them by 45°, see Figure 12.

$$f(x, y) := (f_1(x, y), f_2(x, y)) := \Big((\lambda + 1)y + x,\ (\lambda + 1)y - x\Big)$$

We define a coordinate assignment (x', y') with $(x'(v), y'(v)) := f(\chi(v), \gamma(v))$ for each vertex $v \in V$. The resulting Bi-Monotonicity instance is $\mathcal{G}' = (G, x', y')$ with $x'(v) = (\lambda + 1)\gamma(v) + \chi(v)$ and $y'(v) = (\lambda + 1)\gamma(v) - \chi(v)$.

Recall that L_i denotes the horizontal line with y-coordinate i, which passes through all the vertices of level V_i. We use $S_i \subset L_i$ to denote the *open* line segment between the points $(-1, i)$ and (λ, i). The correctness of our reduction relies on the following property:

Observation 10. *Let $p_i \in f(S_i)$ and $p_{i+1} \in f(S_{i+1})$ for some $0 \leq i < \lambda$. Then $p_i < p_{i+1}$, componentwise.*

The correctness of Observation 10 follows from the fact that for $(j, i) = f^{-1}(p_i)$

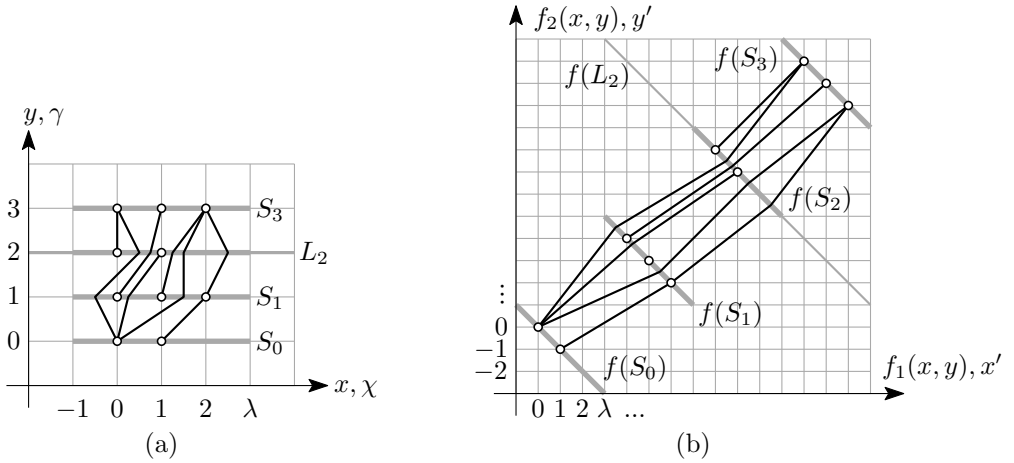

Figure 12: (a) An ordered level planar drawing of $\mathcal{G}$, and (b) the corresponding bi-monotone drawing of $\mathcal{G}'$.

and $(j', i+1) = f^{-1}(p_{i+1})$ we have:

$$\begin{aligned} p_i &= f(j,i) \\ &< \Big((\lambda+1)i + \lambda,\ (\lambda+1)i + 1\Big) \\ &= \Big((\lambda+1)(i+1) - 1,\ (\lambda+1)(i+1) - \lambda\Big) \\ &< f(j', i+1) \\ &= p_{i+1} \end{aligned}$$

Let Γ be an ordered level planar drawing of $\mathcal{G}$. Without loss of generality, we can assume that in Γ all edges are realized as polygonal paths in which bend-points occur only on the horizontal segments S_i, see Figure 12(a). Applying f to all the bend-points yields a drawing $f(\Gamma)$ of $\mathcal{G}'$, see Figure 12(b). Since f is linear, $f(\Gamma)$ is planar. By Observation 10, every edge in $f(\Gamma)$ is realized as a polygonal path whose segments have positive slopes. Therefore $f(\Gamma)$ is bi-monotone.

For the other direction, let Γ' be a planar bi-monotone drawing of $\mathcal{G}'$. The lines $f(L_i) \supset f(S_i)$ have a negative slope (of -1) and, by Observation 10, every edge is realized as a curve that is simultaneously increasing in the x- and y-directions. Therefore, every edge may intersect each line $f(L_i)$ at most once. More precisely, an edge (v_j, v_k) with $v_j \in V_j$, $v_k \in V_k$ and $j < k$ crosses each of the consecutive lines $f(L_{j+1}), ..., f(L_{k-1})$ exactly once. Further, all vertices of level V_i have been mapped to $f(S_i) \subset f(L_i)$. Thus, we can leave the intersection of each edge with each line $f(L_i)$ fixed and replace the intermediate pieces by line-segments. This does not introduce any crossings and turns all edges into x- and y-monotone polygonal paths in which bend-points occur only on the lines $f(L_i)$, see Figure 12(b). Applying f^{-1} yields an ordered level planar drawing $f^{-1}(\Gamma')$ of $\mathcal{G}$, see Figure 12(a). □

3.3 Reductions to variations of Level Planarity

In this section we explore the connection between Ordered Level Planarity and other variants of Level Planarity. We prove the following theorem.

Theorem 5. Ordered Level Planarity *reduces in linear time to* Constrained Level Planarity *and* T-Level Planarity*, and in quadratic time to* Clustered Level Planarity*.*

3.3.1 Reducing to Constrained Level Planarity

The reduction to Constrained Level Planarity is immediate, which together with Theorem 1 also yields:

Theorem 6. Constrained Level Planarity *is $\mathcal{NP}$-hard even for acyclic level graphs with maximum degree $\Delta = 2$ and level-width $\lambda = 2$ and prescribed* total *orderings.* □

3.3.2 Reducing to T-Level Planarity

In this section, we reduce from Ordered Level Planarity to T-Level Planarity. We restrict our attention to ordered level graphs with level-width $\lambda = 2$. As we will see in Section 3.4, this restriction is no loss of generality (Lemma 14).

Lemma 11. Ordered Level Planarity *with maximum degree Δ and level-width $\lambda = 2$ reduces in linear time to* T-Level Planarity *with maximum degree $\Delta' = \max(\Delta, 2)$ and level-width $\lambda' = 4$.*

Proof. Let $\mathcal{G} = (G = (V, E), \gamma, \chi)$ be an ordered level graph with maximum degree Δ and level-width $\lambda = 2$. We augment each level V_i with $|V_i| = 1$ by adding an isolated dummy vertex v with $\gamma(v) = i$ and $\chi(v) = 1$ to avoid having to treat special cases. Thus, each level V_i has a vertex v_i^0 with $\chi(v_i^0) = 0$ and a vertex v_i^1 with $\chi(v_i^1) = 1$. The following steps are illustrated in Figure 13. For each level V_i we create two new vertices v_i^l and v_i^r. We add edges (v_i^l, v_{i+1}^l) and (v_i^r, v_{i+1}^r) for $i = 0, \dots, h-1$, where h is the height of $\mathcal{G}$. Hence, we obtain a path p_l from v_0^l to v_h^l and a path p_r from v_0^r to v_h^r. The root r_i of each tree T_i has two children u_i^l and u_i^r. The two children of u_i^l are v_i^l and v_i^0. The two children of u_i^r are v_i^r and v_i^1. Let $\mathcal{G}'$ denote the resulting T-level graph. The construction of $\mathcal{G}'$ can be carried out in linear time.

Clearly, an ordered level planar drawing Γ of $\mathcal{G}$ can be augmented to a T-level planar drawing of $\mathcal{G}'$ by drawing p_l to the left of Γ and by drawing p_r to the right of Γ. For the other direction, let Γ' be a T-level-planar drawing of $\mathcal{G}'$. We can assume without loss of generality that all vertices are placed on vertical lines with x-coordinates -1, 0, 1 or 2. The paths p_l and p_r are vertex-disjoint and drawn without crossing. Thus, p_l is drawn either to the left or to the right of p_r. By

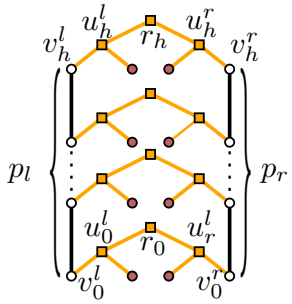

Figure 13: Reduction from Ordered Level Planarity to T-level Planarity. The square vertices illustrate each level's tree.

reflecting horizontally at the line $x = 1/2$ we can assume without loss of generality that p_l is drawn to the left of p_r. Consequently, for each level V_i the vertex v_i^0 has to be drawn to the left of the vertex v_i^1 since v_i^l and v_i^0 are the children of u_i^l and since v_i^r and v_i^1 are the children of u_i^r. Therefore, the subdrawing of G or its mirror image is an ordered level planar drawing of $\mathcal{G}$. □

Together with Theorem 1 this shows the $\mathcal{NP}$-hardness of T-Level Planarity for instances with maximum degree $\Delta = 2$ and level-width $\lambda = 4$. However, a stronger statement was already given by Angelini et al. [9], who show $\mathcal{NP}$-hardness for instances with $\Delta = 2$ and $\lambda = 3$.

3.3.3 Reducing to Clustered Level Planarity

We proceed with a reduction to Clustered Level Planarity. Like in the previous section, we restrict our attention to ordered level graphs with level-width $\lambda = 2$, which is no loss of generality due to Lemma 14.

Lemma 12. Ordered Level Planarity *with maximum degree* Δ *and level-width* $\lambda = 2$ *reduces in quadratic time to* Clustered Level Planarity *with maximum degree* $\Delta' = \max(\Delta, 2)$*, level-width* $\lambda' = 2$*, and a clustering hierarchy that partitions the vertices into only two non-trivial clusters.*

Proof. Let $\mathcal{G} = (G = (V, E), \gamma, \chi)$ be an ordered level graph with maximum degree Δ and level-width $\lambda = 2$. As in the proof of the previous Lemma 11, we augment each level V_i with $|V_i| = 1$ by adding an isolated dummy vertex v with $\gamma(v) = i$ and $\chi(v) = 1$. Thus, each level V_i has a vertex v_i^0 with $\chi(v_i^0) = 0$ and a vertex v_i^1 with $\chi(v_i^1) = 1$. In addition to the trivial cluster that contains all vertices, we create two clusters $c_0 = \{v_0^0, \dots, v_h^0\}$ and $c_1 = \{v_0^1, \dots, v_h^1\}$, where h is the height of $\mathcal{G}$. Now we see the close correspondence between clustered level planar drawings and ordered level planar drawings: the two clusters pass through every level, their boundaries are not allowed to intersect, and they cannot be nested. Thus, by reflecting horizontally if necessary, we can assume without loss of generality that c_0 intersects each level to the

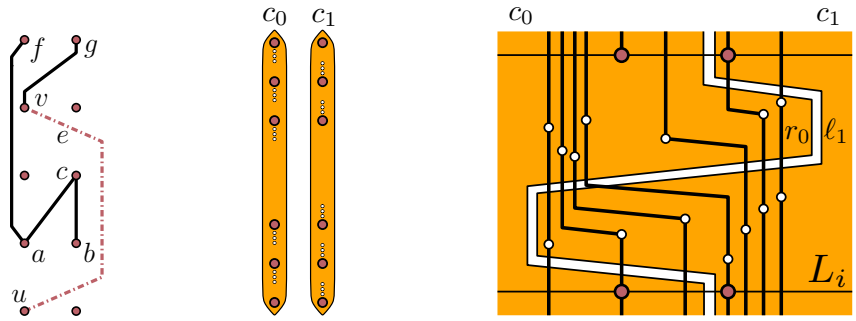

Figure 14: (a) In an ordered level planar drawing, the edge $e = (u, v)$ has to pass the level of b to the right of b: due to the edge (v, g), the edge (a, f) passes to the left of v. As a consequence, e cannot pass the level of b to the left of a. Further, due to (a, c) and (b, c), it can also not pass between a and b. The reduction from Ordered Level Planarity to Clustered Level Planarity is shown in (b) and (c). Big black vertices are the vertices of the Ordered Level Planarity instance. The small vertices are subdivision vertices. (b) Schematic view of the entire clustered level graph. (c) The clustering boundaries can be drawn such that they cross each subdivision edge at most once.

left of c_1 as depicted in Figure 14b. Consequently, on each level V_i the vertex $v_i^0 \in c_0$ is placed to the left of $v_i^1 \in c_1$, just as in an ordered level planar drawing.

To make the reduction work, we have to subdivide each edge several times. Otherwise, an edge might be forced to cross a cluster boundary more than once: consider an edge $e = (u, v)$ with $u, v \in c_0$ that has to pass the level of some vertex $b \in c_1$ with $\gamma(u) < \gamma(b) < \gamma(v)$ to the right of b, see Figure 14a. In this situation, e must cross the right boundary r_0 of c_0 at least twice, as r_0 has to be drawn to the right of $u, v \in c_0$, and to the left of $b \in c_1$. This example can be blown up to enforce arbitrarily many crossings between e and r_0.

To avoid this situation, we subdivide the edges of $\mathcal{G}$ as follows. Each edge from some level i to some level $j > i$ is transformed into a path of $2(j - i) + 1$ edges whose inner vertices alternate between the clusters c_1 and c_0. More precisely, for each pair of consecutive levels V_i and V_{i+1} we add two new subdivision vertices on each edge $e = (u, v) \in E$ with $\gamma(u) \leq i$ and $\gamma(v) \geq i + 1$. The lower one of the resulting subdivision vertices for e is added to c_1, the upper one is added to c_0. We place each subdivision vertex that was added to c_1 on a new separate level between the levels V_i and V_{i+1}. The relative order of these new levels is arbitrary. Above these new levels but below V_{i+1} we place all the subdivision vertices added to c_0, again each on a new separate level, see Figures 14b–14c.

Let $\mathcal{G}^s = (G^s, \gamma^s, \chi^s)$ denote the ordered level graph resulting from applying the subdivision to $\mathcal{G}$. The output of our reduction is the clustered level graph

$\mathcal{G}^{\text{cl}} = (G^s, \gamma^s, T)$ where T is the described hierarchy, with the clusters c_0 and c_1. Since edges may stretch over a linear number of levels, the size of G^s can be quadratic in the size of G and, therefore, the construction of $\mathcal{G}^{\text{cl}}$ might require quadratic time.

Correctness. The subdivision does not affect the realizability of $\mathcal{G}$ as an ordered level planar drawing, since every subdivision vertex in $\mathcal{G}^s$ is the singleton vertex of some new level. Therefore, to prove correctness, it suffices to argue that $\mathcal{G}^{\text{cl}}$ is realizable as a clustered level planar drawing if and only if $\mathcal{G}^s$ is realizable as an ordered level planar drawing.

For the easy direction, let Γ^{cl} be a clustered level planar drawing of $\mathcal{G}^{\text{cl}}$. As discussed above, we may assume that c_0 is drawn to the left of c_1. Further, we may assume without loss of generality that all vertices are placed on vertical lines with x-coordinates 0 and 1, and moreover, all subdivision vertices, being singleton vertices of their levels, are placed on $x = 0$. Recall that each vertex v of the original graph is contained in c_0 if $\chi(v) = 0$; and it is contained in c_1 if $\chi(v) = 1$. Thus, by the above assumptions, $v \in V$ is placed on $x = 0$ if $\chi(v) = 0$; and it is placed on $x = 1$ if $\chi(v) = 1$. Therefore, the drawing Γ^{cl} (without the cluster boundaries) is an ordered level planar drawing of $\mathcal{G}^s$.

For the other direction, let Γ be an ordered level planar drawing of the ordered level graph $\mathcal{G}^s$. We create a clustered level planar drawing of $\mathcal{G}^{\text{cl}}$ by adding the cluster boundaries of c_0 and c_1 to Γ. The left boundary ℓ_0 of c_0 is drawn as a vertical line segment to the left of Γ. Analogously, the right boundary r_1 of c_1 is drawn as a vertical line segment to the right of Γ.

It remains to draw the right boundary r_0 of c_0 and the left boundary ℓ_1 of c_1. We draw them from bottom to top. We keep them close together, and they will always cross the same edge in direct succession, see Figure 14c. Assume inductively that r_0 and ℓ_1 have already been drawn in the closed half-plane H_i below the line L_i through the vertices V_i of $\mathcal{G}$, and this subdrawing violates none of the conditions from the definition of a clustered level planar drawing. In particular, r_0 and ℓ_1 are realized as non-crossing y-monotone curves with all vertices of c_0 to the left of r_0, and with all vertices of c_1 to the right of ℓ_1. Moreover, no edge is intersected more than once by any of r_0 or ℓ_1. Further, let E_i be the set of edges of G^s that are intersected by L_i including the edges having their lower endpoint on L_i, but without the edges having their upper endpoint on L_i. We maintain the following two additional inductive assumptions: (a) L_i intersects the edges in E_i and the boundaries r_0 and ℓ_1 in the following left-to-right order (see Figure 14c): (1) all edges $E_\ell \subseteq E_i$ that intersect L_i to the left of v_i^1, (2) the boundary r_0, (3) the boundary ℓ_1, and (4) the remaining edges $E_r = E_i \setminus E_\ell$, i.e. the edges incident to v_i^1, or passing v_i^1 to its right. (b) No edge of E_ℓ has already been crossed by r_0 or ℓ_1 below L_i. Note that these conditions are easily met for $i = 0$.

We describe how the partial drawings of r_0 and ℓ_1 are extended upwards from L_i. For an illustration, see Figure 14c. Each edge in E_i is part of a path that has two

subdivision vertices between L_i and L_{i+1}. The lower of these vertices belongs to c_1, and the upper one belongs to c_0. We draw r_0 and ℓ_1 in a very schematic and simple way. First we cross all edges in E_ℓ from right to left. By assumption (b), this is permitted. We then pass to the left of all the lower subdivision vertices, ensuring that they lie within the cluster boundaries of c_1. We then cross all edges between their two subdivision vertices from left to right, and pass to the right of all the subdivision vertices in c_0. Finally, we cross from right to left all edges which pass L_{i+1} to the right of v^1_{i+1}, and those whose upper endpoint is v^1_{i+1}. It is easy to check that the inductive assumptions hold again for L_{i+1}. Thus, we may iterate this procedure to obtain a clustered level planar drawing of $\mathcal{G}^{\text{cl}}$. □

Together with Theorem 1 we obtain the following.

Theorem 7. Clustered Level Planarity *is $\mathcal{NP}$-hard even for acyclic clustered level graphs with maximum degree $\Delta = 2$, level-width $\lambda = 2$ and a flat cluster hierarchy that partitions the vertices into two non-trivial clusters.* □

The previous $\mathcal{NP}$-hardness result by Angelini et al. [9] holds for instances with $\Delta = 2$ and $\lambda = 3$. Their cluster hierarchies have linear depths. The authors pose the complexity of Clustered Level Planarity for instances with flat cluster hierarchies as an open problem. Theorem 7 gives an answer to this question and improves the previous result by Angelini et al.

3.4 Results for Ordered Level Planarity

In this section we study Ordered Level Planarity. We begin with the $\mathcal{NP}$-hardness proof, which is by reduction from a variant of 3-Satisfiability described in the following.

Planar Monotone 3-Satisfiability. A *monotone* 3-Satisfiability formula is a Boolean 3-Satisfiability formula in which each clause is either *positive* or *negative*, that is, each clause contains either exclusively positive or exclusively negative literals, respectively (we remark that, in the literature, there exist other meanings of *monotone* in the context of Boolean formulas). A *planar* 3-Satisfiability formula $\varphi = (\mathcal{U}, \mathcal{C})$ is a Boolean 3-Satisfiability formula with a set $\mathcal{U}$ of variables and a set $\mathcal{C}$ of clauses such that its *variable–clause graph* $G_\varphi = (\mathcal{U} \uplus \mathcal{C}, E)$ is planar. The graph G_φ is bipartite, i.e. every edge in E is incident to a *clause* vertex from $\mathcal{C}$ and a *variable* vertex from $\mathcal{U}$. Furthermore, we have that $\{c, u\} \in E$ if and only if a literal of variable $u \in \mathcal{U}$ occurs in $c \in \mathcal{C}$. Planar Monotone 3-Satisfiability is a special case of 3-Satisfiability where we are given a planar and monotone 3-Satisfiability formula φ and a *monotone rectilinear representation* $\mathcal{R}$ of the variable-clause graph of φ. The representation $\mathcal{R}$ is a contact representation on an integer grid in which

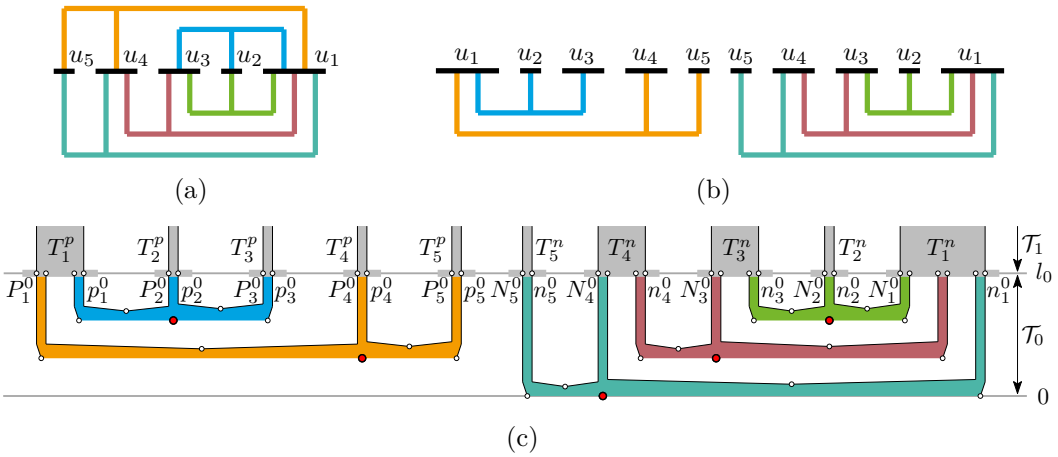

Figure 15: (a) Representation $\mathcal{R}$ of φ with negative clauses $(\overline{u}_1 \vee \overline{u}_4 \vee \overline{u}_5)$, $(\overline{u}_1 \vee \overline{u}_3 \vee \overline{u}_4)$, and $(\overline{u}_1 \vee \overline{u}_2 \vee \overline{u}_3)$ and positive clauses $(u_1 \vee u_4 \vee u_5)$ and $(u_1 \vee u_2 \vee u_3)$ and (b) its modified version $\mathcal{R}'$ in Lemma 13. (c) Tier $\mathcal{T}_0$.

the variables are represented by horizontal line segments arranged on a common horizontal line ℓ. The clauses are represented by E-shapes turned by 90° such that all positive clauses are placed above ℓ and all negative clauses are placed below ℓ, see Figure 15a. Planar Monotone 3-Satisfiability is $\mathcal{NP}$-complete [41].

We are now equipped to prove the core lemma of this section.

Lemma 13. Planar Monotone 3-Satisfiability *reduces in polynomial time to* Ordered Level Planarity. *The resulting instances have maximum degree* $\Delta = 2$ *and contain no source or sink with degree* Δ *on a level* V_i *with width* $\lambda_i > 2$.

Proof. Let $\varphi = (\mathcal{U}, \mathcal{C})$ be a planar and monotone 3-Satisfiability formula with clause set $\mathcal{C} = \{c_1, \ldots, c_{|\mathcal{C}|}\}$. Let G_φ be the variable-clause graph of φ. Let $\mathcal{R}$ be a monotone rectilinear representation of G_φ. We construct an ordered level graph $\mathcal{G} = (G, \gamma, \chi)$ such that $\mathcal{G}$ has an ordered level planar drawing if and only if φ is satisfiable.

Overview. The ordered level graph $\mathcal{G}$ has $l_3 + 1$ levels that are partitioned into four *tiers* $\mathcal{T}_0 = \{0, \ldots, l_0\}$, $\mathcal{T}_1 = \{l_0 + 1, \ldots, l_1\}$, $\mathcal{T}_2 = \{l_1 + 1, \ldots, l_2\}$, and $\mathcal{T}_3 = \{l_2 + 1, \ldots, l_3\}$. Each clause $c_i \in \mathcal{C}$ is associated with a *clause* edge $e_i = (c_i^s, c_i^t)$ starting with c_i^s in tier $\mathcal{T}_0$ and ending with c_i^t in tier $\mathcal{T}_2$. The clause edges have to be drawn in a system of tunnels that encodes the 3-Satisfiability formula φ. In $\mathcal{T}_0$ the layout of the tunnels corresponds directly to the rectilinear representation $\mathcal{R}$, see Figure 15c. For each E-shape there are three tunnels corresponding to the three literals of the associated clause. The bottom vertex c_i^s of each clause edge e_i is placed such that e_i has to be drawn inside one of the three tunnels of the E-shape corresponding to c_i. This corresponds to the fact that in a satisfying truth assignment

every clause has at least one satisfied literal. In tier $\mathcal{T}_1$ we merge all the tunnels corresponding to the same literal. We create variable gadgets that ensure that for each variable u, the edges of clauses containing u can be drawn in the tunnel associated with either the negative or the positive literal of u but not in both. This corresponds to the fact that every variable is set to either true or false. Tiers $\mathcal{T}_2$ and $\mathcal{T}_3$ have a technical purpose.

We proceed by describing the different tiers in detail. Recall that in terms of realizability, Ordered Level Planarity is equivalent to the generalized version where γ and χ map to the reals, cf. Section 3.1.3. For the sake of convenience we will begin by designing $\mathcal{G}$ in this generalized setting. It is easy to transform $\mathcal{G}$ such that it satisfies the standard definition in a polynomial-time post processing step.

Tier 0, tier 2, and clause gadgets. Each clause edge $e_i = (c_i^s, c_i^t)$ ends in tier $\mathcal{T}_2$, which is composed of $l_2 - l_1 = |\mathcal{C}|$ levels each of which contains precisely one vertex. We assign $\gamma(c_i^t) = l_1 + i$. Recall that for levels with width 1, the assigned x-coordinates are irrelevant. Hence, we set $\chi(c_i^t) = 0$. Observe that the positions of the vertices c_i^t impose no constraints on the order in which the incident edges enter $\mathcal{T}_2$.

Tier $\mathcal{T}_0$ consists of a system of tunnels that resembles the monotone rectilinear representation $\mathcal{R}$ of $G_\varphi = (\mathcal{U} \uplus \mathcal{C}, E)$, see Figure 15c. Intuitively it is constructed as follows: We take the top part of $\mathcal{R}$, rotate it by 180° and place it to the left of the bottom part such that the variables' line segments align, see Figure 15b. We call the resulting representation $\mathcal{R}'$. For each E-shape in $\mathcal{R}'$ we create a *clause gadget*, which is a subgraph composed of 11 vertices that are placed on a grid close to the E-shape, see Figure 16. The enlarged vertex at the bottom is the lower vertex c_i^s of the clause edge e_i of the clause c_i corresponding to the E-shape. Without loss of generality we assume the grid to be fine enough such that the resulting ordered level graph can be drawn as in Figure 15c without crossings. Further, we assume that the y-coordinates of every pair of horizontal segments belonging to distinct E-shapes differ by at least 3. This ensures that there are no sources or sinks with degree Δ on levels with width larger than 2.

Construction details. In the following two paragraphs, we describe the construction of the clause gadgets in detail.

For every $i = 1, \dots, |\mathcal{C}|$ where c_i is negative, we create its 11-vertex clause gadget as follows, see Figure 16. Let s_1, s_2, s_3 be the three vertical line segments of the E-shape representing c_i in $\mathcal{R}'$ where s_1 is left-most and s_3 right-most. Let v_1, v_2, v_3 be the lower endpoints and v_1', v_2', v_3' be the upper endpoints of s_1, s_2, s_3, respectively. We place the tail c_i^s of the clause edge e_i of c_i at v_2. We create new vertices at v_1, v_3, v_1', v_2', v_3', $v_4 = v_1 + (1, 1)$, $v_5 = v_2 + (1, 2)$, and at v_6, v_7, v_8 which are the lattice points one unit to the right of v_1', v_2', v_3', respectively. To simplify notation, we identify these new vertices with their locations on the grid. We add edges (v_1, v_1'), (v_3, v_8), (v_4, v_6), (v_4, v_2'), (v_5, v_7) and (v_5, v_3') to G.

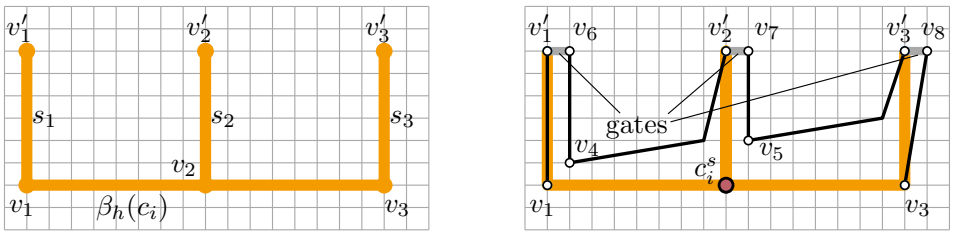

(a) (b)

Figure 16: (a) The E-shape and (b) the clause gadget of clause c_i.

As stated above, we can assume without loss of generality that the grid is fine enough such that the resulting ordered level graph can be drawn as in Figure 15c without crossings. It suffices to assume that the horizontal and vertical distance between any two segment endpoints of $\mathcal{R}'$ is at least 3 (unless the endpoints lie on a common horizontal or vertical line).

Gates and tunnels. The clause gadget (without the clause edge) has a *unique* ordered level planar drawing in the sense that for every level V_i the left-to-right sequence of vertices and edges intersected by the horizontal line L_i through V_i is identical in every ordered level planar drawing. This is due to the fact that the order of the top-most vertices v'_1, v_6, v'_2, v_7, v'_3 and v_8 is fixed and every edge of the gadget is incident to precisely one of these vertices. With the same reasoning, it follows that the subgraph G_0 induced by $\mathcal{T}_0$ (without the clause edges) has a unique ordered level planar drawing, see Figure 15c.

Consider the clause gadget of some clause c_i. We call the line segments v'_1v_6, v'_2v_7 and v'_3v_8 the *gates* of c_i, see Figure 16b. Note that the clause edge e_i has to intersect one of the gates of c_i. This corresponds to the fact the at least one literal of every clause has to be satisfied. In tier $\mathcal{T}_1$ we bundle all gates that belong to the same literal together by creating two long paths for each literal. These two paths form the *tunnel* of the corresponding literal. All clause edges intersecting a gate of some literal have to be drawn inside the literal's tunnel, see Figure 15c. More precisely, for $j = 1, \ldots, |\mathcal{U}|$ we use N_j^0 (n_j^0) to refer to the left-most (right-most) vertex of a negative clause gadget placed on a line segment of $\mathcal{R}'$ representing $u_j \in \mathcal{U}$. The vertices N_j^0 and n_j^0 are the first vertices of the paths forming the *negative* tunnel T_j^n of the negative literal of variable u_j. Analogously, we use P_j^0 (p_j^0) to refer to the left-most (right-most) vertex of a positive clause gadget placed on a line segment of $\mathcal{R}'$ representing u_j. The vertices P_j^0 and p_j^0 are the first vertices of the paths forming the *positive* tunnel T_j^p of the positive literal of variable u_j. If for some j the variable u_j is not contained both in negative and positive clauses, we artificially add two vertices N_j^0 and n_j^0 or P_j^0 and p_j^0 on the corresponding line segments to avoid having to treat special cases in the remainder of the construction.

Tier 1, tier 3, and variable gadgets. Recall that every clause edge has to pass through a gate that is associated with some literal of the clause, and, thus, every edge is drawn in the tunnel of some literal. We need to ensure that for no variable it is possible to use both the tunnel associated with its positive literal and the tunnel associated with its negative literal simultaneously. To this end, we create a *variable gadget* with vertices in tiers $\mathcal{T}_1$ and T_3 for each variable. The variable gadget of variable u_j is illustrated in Figure 17a. The variable gadgets are nested in the sense that they start in $\mathcal{T}_1$ in the order $u_1, u_2, ..., u_{|\mathcal{U}|}$, from bottom to top and they end in the reverse order in $\mathcal{T}_3$, see Figure 18. We force each tunnel with index at least j to be drawn between the vertices u_j^a and u_j^b. This is done by subdividing the tunnel edges on this level, see Figure 17b. The *long edge* (u_j^s, u_j^t) has to be drawn to the left or right of u_j^c in $\mathcal{T}_3$. Accordingly, it is drawn to the left of u_j^a or to the right of u_j^b in $\mathcal{T}_1$. Thus, it is drawn either to the right (Figure 17b) of all the tunnels or to the left (Figure 17c) of all the tunnels. As a consequence, the *blocking edge* (u_j^s, u_j^p) is also drawn either to the right or the left of all the tunnels. Together with the edge (u_j^q, u_j^p) it prevents clause edges from being drawn either in the positive tunnel T_j^p or negative tunnel T_j^n of variable u_j, which end at level $\gamma(u_j^q)$, because they cannot reach their endpoints in $\mathcal{T}_2$ without crossings. We say T_j^p or T_j^n are *blocked*, respectively.

Construction details. In the following two paragraphs, we describe the construction of the variable gadgets in detail, for illustrations refer to Figures 17 and 18.

Tier $\mathcal{T}_3$ has $l_3 - l_2 = 2 \cdot |\mathcal{U}|$ layers each of which contains precisely one vertex. We refer to the vertex in layer $(l_3 - 2j + 1)$ as u_j^t and to the vertex in layer $(l_3 - 2j)$ as u_j^c for $j = 1, \ldots, |\mathcal{U}|$. Tier $\mathcal{T}_1$ has $l_1 - l_0 = 4 \cdot |\mathcal{U}|$ levels. In each of the levels $(l_0 + 4j - 3)$, $(l_0 + 4j - 1)$, and $(l_0 + 4j)$ where $j = 1, \ldots, |\mathcal{U}|$ we create one vertex. These vertices are called u_j^s, u_j^q, and u_j^p, respectively. In level $(l_0 + 4j - 2)$ we create two vertices u_j^a and u_j^b in this order. We add the edges (u_j^s, u_j^t), (u_j^s, u_j^p), (u_j^a, u_j^c), (u_j^b, u_j^c), and (u_j^q, u_j^p).

Finally, for $j = 1, \ldots, |\mathcal{U}|$ we do the following, see Figure 17b or Figure 17c. In level (l_0+4j-2) we create vertices $P_j^j, p_j^j, \ldots, P_{|\mathcal{U}|}^j, p_{|\mathcal{U}|}^j, N_{|\mathcal{U}|}^j, n_{|\mathcal{U}|}^j, \ldots, N_j^j, n_j^j$ and add them in this order between u_j^a and u_j^b. In level $(l_0 + 4j - 1)$ we create vertices P_j^{j+1} and p_j^{j+1} in this order before u_j^q and we create vertices N_j^{j+1} and n_j^{j+1} in this order after u_j^q. We create edges realizing the paths $t_j^P = (P_j^0, \ldots, P_j^{j+1})$, $t_j^p = (p_j^0, \ldots, p_j^{j+1})$, $t_j^N = (N_j^0, \ldots, N_j^{j+1})$ and $t_j^n = (n_j^0, \ldots, n_j^{j+1})$. The pair of paths $T_j^p = (t_j^P, t_j^p)$ is the positive tunnel of variable u_j. The pair of paths $T_j^n = (t_j^N, t_j^n)$ is the negative tunnel of variable u_j. If an edge e is drawn in the region between the two paths of a tunnel T, we say it is drawn *in* T.

Runtime and Properties. The construction of the ordered level graph $\mathcal{G}$ can be carried out in polynomial time. Note that its maximum degree is $\Delta = 2$. Moreover, no source or sink with degree Δ is located on a level V_i with width $\lambda_i > 2$: each

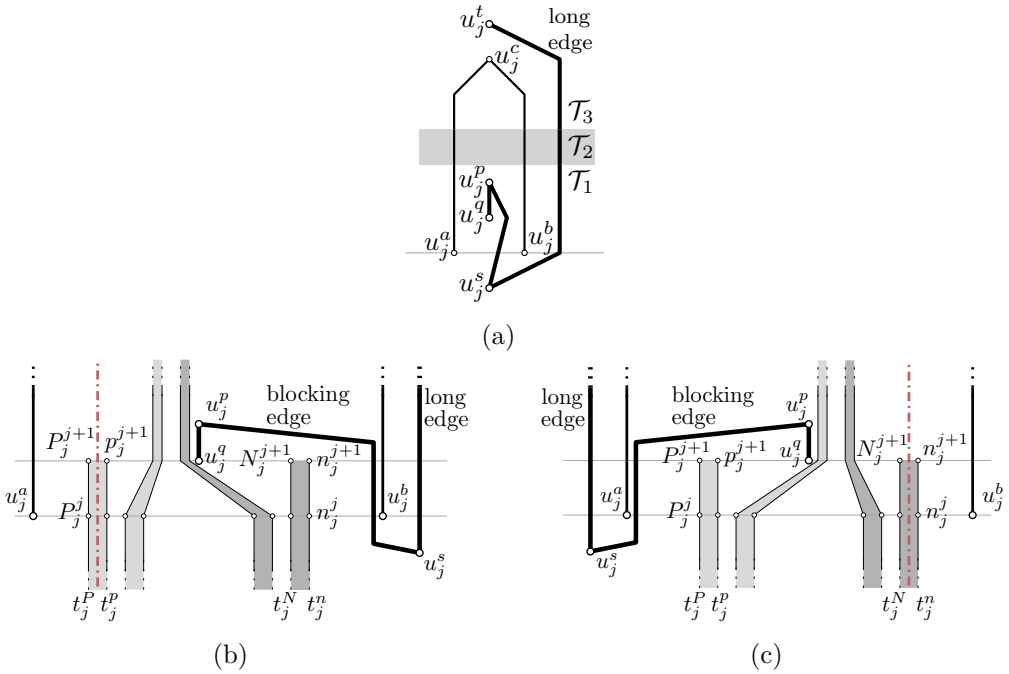

Figure 17: (a) The variable gadget of u_j in (b) positive and (c) negative state. For the sake of visual clarity, these figures make use of the relaxed but equivalent version of Ordered Level Planarity that only requires that the vertices of each level appear according to the total ordering corresponding to χ, cf. Section 3.1.3. In particular, a vertex v of a level V_i with width $\lambda_i = 1$ may appear *anywhere* on the horizontal line L_i. The dash-dotted edges are clause edges.

variable gadget has a total of three sinks or sources with degree Δ (in tiers $\mathcal{T}_1$ and $\mathcal{T}_3$). Each of these vertices is located on a level with width 1. Additionally, each clause gadget has two sources with degree Δ (in tier $\mathcal{T}_0$). These sources are also located on levels with width 1 due to the assumption that, in the representation $\mathcal{R}'$, the y-coordinates of each pair of horizontal segments belonging to distinct E-shapes differ by at least 3.

Correctness. It remains to show that $\mathcal{G}$ has an ordered level planar drawing if and only if φ is satisfiable. Assume that $\mathcal{G}$ has an ordered level planar drawing Γ. We create a satisfying truth assignment for φ. If T_j^n is blocked we set u_j to true, otherwise we set u_j to false for $j \in 1, \ldots, |\mathcal{U}|$. Recall that the subgraph G_0 induced by the vertices in tier $\mathcal{T}_0$ has a unique ordered level planar drawing. Consider a clause c_i and let f, g, j be the indices of the variables whose literals are contained in c_i. Clause edge $e_i = (e_i^s, e_i^t)$ has to pass level l_0 through one of the gates of c_i. More precisely, e_i has to be drawn between either N_f^0 and n_f^0, N_g^0 and n_g^0, or N_j^0 and n_j^0 (if c_i is negative), or between either P_f^0 and p_f^0, P_g^0 and p_g^0, or P_j^0 and p_j^0 (if c_i is

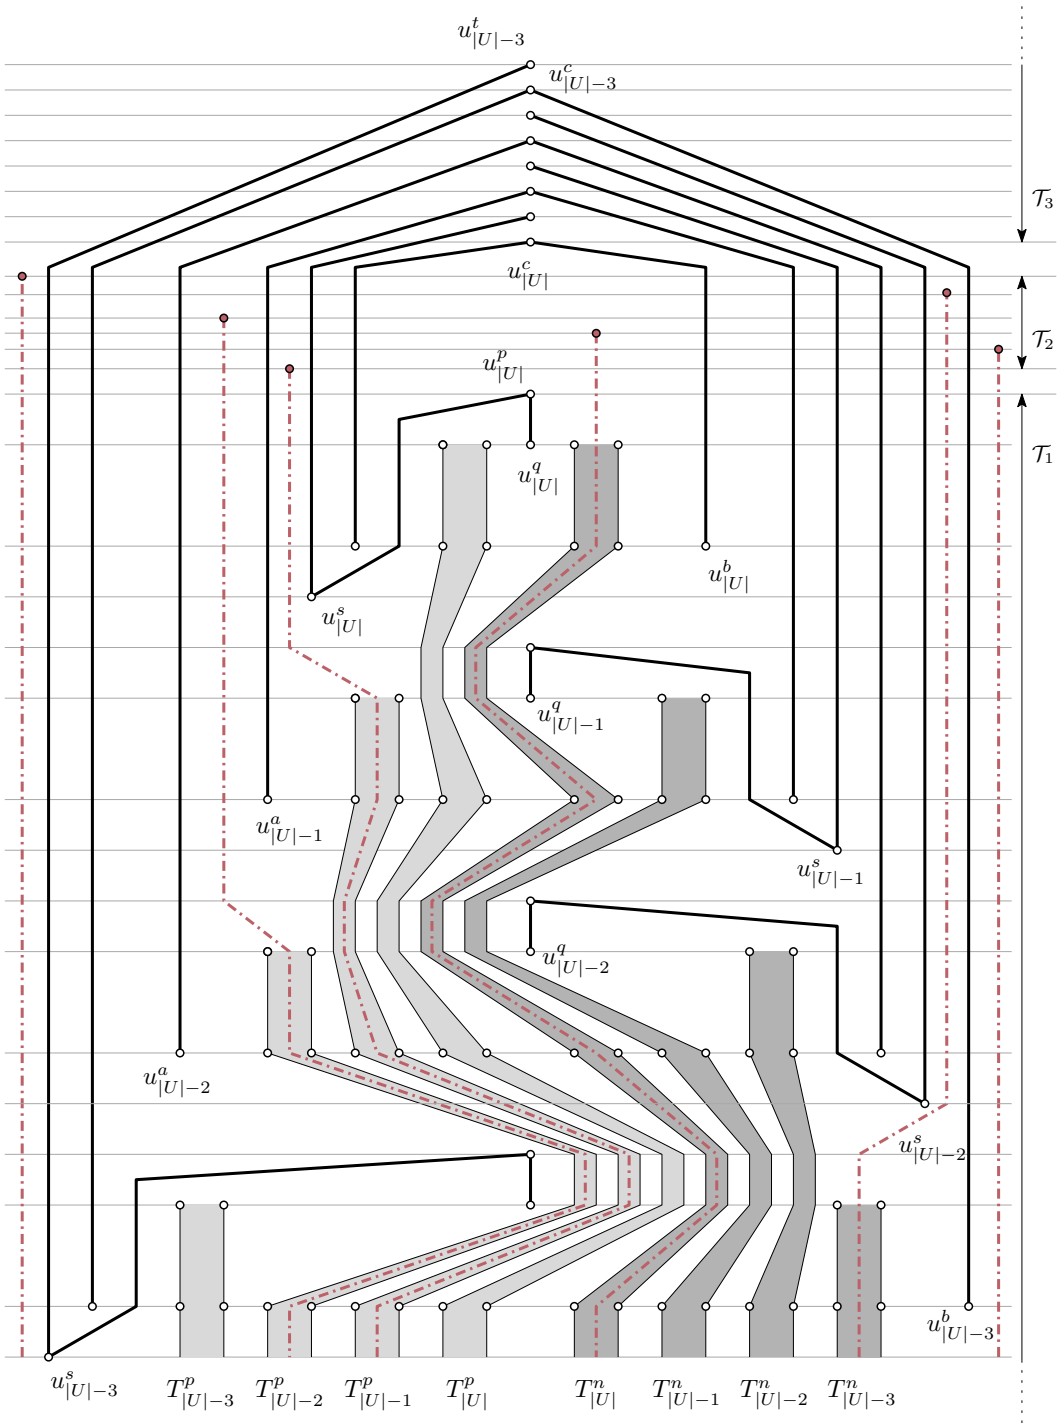

Figure 18: The nesting structure of the variable gadgets. Only the gadgets of the variables with the four largest indices are shown. They are nested within the remaining variable gadgets. Tier $\mathcal{T}_0$ is located below *all* these gadgets. As in Figure 17, this figure makes use of the version of Ordered Level Planarity that uses relative x-coordinates on each level. The dash-dotted edges are clause edges.

positive), see Figure 15c. First, assume that c_i is negative and assume without loss of generality that e_i traverses l_0 between N_j^0 and n_j^0. In this case e_i has to be drawn in T_j^n. Recall that this is only possible if T_j^n is not blocked, which is the case if u_j is false, see Figure 17c. Analogously, if c_i is positive and e_i traverses w.l.o.g. between p_j^P and p_j^p, then u_j is true, Figure 17b. Thus, we have established that one literal of each clause in $\mathcal{C}$ evaluates to true for our truth assignment and, hence, formula φ is satisfiable.

Now assume that φ is satisfiable and consider a satisfying truth assignment. We create an ordered level planar drawing Γ of $\mathcal{G}$. It is clear how to create the unique subdrawing of G_0. The variable gadgets are drawn in a nested fashion, see Figure 18. For $j = 1, \dots, |\mathcal{U}| - 1$ we draw edge (u_j^a, u_j^c) to the left of u_{j+1}^a and u_{j+1}^s, and edge (u_j^b, u_j^c) to the right of u_{j+1}^b and u_{j+1}^s. In other words, the pair $((u_j^a, u_j^c), (u_j^b, u_j^c))$ is drawn between all such pairs with index smaller than j. Recall that the vertices u_j^a, u_j^b, u_j^s, u_j^p and u_j^q are located on higher levels than the according vertices of variables with index smaller than j and that u_j^t and u_j^c are located on lower levels than the according vertices of variables with index smaller than j.

For $j = 1, \dots, |\mathcal{U}|$ if u_j is positive, we draw the long edge (u_j^s, u_j^t) to the right of u_j^b and u_j^c and, accordingly, we have to draw all tunnels left of u_j^s and u_j^q (except for T_j^n, which has to be drawn to the left of u_j^s and must end to the right of u_j^q), see Figure 17b. If u_j is negative we draw the long edge (u_j^s, u_j^t) to the left of u_j^b and u_j^c and, accordingly, we have to draw all tunnels right of u_j^s and u_j^q (except for T_j^p, which has to be drawn to the right of u_j^s and end to the left of u_j^q), see Figure 17c. We have to draw the blocking edge (u_j^s, u_j^p) to the right of n_j^{j+1} if u_j is positive and to the left of P_j^{j+1} if u_j is negative.

It remains to describe how to draw the clause edges. Let c_i be a clause. There is at least one true literal in c_i. Let k be the index of the corresponding variable. We describe the drawing of clause edge $e_i = (c_i^s, c_i^t)$ from bottom to top. We start by drawing e_i in the tunnel T_k^p (T_k^n) if c_i is positive (negative). We exit the tunnel when passing through level $\gamma(u_k^q)$, after which we end up to the left (right) of all tunnels with index larger than k, see Figure 17b (Figure 17c). Note that since T_k^p (T_k^n) is not blocked, we can continue without having to cross blocking edge (u_k^s, u_k^p) or (u_k^q, u_k^p). We draw e_i to the left (right) of all vertices belonging to variable gadgets with index larger than k, see Figure 18. This introduces no crossings since above level $\gamma(u_k^q)$ all tunnels with index larger than k are drawn to the right of $u_{k+1}^a, \dots, u_{|\mathcal{U}|}^a$ and the left of $u_{k+1}^b, \dots, u_{|\mathcal{U}|}^b$. Connecting to c_i^t in tier $\mathcal{T}_2$ is straight-forward since each level of this tier contains only one vertex. □

We obtain $\mathcal{NP}$-hardness for instances with maximum degree $\Delta = 2$.

In fact, we can restrict our attention to instances with level-width $\lambda = 2$. To this end, we split levels with width $\lambda_i > 2$ into $\lambda_i - 1$ levels containing exactly two vertices each.

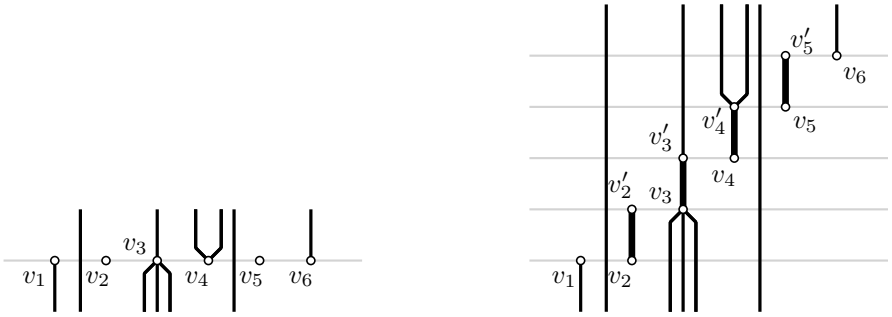

Figure 19: (a) A level V_i with width $\lambda_i > 2$. (b) To reduce the level-width, we replace V_i with $\lambda_i - 1$ levels. Thick edges are the stretch edges.

Lemma 14. *An instance $\mathcal{G} = (G = (V, E), \gamma, \chi)$ of* Ordered Level Planarity *with maximum degree Δ and level-width $\lambda > 2$ can be transformed in linear time into an equivalent instance $\mathcal{G}' = (G' = (V', E'), \gamma', \chi')$ of* Ordered Level Planarity *with maximum degree $\Delta' \leq \Delta + 1$ and level-width $\lambda' = 2$. Further, if $\mathcal{G}$ contains no source or sink with degree Δ on a level V_i with width $\lambda_i > 2$, then $\Delta' \leq \Delta$.*

Proof. We replace each level V_i with width $|V_i| = \lambda_i > 2$ by $\lambda_i - 1$ levels with 2 vertices each, as illustrated in Figure 19. Accordingly, vertices on levels above V_i are shifted upwards by $\lambda_i - 2$ levels. Formally, let $V_i = \{v_1, \ldots, v_{\lambda_i}\}$ with $\chi(v_1) < \cdots < \chi(v_{\lambda_i})$. We increase the level of vertex v_j by $j - 2$ for $j = 3, \ldots, \lambda_i$. For $j = 2, \ldots, \lambda_i - 1$ we create a vertex v_j' one level above v_j with $\chi(v_j') = 0$ and we create a new *stretch* edge (v_j, v_j'). For $j = 2, \ldots, \lambda_i$ we set $\chi(v_j) = 1$.

For all the vertices v that have been split in this way into v and v', the bottom vertex v inherits all the incoming edges and the top vertex v' inherits all the outgoing edges. Let $\mathcal{G}'$ denote the resulting instance, which can be constructed in linear time. It is easy to verify that the vertex degrees behave as desired.

An ordered level planar drawing of $\mathcal{G}$ can easily be converted to a drawing of $\mathcal{G}'$. For the conversion in the other direction, we successively contract each stretch edge (v_i, v_i') back into a single vertex, thereby merging two consecutive levels of $\mathcal{G}'$. Apart from the edge (v_i, v_i'), the vertex v_i has incident edges from below and the vertex v_i' has incident edges from above only. Therefore, such a contraction cannot cause any problems. The stretch edges ensure that the vertices of each level of $\mathcal{G}$ end up in the correct order. □

Corollary 15. Ordered Level Planarity *is $\mathcal{NP}$-hard, even for acyclic ordered level graphs with maximum degree $\Delta = 2$ and level-width $\lambda = 2$.* □

The reduction in Lemma 13 requires degree-2 vertices. With $\Delta = 1$, the problem becomes polynomial-time solvable. In fact, one can easily solve it as long as the maximum in-degree and the maximum out-degree are both bounded by 1.

Lemma 16. Ordered Level Planarity *restricted to instances with maximum in-degree* $\Delta^- = 1$ *and maximum out-degree* $\Delta^+ = 1$ *can be solved in linear time.*

Proof. Let $\mathcal{G} = (G = (V, E), \gamma, \chi)$ be an ordered level graph with maximum indegree $\Delta^- = 1$ and maximum outdegree $\Delta^+ = 1$. Such a graph $\mathcal{G}$ consists of a set P of y-monotone paths. Each path $p \in P$ has vertices on some sequence of levels, possibly skipping intermediate levels.

We define the following relation on P: We write $p \prec q$, meaning that p must be drawn to the left of q, if p and q have vertices v_p and v_q that lie adjacent on a common level, i.e. $\gamma(v_p) = \gamma(v_q)$ and $\chi(v_q) = \chi(v_p) + 1$. This relation has at most $|V|$ pairs, and by topological sorting, we can find in $O(|V|)$ time a linear ordering that is consistent with the relation $\prec$, unless this relation has a cycle. The former case implies the existence of an ordered level drawing while the latter case implies that the problem has no solution.

This follows from considerations about horizontal separability of y-monotone sets by translations, cf. [14, 43]. An easy proof can be given following Guibas and Yao [70, 71]: consider an ordered level planar drawing of $\mathcal{G}$. We say that a vertex is *visible from the left* if the infinite horizontal ray emanating from that vertex to the left does not intersect the drawing. Among the paths whose lower endpoint is visible from the left, the one with the topmost lower endpoint must precede all other paths to which it is related in the $\prec$-relation. Removing this path and iterating the procedure leads to a linear order that extends $\prec$. Regarding the other direction, if we have such a linear order $x\colon P \to \{1, \ldots, |P|\}$, we can simply draw each path p straight at x-coordinate $x(p)$, subdivide all edges properly and, finally, shift the vertices on each level such that the vertices of V are placed according to χ while maintaining the order x. □

For $\lambda = 1$, Ordered Level Planarity is solvable in linear time since Level Planarity can be solved in linear time [83]. Proper instances have a unique drawing (if it exists). The existence can be checked with a simple linear-time sweep through every level. The problem is obviously contained in $\mathcal{NP}$. The results of this section establish Theorem 1.

Theorem 1. Ordered Level Planarity *is* $\mathcal{NP}$*-complete, even for acyclic ordered level graphs with maximum degree* $\Delta = 2$ *and level-width* $\lambda = 2$*. The problem can be solved in linear time if the given level graph is proper, or if the level-width is* $\lambda = 1$*, or if* $\Delta^+ = \Delta^- = 1$*, where* Δ^+ *and* Δ^- *are the maximum in-degree and out-degree respectively.* □

3.5 Connected instances

To reduce from Ordered Level Planarity to Geodesic Planarity, our main reduction (to Ordered Level Planarity) is tailored to achieve a small maximum

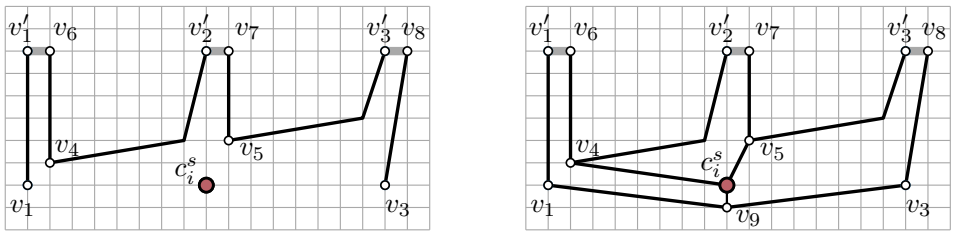

(a) (b)

Figure 20: (a) The original clause gadget and (b) the augmented version for the connected case. The clause edge starting at c_i^s is not shown.

degree of $\Delta = 2$. As a consequence, the resulting graphs are not connected. At the cost of an increased maximum degree, it is possible to make our instances connected by inserting additional edges. In this section, we discuss the necessary adaptations to obtain the following theorem.

Theorem 8. *The following problems are $\mathcal{NP}$-hard even for connected instances with maximum degree $\Delta = 4$:*

- Ordered Level Planarity*, even for level-width $\lambda = 2$;*
- Constrained Level Planarity*, even for level-width $\lambda = 2$ and prescribed* total *orderings;*
- Clustered Level Planarity*, even for level-width $\lambda = 2$ and flat cluster hierarchies that partition the vertices into two non-trivial clusters; and*
- Bi-Monotonicity*.*

We begin by showing the $\mathcal{NP}$-hardness of Ordered Level Planarity for connected instances.

Lemma 17. Planar Monotone 3-Satisfiability *reduces in polynomial time to* Ordered Level Planarity*. The resulting instances are connected and have maximum degree $\Delta = 4$. The maximum in-degree Δ^- and maximum out-degree Δ^+ are both 3.*

Proof. We proceed exactly as in Lemma 13. We augment the resulting instances such that they become connected. During this augmentation step, we need to make sure that the degree constraints remain satisfied.

Recall that $\mathcal{U}$ is the set of variables and that tier $\mathcal{T}_3$ contains precisely $2|\mathcal{U}|$ vertices each of which is the only vertex of its level, see Figure 17a and Figure 18. We connect all these vertices with a directed path, that is, we insert the edges (u_j^c, u_j^t) for $j = 1, \ldots, |\mathcal{U}|$ and the edges (u_j^t, u_{j+1}^c) for $j = 2, \ldots, |\mathcal{U}|$. One can easily check

that the degree constraints are satisfied: the degree of the vertices u_j^t is now 3 (except for u_1^t, which has degree 2). The degree of the vertices u_j^c is now $4 = \Delta$ (except for $u_{|\mathcal{U}|}^c$, which has degree 3). The largest out-degree of all these vertices is $1 < \Delta^+$, while the largest in-degree is $3 = \Delta^-$.

Recall that for each clause c_i we have created a clause gadget as depicted in Figure 20a. We replace this graph with the graph shown in Figure 20b. Precisely, we do the following: We add a new vertex v_9 one unit below c_i^s and we add the edges (v_9, c_i^s), (c_i^s, v_4), (c_i^s, v_5). Again, the degree bounds are easily verified: Vertex c_i^s now has degree $4 = \Delta$ (including the clause edge); vertices v_4, v_5 and v_9 have degree 3 and vertices v_1 and v_3 have degree 2. The overall maximum out-degree is $3 = \Delta^+$, while the maximum in-degree is 1.

Recall that the segments $v_1'v_6$, $v_2'v_7$, and $v_3'v_8$ of each clause gadget are called the gates of c_i. All gates (of all clauses) are located on the same level V_g, see Figure 15c. We now ensure that all vertices of V_g become connected to each other. The two vertices that bound each gate are already connected through the augmented clause gadgets. We connect two consecutive vertices u, v from different gates by adding for each such pair u, v a new vertex w one level below V_g with two edges (w, u) and (w, v).

The resulting instance has two connected components: one component contains all the clause gadgets, clause edges, and tunnels, and the other contains all the variable gadgets. We can connect these components by adding a path P between the top-most vertex v_t and bottom-most vertex v_b of the instance. Note that $v_t = u_{|\mathcal{U}|}^t$. The bottom-most vertex is vertex v_9 of the clause gadget corresponding to the (unique) E-shape with the lowest horizontal line segment. Simply choosing $P = (v_b, v_t)$ would result in an increased maximum out-degree of 4. Instead, we choose the (undirected) path $P = (v_b, v_b', v_t', v_t)$, where v_b' and v_t' are new vertices placed below v_b and above v_t respectively. This way, the out-degree of v_b remains 3.

The new connected instance is equivalent to the original one as the clause edge (c_i^s, c_i^t) can still reach each of the three gates of c_i by choosing the corresponding embedding. Aside from the edges incident to the vertices c_i^s, no new edge impairs the realizability of the instance in any way. □

We remark that it is possible to decrease the maximum in-degree guaranteed in Lemma 17 to $\Delta^- = 2$ by splitting the vertices u_j^c before the augmentation step.

Lemma 17 produces instances in which the maximum out-degree and the maximum in-degree are strictly smaller than the maximum degree $\Delta = 4$. It follows that no source or sink has degree Δ. Thus, Lemma 14 implies the statement about Ordered Level Planarity and Constrained Level Planarity in Theorem 8. The statement about Bi-Monotonicity follows from Theorem 4. Finally, the statement about Clustered Level Planarity follows from the fact that the reduction given in Lemma 12 does not change the graph except for the subdivisions of the edges and the addition of isolated vertices, which can be removed without affecting the realizability of the instance. This concludes the proof of Theorem 8.

3.6 Conclusion

We introduced and studied the problem Ordered Level Planarity. Our main result is an $\mathcal{NP}$-hardness statement that cannot be strengthened. We demonstrated the relevance of our result by stating reductions to several other graph drawing problems. These reductions answer multiple questions posed by the graph drawing community and establish connections between problems that (to the best of our knowledge) have not been considered in the same context before. Recently, Da Lozzo, Di Battista, and Frati [39] used Theorem 1 to show the $\mathcal{NP}$-hardness of another generalization of Ordered Level Planarity. We expect that Theorem 1 will serve as a useful tool for further reductions.

In Section 3.5, we extended most of our reductions to produce problem instances that are connected. We did not provide such a modification for our reduction to Geodesic Planarity. Due to the increased vertex degrees in Ordered Level Planarity instances generated by Theorem 8, our reduction to Geodesic Planarity in Step (i) of Lemma 9 breaks down, as there is not enough space anymore to attach all the edges around each vertex. It does not seem straight-forward to modify our construction to obtain a reduction to Geodesic Planarity that produces connected instances. Thus, we leave it as an open question whether $\mathcal{NP}$-hardness still holds for connected instances of (Manhattan) Geodesic Planarity.

Chapter 4

Two-page book embeddings of triconnected planar graphs

4.1 Introduction

A *book embedding* [22, 75, 131] can be thought of as a drawing of a graph in a book. Formally, the corresponding drawing convention requires that all vertices are embedded on a line in $\mathbb{R}^3$ called *spine*, and every edge is embedded in a half-plane, called *page*, bounded by the spine. No two edges (on the same page) are allowed to cross. If k pages are used, then the corresponding embedding is a k-page book embedding. Applications for book embeddings include VLSI design [35], bioinformatics [74], and other Graph Drawing problems [11, 128].

Obviously, every graph has a book embedding: we can simply realize each edge on a separate page. The *book thickness* of a graph is the smallest number of pages that suffices to realize the graph in form of a book embedding. Graphs with book thickness at most two are planar since a k-page book embedding with $k \leq 2$ corresponds to a planar drawing of the graph: each page by itself is crossing-free and two pages can be arranged on a common plane such that they only intersect along the spine.

Bernhart and Kainen [22] characterized the graphs that can be embedded on k pages, for $k \leq 2$. For $k = 2$ these are the subhamiltonian planar graphs, that is, subgraphs of Hamiltonian planar graphs, cf. Figure 21. This turns the problem of embedding a graph on two pages into a graph augmentation problem: a planar graph admits a 2-page book embedding if and only if one can add edges[1] to make it Hamiltonian while maintaining planarity. It follows that a *maximal* planar graph is subhamiltonian if and only if it is Hamiltonian, and, hence, the recognition of graphs with book thickness at most two is $\mathcal{NP}$-hard [127]. However, no planar graph is too far away from being subhamiltonian: subdividing at most $n/2$ of the up to $3n - 6$

[1] Kainen and Overbay [84] observe that it suffices to add edges (rather than vertices and edges) since by removing the added vertices of a Hamiltonian cycle in a plane augmentation, one obtains a Hamiltonian cycle in a plane augmentation that only uses new edges.

edges of a planar graph on n vertices yields a subhamiltonian planar graph [29].

The structure of Hamiltonian paths and cycles in graphs has been a fruitful subject of intense research over many decades, both from a combinatorial and from an algorithmic point of view. For general graphs, sufficient conditions for the existence of a Hamiltonian cycle typically involve rather strong assumptions on the degree, such as in Dirac's Theorem [49] (minimum degree $\geq n/2$), Ore's Theorem [103] (degree sum of every nonadjacent vertex pair $\geq n$), or Asratian and Khachatrian's Theorem [15] ($\deg(u)+\deg(w) = |\mathrm{N}(u)|+|\mathrm{N}(v)|+|\mathrm{N}(w)|$, for every induced path uvw). Planar graphs provide a lot more structure so that by a famous theorem of Tutte, 4-connectivity suffices to guarantee the existence of a Hamiltonian cycle [122], which can be computed in linear time [30]. In contrast, deciding Hamiltonicity is $\mathcal{NP}$-complete for 3-connected cubic planar graphs [61], and for maximal planar graphs [127]. Finally, maximal planar graphs of degree at most six are Hamiltonian [52]. Hence, both vertex degree and connectivity appear to be crucial parameters concerning Hamiltonicity.

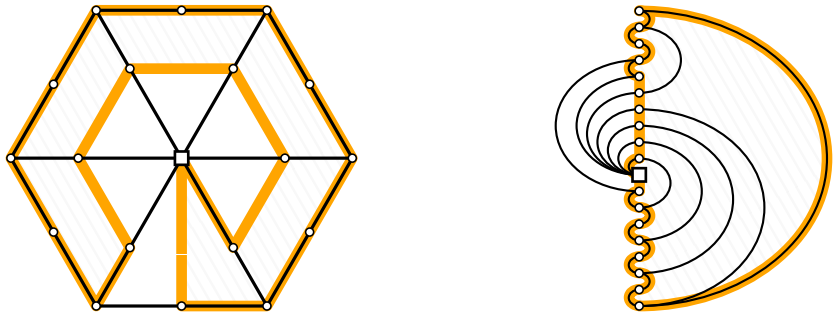

(a) (b)

Figure 21: A nonhamiltonian graph with a subhamiltonian cycle (a Hamiltonian cycle in a plane augmentation) and a corresponding two-page book embedding. In particular, the order of the vertices along the cycle corresponds to the ordering of the vertices along the spine. The interior of the cycle corresponds to the right page and its exterior corresponds to the left page.

Despite a plethora of results concerned with Hamiltonian cycles in planar graphs and book embeddings, several fundamental questions are still open. Let us give just two prominent examples to illustrate this point. For once, there is Barnette's Conjecture: "Every 3-connected cubic bipartite planar graph is Hamiltonian." And then there is the question if every planar graph can be embedded on three pages. Yannakakis showed, improving a series of earlier results, that four pages are sufficient for every planar graph [131]. However, a corresponding lower bound is still elusive, in spite of initial claims [130].

In this chapter, we are interested in sufficient conditions for two-page book embeddability. Previously established sufficient conditions for a planar graph to be subhamiltonian involve assumptions about vertex degrees [20, 21, 75], as is the case

with sufficient conditions for classic Hamiltonicity. However, in the subhamiltonian case, these assumptions require small vertex degrees rather than large vertex degrees. Intuitively, the reason is that when looking for a Hamiltonian cycle in some augmentation of a given planar graph G, we may create any desired edges that are not part of G. In contrast, unwanted edges that belong to G may act as obstructions.

Heath [75] showed that planar graphs of maximum degree three are always subhamiltonian. Later, Bauernöppel [20] and, independently, Bekos, Gronemann, and Raftopoulou [21] showed that maximum degree four is also a sufficient condition for a planar graph to be subhamiltonian. On the negative side, Bauernöppel [20] showed that recognizing subhamiltonian planar graphs is $\mathcal{NP}$-hard even if the maximum degree is seven (the graphs in his reduction are biconnected, but not 3-connected). Guan and Yang [69] showed that planar graphs of maximum degree five can be embedded on *three* pages.

When sticking to two pages, the ultimate goal is to determine the largest k such that every planar graph with maximum degree at most k is subhamiltonian. We take a considerable next step towards this goal by showing that all triconnected planar graphs of maximum degree five are subhamiltonian planar. In fact, we prove the following more general statement where the degree restriction applies to vertices of separating 3-cycles only.

Theorem 18. *Let G be a 3-connected simple planar graph on n vertices where every vertex that belongs to a separating 3-cycle has degree at most five. Then G is subhamiltonian planar. Moreover, a subhamiltonian plane cycle for G can be computed in $O(n^2)$ time.*

Corollary 19. *Every 3-connected simple planar graph with maximum vertex degree five can be embedded on two pages, and such an embedding can be computed in quadratic time.*

We also show that the degree bound in Theorem 18 is tight.

Theorem 20. *There exists an infinite family of 3-connected simple planar graphs that are not subhamiltonian planar and where every vertex of a separating 3-cycle has degree at most six.*

We remark that the condition of Theorem 18 is easy to test in $O(n^2)$ time: start by testing planarity [25, 81] and 3-connectivity [72, 80] in linear time. If the graph is indeed 3-connected and planar, then it has a unique combinatorial embedding (up to its orientation) and its separating 3-cycles are separating triangles in this embedding (see Section 4.2 for the formal definition). Then, for each edge that is incident to a vertex of degree at least six we check whether its endpoints have a common neighbor, which is easily done in linear time per edge. It remains to check whether the thereby found triangles are separating. This can be accomplished in constant time per triangle by checking whether there is at least one vertex on each of its sides in the combinatorial embedding.

Organization. We begin by introducing some terminology and notation in Section 4.2. In Section 4.3, we study three special cases for which Theorem 18 is easily proven. In particular, we show how to deal with the case that the given graph contains no separating 3-cycle.

In Section 4.4, we proceed with a high-level overview of the proof of Theorem 18, which is then developed in detail in Sections 4.5–4.8. In short, it consists of three main steps: in Section 4.5 we describe a recursive algorithm that gets rid of all separating 3-cycles using edge contractions. In Section 4.6, we then augment the graph (by adding new vertices and edges) without creating any new separating 3-cycles; and apply the results from Section 4.3 to obtain a subhamiltonian cycle in this contracted, augmented graph. Finally, in Section 4.7, we reconstruct the collapsed 3-cycles, one-by-one. Due to the augmentation from Section 4.6, it is possible to maintain a subhamiltonian cycle throughout the procedure. This is done by applying small local modifications to the current cycle at every reconstruction step. We conclude the proof in Section 4.8, in which we deal with the remaining special cases. In this section we also formally summarize the proof. In Section 4.9, we provide a high-level summary of the algorithm.

The infinite family of nonsubhamiltonian 3-connected planar graphs described by Theorem 20 is constructed in Section 4.10. We conclude by discussing future directions and open questions in Section 4.11.

4.2 Notation

In this chapter, we use the shorthand uv to denote an edge between u and v. We also use shorthands such as $stuvw$ to denote a path (or cycle) (s, t, u, v, w) (it will be clear from the context whether a path or cycle is meant). The reason for this inconsistency in notation is that in this chapter, we often have to argue about small constant sized configurations, where the shorthands help to remove clutter from the notation. In contrast, in Chapters 3 and 5, we encounter large graphs with heavily indexed labeling, where the classic notation improves readability.

A *Hamiltonian* cycle for a graph is a simple cycle through all vertices and a graph is *Hamiltonian* if it contains a Hamiltonian cycle. An *augmentation* of a graph $G = (V, E)$ is a supergraph $A = (V, E')$ with $E' \supseteq E$. If G is a plane graph, then a *plane augmentation* H of G is an augmentation that is plane and agrees with G, that is, the plane subgraph of H the corresponds to the underlying abstract graph of G is G. A *(plane) subhamiltonian cycle*[2] for a (plane) graph G is a Hamiltonian cycle in some (plane) augmentation of G.

A set of vertices is *separating* if its removal disconnects the graph. We distinguish between separating 3-cycles as a notion for both abstract and plane graphs, and

[2] This term has been established in the related literature. Arguably, the term *super*hamiltonian would be more fitting.

separating triangles in plane graphs. A *separating* 3*-cycle* is a 3-cycle whose removal disconnects the graph. A *separating triangle* is a 3-cycle C of a plane graph G such that both the interior and the exterior region bounded by C contain some vertex of G. Note that every separating triangle is a separating 3-cycle, whereas the converse is not true in general. Distinguishing between the two concepts is important in the context of graph augmentation problems. For instance, if a plane graph G has a separating triangle abc, then a, b, c form a separating triple in *every* plane augmentation. Hence, it is not possible to make G 4-connected by adding edges. On the other hand, if uvw is a separating 3-cycle of G that is not a separating triangle, then it is easy to find a plane augmentation in which u, v, w are not separating.

For a cycle C in a plane graph G denote by G_C^+ the plane subgraph of G that contains all vertices and edges on C and exterior to C. Similarly denote by G_C^- the plane subgraph of G that contains all vertices and edges on C and interior to C. The *inside* of C refers to the interior of the bounded region enclosed by C. So a vertex of G *inside of* C is a vertex of $G_C^- \setminus C$. Analogously a vertex *outside of* C is a vertex of $G_C^+ \setminus C$. A separating triangle T is called *trivial* if $G_T^- \simeq K_4$, that is, if G_T^- is isomorphic to K_4. Otherwise, T is called *nontrivial*. A separating triangle T is *maximal* if no vertex of T is inside a separating triangle of G_T^+.

Triconnected components. Our algorithm described in Section 4.5 uses a decomposition of a (biconnected) graph into its triconnected components. There are several distinct, but equivalent definitions of this concept. Here we adopt the notation by Gutwenger and Mutzel [72].

Let $G = (V, E)$ be a biconnected multigraph and $a, b \in V$. The edge set E can be partitioned into equivalence classes $E_1, \ldots, E_k$ such that two edges belong to the same set E_i if and only if they belong to a common path that does not contain a or b as an interior vertex. These sets E_i are called *separation classes* of G with respect to the pair a, b. Moreover, if $k \geq 2$, then a, b is a *separation pair* unless $k = 2$ and one class consists of a single edge; or $k = 3$ and each class consists of a single edge. Let $E' = \bigcup_{i \in I} E_i$ and $E'' = E \setminus E'$ with $I \subset \{1, \ldots, k\}$ such that $|E'| \geq 2$ and $|E''| \geq 2$. The two graphs $G' = (\mathrm{V}(E'), E' \cup \{e\})$ and $G'' = (\mathrm{V}(E''), E'' \cup \{e\})$, where $e = ab$ is a new edge, are called *split graphs* of G with respect to a, b. Replacing G by two split graphs is called *splitting* G. The edge e, which occurs in both G' and G'', is called *virtual* edge and it identifies the split operation. Split graphs are again biconnected [72].

Iteratively splitting G and the resulting splits graphs until no further split is possible yields a set of multigraphs that are called *split components* of G. Each split component is of one of three types: (1) a set of three parallel edges between two common vertices, called *triple bond*; (2) a cycle of length 3, called *triangle*; (3) a simple triconnected graph. Consider two split components $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ of G such that E_1 and E_2 both contain the same virtual edge e'. Replacing G_1 and G_2 with the graph $(V_1 \cup V_2, (E_1 \cup E_2) \setminus \{e'\})$ is called *merging* G_1 and G_2.

The *triconnected components* of G are obtained from its split components by merging all triple bonds of G into maximal sets of parallel edges; and by merging all triangles into maximal simple cycles. Each triconnected component is of one of three types: (1) a set of at least three parallel edges between two common vertices, called *parallel* triconnected component; (2) a simple cycle of length at least three, called *series* triconnected component; (3) a simple triconnected graph, called *rigid* triconnected component. While the split components of a (multi–) graph are not necessarily unique; its triconnected components are [80, 100]. Each edge in E occurs in exactly one triconnected component and is called *real*; and each virtual edge e'' occurs in exactly two triconnected components N_1 and N_2 [72]. In N_1, the edge e'' acts as a placeholder for N_2; we say it *corresponds* to N_2.

The triconnected components of a biconnected graph can be efficiently computed [72] and maintained [78] via the SPQR-tree data structure.

4.3 Three simple cases

It suffices to prove Theorem 18 for an arbitrary combinatorial embedding of the given graph (recall that we assume combinatorial embeddings to be given as doubly-connected edge lists (DCEL), cf. Chapter 2). In fact, by 3-connectivity, the combinatorial embedding is unique up to its orientation (cf. Chapter 2), and there is no difference between separating 3-cycles and separating triangles. So let G be a plane 3-connected simple graph where every vertex that belongs to a separating triangle has degree at most five.

In this section, we deal with three special cases for which Theorem 18 is easily proven. In particular, we discuss the case of graphs without separating triangles. Then, we show that we can assume the separating triangles of G to be trivial and pairwise vertex-disjoint.

We remark that the notation G retains its meaning throughout the remainder of Chapter 4. We refer to G as our *general assumption*. We emphasize that many of the following theorems are stated for more general graph classes. To avoid mix-ups with G, the graphs in the preconditions of these theorems are usually labeled $\mathcal{G}$.

4.3.1 Graphs without separating triangles

By combining known results, it is easy to deal with the case that G has no separating triangle. In fact, a plane graph without separating triangles is subhamiltonian even if it is not (3-)connected or degree bounded. This was observed already by Bauernöppel [20], and by Kainen and Overbay [85]. The statement is useful even if G has separating triangles, in which case we plan to remove them by using edge contractions. To make use of the thereby obtained subhamiltonian cycle for the reduced graph, it will be useful to prescribe edges to be part of the cycle.

Bauernöppel [20] states the fact that one edge may be prescribed. Here, we need a stronger statement that allows two edges of the graph to be prescribed:

Theorem 21. *If a simple plane graph $\mathcal{G} = (V, E)$ does not contain any separating triangle, then for any two distinct edges $e_1, e_2 \in E$ there is a plane augmentation $\mathcal{H}$ of $\mathcal{G}$ that contains a Hamiltonian cycle C using e_1 and e_2. Moreover, the cycle C can be computed in $O(n^2)$ time, where n is the number of vertices in $\mathcal{G}$. If both e_1 and e_2 belong to the outer face of $\mathcal{G}$ and the outer face is a triangle, then C can be computed in linear time.*

Proof. We may assume that $\mathcal{G}$ is biconnected. If not, use a linear time algorithm of Read [110] to make the graph biconnected by adding edges without introducing separating triangles. As $\mathcal{G}$ is biconnected, all faces are bounded by simple cycles.

Biedl, Kant, and Kaufmann [23] have shown that if $\mathcal{G}$ does not contain a star, then $\mathcal{G}$ can be augmented (by adding edges) in linear time and space to a maximal plane graph, without introducing separating triangles. Here, $\mathcal{G}$ *contains a star* (f, v) if there is a face f and a vertex v such that

(1) either ∂f is a cycle on at least 4 vertices such that $v \notin \partial f$, and all vertices of ∂f are neighbors of v in $\mathcal{G}$;

(2) or ∂f is a cycle on at least 5 vertices such that $v \in \partial f$, and all vertices of $\partial f \setminus \{v\}$ are neighbors of v in $\mathcal{G}$.

The algorithm by Biedl et al. either returns a star (f, s) in $\mathcal{G}$, or it returns a maximal plane graph $\mathcal{H} = (V, E')$ without separating triangles.

Hence, if $\mathcal{G}$ does not contain a star, then we obtain a 4-connected plane augmentation $\mathcal{H}$ of $\mathcal{G}$, in which the desired cycle exists by a theorem of Sanders [114, Corollary 2]. Algorithmically, such a cycle may be obtained in quadratic time using a recently developed algorithmic version of Sanders' theorem by Schmid and Schmidt [115]. For the special case where both prescribed edges are on the outer face of $\mathcal{G}$ and the outer face is triangular, we can use a linear time algorithm by Chiba and Nishizeki [30].

Otherwise, $\mathcal{G}$ contains a star (f, s). Let $S = V(\partial f) \cup \{s\}$ and consider the induced plane graph $\mathcal{G}[S]$. As s is adjacent to all vertices in $S \setminus \{s\}$, we know that all faces of $\mathcal{G}[S]$ other then f are triangles. Moreover, these triangles are also faces of $\mathcal{G}$ because $\mathcal{G}$ does not have a separating triangle by assumption. It follows that $\mathcal{G} = \mathcal{G}[S]$.

If (f, s) is of type (1), then set $\mathcal{G}' := \mathcal{G}$; otherwise, (f, s) is of type (2) and we derive $\mathcal{G}'$ from $\mathcal{G}$ by adding an edge between the two neighbors of s on ∂f. Then $\mathcal{G}'$ is a wheel with at least four spokes. Let f' denote the (only) nontriangular face of $\mathcal{G}'$, and let $(u_0, \ldots, u_k = u_0)$ denote the circular sequence of vertices along $\partial f'$. Unless both e_1 and e_2 are incident to s, the desired Hamiltonian cycle is easily found in $\mathcal{H} := \mathcal{G}'$. Otherwise, $e_1 = \{s, u_i\}$ and $e_2 = \{s, u_j\}$, with $0 \leq i < j < k$. If u_i and u_j are adjacent along ∂f, then again we find the desired Hamiltonian cycle in $\mathcal{H} := \mathcal{G}'$.

Otherwise, let $\mathcal{H} := \mathcal{G}' \cup e$, where $e = \{u_{i+1}, u_{(j+1) \bmod k}\}$. The desired Hamiltonian cycle in $\mathcal{H}$ is $(u_{(j+1) \bmod k}, \ldots, u_i, s, u_j, \ldots, u_{i+1})$, see Figure 22. Clearly, this case can be handled in linear time. □

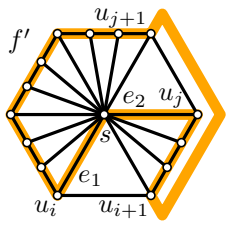

Figure 22: A plane subhamiltonian cycle using e_1 and e_2 in a wheel with at least four spokes.

4.3.2 Graphs with nontrivial separating triangles

To be able to argue inductively, we prove a stronger statement than necessary, namely a version of Theorem 18 where, similar to the statement in Theorem 21, two edges of the desired plane subhamiltonian cycle may be prescribed.

Theorem 22. *Let $\mathcal{G} = (V, E)$ be a 3-connected simple plane graph on n vertices where every vertex that belongs to a separating triangle has degree at most five. Then there is a plane augmentation of $\mathcal{G}$ that contains a Hamiltonian cycle C, which can be computed in $O(n^2)$ time.*

Moreover, for certain graphs, we may prescribe two edges to be part of the cycle C. Let $F \subset E$ be a set of up to two edges such that if $F \neq \emptyset$, the following conditions are satisfied:

(P1) the edges of F belong to the outer face $T_\circ$ of $\mathcal{G}$;

(P2) $T_\circ$ is a triangle;

(P3) no vertex of $T_\circ$ belongs to a separating triangle of $\mathcal{G}$; and

(P4) either at least one vertex of $T_\circ$ has degree three in $\mathcal{G}$, or all vertices of $T_\circ$ have degree four in $\mathcal{G}$.

The cycle C uses all edges in F.

The proof, described in Sections 4.5–4.8, is carried out by induction on the number of vertices. Due to the ability to prescribe edges, it is easy to recursively deal with the case that G contains at least one nontrivial separating triangle.

Lemma 23. *Suppose that the statement of Theorem 22 holds for all graphs with at most $n - 1$ vertices, where $n \geq 6$. Then it also holds for every graph on n vertices that contains at least one nontrivial separating triangle.*

Proof. As a first step we determine all maximal separating triangles, which is easy to accomplish in linear time, given that the vertices of separating triangles have degree at most five. If there exists a nontrivial maximal separating triangle T, then consider the graph $\mathcal{G}'$ obtained from $\mathcal{G}$ by replacing $\mathcal{G}_T^-$ by K_4. As $\mathcal{G}'$ has less vertices than $\mathcal{G}$, by the inductive hypothesis we obtain in $O(n^2)$ time a plane augmentation of $\mathcal{G}'$ that contains a Hamiltonian cycle $\mathcal{H}'$ using all edges from F (note that the outer edges of $\mathcal{G}$ remain in $\mathcal{G}'$). Denoting $T = uvw$, the cycle $\mathcal{H}'$ traverses the unique vertex t inside T via two neighbors, without loss of generality $\mathcal{H}' = \dots utv \dots$.

We will show that there is a plane subhamiltonian cycle $\mathcal{H}''$ for $\mathcal{G}_T^-$ that uses the edges uw and wv. The path $P_1 = \mathcal{H}'' \setminus \{w\}$ is a path from u to v in a plane augmentation of G_T^- that visits all vertices of G_T^- except w. On the other hand, $P_2 = \mathcal{H}' \setminus \{t\}$ is a path from u to v in a plane augmentation of $\mathcal{G}'$ that visits all vertices of $\mathcal{G}'$ except t and uses the edges of F. The combination of P_1 and P_2 yields the desired plane subhamiltonian cycle $\mathcal{H}$ for $\mathcal{G}$ that passes through the edges from F.

It remains to prove the existence of $\mathcal{H}''$. If $\mathcal{G}_T^-$ satisfies the preconditions of Theorem 22 for the prescribed edges uw and wv, then the cycle $\mathcal{H}''$ may be obtained in $O(n^2)$ time due to the assumption that Theorem 22 holds for graphs with at most $n-1$ vertices. So suppose that the preconditions are not satisfied. We claim that the only violated property is (P3), which can be recognized in constant time. Indeed, clearly $\mathcal{G}_T^-$ is 3-connected and satisfies the degree bounds. Moreover, the Properties (P1) and (P2) are satisfied. Property (P4) also holds: due to the degree bounds of $\mathcal{G}$, the vertices u, v, w have degree at most four in $\mathcal{G}_T^-$. Consequently, Property (P3) is violated and, so, one of the vertices of u, v, w belongs to a separating triangle in $\mathcal{G}_T^-$. In fact, since the degrees of u, v, and w are bounded by four and by 3-connectivity, it follows that an *edge* of uvw belongs to a separating triangle abc, say w.l.o.g. $\{a, b\} \subset \{u, v, w\}$. Note that $\mathrm{V}(\mathcal{G}_T^-) \setminus \mathrm{V}(\mathcal{G}_{abc}^-) = \{u, v, w\} \setminus \{a, b\}$ by 3-connectivity and the degree bounds of $\{a, b\} \cap \{u, v, w\}$, for illustrations see Figure 23. Hence, our plan is to recurse on $\mathcal{G}_{abc}^-$ and to then incorporate the missing vertex such the two desired edges uw and vw are part of the resulting cycle.

Just like $\mathcal{G}_T^-$, for any two prescribed outer edges the graph $\mathcal{G}_{abc}^-$ satisfies all preconditions except, possibly, Property (P3). In fact, Property (P3) also holds: since the degrees of a, b, c in $\mathcal{G}_{abc}^-$ are bounded by four, a violation of Property (P3) would imply that an edge of abc belongs to a separating triangle. However, given that the vertices a, b have degree exactly four in $\mathcal{G}_T^-$, their degrees in $\mathcal{G}_{abc}^-$ are exactly three. Consequently, no edge of abc is part of a separating triangle since in a 3-connected graph, every vertex of a separating triangle has degree at least four. Thus, all preconditions are satisfied and, by induction, we obtain in $O(n^2)$ time

- a plane subhamiltonian cycle $\mathcal{H}_{ab}$ for $\mathcal{G}_{abc}^-$ that uses the edges ca and cb;
- a plane subhamiltonian cycle $\mathcal{H}_{bc}$ for $\mathcal{G}_{abc}^-$ that uses the edges ab and ac; and
- a plane subhamiltonian cycle $\mathcal{H}_{ca}$ for $\mathcal{G}_{abc}^-$ that uses the edges ba and bc.

Using these cycles, we obtain the desired plane subhamiltonian cycle $\mathcal{H}''$ for $\mathcal{G}_T^-$. We distinguish three cases regarding the identity of a, b.

Case 1: $\{a, b\} = \{u, v\}$. Refer to Figure 23a. We set $\mathcal{H}'' = (\mathcal{H}_{cu} \setminus \{v\}) \cup cvwu$. ◁

Case 2: $\{a, b\} = \{u, w\}$. Refer to Figure 23b. We set $\mathcal{H}'' = (\mathcal{H}_{cu} \setminus \{w\}) \cup cvwu$. ◁

Case 3: $\{a, b\} = \{v, w\}$. Refer to Figure 23c. We set $\mathcal{H}'' = (\mathcal{H}_{cv} \setminus \{w\}) \cup cuwv$. ◁

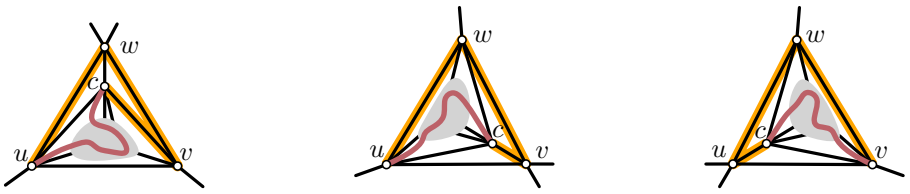

(a) $\{a, b\} = \{u, v\}$. (b) $\{a, b\} = \{u, w\}$. (c) $\{a, b\} = \{v, w\}$.

Figure 23: An edge of the triangle uvw may be part of another separating triangle abc.

As discussed above, these cycles are indeed plane subhamiltonian for $\mathcal{G}_T^-$ and they use the edges uw and vw, as desired. The overall runtime is quadratic because we use a constant number of recursive calls and merge the obtained cycles in constant time. □

4.3.3 Graphs with nondisjoint separating triangles

The degree restriction basically enforces that the separating triangles of G are pairwise vertex-disjoint. However, there are some exceptional configurations where two separating triangles share an edge (we encountered such a situation already in the proof of Lemma 23, see Figure 23). Our next goal is to classify these configurations precisely.

A *double kite* is a subgraph $U \simeq K_4$ of a plane graph $\mathcal{G}$ so that exactly two of the four triangles of U are separating in $\mathcal{G}$, see Figures 23 and 24a. The two separating triangles are said to *define* the double kite. Note that $\mathcal{G}$ may contain multiple double kites, see Figure 24c. We refer $\mathcal{G}$ itself as a *trivial double kite* if it is 3-connected, contains a double kite, and has precisely 6 vertices, see Figure 24b.

The following lemma shows that two separating triangles of graphs such as our general assumption G never share a single vertex, and they share an edge if and only if they define a double kite. In fact, the statement is more general since the degree constraint is only required for one of the two separating triangles. This plays an important role later on: in Section 4.5, we describe a recursive algorithm to get rid of the trivial separating triangles in G. The recursive calls are applied to rigid triconnected components of a graph that is the result of contracting some edges in G. These components may contain separating triangles that are not part of the original graph G. In particular, these triangles do not need to satisfy the degree bound.

Lemma 24. *Let $\mathcal{G}$ be a 3-connected simple plane graph and let T_1, T_2 be two distinct separating triangles of $\mathcal{G}$ such that every vertex that belongs to T_1 has degree at most five. Then if T_1 and T_2 share a vertex, the triangles T_1 and T_2 define a double kite.*

Proof. Let v denote a common vertex of T_1 and T_2. As $\mathcal{G}$ is 3-connected, every vertex of a separating triangle is adjacent to at least one vertex on each of its sides, as otherwise, the other two vertices of the triangle would be separating. We distinguish two cases.

Case 1: v is the only common vertex of T_1 and T_2. Then v is incident to four pairwise distinct edges from T_1 and T_2. Moreover, v has some neighbor v' on the side of T_1 that does not contain T_2, bringing the degree of v up to five. However, v also has some neighbor on the side of T_2 that does not contain T_1 and v'. This brings the degree of v up to six; contrary to our assumption. ◁

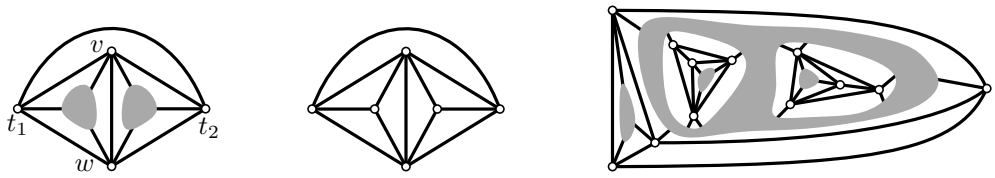

(a) (b) (c)

Figure 24: (a) The degree bounds imply that two nondisjoint separating triangles form a double kite. (b) A trivial double kite. (c) Schematic of a graph with three double kites.

Case 2: T_1 and T_2 share two vertices v and w. For an illustration refer to Figure 24a. Then v and w each are incident to three pairwise distinct edges from T_1 and T_2. Let t_1 and t_2 denote the remaining vertex ($\neq v, w$) of T_1 and T_2, respectively.

Both v and w have some neighbor on the side of T_1 that does not contain t_2. These neighbors belong to the same side of T_2 as t_1. Hence, both v and w have some additional neighbor on the other side of T_2, bringing their degrees up to the maximum of five. This implies that there is no vertex located in the intersection of the side of T_1 that contains t_2 and the side of T_2 that contains t_1, since otherwise t_1 and t_2 would form a separation pair; in contradiction to $\mathcal{G}$ being 3-connected. It also follows that t_1 and t_2 are adjacent because otherwise v, w is a separation pair of $\mathcal{G}$. Thus, T_1 and T_2 define a double kite. □

Observation 25. *Let $\mathcal{G}$ be a 3-connected simple plane graph and let T_1, T_2 be two distinct trivial separating triangles of $\mathcal{G}$ such that all vertices of T_1 and T_2 have degree at most five. Then if T_1 and T_2 share a vertex, the graph $\mathcal{G}$ is a trivial double kite.*

Proof. If T_1 and T_2 share a vertex, Lemma 24 implies that they define a double kite. The claim follows since both triangles are trivial. □

For a trivial double kite G, the statement of Theorem 22 is easily verified. Hence, from now on we may assume that the separating triangles of G are trivial (by Lemma 23) and pairwise vertex-disjoint (by Observation 25).

4.4 Proof overview

To prove Theorem 22 we proceed in three main steps. We will first remove all separating triangles of G by using edge contractions. We then obtain a plane subhamiltonian cycle in an augmented version of the reduced graph. Finally, we undo the contractions while maintaining a plane subhamiltonian cycle. In this section, we give a high-level overview of these three steps.

Collapsing edges. In the first step (Section 4.5), we describe a recursive algorithm that destroys all separating triangles of G by using an operation that replaces a trivial separating triangle with a single edge as illustrated in Figure 25.

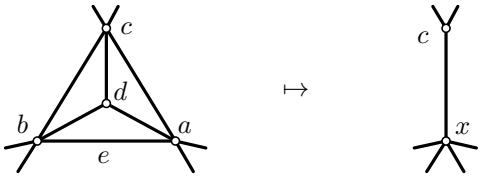

Figure 25: Collapsing the edge $e = ab$ of the separating triangle $T = abc$.

To define this operation, let $\mathcal{G}$ be a simple plane graph. We emphasize that we do not require the separating triangles of $\mathcal{G}$ to be trivial or pairwise vertex-disjoint. Moreover, the vertices of separating triangles can have arbitrary degree.

Definition 26 (Collapsible edge)**.** An edge $e = ab$ of $\mathcal{G}$ is *collapsible* if it belongs to exactly one separating triangle $T = abc$ of $\mathcal{G}$ and this triangle is trivial and its vertices have degree at most five.

We remark that in our general assumption G *each* edge of a separating triangle is collapsible. However, to make our recursive strategy work, we need our operation to be applicable to the more general case of graphs such as $\mathcal{G}$, where this might not be the case.

The edge e of $\mathcal{G}$ is *collapsed* by contracting (the three edges of) the triangle formed by a, b, and the single vertex d inside T into a new vertex x and merging the two edges ac and bc to a new edge cx, see Figure 25. Let $\mathcal{G}'$ denote the resulting graph. If a and b have a common neighbor $z \notin \{c, d\}$ in $\mathcal{G}$, then, in $\mathcal{G}'$, we merge the two parallel edges between z and x into a singleton edge zx. Since e is collapsible, the triangle abz is nonseparating and, hence, the two parallel edges form a 2-cycle that has all vertices of $\mathrm{V}(\mathcal{G}') \setminus \{x, z\}$ on one of its two sides. Therefore, the embedding of the merged edge zx is uniquely determined. Due to this merging step, the result

of an edge collapse is always a simple graph $\mathcal{G}''$. The collapse operation does not increase the degree of any vertex. In particular, the degree of c decreases by two. The degree of the new vertex x is at most five. Hence, the graph $\mathcal{G}''$ satisfies the degree constraints, unless the collapse creates a *new* separating triangle that has a vertex of degree larger than five. Hence, we need to ensure that this does not happen.

An edge collapse can be performed in constant time in a DCEL.

In Section 4.5, we describe a procedure to find a set $K \subset \mathrm{E}(G)$ of collapsible edges such that simultaneously collapsing all edges of K in G results in a graph G' that does not contain any separating triangle. By *simultaneously collapsing* the edges of K, we refer to the process of collapsing the edges of K iteratively, in an arbitrary order. This is not always a well defined operation since the collapse of an edge $e \in K$ might make some other edge $e' \in K$ noncollapsible. We discuss in Section 4.5 how to ensure that G' is a well defined graph.

Stellation. In the second step (Section 4.6), we augment G' by *stellating* every nontriangular face, that is, for each such face f of G' we insert a new vertex v_f into f and we add an edge between v_f and each vertex on the boundary of f. The result is a triangulation G''. By choosing the set K suitably in the previous step, we ensure that G'' does not contain any separating triangle. Thus, using Theorem 21 we obtain a plane subhamiltonian cycle C'' for G''.

Reconstruction. By removing the new vertices v_f from C'', we obtain a plane subhamiltonian cycle C' for G' that passes through every face at most once, that is, for each face f of G', the cycle C' contains at most one chord that is drawn in f. We say that C' is unichordal.

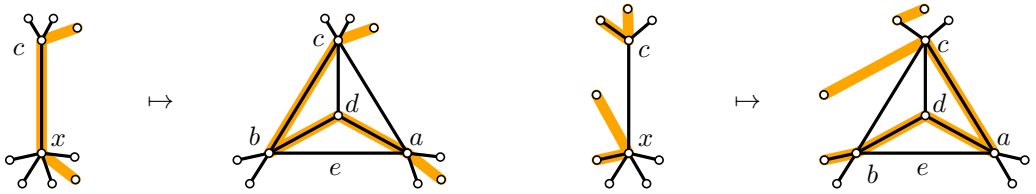

Figure 26: Examples of two reconstructions. In (a) we can simply replace the edge cx with a path that is spanning for the closed interior of the separating triangle. In (b) we first shortcut the current cycle at c in order to visit a and d.

In the third step of our algorithm (Section 4.7), we iteratively revert the edge collapses while maintaining a plane subhamiltonian cycle. This is done by applying small local modifications to the current cycle at every reconstruction step; Figure 26 depicts some examples. Such modifications are not always possible: Figure 27 shows two examples for which there is no easy way to locally extend the current cycle.

However, since the initial cycle C' is unichordal, it cannot contain any of these two forbidden configurations. The challenging aspect is, that applying local modifications to one of the triangles might introduce one of the forbidden configurations at another triangle. Hence, it is important to control the interaction of the cycle with the remaining triangles at each reconstruction step.

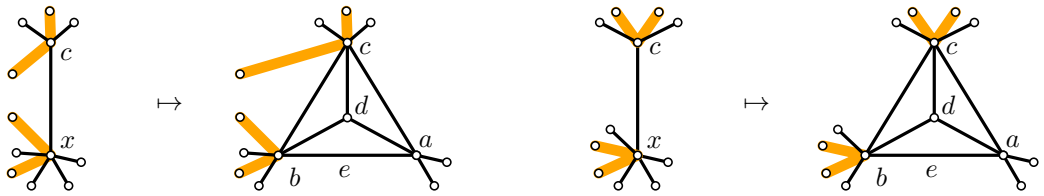

Figure 27: Two types of remaining triangles to avoid: there is no easy way to extend the cycle. Note that in (a) there is no way to shortcut the current cycle. On the other hand, in (b) we can shortcut the cycle at c (and, symmetrically, at b), however, this shortcut is not beneficial since the cycle edges incident to x do not pass through a face with c on its boundary as in Figure 26b.

Special cases. In some special cases, the three step procedure described above cannot be applied. In Section 4.8, we deal with these special cases and we formally summarize the proof. We proceed by describing the three steps for the general case in more detail.

We remark that Bauernöppel [20] also used edge contractions and local modifications as in Figure 26 to show that every planar graph of maximum degree four is subhamltonian. However, in the degree four case, it suffices to contract a *single* separating triangle T and inductively obtain some arbitrary subhamiltonian cycle. Challenging cases as in Figure 27 do not exist, so the reconstruction can be easily performed by distinguishing a small number of cases. This allows Bauernöppel to describe a very short (one page) proof of his result. Our proof for the degree five case is significantly more involved since, to avoid the forbidden configurations, we need to augment the graph and, thus, cannot proceed inductively. Instead, we rely on Theorem 21, which requires us to get rid of *every* separating triangle. In particular, we also have to avoid creating new separating triangles when performing a collapse.

4.5 Collapsing edges

Recall that our general assumption G is a 3-connected simple plane graph such that every vertex that belongs to a separating triangle has degree at most five. Moreover, the separating triangles of G are trivial and pairwise vertex-disjoint (by Lemma 23 and Observation 25). In particular, this implies that G does not contain a double

kite. Let S denote the set of separating triangles of G. Our goal is to get rid of all triangles in S by using edge collapses. To this end, we want to find a set $K \subset E$ of edges so that (1) every edge in K belongs to a triangle in S; (2) every triangle in S has exactly one edge of K; and (3) the graph obtained by simultaneously collapsing the edges of K is well defined and does not contain any separating triangle.

Obviously, there exists a set $\hat{K}$ of edges that satisfies (1) and (2). Suppose that the graph $\hat{G}$ obtained from G by collapsing the edges in $\hat{K}$ is well defined, but contains a separating triangle T. By (2) we know that T is not a triangle in G, but T corresponds to a separating k-cycle in G, for $k \geq 4$. By (1), (2), and since the triangles in S are pairwise vertex-disjoint, at most every other edge of a cycle in G is in $\hat{K}$ and, therefore, we have $k \leq 6$. In other words, every separating triangle in $\hat{G}$ corresponds to a separating k-cycle in G where $k \in \{4, 5, 6\}$ and exactly $k - 3$ edges are in $\hat{K}$. Moreover, for any such separating k-cycle in G, both the interior and the exterior must contain at least one vertex that is not the interior vertex of a triangle from S because by (2) every interior vertex of a triangle from S disappears when collapsing the edges in $\hat{K}$.

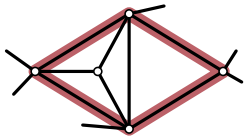

Figure 28: A separating 4-cycle that is not hyperseparating with respect to the depicted separating triangle.

We now formally define these special separating cycles and introduce some related terminology. The set K will be obtained with a recursive algorithm. For this reason, we need the definition to apply not only to graphs such as G, but to a more general case: let $\mathcal{G}$ be a simple plane graph and let $\mathcal{S}$ be a set of trivial separating triangles of $\mathcal{G}$ where every vertex of a triangle in $\mathcal{S}$ has degree at most five. We emphasize that not all separating triangles of $\mathcal{G}$ need to be trivial, and the set $\mathcal{S}$ does not have to contain all trivial separating triangles of $\mathcal{G}$. Inspired by the observations in the previous paragraph, we call a cycle C of $\mathcal{G}$ *hyperseparating* with respect to $\mathcal{S}$ if both the interior and the exterior contain at least one vertex that is not the interior vertex of a triangle from $\mathcal{S}$. Every hyperseparating cycle is separating, but the converse is not true, see Figure 28. We define an *inhibitor* to be a hyperseparating (with respect to $\mathcal{S}$) k-cycle I, where $k \in \{4, 5, 6\}$, for which at least $k - 3$ edges belong to a trivial separating triangle of $\mathcal{S}$, for an illustration refer to Figure 30d. We refer to these at least $k - 3$ edges as *constrained*. The idea behind this definition is that an inhibitor inhibits us from collapsing all of its constrained edges simultaneously since the result would be a separating triangle (or even a separating 2-cycle if $k = 4$ and two edges of I are constrained). An edge e of $\mathcal{G}$ is *constrained* with respect to $\mathcal{S}$ if there exists an inhibitor I with respect to $\mathcal{S}$ so that e is a constrained edge of I and *unconstrained*

(with respect to $\mathcal{S}$), otherwise. An inhibitor of length $k \in \{4, 5, 6\}$ is also called a *k-inhibitor*.

For brevity, when we speak about inhibitors in the context of our general assumption G, we always refer to inhibitors with respect to S. Note that a set $\hat{K}$ of edges from G that satisfies (1) and (2) also satisfies (3) if for every k-inhibitor, no more than $k-4$ of its constrained edges are in $\hat{K}$. In particular, if there is no k-inhibitor such that $\hat{K}$ has more than $k-3$ of its constrained edges, then collapsing the edges of $\hat{K}$ is a well defined operation.

Our recursive algorithm can be thought of as an iterative process that selects the edges of K one by one and performs a collapse immediately after each edge is added to K. Hence, the most important case of inhibitors are those of length four since only these can be turned into a separating triangle by a single edge collapse. For this reason, the next section is devoted to studying the structure of 4-inhibitors.

4.5.1 Structure of 4-inhibitors

First, let us observe that 4-inhibitors of G are chordless. For the proof we require the following basic observation.

Observation 27. *Let $\mathcal{G}$ be a 3-connected simple plane graph and let $T = abc$ be a trivial separating triangle of $\mathcal{G}$ with inner vertex d such that all vertices of T have degree at most five. Further, let $I_{ab} = abxy$ be a 4-inhibitor with respect to some set $\mathcal{S} \supseteq \{T\}$ that constrains ab. Then $\{x, y\} \cap \{c, d\} = \emptyset$.*

Proof. Clearly, $x \neq d$ and $y \neq d$, as otherwise I_{ab} is uniquely determined and not separating.

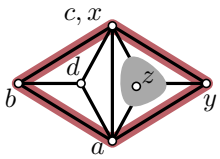

Figure 29: An inhibitor constraining the edge ab of the triangle abc cannot pass through c or d.

So assume without loss of generality that $x = c$. For an illustration refer to Figure 29. By 3-connectivity and since I_{ab} is *hyper*separating with respect to $\mathcal{S}$, there exists a vertex $z \neq d$ on the side of I_{ab} that contains d. Since d is the only vertex on its side of abc, the vertex z is located on the side of acy that does not contain d. By 3-connectivity, this saturates the degree bound of both a and c. However, by 3-connectivity and since I_{ab} is separating, at least one of a, c has an additional neighbor on the side of I_{ab} that does not contain d and z; a contradiction. □

Lemma 28. *Let $\mathcal{G}$ be a 3-connected simple plane graph and let $T = abc$ be a trivial separating triangle of $\mathcal{G}$ with inner vertex d such that all vertices of T have degree at most five. Suppose that ab is constrained by a 4-inhibitor $I_{ab} = abxy$ with respect to some set $\mathcal{S} \supseteq \{T\}$. Then I_{ab} is chordless, unless ab belongs to two separating triangles that define a double kite.*

Proof. Assume without loss of generality that by is a chord of I_{ab}. By Observation 27, $x, y \notin \{c, d\}$ and, so, the degree of b is saturated, see Figure 30a. We claim that by is on the side of I_{ab} that contains c and d. To see this, assume the contrary. Since I_{ab} is separating, there is some vertex z on the side of I_{ab} that does not contain c, d. The chord by splits this side of I_{ab} into two triangles, one of which contains z, see Figure 30b and 30c. However, by 3-connectivity, this implies that b has a neighbor in this triangle, in contradiction to the degree bound of b. So the claim holds. Then aby is a triangle that separates x from c, d, see Figure 30d. By Lemma 24, the triangles abc and aby define a double kite. □

Since G contains no double kite, it follows that all its 4-inhibitors are chordless.

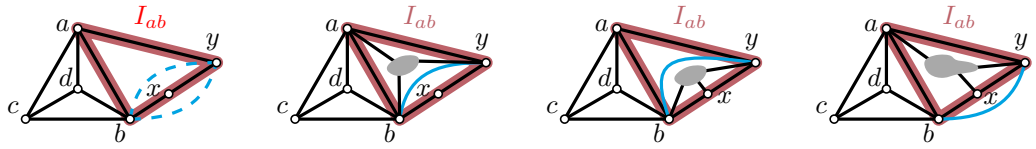

(a) (b) (c) (d)

Figure 30: A 4-inhibitor $I_{ab} = abxy$ that constrains the edge ab of the separating triangle abc.

A necessary requirement for the existence of a set K of edges to collapse is that our graph does not contain a separating triangle for which every edge is constrained by a 4-inhibitor. We show that this requirement is met, except for some specific special cases.

Lemma 29. *Let $\mathcal{G}$ be a 3-connected simple plane graph and let T be a trivial separating triangle of $\mathcal{G}$ such that all vertices of T have degree at most five. Further, let $\mathcal{I}$ be a set of 4-inhibitors of $\mathcal{G}$ with respect to some set $\mathcal{S} \supseteq \{T\}$. Finally, assume that $\mathcal{G}$ contains no separating triangle that together with T defines a double kite.*

Then either (1) T has at least one edge that is not constrained by a 4-inhibitor of $\mathcal{I}$; or (2) $\mathcal{G}$ contains a subgraph $\mathcal{G}'$ that is isomorphic to $\mathcal{G}_2$ or one of the graphs of the family $\mathcal{G}_1$ defined in Figure 31 such that each thick (colored) edge of $\mathcal{G}'$ belongs to some 4-inhibitor of $\mathcal{I}$ that constrains an edge of T.

Proof. For the sake of contradiction, let us assume that each edge of $T = abc$ is constrained by a 4-inhibitor of $\mathcal{I}$, and that $\mathcal{G}$ contains no subgraph $\mathcal{G}'$ as in the claim. We denote the inhibitors of $\mathcal{I}$ constraining ab, bc, ac by $I_{ab} = abxy$, $I_{bc} = bcst$, and

(a) $\mathcal{G}_1$ (b) $\mathcal{G}_2$

Figure 31: If all three edges of a separating triangle in the graph $\mathcal{G}$ are constrained by a 4-inhibitor, there is a subgraph $\mathcal{G}'$ isomorphic to $\mathcal{G}_2$, depicted in (b), or one of the graphs of the family $\mathcal{G}_1$, depicted in (a). The gray parts in (a) represent arbitrary nonempty subgraphs (possibly consisting of a single vertex) that result in an overall 3-connected graph. Each black edge is part of the graph. The vertices s, x, y might have additional neighbors, whereas the neighborhood of the vertices a, b, c, d is exactly as depicted.

$I_{ac} = acvw$, respectively. By Observation 27, we have $c, d \notin I_{ab}$, $a, d \notin I_{bc}$, and $b, d \notin I_{ac}$. By 3-connectivity and since I_{ab} is separating, we may assume without loss of generality that b has a neighbor z which is located on the side of I_{ab} that does not contain c, see Figure 32b. Due to the degree bound of b, we have $t \in \{x, z\}$. Accordingly, we distinguish two cases.

Case 1: $t = x$. We proceed by studying the identity of s, the fourth vertex of I_{bc}. We distinguish several cases. By planarity, we have $s \neq z$, as z is located on the side of I_{ab} that does not contain c. Next, assume $s = y$ and, thus, $I_{bc} = bcyx$, for an illustration see Figure 32a. By 3-connectivity, the degree bound of b, and since I_{bc} is separating, each of the vertices c, x, y has a neighbor on the side of I_{bc} that does not contain a. Let r denote the according neighbor of c. We study the third and final inhibitor $I_{ac} = acvw$. By Lemma 28, we have $v \neq y$, as otherwise ay would be a chord of I_{ac}. Due to the degree bound of c, it remains to consider the case $v = r$. By planarity, the vertex w has to belong to I_{bc} since r and a are located on distinct sides of I_{bc}. By Lemma 28, we have $w \neq y$, as otherwise cy would be a chord of I_{ac}. Therefore, the only remaining option is $w = x$. However, this implies the existence of the edge ax, which is a chord of I_{ab}; a contradiction to Lemma 28.

So far, we have established that $s \notin \{a, b, c, d, x, y, z\}$. Therefore, we have $I_{bc} = bcsx$ where s is distinct from all the previously considered vertices. By planarity, s is located on the side of I_{ab} that contains c, for an illustration see Figure 32b. By 3-connectivity, the degree bound of b, and since I_{bc} is separating, each of the vertices c, x, s has a neighbor on the side of I_{bc} that does not contain a. Let r' denote the according neighbor of c. We study the third inhibitor $I_{ac} = acvw$. Due to the degree bound of c, we have $v \in \{s, r'\}$. Accordingly, we distinguish two cases.

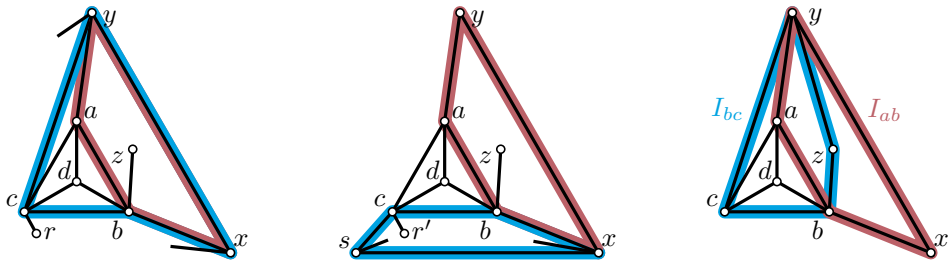

(a) $I_{bc} = bcyx$ (b) $I_{bc} = bcsx$ (c) $I_{bc} = bcyz$

Figure 32: Illustrations for the Cases 1 (a)–(b) and 2 (c) considered in Lemma 29.

First, assume that $v = r'$. By planarity, we have $w \in I_{bc}$. Further, by Lemma 28 applied to I_{ac}, we have $w = x$. This implies that ax is a chord of I_{ab}; a contradiction to Lemma 28. It remains to consider the case $v = s$. We study the identity of w. The case $w = x$ can be excluded by Lemma 28 applied to I_{ab}. Due to the degree bounds of a, b, and c, by 3-connectivity, and since I_{ac} is separating, we also have that $w \neq y$; otherwise, by 3-connectivity and the degree bounds, $\mathcal{G}$ would contain a subgraph $\mathcal{G}'$ isomorphic to a $\mathcal{G}_1$ as in the claim, which contradicts our assumptions. This establishes that w does not belong to I_{ab} or I_{bc}. Thus, by planarity, w it is located on the sides of I_{ab} and I_{bc} that contain c and a, respectively, and, therefore, it is distinct from all the previously considered vertices. This brings the degree of a up to the maximum of five. By 3-connectivity and since I_{ac} is separating, three vertices of I_{ac} have a neighbor on the side of I_{ac} that does not contain b. This is a contradiction to the degree bound of a or c. ◁

Case 2: $t = z$. Recall that $I_{bc} = bcst = bcsz$ where z is on the side of I_{ab} that does not contain c, for an illustration see Figure 32c. By planarity, it follows that s belongs to I_{ab} and by Lemma 28 applied to I_{bc} we have $s = y$ and, thus, $I_{bc} = bcyz$.

We study the third inhibitor $I_{ac} = acvw$. The vertex v has to be distinct from y, as otherwise ay would be a chord of I_{ac} in contradiction to Lemma 28. Note that the triangle acy is nonseparating by Lemma 24. Thus, v is located on the side of I_{bc} that does not contain a and, hence, by planarity, v or w has to belong to I_{bc}. Since $c \in I_{ac} \cap I_{bc}$, Lemma 28 implies that the only other vertex of I_{ac} in this intersection is z. As z and c are located on distinct sides of I_{ab}, the vertex v must belong to I_{ab} and, thus, $v = x$ and $I_{ac} = acxz$. Altogether, this establishes that $\mathcal{G}$ contains a subgraph $\mathcal{G}'$ isomorphic to $\mathcal{G}_2$ as in the claim, which contradicts our assumption. ◁

Altogether, we have obtained a contradiction. Hence, at least one of the statements (1) and (2) is true. Moreover, it is easy to see that not both of them can be true simultaneously. □

In Lemma 23, we have used the inductive framework of the proof of Theorem 22 to deal with the case that G contains a nontrivial separating triangle. In Section 4.8, we

show that a similar but more involved idea can be used to deal with the case that G contains a subgraph isomorphic to a $\mathcal{G}_1$ as in Lemma 29. In the same section, we also show that if G contains a subgraph isomorphic to $\mathcal{G}_2$ as in Lemma 29, then it has constant size, which makes it easy to directly find the desired plane subhamiltonian cycle.

For now, let us focus on the idea for the general case and, hence, assume from now on that G does not contain a subgraph isomorphic to a $\mathcal{G}_1$ or $\mathcal{G}_2$ as in Lemma 29. The idea to obtain the desired edge set K is to collapse some unconstrained edge of a separating triangle in G. The existence of such an edge is guaranteed by Lemma 29. This process may be repeated if the resulting graph is 3-connected. However, in general an edge collapse may reduce the connectivity of the graph. In this case, we plan to recurse on the triconnected components of the graph. To make such a recursion work in the context of our overall proof strategy, we must take special care concerning the vertices of separation pairs. Specifically, as we will discuss in the following section, we should never create a separation pair whose vertices are adjacent.

4.5.2 Avoiding adjacent separation pairs

Recall that we plan to stellate each nontriangular face of the graph G' that is obtained by simultaneously collapsing all edges in K, and that we need to ensure that the resulting graph G'' does not contain a separating triangle. Consider a face f of G' and assume that its stellation creates a separating triangle $s = abv_f$ where v_f is the new vertex inserted into f. Note that the vertices a and b belong to f. Therefore, the edge $ab \in \mathrm{E}(G')$ is a chord of ∂f and, moreover, a, b is a separation pair of G'.

In order to avoid this situation, it suffices to choose the set K subject to the following additional constraint: (4) the graph obtained by simultaneously collapsing the edges of K does not create an *adjacent* separation pair, i.e., a separation pair a, b where a and b are adjacent. Inspired by this observation, we devise a strengthened version (Lemma 32) of Lemma 29. For its proof, we require the two observations. The first of them is a well-known fact:

Observation 30. *Let $\mathcal{G}$ be a k-connected simple graph and $S \subseteq \mathrm{V}(\mathcal{G})$ with $|S| = k$ such that $\mathcal{G} \setminus S$ is disconnected. Then for each $x \in S$ the set $\mathrm{N}_{\mathcal{G}}(x) \setminus S$ contains at least one vertex of each connected component of $\mathcal{G} \setminus S$.*

Proof. Let $x \in S$ and let K_u and K_v be two distinct connected components of $\mathcal{G} \setminus S$. Further, let $u \in \mathrm{V}(K_u)$ and $v \in \mathrm{V}(K_v)$. By k-connectivity, there exist k interior disjoint paths between u and v in $\mathcal{G}$. Since S is separating and of size k, each of these paths contains precisely one vertex of S. In particular, one of the paths passes through x and, hence, x has at least one neighbor in K_u and at least one neighbor in K_v. □

Observation 31. *Let $\mathcal{G}$ be a 3-connected simple plane graph and let $T = abc$ be a trivial separating triangle of $\mathcal{G}$ with inner vertex d. Further, assume that collapsing ab results in a graph $\mathcal{G}'$ that is 2-connected but not 3-connected; and let p_{ab}, q_{ab} be a separation pair of $\mathcal{G}'$. Then we may assume (by exchanging the roles of p_{ab}, q_{ab} if necessary) that p_{ab} is the new vertex that is created by the collapse of ab and that $q_{ab} \notin \{c, d\}$.*

Proof. Let x denote the new vertex resulting from collapsing ab and assume, for the sake of contradiction, that there is some separation pair p_{ab}, q_{ab} of $\mathcal{G}'$ where $x \notin \{p_{ab}, q_{ab}\}$. All vertices in the neighborhood of x in $\mathcal{G}' \setminus \{p_{ab}, q_{ab}\}$ belong to the same connected component. Consequently, undoing the collapse of ab does not reduce the number of connected components. This implies that p_{ab}, q_{ab} is a separation pair in $\mathcal{G}$, which contradicts the 3-connectivity of $\mathcal{G}$. Hence, we may assume without loss of generality that $p_{ab} = x$.

Obviously $q_{ab} \neq d$ since $d \notin \mathcal{G}'$. Assume that $q_{ab} = c$. By planarity and 3-connectivity, it follows that the graph $\mathcal{G} \setminus \{a, b, c\}$ has exactly two connected components. One of these components is the graph $(d, \emptyset)$, the other is the graph $\mathcal{G}''$ induced by $\mathrm{V}(\mathcal{G}) \setminus \{a, b, c, d\}$. This implies that $\mathcal{G}' \setminus \{p_{ab}, q_{ab}\} = \mathcal{G}''$ is connected, which is a contradiction to the fact that p_{ab}, q_{ab} are separating in $\mathcal{G}'$. □

We are now prepared to prove the strengthened version of Lemma 29.

Lemma 32. *Let $\mathcal{G}$ be a 3-connected simple plane graph and let $T = abc$ be a trivial separating triangle of $\mathcal{G}$ with inner vertex d such that all vertices of T have degree at most five. Further, let $\mathcal{I}$ be a set of 4-inhibitors of $\mathcal{G}$ with respect to some set $\mathcal{S} \supseteq \{T\}$. Finally, assume that $\mathcal{G}$ contains no separating triangle that together with T defines a double kite.*

Then either (1) T has at least one edge e such that (i) e is not constrained by a 4-inhibitor of $\mathcal{I}$; and (ii) collapsing e does not create an adjacent separation pair; or (2) $\mathcal{G}$ contains a subgraph $\mathcal{G}'$ that isomorphic to $\mathcal{G}_2$ or one of the graphs of the family $\mathcal{G}_1$ defined in Figure 31 such that each thick (colored) edge of $\mathcal{G}'$ belongs to some 4-inhibitor of $\mathcal{I}$ that constrains an edge of T.

Proof. For the sake of contradiction, assume that $\mathcal{G}$ contains no subgraph $\mathcal{G}'$ as in the claim and that every edge of T violates Property (i) or (ii). By Lemma 29, we may assume without loss of generality that ab is not constrained by a 4-inhibitor of $\mathcal{I}$ and, consequently, collapsing ab results in a graph containing an adjacent separation pair pq. Further, by Observation 31, we may assume without loss of generality that p is the new vertex resulting from collapsing ab and that $q \notin \{c, d\}$. In other words, a, b, q is a separating triple of $\mathcal{G}$ and, without loss of generality, we may assume that b and q are adjacent. By planarity and 3-connectivity, the graph $\mathcal{G} \setminus \{a, b, q\}$ has exactly two connected components K, K' and each of the vertices a, b, q is adjacent to at least one vertex in each of the two components. Assume w.l.o.g. that $c, d \in K$ and

let b' denote the singleton neighbor of b in K' (recall that the degree of b is bounded by five), for an illustration see Figure 33a.

We distinguish three main cases and show that none of them can occur. First, we assume that both bc and ac violate Property (i). Next, we consider the case that bc does not violate Property (i). Finally, we consider the case that bc violates Property (i), but ac does not. Overall, this yields the desired contradiction.

Case 1: Each of bc and ac is constrained by 4-inhibitor of $\mathcal{I}$. Let $I_{bc} = bcxy$ and $I_{ac} = acst$ denote the 4-inhibitors of $\mathcal{I}$ constraining bc and ac, respectively. Due to the degree bound of b and by Observation 27, we have $y \in \{b', q\}$. First, assume that $y = b'$. Since a, b, q is a separating triple, it follows that $x \in \{a, q\}$. If $x = a$, then ab is a chord of I_{bc}, and if $x = q$, then bq is a chord of I_{bc}. Hence, in both cases we obtain a contradiction to Lemma 28. It remains to consider the case $y = q$. By Observation 27 and since a, b, q is separating, the vertex x belongs to $K \setminus \{c, d\}$. By 3-connectivity and the degree bound of b, each of the vertices c, q, x has a neighbor on the side of I_{bc} that does not contain a. Let z denote the according neighbor of c, for an illustration see Figure 33b.

We study the second inhibitor $I_{ac} = acst$. By Observation 27, we have $\{s, t\} \cap \{b, d\} = \emptyset$. Moreover, we observe that $t \neq q$ since the contrary would imply that abq is a separating *triangle* that is not vertex-disjoint from abc; a contradiction to Lemma 24 and the assumption that T does not define a double kite. Due to the degree bound of c, we have $s \in \{x, z\}$. First, assume $s = z$. Since a and z are on distinct sides of I_{bc}, it follows that $t \in I_{bc}$. Further, by Lemma 28 applied to I_{ac}, we have $t = q$; a contradiction (recall that we have already established that $t \neq q$). Therefore, $s = x$. Since $t \neq q$, it follows that t belongs to K. Moreover, by planarity, t belongs to the side of I_{bc} that contains a. Since $t \neq d$, it follows that t is distinct from all previously considered vertices. Note that this brings the degree of a up to the maximum of five (recall that a has a neighbor in K'). By 3-connectivity, three vertices of $I_{ac} = acxt$ need a neighbor on the side of I_{ac} that does not contain b. However, this is a contradiction to the degree bound of c or a. ◁

Figure 33: Illustrations for Lemma 32.

Case 2: bc is not constrained by a 4-inhibitor of $\mathcal{I}$. Hence, by assumption, collapsing bc creates an adjacent separation pair sr. By Observation 31 we may assume without loss of generality that s is the new vertex created by collapsing bc and that $r \notin \{a, d\}$. Thus, b, c, r is a separating triple in $\mathcal{G}$ such that $br \in \mathrm{E}(\mathcal{G})$ or $cr \in \mathrm{E}(\mathcal{G})$. By planarity and 3-connectivity, the graph $\mathcal{G} \setminus \{b, c, r\}$ has two connected components $\hat{K}, \hat{K}'$ and each of the vertices b, c, r is adjacent to at least one vertex in each of these components. Assume w.l.o.g. that $a \in \hat{K}$.

Our plan is to show that $\mathcal{G}$ is isomorphic to the graph $\mathcal{G}'$ depicted in Figure 34, from which we can easily derive the desired contradiction. To show that $\mathcal{G} \simeq \mathcal{G}'$, we proceed in four steps. First, we show that $r \neq q$. From this we conclude that c and q are adjacent. We then prove that the neighborhood of a is $\mathrm{N}(a) = \{b, c, d, b'\}$. Finally, we study the neighborhoods of c and q. Let us carry out this plan.

The identity of r. We start by showing that $r \neq q$. Assume otherwise. By 3-connectivity and since a, b, q is a separating triple of $\mathcal{G}$, there exists a path from b' to a that does not visit b or q $(= r)$ and whose inner vertices belong to K'. This implies that the vertices $\mathrm{N}_{\mathcal{G}}(b) \setminus \{b, c, r\} = \{a, b', d\}$ are connected in $\mathcal{G} \setminus \{b, c, r\}$, which yields a contradiction to Observation 30. Thus, $r \neq q$ as claimed.

Towards a contradiction, assume that $cr \in \mathrm{E}(\mathcal{G})$, in which case r is some vertex of $K \setminus \{c, d\}$. By 3-connectivity and since a, b, q is a separating triple of $\mathcal{G}$, there exists a path from q to a via b' whose internal vertices belong to K'. Since $c, r \in K$, they do not belong to this path and, hence, we have that $\mathrm{N}_{\mathcal{G}}(b) \setminus \{b, c, r\} = \{a, b', d, q\}$ are connected in $\mathcal{G} \setminus \{b, c, r\}$), which yields a contradiction to Observation 30. Hence, $cr \notin \mathrm{E}(\mathcal{G})$.

It follows that $br \in \mathrm{E}(\mathcal{G})$. Moreover, due to the degree constraint of b and since $r \neq q$, we have $r = b'$.

The existence of the edge cq. Assume for the sake of contradiction that c and q are not adjacent. Since, by Observation 30, the vertices $\mathrm{N}_{\mathcal{G}}(b) \setminus \{b, c, r\} = \{a, d, q\}$ are not connected in $\mathcal{G} \setminus \{b, c, r\}$) and due to the edge ad, it follows that $q \in \hat{K}'$ and $a \in \hat{K}$. Hence, by 3-connectivity, it follows that there exist three simple internally vertex-disjoint paths P_b, $P_{b'}$, P_c from q to a where b, b', c is an interior vertex of P_b, P_b', P_c, respectively.

By assumption, $cq \notin \mathrm{E}(\mathcal{G})$. Hence, there exists a vertex c' that belongs to the subpath of P_c between q and c. Since P_c is simple and does not pass through b, it follows that $c' \in K \setminus \{c, d\}$. By 3-connectivity, c, q is not a separation pair and, thus, there must be some simple path $P_{c'd}$ from c' to d that does not pass through any of c or q. This path must visit a or b. Accordingly, we distinguish two cases.

First, assume that $P_{c'd}$ visits a before potentially also visiting b, that is, if b is visited, then it belongs to the subpath between a and d. Then the subpath of P_c from q to c' together with the subpath of $P_{c'd}$ from c' to a forms a path between a

and q that visits neither of b, c, and b' $(= r)$; a contradiction to the fact that $a \in \hat{K}$ and $q \in \hat{K}'$.

On the other hand, if $P_{c'd}$ visits b before potentially visiting a, then it has to visit b' before b due to the degree bound of b and since $P_{c'd}$ does not pass through any of c or q. However, this implies that $P_{c'd}$ contains a subpath that connects $b' \in K'$ and $c' \in K$ without passing through any of a, b, q, which yields the desired contradiction. Altogether, this shows that c and q are adjacent, as claimed.

The neighborhood of a. Our goal is to show that $\mathrm{N}(a) = \{b, c, d, b'\}$. First, observe that $aq \notin \mathrm{E}(\mathcal{G})$, as otherwise abq would be a separating triangle that is non vertex-disjoint from abc; a contradiction to Lemma 24 and the assumption that T does not define a double kite.

Next, we will show that $\mathrm{N}(a) \cap K = \{c, d\}$. Assume otherwise, and let $z \in K \setminus \{c, d\}$ such that a and z are adjacent. By planarity z is contained on the side of bcq that contains a. By planarity and since a, b, q is separating, this implies that a, c, q is separating z from b, d, and all vertices in K'. By 3-connectivity, this implies that there exist three path from z to a, c, and q, respectively, that are vertex-disjoint except for the common endpoint z, and that do not pass through b, d or any vertex of K'. However, the path to a together with the path to q forms a path between a and q that does not visit any of b, c, and b' $(= r)$; a contradiction to the fact that $a \in \hat{K}$ and $q \in \hat{K}'$.

Finally, we show that $\mathrm{N}(a) \cap K' = b'$. Assume otherwise, and let $a' \in \mathrm{N}(a) \cap K'$ such that $a' \neq b'$. By 3-connectivity and since a, b, q is a separating triple, there exist three paths from a' to a, b, q, respectively, that are pairwise vertex-disjoint except for the common endpoint a'. The path to b has to contain b', as b' is its only neighbor in K'. This implies that in $\mathcal{G} \setminus \{b, c, b'\}$ the vertices a and q belong to the same connected component, as they are connected via the paths that connect a' with a and with q. Once again, this yields a contradiction to the fact that $a \in \hat{K}$ and $q \in \hat{K}'$. Altogether, we obtain $\mathrm{N}(a) = \{b, c, d, b'\}$, as claimed.

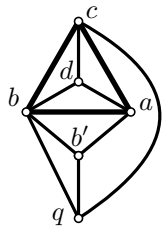

Figure 34: The graph $\mathcal{G}'$ in Case 2 of Lemma 32.

The neighborhoods of c and q. Note that the fact that $\mathrm{N}(a) = \{b, c, d, b'\}$ implies that that $K' = \{b'\}$, as otherwise q, b' would be a separation pair in $\mathcal{G}$. Moreover,

the side of the triangle bcq that does not contain a is empty by Lemma 24 and the assumption that T does not define a double kite.

Towards a contradiction, assume that there exists some vertex t on the side of the cycle $cab'q$ that does not contain b. By 3-connectivity, there exist three paths from t to $cab'q$ that are pairwise vertex-disjoint except for the common endpoint t. Since $\mathrm{N}(a) = \{b, c, d, b'\}$, none of these paths ends at a, and so the three endpoints on $cab'q$ are c, q, b'. However, this implies that $b' \in K'$ and $c \in K$ are connected in $\mathcal{G} \setminus \{a, b, q\}$; a contradiction. It follows that the side of the cycle $cab'q$ that does not contain b does not contain any vertex. It may also not contain the chord cb' since, again, this would imply that $b' \in K'$ and $c \in K$ are connected in $\mathcal{G} \setminus \{a, b, q\}$. Altogether, this shows that $\mathcal{G}$ is isomorphic to $\mathcal{G}'$, as claimed.

Wrapping up. Now that we have shown that $\mathcal{G} \simeq \mathcal{G}'$, it is easy to derive the desired contradiction: By our assumption, the edge ac violates property (i) or (ii).

First, assume the former is true, and let $I_{ac} = acst$ be a 4-inhibitor of $\mathcal{I}$ constraining ac. By Observation 27, we have $I_{ac} = acqb'$. However, this cycle is nonseparating; a contradiction.

Therefore ac violates property (ii), i.e., collapsing ac creates an adjacent separation pair. However, collapsing ac creates a graph isomorphic to K_4, which does not contain any separation pair. Thus, we obtain the desired contraction. ◁

Case 3: bc is constrained by a 4-inhibitor $I_{bc} = bcxy$ of $\mathcal{I}$ and ac is not constrained by a 4-inhibitor of $\mathcal{I}$. For illustrations, refer to Figure 35a. By assumption, collapsing ac creates an adjacent separation pair vw. By Observation 31, we may assume without loss of generality that v is the new vertex created by collapsing ac and that $w \notin \{b, d\}$. Thus, there exists a separating triple a, c, w in $\mathcal{G}$ such that $aw \in \mathrm{E}(\mathcal{G})$ or $cw \in \mathrm{E}(\mathcal{G})$. By planarity and 3-connectivity, the graph $\mathcal{G} \setminus \{a, c, w\}$ has exactly two connected components $\tilde{K}, \tilde{K}'$ and each of the vertices a, c, w is adjacent to at least one vertex in each of these components. Assume w.l.o.g. that $b \in \tilde{K}$.

We proceed in four steps. First, we show that a and w cannot be adjacent. It follows that $cw \in \mathrm{E}(\mathcal{G})$. We then show that this implies that $w \in K \setminus \{c, d\}$. From this, we conclude that $I_{bc} = bcwq$, as shown in Figure 35b. Finally, we use the fact that I_{bc} is separating and the degree bound of c to show that c cannot actually have a neighbor in $\tilde{K}'$, which yields the desired contradiction. Let us proceed by carrying out this plan.

Vertices a and w are nonadjacent. Towards a contradiction, assume $aw \in \mathrm{E}(\mathcal{G})$. Then $w \neq q$, as otherwise a, b, q would be a separating triangle that is not vertex-disjoint from abc; a contradiction to Lemma 24 and the assumption that T does not define a double kite. We claim that $w \in K'$. Assume otherwise, that is, $w \in K \setminus \{c, d\}$. Let a' denote the (unique) neighbor of a in K'. By 3-connectivity and since the triple a, b, q separates K from K', there is a path between a' and b whose interior

vertices belong to K'. This path cannot pass through $c, w \in K$. It follows that the vertices $\mathrm{N}_{\mathcal{G}}(a) \setminus \{a, c, w\} = \{a', b, d\}$ are connected in $\mathcal{G} \setminus \{a, c, w\}$, which yields a contradiction to Observation 30. Therefore $w \in K'$ as claimed.

If a has no neighbor distinct from b, c, d, w, then, once again, $\mathrm{N}_{\mathcal{G}}(a) \setminus \{a, c, w\} = \{b, d\} \subseteq \tilde{K}$ and we obtain a contradiction to Observation 30. So let $a' \in \mathrm{N}_{\mathcal{G}}(a) \setminus \{b, c, d, w\}$. If $a' \in K'$, then by 3-connectivity there exist three paths from a' to the vertices a, b, q that are pairwise vertex-disjoint except for the common endpoint a' and whose interior vertices belong to K'. Only one of these paths may pass through w and, thus, at least one of the two paths that connect a' to b and to q is still present in $\mathcal{G} \setminus \{a, c, w\}$. Due to the edge bq, this implies that, once again, $(\mathrm{N}_{\mathcal{G}}(a) \setminus \{a, c, w\}) \subseteq \tilde{K}$, which yields a contradiction to Observation 30.

It remains to consider the case $a' \in K$. By 3-connectivity, there exist three simple paths from a' to the vertices a, b, q that are pairwise vertex-disjoint except for the common endpoint a' and whose interior vertices belong to K. Deleting the vertices a and c may only destroy two of these paths, one of which is the one to a. Therefore the path to b or the path to q still exists. Once again, this implies that $(\mathrm{N}_{\mathcal{G}}(a) \setminus \{a, c, w\}) \subseteq \tilde{K}$, which yields a contradiction to Observation 30. Altogether, this shows that $aw \notin \mathrm{E}(\mathcal{G})$.

Figure 35: Illustrations for Lemma 32.

The identity of w. The discussion in the previous paragraph shows that $cw \in \mathrm{E}(\mathcal{G})$ (since vw is an adjacent separation pair). Let us study the identity of w.

First, assume $w = q$. Recall that bc is constrained by a 4-inhibitor $I_{bc} = bcxy$. Due to the degree bound of b and by Observation 27, we have $y \in \{q, b'\}$. If $y = b'$, then $x \in \{a, q\}$ since a, b, q is a separating triple. If $x = a$, we obtain a contradiction to Observation 27. On the other hand, if $x = q$, we obtain a contradiction to Lemma 28 due the edge bq. It follows that $y = q$. In this case, the edge cq is a chord of I_{bc}, again contradicting Lemma 28. Altogether, this shows that $w \neq q$.

Since a, b, q is separating and c and w are adjacent, it follows that $w \notin K'$. Consequently, $w \in K \setminus \{c, d\}$.

The identity of I_{bc}. Let us study the identity of $I_{bc} = bcxy$. As in the previous paragraph, we have $y = q$ due to the degree bound of b and by Observation 27 and Lemma 28. Recall that c has a neighbor c' in $\tilde{K}'$. Hence, due the degree bound of c, its neighborhood is $\mathrm{N}_{\mathcal{G}}(c) = \{a, b, d, w, c'\}$. By Observation 27 it follows that $x \in \{w, c'\}$.

We have $x \neq c'$, since otherwise $c' \in \tilde{K}'$ is connected to $b \in \tilde{K}$ in $\mathcal{G} \setminus \{a, c, w\}$ via the path $c'qb$ $(= xyb)$. Consequently, we obtain $x = w$. Therefore, $I_{bc} = bcwq$, for an illustration see Figure 35b.

Wrapping up. Since I_{bc} is separating, there exists a vertex z on the side of $I_{bc} = bcwq$ that does not contain a. By 3-connectivity, there exist three paths from z to I_{bc} that are disjoint except for the common endpoint z. Due to the degree bound of b, it follows that the endpoints of these paths are c, w, q, respectively. We denote that path that connects z with q by P_q. Due to the degree bound of c, we may assume without loss of generality that $z = c'$. However, this implies that $P_q \cup qb$ is a path between $c' \in \tilde{K}'$ and $b \in \tilde{K}$, that does not pass through any of a, c, w. This yields a contradiction to the definition of $\tilde{K}$ and $\tilde{K}'$.

Altogether, this shows that collapsing ac does not create an adjacent separation pair. ◁

We have shown that the statements (1) and (2) cannot simultaneously be violated. By Lemma 29 they also cannot simultaneously be satisfied, which proves the claim. □

4.5.3 Strategy to choose the set of edges to collapse

We are now prepared to discuss how the desired edge set K is obtained. More precisely, we prove the following statement.

Theorem 33. *Let $\mathcal{H}$ be a 3-connected simple plane graph on n vertices where every vertex that belongs to a separating triangle has degree at most five and where every separating triangle is trivial. Further, let $\mathcal{S}$ denote the set of separating triangles in $\mathcal{H}$. Assume that $\mathcal{H}$ is not a trivial double kite and that $\mathcal{H}$ contains no subgraph isomorphic to a $\mathcal{G}_1$ or $\mathcal{G}_2$. Then we can compute in $O(n^2)$ time a set $\mathcal{K} \subset \mathrm{E}(\mathcal{H})$ of edges so that:*

(I1) every edge in $\mathcal{K}$ is collapsible and belongs to a triangle in $\mathcal{S}$; and

(I2) every triangle in $\mathcal{S}$ has exactly one edge of $\mathcal{K}$.

(I3) The graph $\mathcal{H}$ contains no k-inhibitor I with respect to $\mathcal{S}$ such that $\mathcal{K}$ contains more than $k-3$ edges of I. Let $\mathcal{H}''$ denote the graph obtained by simultaneously collapsing the edges of $\mathcal{K}$. If $\mathcal{H}$ contains a k-inhibitor I with respect to $\mathcal{S}$ such that $\mathcal{K}$ contains exactly $k-3$ edges of I, then the triangle corresponding to I in $\mathcal{H}''$ is nonseparating.

(I4) The graph $\mathcal{H}''$ is biconnected and does not contain an adjacent separation pair.

The general idea for the proof of Theorem 33 is to collapse some edge e that satisfies Property (1) of Lemma 32. If the resulting graph $\mathcal{H}'$ is 3-connected, the process may be repeated. On the other hand, if $\mathcal{H}'$ is biconnected only, we plan to recurse on the rigid triconnected components of $\mathcal{H}'$, making use of the fact that each separating triangle still present in $\mathcal{H}'$ appears in one of these components:

Lemma 34. *Let $\mathcal{G}$ be a 3-connected simple plane graph and let $T_1 = abc$ and $T_2 = fgh$ be two vertex-disjoint trivial separating triangles of $\mathcal{G}$. Assume that collapsing the edge $e = ab$ results in a graph $\mathcal{G}'$ that is biconnected, but not 3-connected. Then the graph $\mathcal{G}^-_{T_2}$ appears with all its real edges in a common rigid triconnected component of $\mathcal{G}'$.*

Proof. Let d and d' denote the unique interior vertices of T_1 and T_2, respectively. Since T_1 and T_2 are vertex-disjoint, the graph $\mathcal{G}^-_{T_2}$ is a subgraph of $\mathcal{G}'$. Moreover, since $\mathcal{G}^-_{T_2}$ is 3-connected, there is a rigid triconnected component R of $\mathcal{G}'$ that contains a induced subgraph T_2' isomorphic to $\mathcal{G}^-_{T_2}$ whose vertices are f, g, h, d' (copies of some of these vertices may additionally appear in other triconnected components).

It remains to show that all edges of T_2' are real. Since virtual edges of R correspond to separation pairs of $\mathcal{G}'$, it suffices to show that there is no separation pair p, q of $\mathcal{G}'$ such that $p, q \subset \{f, g, h, d'\}$. So let p, q be a separation pair of G'. By Observation 31, we may assume without loss of generality p is the vertex resulting from the collapse of e. Since T_1 and T_2 are vertex-disjoint, we have $p \notin \{f, g, h, d'\}$, which proves the claim. □

In order to make this recursive strategy work, we actually prove a generalized version of Theorem 33, which is formulated as the upcoming Lemma 38. In preparation, let us discuss the intuition behind its statement: Let $\mathcal{G}$ be a 3-connected simple plane graph in which each edge is labeled as either real or virtual. We refer to a separating triangle as *real* if its three edges are real. Assume that in $\mathcal{G}$ every vertex that belongs to a real separating triangle has degree at most five and that each real separating triangle is trivial and its three interior edges are real. Such a graph $\mathcal{G}$ can arise as a rigid triconnected component of a biconnected simple graph that is the result of iteratively applying some number of edge collapses to a graph $\mathcal{H}$ that satisfies the preconditions of Theorem 33.

Our goal is to find a set $\mathcal{K} \subset \mathrm{E}(\mathcal{G})$ of edges that satisfies relaxed versions of the Properties (I1)–(I4). Property (I4) remains unchanged and in order to maintain it, our plan is to never perform edge collapses that result in an adjacent separation pair. Therefore, no (newly created) virtual edge e will ever correspond to a (parallel) triconnected component with a real edge parallel to e. Hence, virtual edges may be thought of as paths of length at least two and, thus, we may ignore separating triangles that use virtual edges. Similarly, we ignore inhibitors that use virtual edges. To this end, in $\mathcal{G}$ (and, more generally, in graphs that are not 3-connected but otherwise satisfy all properties of $\mathcal{G}$), we refer to an inhibitor with respect to the

set of real separating triangles as *real* if all of its edges are real. Further, we call an edge e of a real separating triangle a *candidate* edge if (i) e is not constrained by a *real* 4-inhibitor; and (ii) collapsing e does not create an adjacent separation pair. Note that collapsing a candidate edge cannot create new real separating triangles due to Property (i), but it may create separating triangles that are not real. Lemma 32 specializes to:

Lemma 35. *Let $\mathcal{G}$ be a 3-connected simple plane graph in which each edge is labeled as either real or virtual. Assume that in $\mathcal{G}$ every vertex that belongs to a real separating triangle has degree at most five and that each real separating triangle is trivial and its three interior edges are real. Further, let T be such a real separating triangle. Finally, assume that $\mathcal{G}$ contains no separating triangle (real or not) that together with T defines a double kite.*

Then either (1) T has at least one candidate edge; or (2) $\mathcal{G}$ contains a subgraph $\mathcal{G}'$ that isomorphic to $\mathcal{G}_2$ or one of the graphs of the family $\mathcal{G}_1$ depicted in Figure 31 such that each thick (colored) edge of $\mathcal{G}'$ belongs to some real 4-inhibitor that constrains an edge of T.

Proof. Follows from Lemma 32 by choosing $\mathcal{I}$ as the set of real 4-inhibitors. □

By the preconditions of Theorem 33, the initial graph $\mathcal{H}$ does not satisfy Property (2) of Lemma 35. Consequently, by Property (1), we find a first candidate edge for inclusion in the desired edge set $\mathcal{K}$. Given that our plan is to iterate this process, we need to ensure that our edge collapses never result in a situation where Property (2) is satisfied. Specifically, we need to avoid creating one of the following:

Definition 36 (Real subgraphs)**.** Let $\mathcal{G}$ be a simple plane graph in which each edge is labeled as either real or virtual. Then,

- $\mathcal{G}$ is a *real* trivial double kite if it is a trivial double kite whose edges are real.

Moreover, we say that

- $\mathcal{G}$ contains a *real* $\mathcal{G}_1$ subgraph if it contains a subgraph isomorphic to a $\mathcal{G}_1$ such that the thick (colored) edges and the three edges incident to d of the $\mathcal{G}_1$ are real; and

- $\mathcal{G}$ contains a *real* $\mathcal{G}_2$ subgraph, if it contains a subgraph isomorphic to $\mathcal{G}_2$ whose edges are real.

Lemma 37. *Let $\mathcal{G}$ be a 3-connected simple plane graph in which each edge is labeled as either real or virtual. Assume that in $\mathcal{G}$ every vertex that belongs to a real separating triangle has degree at most five and that each real separating triangle is trivial and its three interior edges are real. Let $\mathcal{G}'$ denote the graph resulting from collapsing some candidate edge e of a real separating triangle T in $\mathcal{G}$. Then:*

– *$\mathcal{G}'$ is not a real trivial double kite.*

Moreover, under the assumption that in $\mathcal{G}$ no real separating triangle defines a double kite with some other separating triangle (real or not), the following properties hold:

– *if $\mathcal{G}$ contains no real $\mathcal{G}_2$ subgraph, then $\mathcal{G}'$ contains no real $\mathcal{G}_2$ subgraph; and*

– *if $\mathcal{G}$ contains no real $\mathcal{G}_1$ subgraph, but $\mathcal{G}'$ contains a real $\mathcal{G}_1$ subgraph, then there is a candidate edge e' of T such that collapsing e' in $\mathcal{G}$ results in a 3-connected graph that does not contain a real $\mathcal{G}_1$ subgraph.*

Proof. We start with some basic observations and then prove the three conclusions of the lemma individually.

Preliminary considerations. Since e is a candidate edge, it is not constrained by a real 4-inhibitor in $\mathcal{G}$. Consequently, each real separating triangle $T_{\mathcal{G}'}$ ($\neq T$) of $\mathcal{G}'$ corresponds to a real separating triangle $T_{\mathcal{G}}$ ($\neq T$) of $\mathcal{G}$. The triangle $T_{\mathcal{G}}$ is trivial and its three interior edges are real and none of these edges can be e since they do not belong to separating triangles. Consequently, the triangle $T_{\mathcal{G}'}$ is also trivial and its three interior edges are real and correspond to the interior edges of $T_{\mathcal{G}}$.

Avoiding real trivial double kites. Towards a contradiction, assume that $\mathcal{G}'$ is a real trivial double kite. Then, by the preliminary considerations, $\mathcal{G}$ contains two real trivial separating triangles $T_1 \neq T$ and $T_2 \neq T$ corresponding to the triangles of the double kite. If the trivial triangle T shares a vertex with one of these triangles, then, by Observation 25, the graph $\mathcal{G}$ is a trivial double kite (that is not necessarily real). We obtain a contradiction due to the fact that edge collapses always decrease the number of vertices.

So assume that T does not share a vertex with T_1 or T_2. In particular, this implies that these triangles do not share a vertex with e. Hence, since their corresponding triangles in $\mathcal{G}'$ share an edge, the triangles T_1 and T_2 also share an edge. Again, by Observation 25, the graph $\mathcal{G}$ is a trivial double kite (that is not necessarily real) and we obtain a contradiction due to the fact that edge collapses always decrease the number of vertices.

Additional assumption. From now on, assume that $\mathcal{G}$ contains no real separating triangle that defines a double kite with some other separating triangle (real or not). Lemma 24 implies that

(F1) the edge e is not incident to a vertex that belongs to a real separating triangle of $\mathcal{G}$ other than T.

Avoiding real $\mathcal{G}_2$ subgraphs. Towards a contradiction assume that $\mathcal{G}$ contains no real $\mathcal{G}_2$ subgraph, but that $\mathcal{G}'$ contains a real $\mathcal{G}_2$ subgraph whose vertices are denoted by a, b, c, d, x, y, z as in Figure 31a on page 64. One of these vertices has to be the vertex $\hat{e}$ created by collapsing e, since otherwise $\mathcal{G}$ also contains a real $\mathcal{G}_2$ subgraph. Let T' be the real separating triangle of $\mathcal{G}$ that corresponds to abc. By (F1), the edge e and T' are vertex-disjoint. Consequently, $\hat{e} \notin \{a, b, c\}$. Moreover, by the preliminary considerations, we have $\hat{e} \neq d$. Hence, $\hat{e} \in \{x, y, z\}$.

Collapsing e replaces T and its interior vertex with an edge incident to $\hat{e}$ and some other vertex e_3 whose degree is at most three. Since all neighbors of $\hat{e}$ in the $\mathcal{G}_2$ subgraph of $\mathcal{G}'$ have degree at least four, it follows, that e_3 is located in one of the triangular faces of the $\mathcal{G}_2$ subgraph of $\mathcal{G}'$ with $\hat{e}$ on its boundary. However, this implies that the boundary ∂f of this face f is a real separating triangle in $\mathcal{G}'$. By the preliminary considerations, a corresponding real separating triangle exists in $\mathcal{G}$ and we obtain a contradiction to (F1).

Avoiding real $\mathcal{G}_1$ subgraphs. Assume that $\mathcal{G}$ contains no real $\mathcal{G}_1$ subgraph, but that $\mathcal{G}'$ contains a real $\mathcal{G}_1$ subgraph whose vertices are denoted by a, b, c, d, x, y, s as depicted in Figure 36a. One of these vertices has to be the vertex $\hat{e}$ created by collapsing e, since otherwise $\mathcal{G}$ also contains a real $\mathcal{G}_1$ subgraph. Let T' be the real separating triangle of $\mathcal{G}$ that corresponds to abc. By (F1), the edge e and T' are vertex-disjoint. Consequently, $\hat{e} \notin \{a, b, c\}$. Moreover, by the preliminary considerations, we have $\hat{e} \neq d$. Therefore, we may assume without loss of generality that $\hat{e} = y$. Collapsing e replaces the triangle T and its interior vertex with an edge yz $(= \hat{e}z)$, which has one endpoint with degree at most three and another with degree at most five. Since y has degree at least five in $\mathcal{G}'$, the vertex z has degree at most three. Moreover, since x and s have degree at least five in $\mathcal{G}'$, it follows that $z \notin \{x, s\}$. Moreover, the vertex z cannot lie on the side of xys in $\mathcal{G}'$ that does not contain d since this would imply that y has degree hat least six. Consequently, the vertex z is located on the side of $C'_{ab} = abxy$ (see Figure 36a) or $C'_{ca} = cays$ (see Figure 36b), respectively, that does not contain d.

Let $T = vwz$ with interior vertex d' such that v is adjacent to a in $\mathcal{G}$, see Figure 37. The five neighbors of y in $\mathcal{G}'$ are s, x, z, a, u. If z is located on the side of C'_{ab} that does not contain d, then the vertex u is located on the side of C'_{ca} that does not contain d. Otherwise, that is, if z is located on the side of C'_{ca} that does not contain d, the vertex u is located on the side of C'_{ab} that does not contain d. By planarity and due to the degree bounds, the edges of $\mathcal{G}$ that correspond to yx and ys of $\mathcal{G}'$ are both incident to w, and the edges of $\mathcal{G}$ that correspond to ya and yu are both incident to v. The neighborhood of v is $\mathrm{N}_{\mathcal{G}}(v) = \{w, z, d', a, u\}$ where u lies in the side of $C_{ca} = cavws$ that does not contain d if z lies on the side of $C_{ab} = abxwv$ that does not contain d (see Figure 37a); and otherwise u lies on the side of C_{ab} that does not contain d (Figure 37b). We will show that the edge vz is the desired candidate edge e'. Hence, we need to show that vz is not constrained by a real 4-inhibitor,

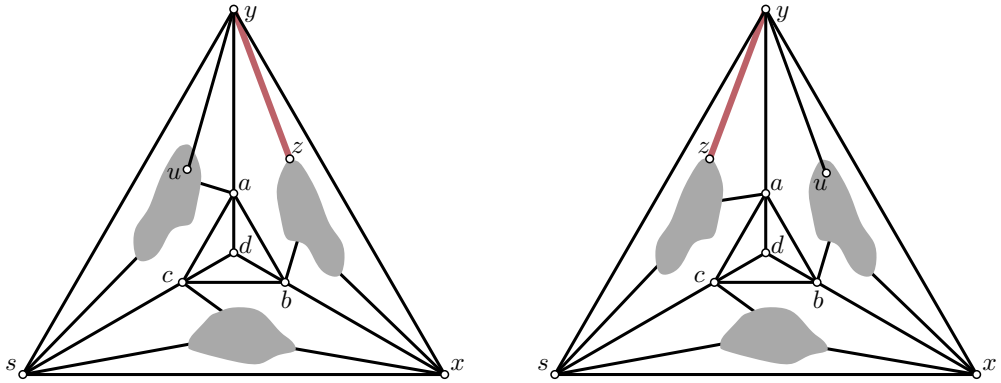

Figure 36: The edge yz corresponds to the collapsed triangle T.

that vz can be collapsed without creating a separation pair, and that collapsing vz does not result in a graph that contains a real $\mathcal{G}_1$ subgraph.

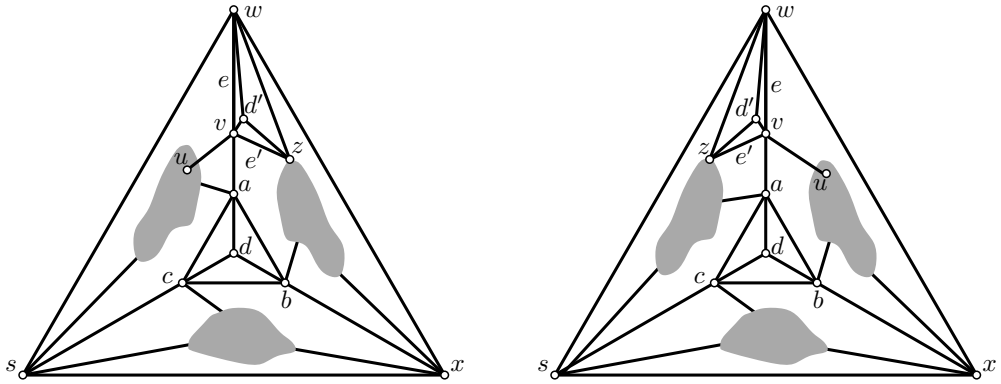

Figure 37: Collapsing the edge $e = vw$ creates a $\mathcal{G}_1$.

First, we show that vz is not constrained by a real 4-inhibitor in $\mathcal{G}$. For the sake of contradiction, assume that $vzij$ is a real 4-inhibitor constraining vz. Let us study the identity of j. By Observation 27, we have $j \in \{a, u\}$. Assume first that $j = u$. By planarity, we have that $i \in \{w, a\}$. However, $i = w$ can be excluded by Observation 27 and $i = a$ contradicts Lemma 28. So assume that $j = a$. We distinguish two cases regarding the location of z. If z belongs to the side of $C_{ab} = abxwv$ that does not contain d (Figure 37a), then, by planarity and the degree bound of a, we have $i = b$. However, since the inhibitor $vzij$ is separating and $\mathcal{G}$ is 3-connected, this implies a contradiction to the degree bound of v or a. On the other hand, if z belongs to the side of $C_{ca} = cavws$ that does not contain d (Figure 37b), then, by planarity and the degree bound of a, we have that $i = c$ or i is the unique neighbor of a on

the side of $C_{ca} = cavws$ that does not contain d. The former case contradicts the degree bound of c. The latter case contradicts the degree bound of v or a due to the 3-connectivity of $\mathcal{G}$ and since $vzij$ is separating. So overall, it follows that vz is not constrained by a real 4-inhibitor.

Next, assume for the sake of contradiction that collapsing vz creates a separation pair p, q. By Observation 31, we may assume without loss of generality that p is the new vertex resulting from collapsing vz in $\mathcal{G}$ and that $q \notin \{w, d'\}$. We will show that the vertices $\mathrm{N}_{\mathcal{G}}(v) \setminus \{v, z, q\} = \{a, u, w, d'\} \setminus \{q\}$ are connected in $\mathcal{G} \setminus \{v, z, q\}$, which yields a contradiction to Observation 30. There exist two interior vertex-disjoint paths between a and w in $\mathcal{G} \setminus \{v, z\}$, namely, $acsw$ and $abxw$. Consequently, the vertices $\{a, w, d'\} \setminus \{q\}$ are connected in $\mathcal{G} \setminus \{v, z, q\}$. It remains to show that u also remains connected to these vertices (if $u \neq q$). If u lies on the side of C_{ca} that does not contain d (Figure 37a), there exist three paths from u to v, a, s, respectively, that are pairwise vertex-disjoint except for the common endpoint u. Given that z is located on the side of C_{ab} that does not contain d, it follows that at least one of the paths to a or to s, respectively, remains in $\mathcal{G} \setminus \{v, z, q\}$. If the path to a remains, there is nothing to show. On the other hand, if the path to s remains, then u is also connected to $\{a, w, d'\} \setminus \{q\}$ due to the edge sw. Similarly, if u lies on the side of C_{ab} that does not contain d (Figure 37b), there exist three paths from u to v, b, x, respectively, that are pairwise vertex-disjoint except for the common endpoint u. Given that z is located on the side of C_{ca} that does not contain d, it follows that at least one of the paths to b or to x, respectively, remains in $\mathcal{G} \setminus \{v, z, q\}$. If the path to x remains, then the edge xw implies the claim. On the other hand, if the path to b remains, the two interior vertex-disjoint paths $bcsw$ and bxw imply the claim. Altogether this shows that collapsing vz does not create a separation pair.

Finally, assume towards a contradiction that collapsing vz in $\mathcal{G}$ results in a graph $\tilde{\mathcal{G}}$ that contains a real $\mathcal{G}_1$ subgraph. We denote the vertices of this subgraph by $\tilde{a}, \tilde{b}, \tilde{c}, \tilde{d}, \tilde{x}, \tilde{y}, \tilde{s}$ where $\tilde{a}$ corresponds to its pendant a in Figure 36a and so on. As above, we may assume without loss of generality, that $\tilde{y}$ is the vertex generated be collapsing vz and we use $\tilde{v}, \tilde{z}, \tilde{w}$ to denote the vertices of $\mathcal{G}$ that correspond to their pendants v, z, w in Figure 37a and 37b, i.e., we have $\{\tilde{v}, \tilde{w}\} = \{v, z\}$ such that $\tilde{w}$ is the vertex incident to the two edges of $\mathcal{G}$ corresponding to the edges $\tilde{y}\tilde{x}$ and $\tilde{y}\tilde{s}$ of $\tilde{\mathcal{G}}$, $\tilde{v}$ is the vertex adjacent to $\tilde{a}$, and $\tilde{z} = w$.

Note that due to the degree bounds, none of $\tilde{x}$ and $\tilde{s}$, which are adjacent to $\tilde{w}$, can belong to a real separating triangle in $\mathcal{G}$. Consequently, we have $\tilde{w} \neq v$, as $\tilde{w} = v$ would imply that $a \in \{\tilde{x}, \tilde{s}\}$ due to the degree bound of v. Therefore $\tilde{w} = z$ and $\tilde{v} = v$. Let $C = C_{ab}$ if z is located on the side of C_{ab} that does not contain d; and otherwise let $C = C_{ca}$. By planarity, the vertices $\tilde{x}$ and $\tilde{s}$ (which are adjacent to $\tilde{w}$) are located on C or on the side of C that contains z ($= \tilde{w}$). Note that not both $\tilde{x}$ and $\tilde{s}$ can be located on C simultaneously, since C has only one vertex (namely x or s) that does not belong to a real separating triangle. Hence, at least one of $\tilde{x}$ and $\tilde{s}$ is located on the side of C that contains z. Since the edges $\tilde{v}\tilde{a}, \tilde{x}\tilde{b}, \tilde{s}\tilde{c}$ connect $\tilde{v}, \tilde{x}, \tilde{s}$

to a common real separating triangle $\tilde{a}\tilde{b}\tilde{c}$, it follows that $\{\tilde{a}, \tilde{b}, \tilde{c}\} = \{a, b, c\}$ by planarity and due to the degree bound of $\tilde{v} = v$. Regardless of the position of z, the edges $\tilde{v}\tilde{a}, \tilde{x}\tilde{b}, \tilde{s}\tilde{c}$ cannot be drawn without violating planarity or the degree bounds: if $C = C_{ab}$, then none of $\tilde{v}, \tilde{x}, \tilde{s}$ can be adjacent to $c \in \{\tilde{a}, \tilde{b}, \tilde{c}\}$ due to the definition of $\mathcal{G}_1$. Similarly, if $C = C_{ca}$, then none of $\tilde{v}, \tilde{x}, \tilde{s}$ that can be adjacent to $b \in \{\tilde{a}, \tilde{b}, \tilde{c}\}$ due to the definition of $\mathcal{G}_1$. Overall, it follows that collapsing vz cannot create a real $\mathcal{G}_1$ subgraph and, hence, $e' = vz$ is an edge as in the claim. □

Let us now formally state Lemma 38, which is the generalized version of Theorem 33. Indeed, if all edges of $\mathcal{G}$ are real, the statement of Lemma 38 matches that of Theorem 33.

Lemma 38. *Let $\mathcal{G}$ be a 3-connected simple plane graph on n vertices in which each edge is labeled as either real or virtual. Further let $\mathcal{S}$ denote the set of real separating triangles in $\mathcal{G}$ and assume that the following properties are satisfied.*

(G1) each triangle of $\mathcal{S}$ is trivial and its three interior edges are real;

(G2) each vertex that belongs to a triangle of $\mathcal{S}$ has degree at most five;

(G3) there is no virtual edge incident to vertices of two distinct triangles of $\mathcal{S}$; and

(G4) $\mathcal{G}$ is not a real trivial double kite and it does not contain a real $\mathcal{G}_1$ subgraph or a real $\mathcal{G}_2$ subgraph, cf. Definition 36.

Then we can compute in $O(n^2)$ time a set $\mathcal{K} \subset \mathrm{E}(\mathcal{G})$ of edges so that

(J1) every edge in $\mathcal{K}$ is collapsible and belongs to a triangle in $\mathcal{S}$; and

(J2) every triangle in $\mathcal{S}$ has exactly one edge of $\mathcal{K}$.

(J3) The graph $\mathcal{G}$ contains no k-inhibitor with respect to $\mathcal{S}$ with more than $k-3$ edges of $\mathcal{K}$. Let $\mathcal{G}''$ denote the graph obtained by simultaneously collapsing the edges of $\mathcal{K}$. If $\mathcal{G}$ contains a real k-inhibitor I with exactly $k-3$ edges of $\mathcal{K}$, then the triangle corresponding to I in $\mathcal{G}''$ is nonseparating.

(J4) The graph $\mathcal{G}''$ is biconnected and does not contain an adjacent separation pair.

Some remarks regarding the upcoming proof of Lemma 38, which stretches over the pages 82–130. We collapse edges of $\mathcal{G}$ resulting in simple graphs $\mathcal{G}'$ (after collapsing a single edge of $\mathcal{K}$) and $\mathcal{G}''$ (after collapsing all edges in $\mathcal{K}$) that are possibly biconnected only. As discussed earlier, it is useful to think of $\mathcal{G}$ as a rigid triconnected component of a graph resulting from collapsing some edges in a graph such as $\mathcal{H}$ from the statement of Theorem 33. However, we emphasize that our algorithm and our terminology is actually oblivious to this notion, that is, the endpoints of a virtual edge in $\mathcal{G}$ are *not* considered to be a separation pair. Similarly, the edge-labeling

(real and virtual) of the edges in $\mathcal{G}'$ and $\mathcal{G}''$ is ignored when it comes to determining the separation pairs in these graphs.

In particular, an edge collapse that merges a virtual edge with a real edge corresponds (conceptually) to the creation of an adjacent separation pair in the graph $\mathcal{H}$. However, this is not a violation of Property (J4), which is only concerned with separation pairs in the graph $\mathcal{G}''$ that results from collapsing some edges in the 3-connected graph $\mathcal{G}$. Note that when Lemma 38 is applied to an original graph $\mathcal{H}$ that only has real edges (as is the case in the statement of Theorem 33), Property (J4) ensures that this special type of edge collapse cannot actually occur during one of the recursive calls. For the sake of proving the claim in its full generality, when merging a virtual edge with a real edge, we refer to the resulting singleton edge as real. This is justified since real edges are strictly more restrictive than virtual edges when it comes to the Properties (G1), (G2), (G4), and (J1)–(J4). As an exception, virtual edges are more restrictive, when it comes to Property (G3). However, when the collapse of a candidate edge results in the merge of a real with a virtual edge, the resulting singleton edge is incident to the new vertex created by the collapse, and this vertex does not belong to a real separating triangle since the real separating triangles of our graphs are pairwise vertex-disjoint (by (G4) and Observation 25) and since the collapse of a candidate edge does not create new real separating triangles. Consequently, Property (G3) is maintained regardless of the label that is assigned to the new edge.

Occasionally, it will be important to be able to distinguish between virtual edges of $\mathcal{G}$ and virtual edges created during the decomposition of, say $\mathcal{G}''$, into its triconnected components. In such cases, we refer to the latter as $\mathcal{G}''$-*virtual* edges. Note that the endpoints of a $\mathcal{G}''$-virtual edge uv that appears in some rigid triconnected component R of $\mathcal{G}''$ correspond to a separation pair in $\mathcal{G}''$, but they do not form a separation pair in R.

To prove Lemma 38, we proceed recursively. More specifically, we will identify a collapsible candidate edge e, collapse it, and then recursively determine a set of edges $\mathcal{K}'$ that satisfies the Properties (J1)–(J4) for the reduced graph. This is sufficient to establish the Properties (J1)–(J4) of the set $\mathcal{K} = \{e\} \cup \mathcal{K}'$ for $\mathcal{G}$:

Lemma 39. *Let $\mathcal{G}$ be a graph as in the statement of Lemma 38. Assume that e is a collapsible candidate edge of $\mathcal{G}$. Finally, let $\mathcal{G}'$ denote the graph obtained from collapsing e in $\mathcal{G}$ and assume that there exists a set of edges $\mathcal{K}' \subset \mathrm{E}(\mathcal{G}')$ that satisfies the Properties (J1)–(J4) for $\mathcal{G}'$. Then the set $\mathcal{K} = \{e\} \cup \mathcal{K}'$ satisfies (J1)–(J4) for $\mathcal{G}$.*

Proof. Let $\mathcal{S}$ be the set of real separating triangles of $\mathcal{G}$. The triangles of $\mathcal{S}$ are pairwise vertex-disjoint by Property (G4) and Observation 25. Hence, there is a unique real separating triangle $T = abc$ of $\mathcal{S}$ with interior vertex d that uses the candidate edge $e = ab$.

Since e is a candidate edge, its collapse does not create a new real separating triangle. Hence, the set of real separating triangles of $\mathcal{G}'$ is $\mathcal{S} \setminus \{T\}$. Moreover, since

the triangles of $\mathcal{S}$ are pairwise vertex-disjoint, none of the triangles $\mathcal{S} \setminus \{T\}$ passes through c or the vertex created by the collapse of e in $\mathcal{G}'$. Thus, the Properties (J1) and (J2) for $\mathcal{K}'$ imply (J1) and (J2) for $\mathcal{K}$.

Next, we show that $\mathcal{K}$ satisfies (J3). Let I denote a k-inhibitor of $\mathcal{G}$ with respect to $\mathcal{S}$. We consider two cases regarding whether or not e belongs to I. First, assume that $e \in I$. We consider two subcases. For the first subcase, assume that $k \in \{5, 6\}$. Consequently, I becomes a $(k-1)$-inhibitor with respect to $\mathcal{S} \setminus \{T\}$ in $\mathcal{G}'$. In this case, the Property (J3) of $\mathcal{K}'$ implies Property (J3) of $\mathcal{K}$ for the inhibitor I. For the second subcase, assume that $k = 4$ (in which case I is not real since e is a candidate edge). Hence, I becomes a separating triangle T' in $\mathcal{G}'$, which is not real. Since, by (J1), each edge of $\mathcal{K}'$ is collapsible in $\mathcal{G}'$, it follows that no edge of $\mathcal{K}'$ belongs to T' and, hence, (J3) of $\mathcal{K}$ is established for the inhibitor I.

It remains to consider the case that $e \notin I$. Consequently, I corresponds to a k-cycle I' in $\mathcal{G}'$. If I' is hyperseparating with respect to $\mathcal{S} \setminus \{T\}$, then I' is a k-inhibitor with respect to $\mathcal{S} \setminus \{T\}$ and the Property (J3) of $\mathcal{K}'$ implies Property (J3) of $\mathcal{K}$ for the inhibitor I. So assume that I' is not hyperseparating with respect to $\mathcal{S} \setminus \{T\}$. Since I is hyperseparating with respect to $\mathcal{S}$, both of its sides contain at least one vertex that is not the interior vertex of a real separating triangle of $\mathcal{G}$. The collapse operation decreases the number of vertices that are not interior to a real separating triangle by exactly one. Since one side of I' contains no such vertex, the corresponding side of I contains either a or b, say a, and the edge bc belongs to I. The edge between c and the vertex created by the collapse of ab and its two neighboring edges along I' cannot belong to $\mathcal{K}'$ by (J1) of $\mathcal{K}'$ since the real separating triangles of $\mathcal{G}$ are pairwise vertex-disjoint. Hence, $\mathcal{K}$ cannot possibly contain *more* than $k-3$ edges of I. If I is real and $\mathcal{K}$ contains exactly $k-3$ edges of I, the triangle corresponding to I in the graph $\mathcal{G}''$ created by simultaneously collapsing all edges of $\mathcal{K}$ in $\mathcal{G}$ cannot be separating since all (possibly none) vertices on one of the sides of I' are interior vertices of real separating triangles and none of these vertices are contained in $\mathcal{G}''$ by (J2) of $\mathcal{K}' \subseteq \mathcal{K}$.

Property (J4) for $\mathcal{K}$ follows directly from (J4) for $\mathcal{K}'$. □

Proof of Lemma 38. We prove the claim by induction on the number $|\mathcal{S}|$ of real separating triangles of $\mathcal{G}$. If $|\mathcal{S}| = 0$ there is nothing to show. So assume that $|\mathcal{S}| > 0$.

4.5.4 Data structures

We begin by describing the data structures that are used to efficiently carry out the algorithm corresponding to the proof. As a first step we determine all real separating triangles (which are trivial, by (G1)). This is easy to accomplish in linear time, given that the vertices of these triangles have degree at most five. Due to the degree bound, it is also easy to determine in quadratic time which edges of triangles in $\mathcal{S}$ are constrained by a real 4-inhibitor. For each constrained edge, we store a list of

the (constant number of) real 4-inhibitors that constrain them. We can update this information under an edge collapse in constant time, by checking all real separating triangles in the constant size neighborhood of the collapsed edge.

To maintain the triconnected components, we use a decremental data structure by Holm et al. [78, Theorem 11] that allows to dynamically maintain an SPQR-tree under a sequence of edge contractions or deletions in $O(n \log^2 n)$ total time. The structure can also be used to answer triconnectivity queries (Are two given vertices in the same triconnected component?) in constant time [78, Theorem 1]. Note that an edge collapse can be implemented using a constant number of edge contractions and deletions.

In order to drive the algorithm, we also need to determine all separation pairs created by a potential edge collapse. To this end, we maintain for every edge $e = uv$ of a real separating triangle $uvw \in \mathcal{S}$ a list $\mathcal{P}(e)$ of vertices—other than w—that together with u, v form a separating triple in the unique rigid triconnected component of the current graph that contains e (cf. Lemma 34). For a single edge, this list can be initialized using the following well-known fact:

Observation 40. *In a 3-connected simple planar graph $\mathcal{Q}$, a triple of vertices u, v, x is separating if and only if the 3-connected simple graph $\mathcal{Q} \cup \{uv, vx, xu\}$ is planar and uvx is separating triangle in $\mathcal{Q}$ (in the unique combinatorial embedding).* □

Hence, it suffices to check whether there exist two distinct faces f_u and f_v and a vertex $x \notin \{u, v, w\}$ of $\mathcal{G}$ such that

- neither of $\partial f_u, \partial f_v$ contains the edge uv;
- $u \in \partial f_u$ and $x \in \partial f_u$;
- $v \in \partial f_v$ and $x \in \partial f_v$; and
- if u, v, x form a triangle, then it is separating.

This test can be performed in linear time, given that there is only a constant number of choices for f_u and f_v. As in a 3-connected graph every pair of faces can share at most two vertices,

(F2) each of the lists $\mathcal{P}(e)$ has constant size.

When an edge xy collapses to a vertex $\widehat{xy}$, we update the structure as follows. We want to find those edges $e = uv$ of some triangle in $\mathcal{S}$, for which $u, v, \widehat{xy}$ is a separating triple. Due to the degree bound, there is a constant number of faces with $\widehat{xy}$ on their boundary. For every pair f_u, f_v of such faces, we traverse the boundary ∂f_u and test for every vertex $u \in \partial f_u$ whether it is incident to an edge such as e whose other endpoint v is on ∂f_v such that $u, v, w = \widehat{xy}, f_u, f_v$ satisfy the above properties. Since the endpoints of edges of real separating triangles have degree at

most five, testing a single endpoint can be done in constant time, and the overall update time per collapse is linear.

As a final step of processing the collapse of an edge e, for every edge h that was in the rigid triconnected component of e before the collapse, we go over the list $\mathcal{P}(h)$ and remove all vertices that are not in the same rigid triconnected component as h after the collapse of e. These decisions can be made using triconnectivity queries in constant time per edge and, therefore, by (F2) in linear time per collapse. We also check for each real 4-inhibitor whether its four real edges belong to a common (rigid) triconnected component R and, if so, whether it is hyperseparating in R. If one of these conditions is not satisfied, we remove the inhibitor from the (up to two) lists that it is stored in.

In each step of the algorithm, we select and collapse a collapsible candidate edge. The edge is selected according to certain priorities explained below (in particular, we prioritize candidate edges whose collapse preserves 3-connectivity). To efficiently find the next edge to collapse, we partition the set of edges that belong to triangles of $\mathcal{S}$ into three lists. The first list contains those edges whose collapse creates an adjacent separation pair (the edges in this list are never considered for inclusion in the set $\mathcal{K}$). Among the edges whose collapse does not create an adjacent separation pair, the second list contains those that are constrained by real 4-inhibitors. The third list contains the remaining edges (which are not constrained by real 4-inhibitors). Both the second and third list are further partitioned into two sublists: the first sublist contains those edges e whose collapse preserves 3-connectivity of the rigid triconnected component that contains e. Note that these edges are easy to identify since their lists $\mathcal{P}(e)$ are empty. The other sublist contains the remaining edges, whose collapse decreases the connectivity. As described above, it is easy recognize when an edge should be moved to another list, which can be done in constant time.

Among the candidate edges whose collapse decreases the connectivity, we will prioritize those which are safe. This property is defined below, see Definition 41. Intuitively, the separation pairs resulting from the collapse of a safe edge cannot become adjacent by performing further collapses. To find the next safe candidate efficiently, we further partition each of the two lists of edges whose collapse decreases connectivity into a sublist of safe edges and a sublist of unsafe edges.

4.5.5 Handling triangles of double kites

We start by taking care of the special case that $\mathcal{S}$ contains a real separating triangle T_1 that together with some other separating triangle T_2 (real or not) defines a double kite D. Note that we can find all such double kites in linear time because T_1 has bounded degree and a vertex of T_2 must be adjacent to all three vertices of T_1. We distinguish two cases regarding T_2.

First, towards a contradiction, assume that there is such a triangle T_2 that is real, i.e., $T_2 \in \mathcal{S}$. Then, by Observation 25, the graph $\mathcal{G}$ is a trivial double kite. The edges

of T_1 and T_2 are real by assumption. The edges interior to T_1 and T_2 are real by Property (G1). The remaining edge of D, which is exterior to both T_1 and T_2, is also real by Property (G3). However, this implies that $\mathcal{G}$ is a real trivial double kite, contradicting Condition (G4).

It remains to consider the case that there is no real separating triangle that defines a double kite with T_1. In particular, the triangle T_2 is not real. Note that by Lemma 24, the triangle T_1 is vertex-disjoint from all other separating triangles in $\mathcal{S}$. Let $T_1 = uvw$ where uv is the edge that also belongs to T_2, for an illustration see Figure 38. The edge uv is not collapsible, but uw and vw are collapsible. In fact, both of these edges are candidate edges; we show this for the edge $e = uw$: Let $\mathcal{G}'$ denote the graph that results from collapsing e. Clearly, the graph $\mathcal{G}'$ is 3-connected, and so collapsing e does not create an adjacent separation pair. So assume that e is constrained by a real 4-inhibitor I_{uw} in $\mathcal{G}$. By Observation 27, we may assume that $I_{uw} = uwxy$ where wx is the edge of D that belongs to the two nonseparating triangles of D. But then, xu is a chord of I_{uw}; a contradiction to Lemma 28. So e is a candidate edge and $\mathcal{G}'$ clearly satisfies the Properties (G1)–(G3).

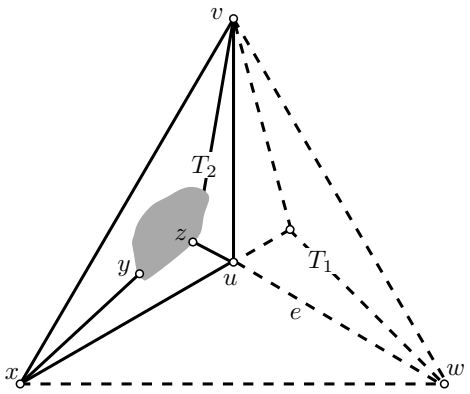

Figure 38: A double kite defined by the real separating triangle T_1 and the separating triangle T_2. The solid edges belong to the subgraph $\hat{\mathcal{G}}$, which is isomorphic to the graph $\mathcal{G}'$ that results from collapsing e.

To see that Condition (G4) also holds, observe that $\mathcal{G}$ contains a subgraph $\hat{\mathcal{G}}$ isomorphic to $\mathcal{G}'$. The only possible difference between $\hat{\mathcal{G}}$ and $\mathcal{G}'$ is that the edge corresponding to ux might be real in $\mathcal{G}'$, but virtual in $\hat{\mathcal{G}}$ (in case ux is virtual and wx is real in $\mathcal{G}$). By Lemma 37, the graph $\mathcal{G}'$ is not a real trivial double kite. Towards a contradiction, assume that $\mathcal{G}'$ contains a real $\mathcal{G}_1$ subgraph or a real $\mathcal{G}_2$ subgraph. By Condition (G4) of $\mathcal{G}$, the graph $\hat{\mathcal{G}}$ does not contain a real $\mathcal{G}_1$ subgraph or a real $\mathcal{G}_2$ subgraph. Consequently, by Definition 36, the edge of $\mathcal{G}'$ corresponding to ux has to correspond to one of the thick (colored) edges or to one of the edges incident to the vertex d of the $\mathcal{G}_1$ or the $\mathcal{G}_2$. In fact, the latter has to be the case given that the degree of u is three in $\hat{\mathcal{G}}$ (by (G2) for $\mathcal{G}$), i.e, u corresponds to the vertex d of the $\mathcal{G}_1$ or $\mathcal{G}_2$ subgraph. However, this implies that the three neighbors of u in $\hat{\mathcal{G}}$, that is, v, x, and

some other vertex z, form a separating triangle in $\hat{\mathcal{G}}$. Hence, by 3-connectivity, the vertex v has some neighbor on the side of xzv that does not contain u. Consequently, in $\mathcal{G}$ the degree of v is at least six, contradicting (G2). Altogether, this shows that Condition (G4) holds for $\mathcal{G}'$.

Thus, all preconditions of the lemma hold and, by induction, we obtain a set $\mathcal{K}'$ of edges satisfying Properties (J1)–(J4) for $\mathcal{G}'$. Thus, Lemma 39, $\mathcal{K} = \mathcal{K}' \cup \{e\}$ is the desired set for $\mathcal{G}$ that satisfies Properties (J1)–(J4).

4.5.6 Handling the remaining triangles

So from now on, assume that no separating triangle of $\mathcal{S}$ defines a double kite together with some other separating triangle (real or not) and, hence, by Lemma 24

(F3) every (real) separating triangle in $\mathcal{S}$ is vertex-disjoint from every other separating triangle (real or not) of $\mathcal{G}$.

Consequently, every candidate edge is collapsible.

Let e be a candidate edge, which exists by Lemma 35 and Condition (G4). If $|\mathcal{S}| = 1$, then $\mathcal{K} = \{e\}$ is the desired set of edges that satisfies the Properties (J1)–(J4). So assume that $|\mathcal{S}| \geq 2$. Our plan is to proceed inductively and, thus, we need to argue that collapsing e results in a graph $\mathcal{G}'$ that satisfies the Preconditions (G1)–(G4). These conditions hold for the input graph $\mathcal{G}$. Since e is a candidate edge, it is not constrained by a real 4-inhibitor. Hence, $\mathcal{G}'$ does not contain any real separating triangle that is not already present in $\mathcal{G}$ and, by Property (G1) of $\mathcal{G}$, we obtain Property (G1) for $\mathcal{G}'$. Collapsing e creates a new vertex x of degree at most five. The degree of the remaining vertices cannot increase. Hence, we obtain Property (G2) for $\mathcal{G}'$. Observe that the endpoints of the edges of $\mathcal{G}$ remain unchanged in $\mathcal{G}'$, except for the edges that are adjacent to e in $\mathcal{G}$ and, hence, incident to x in $\mathcal{G}'$. The vertex x cannot belong to a real separating triangle since the triangles of $\mathcal{S}$ are pairwise vertex-disjoint. Hence, Property (G3) of $\mathcal{G}$ implies Property (G3) of $\mathcal{G}'$. Finally, Property (G4) follows from Lemma 37, possibly after choosing a different candidate edge e. If, additionally, $\mathcal{G}'$ is 3-connected, it satisfies all preconditions of the lemma. Hence, we obtain by induction a set $\mathcal{K}'$ of edges satisfying Properties (J1)–(J4) for $\mathcal{G}'$ and, thus, by Lemma 39, $\mathcal{K} = \mathcal{K}' \cup \{e\}$ is the desired set for $\mathcal{G}$.

It remains to consider the case that

(F4) for *every* candidate edge e, collapsing e results in a graph $\mathcal{G}'$ that is biconnected, but not 3-connected.

Note that collapsing e creates a separation pair, which is nonadjacent since e is a candidate edge. So let $e = ab$ and let $abc \in \mathcal{S}$ be the real separating triangle that e belongs to. Moreover, let d be the singleton inner vertex of abc. Even though $\mathcal{G}'$ is only biconnected, our plan is to proceed inductively, making use of the decomposition of $\mathcal{G}'$ into its triconnected components.

By Lemma 34, every real separating triangle of $\mathcal{G}'$ appears with all its real edges and its three interior real edges in a common rigid triconnected component of $\mathcal{G}'$. Let $\mathcal{R}$ denote the set of rigid triconnected components of $\mathcal{G}'$. Each $R \in \mathcal{R}$ satisfies the preconditions of the lemma: By definition, R is 3-connected. Recall that R is created by performing some number of split operations. Clearly, the split operation cannot produce any new *real* separating triangle in R since the newly created edges are virtual. Moreover, for each vertex v of R the degree of v in R cannot be larger than its degree in $\mathcal{G}'$. Hence, the Properties (G1) and (G2) of $\mathcal{G}'$ imply the Properties (G1) and (G2) for R. Recall that the new vertex x that is created by collapsing e does not belong to a real separating triangle. Moreover, by Observation 31, every newly created virtual edge of R is incident to x. Thus, we obtain Property (G3) for R by Property (G3) of $\mathcal{G}'$. Finally, Property (G4) holds since every subgraph of R that consists exclusively of real edges also appears in $\mathcal{G}'$, for which Property (G4) is satisfied.

For each $R \in \mathcal{R}$ we will obtain a set of edges $\mathcal{K}_R$ to collapse that satisfies Properties (J1)–(J4) for R. In general these sets are obtained via induction, but in some special cases (in which R has constant size) we will determine $\mathcal{K}_R$ explicitly. Let $\mathcal{K}' = \bigcup_{R \in \mathcal{R}} \mathcal{K}_R$ denote the (disjoint) union of all the edge sets $\mathcal{K}_R$ (the sets $\mathcal{K}_R$ are disjoint since they only contain real edges). We show that $\mathcal{K} = \mathcal{K}' \cup \{e\}$ is an edge set that satisfies Properties (J1)–(J4) for $\mathcal{G}$. In fact, by Lemma 39, it suffices to argue that the set $\mathcal{K}'$ satisfies these properties for $\mathcal{G}'$.

Clearly, Properties (J1) and (J2) are satisfied. The main and quite substantial part of the work is to guarantee Property (J4), which we will discuss next. Eventually, once (J4) is established, we can use it to prove that Property (J3) holds as well.

4.5.7 Ensuring Property (J4)

Recall that simultaneously collapsing the edges of $\mathcal{K}'$ refers to the process of collapsing them iteratively, in an arbitrary order. Given that we have not yet established (the first part of) Property (J3), it is a priori not clear that this simultaneous collapse of $\mathcal{K}'$ in $\mathcal{G}'$ is a well defined operation (since some of the edges of $\mathcal{K}'$, which are initially collapsible by (J1), could become noncollapsible after performing some of the collapses). Therefore, we will prove a slightly generalized version of Property (J4): the statement remains exactly the same, but we will relax the definition of the collapse operation, allowing it to be applied to noncollapsible edges. This ensures that the graph $\mathcal{G}''$ from the statement of Property (J4) is well defined.

Collapsing noncollapsible edges. Let $e' = a'b'$ be a noncollapsible candidate edge of a real separating triangle $a'b'c'$. Since e' is noncollapsible, there is a vertex $z' \neq c'$ that forms a separating triangle with a' and b'. The only problem with performing the collapse as in the collapsible case is that the resulting parallel edges between z' and the new vertex created by the collapse form a separating 2-cycle with vertices

on both of its sides. Hence, when merging these parallel edges, the embedding of the resulting edge is not uniquely determined. To allow the collapse of e', we simply choose the embedding of the merged parallel edges arbitrarily.

With this relaxation in place, simultaneously collapsing the edges of $\mathcal{K}'$ in $\mathcal{G}'$ is always a well defined operation, though the resulting plane graph is a priori not necessarily unique. So let $\mathcal{G}''$ denote a graph resulting from simultaneously collapsing all edges of $\mathcal{K}'$ in $\mathcal{G}'$ (later, once Property (J3) is established, it follows that $\mathcal{G}''$ is unique after all).

Overview. To prove the generalized version of Property (J4), we begin by showing that $\mathcal{G}''$ is biconnected in Section 4.5.7.1. It remains to show that $\mathcal{G}''$ does not contain any adjacent separation pair. We prove this statement in two steps. First, in Section 4.5.7.2 we show (after possibly choosing a different candidate edge e, see (F5) and the subsequent discussion) how to choose the edge sets $\mathcal{K}_R$, for $R \in \mathcal{R}$, such that no separation pair that is already present in $\mathcal{G}'$ becomes adjacent in $\mathcal{G}''$. Then, in Section 4.5.7.3 on page 128 we use the Property (J4) of the sets $\mathcal{K}_R$ in order to show that $\mathcal{G}''$ also does not contain any *new* separation pair that is adjacent.

4.5.7.1 Biconnectivity of $\mathcal{G}''$

The set $\mathcal{K}'$ contains no virtual edges by (J1) (we consider *real* separating triangles only). In particular, none of the $\mathcal{G}'$-virtual edges, created by computing the triconnected components of $\mathcal{G}'$, are contained in $\mathcal{K}'$. It follows that every edge of $\mathcal{K}'$ belongs to a single rigid triconnected component exclusively.

We claim that this implies that to obtain the graph $\mathcal{G}''$, instead of collapsing the edges of $\mathcal{K}'$ directly in $\mathcal{G}'$, we may collapse the edges of $\mathcal{K}'$ in the decomposition of $\mathcal{G}'$ into its triconnected components, and then merge all pairs of identical $\mathcal{G}'$-virtual edges. To see this, it suffices to show that for each separation pair p, q of $\mathcal{G}'$ there are two distinct vertices of $\mathcal{G}''$ corresponding to p and q, respectively (in other words, p and q cannot be contracted to a single vertex). By Observation 31, we may assume without loss of generality that p is the vertex created by collapsing e in $\mathcal{G}$. Moreover, by (F3), this vertex does not belong to a real separating triangle of $\mathcal{G}'$. Since $\mathcal{K}'$ contains only edges of real separating triangles (by (J1)), it follows that simultaneously collapsing the edges of $\mathcal{K}'$ does not merge *any* vertex with p. This proves the claim.

Consequently, to argue about the biconnectivity of $\mathcal{G}''$, we may use the fact that each of the triconnected components of $\mathcal{G}'$ remains biconnected when collapsing its edges of $\mathcal{K}'$. This is discussed in the following paragraphs, in which we also analyze what types of triconnected components appear in the decomposition of $\mathcal{G}'$. For illustrations refer to Figure 39.

Triconnected components of $\mathcal{G}'$. Let N be one of the triconnected components of $\mathcal{G}'$. We observe that N cannot be parallel: it suffices to show that no separation

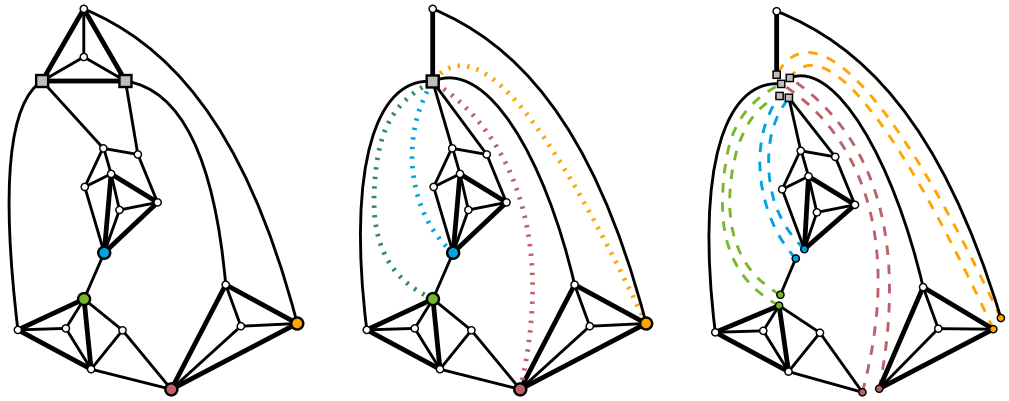

(a) $\mathcal{G}$ (b) $\mathcal{G}'$ (c) Decomposition of $\mathcal{G}'$

Figure 39: The graph in (a) is 3-connected. Each of the large round colored vertices forms a separating triple together with the two square vertices. Collapsing the edge spanned by the square vertices results in the graph in (b), which is biconnected only. Each of its separation pairs involves the vertex created by the collapse. Its decomposition into triconnected components is depicted in (c); dashed edges of the same color correspond to the same virtual edge.

pair in $\mathcal{G}'$ has three or more separation classes. All separation pairs in $\mathcal{G}'$ result from collapsing a single edge $e = ab$ whose endpoints together with some other vertex form a separating triple in the 3-connected planar graph $\mathcal{G}$. Hence, the removal of any separation pair of $\mathcal{G}'$ results in a graph with exactly two connected components. As no separation pair is adjacent, it follows that every separation pair of $\mathcal{G}'$ has exactly two separation classes. Thus, $\mathcal{G}'$ does not have any parallel triconnected component. Therefore, N is series or rigid. We remark that if N is series, it is a 3-cycle. This follows from the facts that all pairs of nonadjacent vertices of a series triconnected component correspond to a separation pair, and that, by Observation 31, every separation pair of $\mathcal{G}'$ involves the vertex created by collapsing e.

Biconnectivity of the collapsed components. Let N' be the graph created by collapsing all edges of $\mathcal{K}' \cap \mathrm{E}(N)$ in N. Then N' is biconnected: this is clear if N is series, in which case $N' = N$ since edge collapses only affect rigid components. If N is rigid, i.e. $N \in \mathcal{R}$, the biconnectivity of N' follows from Property (J4) of $\mathcal{K}_N$.

Biconnectivity of $\mathcal{G}''$. We have shown that each of the triconnected components of $\mathcal{G}'$ remains biconnected when collapsing its edges of $\mathcal{K}'$, which we can now use to prove that $\mathcal{G}''$ is biconnected. To this end, we consider vertices $u, v, w \in \mathrm{V}(\mathcal{G}'')$ with $w \notin \{u, v\}$ and show that u and v are connected in $\mathcal{G}'' \setminus \{w\}$. The collapsed version N' of each triconnected component N of $\mathcal{G}'$ is biconnected. Hence, the graph $N' \setminus \{w\}$ is connected. Therefore, if one of the collapsed triconnected compo-

nents N' contains both u and v, there is nothing to show. So assume that this is not the case and let N'_u and N'_v denote collapsed triconnected components that contain u and v, respectively, and let N_u and N_v denote the triconnected components of $\mathcal{G}'$ that these collapsed graphs originate from. In the decomposition of $\mathcal{G}'$, there is a sequence of triconnected components $N_u = N_1, N_2, \ldots, N_t = N_v$ such that N_i and N_{i+1} share a $\mathcal{G}'$-virtual edge e_i for each $i = 1, 2, \ldots, t-1$. Let $N'_1, N'_2, \ldots, N'_t$ denote the collapsed versions of $N_1, N_2, \ldots, N_t$. Since N'_i is biconnected, the graph $N'_i \setminus \{w\}$ is connected for each $i = 1, 2, \ldots, t$. Since N'_i and N'_{i+1} both contain (an edge corresponding to) e_i, each of the graphs $N'_i \setminus \{w\}$ and $N'_{i+1} \setminus \{w\}$ still contains at least one endpoint of (the edge corresponding to) e_i. Hence, the vertices of $N'_i \setminus \{w\}$ and $N'_{i+1} \setminus \{w\}$ are connected in $\mathcal{G}'' \setminus \{w\}$, and so are the vertices of $N'_1 \setminus \{w\}$ and $N'_t \setminus \{w\}$. Therefore, u and v are connected in $\mathcal{G}'' \setminus \{w\}$.

Altogether, this shows that $\mathcal{G}''$ is biconnected, which proves the first part of the statement in (J4).

4.5.7.2 Separation pairs of $\mathcal{G}'$ remain nonadjacent in $\mathcal{G}''$

Let p, q denote a separation pair of $\mathcal{G}'$. By Observation 31, we may assume without loss of generality that p is the new vertex resulting from collapsing $e = ab$ in $\mathcal{G}$ and that $q \notin \{c, d\}$. Suppose that we have obtained a set $\mathcal{K}_R$ for each $R \in \mathcal{R}$ that satisfies the Properties (J1)–(J4) for R. Assume that simultaneously collapsing all edges in $\mathcal{K}'$ results in an edge between (the vertices corresponding to) p and q. Then there exists some simple path P between p and q in $\mathcal{G}'$ that uses some edges of $\mathcal{K}'$. By (F3), the edges of $\mathcal{K}'$ are pairwise vertex-disjoint. In particular, no two edges of $\mathcal{K}'$ are adjacent along P. Together with the fact that p does not belong to a separating triangle in $\mathcal{G}'$, it follows that the path P has only two edges, exactly one of which belongs to $\mathcal{K}'$, that is, we have $P = psq$, where $sq \in \mathcal{K}'$. The motivates the following definition.

Definition 41. An edge ab of a real separating triangle abc is called *unsafe* if a or b is adjacent to a vertex s of another real separating triangle sqt, and a, b, q is a separating triple. If ab is not unsafe, it is called *safe*.

We may assume that

(F5) *no* candidate edge is safe,

as otherwise we can simply select our candidate e to be safe to ensure that the separation pairs of $\mathcal{G}'$ remain nonadjacent. Algorithmically, whenever we add a new vertex q to one of the lists $\mathcal{P}(e)$, it is easy to determine in constant time whether there is a path asq or bsq that makes e unsafe since the three vertices of such a path have degree at most five. By performing further edge collapses, an unsafe edge may become safe (namely, if for each triangle sqt that is part of a path that makes e unsafe some edge other than sq is collapsed), which is easy to recognize (when performing the collapse). Therefore, the two lists of safe edges whose collapse decreases connectivity

(introduced in Section 4.5.4) are easy to maintain, and, using these lists we can find the next safe candidate edge (if it exists) in constant time.

Often, we will make use of the following weak version of (F5), which follows directly from (F4), (F5), and Observation 31:

(F6) for every candidate edge $e = ab$, collapsing e results in at least one (nonadjacent) separation pair p, q where p is the vertex created by the collapse of e and q belongs to a real separating triangle.

We consider two main cases. In Case 1, we show that if an edge such as sq is constrained by a real 4-inhibitor in $\mathcal{G}$, then for every collection of sets $\mathcal{K}_R$ that satisfy Properties (J1)–(J4), for each $R \in \mathcal{R}$, we have $sq \notin \mathcal{K}' = \bigcup_{R \in \mathcal{R}} \mathcal{K}_R$. In other words, the path P is an obstruction that is already taken into account implicitly when selecting $\mathcal{K}_R$.

In Case 2, the edge sq is not constrained by a real 4-inhibitor. This implies that collapsing sq in $\mathcal{G}$ creates a separation pair (that is not necessarily adjacent) since either sq is a candidate edge, which creates a (nonadjacent) separation pair by assumption (F4), or it is not a candidate edge and its collapse (given that we are not in Case 1) creates an adjacent separation pair. We will show that this is only possible in two very specific local configurations:

- either (in Case 2.2.2.1) the triconnected component R of $\mathcal{G}'$ that contains the separating triangle that uses sq has constant size (in this case we select the set $\mathcal{K}_R$ explicitly, rather than inductively);

- or (in Case 2.2.2.3.2) we can apply a replacement strategy ϱ to obtain a new candidate edge $\varrho(e) \neq e$. We show that ϱ is acyclic, that is, if the Case 2.2.2.3.2 also arises for $\varrho(e)$, we may then consider $\varrho(\varrho(e))$ and so on, until, eventually, we find a candidate edge $\varrho^i(e)$ with $i \in O(n)$ for which we end up in one of the easy Cases 1 or 2.2.2.1.

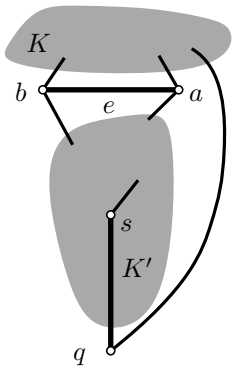

Figure 40: Notation for Section 4.5.7.2 in Lemma 38.

Before we proceed to discuss the two cases, let us introduce some notation; for illustration, refer to Figure 40. The vertices a, b, q form a separating triple in $\mathcal{G}$. The vertex s has to be adjacent to a or b since p and s are adjacent in $\mathcal{G}'$. By planarity and 3-connectivity, the graph $\mathcal{G}\setminus\{a, b, q\}$ has exactly two connected components K, K', and each of the vertices a, b, q is adjacent to at least one vertex in each of the two components. Assume w.l.o.g. that $s \in K'$. Further, let sqt denote the separating triangle that uses sq, and let d' be the unique inner vertex of sqt.

A more detailed overview of Section 4.5.7.2 and its two main cases is illustrated in Figure 41.

Case 1: sq is constrained by a real 4-inhibitor $I_{sq} = sqxy$ in $\mathcal{G}$. We will show that I_{sq} also appears as a 4-cycle of real edges in the triconnected component $R \in \mathcal{R}$ of $\mathcal{G}'$ that contains sqt. Moreover, we will also show that this 4-cycle is hyperseparating (i.e., it is a real 4-inhibitor) in R, or collapsing sq in R creates an *adjacent* separation pair. Thus, any $\mathcal{K}_R$ that satisfies Properties (J3) and (J4) for R cannot contain sq, and so the path P is not an obstruction (that can make the separation pair p, q adjacent).

We start by showing that

(F7) the vertices of I_{sq} belong to $K' \cup \{a, b, q\}$.

Assume otherwise, i.e., some vertex of I_{sq} belongs to K. Since a, b, q are separating, we have $y \in \{a, b\}$, and $x \in K$. We may assume without loss of generality that $y = b$. We will now establish that $x = c$, and that b and q have neighbors on the side of I_{sq} that does not contain a, as illustrated in Figure 42b.

For illustrations of the following steps, refer to Figure 42a. Clearly, we have $x \neq d$ since $dq \notin \mathrm{E}(\mathcal{G})$. Since I_{sq} is separating, there is some vertex z on the side of I_{sq} that does not contain a, d. By 3-connectivity, there are three paths from z to three distinct vertices of $I_{sq} = sqxb$ and these paths are pairwise vertex-disjoint except for the common endpoint z. None of these paths can pass through a given that z and a are on distinct sides of I_{sq}. The paths can also not have b or q as interior vertices since b and q belong to I_{sq}. Hence, it is not possible that one of the three paths leads to $s \in K'$ while another leads to $x \in K$. It follows that two of the paths lead to b and q, respectively. Let b' denote the neighbor of b along its path to z. Hence, $\{a, c, d, s, x, b'\} \subseteq \mathrm{N}_{\mathcal{G}}(b)$. Since the degree of b is bounded by five, two of these vertices have to coincide. The vertices a, c, d, s, b' are pairwise distinct. In particular, $c \neq s$ by (F3), and b' is distinct from a, c, d, s since each of these vertices either belongs to I_{sq} or the side of I_{sq} that does not contain b'. It follows that $x = c$, as claimed.

Recall that sq is an edge of a real separating triangle sqt with unique interior vertex d'. We distinguish two cases regarding the positions of d and t in relation to $I_{sq} = sqcb$. In both cases we show that there exists some candidate edge that can be collapsed without creating a separation pair, thus, obtaining a contradiction to assumption (F4).

Collapsing ab of $abc \in \mathcal{S}$ creates a separation pair p, q. There is a path psq such that sq belongs to a triangle of $\mathcal{S}$. We need to ensure that p, q remain nonadjacent and, hence, that sq is never collapsed.

Case 1: sq is constrained by a real 4-inhibitor I_{sq}.

Implicitly taken care of by (J1)–(J4) for $\mathcal{K}_R$.

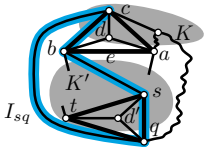

Case 2: Collapsing sq creates a separation pair p_{sq}, q_{sq} (implied by (F4) and the negation of Case 1). Consider a cycle C formed by the path bsq and some path between b and q through K.

Case 2.1: t and a belong to distinct sides of C.

Cannot occur.

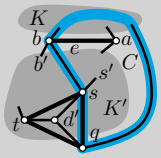

Case 2.2: t and a belong to the same side of C.

Case 2.2.1: st is constrained by a real 4-inhibitor I_{st}.

Cannot occur.

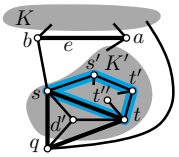

Case 2.2.2: Collapsing st creates a separation pair p_{st}, q_{st} (implied by (F4)).

Case 2.2.2.1: $q_{st} \in K$.

Explicitly choose K_R.

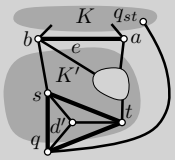

Case 2.2.2.2: $q_{st} = b$.

Cannot occur.

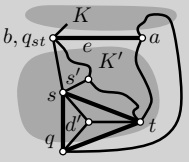

Case 2.2.2.3: $q_{st} \in K' \cup \{a\}$.

Case 2.2.2.3.1: qt is constrained by a real 4-inhibitor I_{qt}.

Cannot occur.

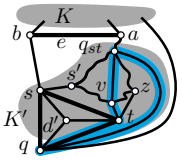

Case 2.2.2.3.2: Collapsing qt creates a separation pair p_{qt}, q_{qt}.

Replace e with $\varrho(e) \in \{st, qt\}$.

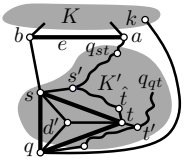

Figure 41: Overview of Section 4.5.7.2 in the proof of Lemma 38.

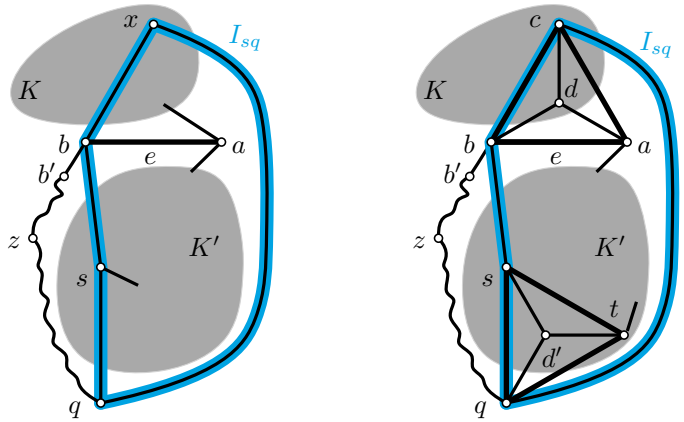

Figure 42: Illustrations for (a) Case 1 and (b) Case 1.1 of Section 4.5.7.2 in Lemma 38.

Case 1.1: d and t are located on the same side of I_{sq}. For an illustration refer to Figure 42b. Note that due to the paths from b and q to z, the degree bounds for b and q are saturated. This implies that

(F8) d is the unique vertex of K on the side of I_{sq} that contains d and t,

since, by 3-connectivity, another vertex in K on this side of I_{sq} would need three disjoint paths via K to a, b, c. However, by planarity, the path to b has to pass through c, and this implies an additional edge incident to q, contradicting its degree bound.

Our plan is to show that ac or st is a candidate edge that can be collapsed without creating a separation pair, which is a contradiction to assumption (F4).

Assume that collapsing ac in $\mathcal{G}$ creates a separation pair p_{ac}, q_{ac}. By Observation 31, we may assume without loss of generality that p_{ac} is the new vertex resulting from collapsing ac in $\mathcal{G}$ and that $q_{ac} \notin \{b, d\}$. We show that the vertices of $\mathrm{N}_{\mathcal{G}}(c) \setminus \{a, c, q_{ac}\}$ are connected in $\mathcal{G} \setminus \{a, c, q_{ac}\}$, which yields a contradiction to Observation 30. The vertices $b, d \in \mathrm{N}_{\mathcal{G}}(c) \setminus \{a, c, q_{ac}\}$ are adjacent. There are two interior vertex-disjoint paths connecting b and q: (1) bsq; and (2) the path via z on the side of I_{sq} that does not contain d and t. Neither of these paths contains a or c as a vertex. Thus, removing a, c, q_{ac} destroys at most one of them. By (F8), there is no neighbor of c in K, other than d, on the side of I_{sq} that contains d and t. Finally, if c has a neighbor $z' \neq q_{ac}$ on the side of I_{sq} that does not contain d and t, then there exist three paths from z' to b, c, q that are pairwise internally vertex-disjoint except for the common endpoint z'. Removing a, c, q_{ac} can destroy either the path to b or the path to q, but not both. Therefore, the neighborhood of c remains connected, as claimed, and we obtain a contradiction to the assumption that ac creates a separation pair. In particular, by (F4) we may assume that ac is not a candidate edge.

Therefore, ac is constrained by a real 4-inhibitor $I_{ac} = acvw$. The vertex v belongs to I_{sq} or to the side of I_{sq} that contains d and t, as otherwise, by Lemma 28 applied to I_{ac} we would obtain that $w = s$ and, thus, an edge between $v \in K$ and $s \in K'$. By Observation 27 and (F8), the only remaining option is $v = q$. As the degree bound for q is saturated with its five neighbors t, d', s, c, and some vertex on the path that leads to z, we conclude that w must be among these neighbors. The vertex c is already used on I_{ac} and d' is nonadjacent to a. Moreover, by planarity, the vertex w cannot be located on the side of I_{sq} that does not contain d and t. This leaves the options $w = t$ and $w = s$. If $w = t$, then I_{ac} cannot be separating: A vertex on the side of I_{ac} that does not contain s is connected to three distinct vertices of I_{ac}. Two of these vertices need to be a and q as otherwise a, b, q would not be separating. This is a contradiction to the degree bound of q. Therefore, $I_{ac} = acqs$, as depicted in Figure 43a.

We claim that in this case the edge st is not constrained by a real 4-inhibitor, and it can be collapsed without creating a separation pair. Note that a, s, q is a separating triple because every path between $c \in K$ and $t \in K'$ must pass through one of a, b, or q, and b and t are on different side of I_{ac}. By 3-connectivity, there is a simple path from t to a that does not pass through s or q. The degree bounds of s and q imply that t and a are adjacent, as otherwise t and a would be separating in $\mathcal{G}$.

Assume that st is constrained by a real 4-inhibitor $I_{st} = stgh$. If $h = a$, the edge at would be a chord of I_{st}, contradicting Lemma 28. Thus, $h = b$ by Observation 27. It remains to study the identity of g. If $g = a$ then sa is a chord of I_{st}, contradicting Lemma 28. We also have $g \neq c$ since otherwise ct is an edge between a vertex of K and a vertex of K'. The final two options, d and b', violate planarity since these vertices belong to the side of I_{ac} that does not contain t. Thus, st is not constrained by a real 4-inhibitor, as claimed.

Therefore, by assumption (F4), collapsing st creates a separation pair p_{st}, q_{st}. By Observation 31, we may assume without loss of generality that p_{st} is the new vertex resulting from collapsing st in $\mathcal{G}$, and that $q_{st} \notin \{q, d'\}$. The degree bounds of a, s, q and 3-connectivity imply that $\mathrm{N}_{\mathcal{G}}(t) = \{s, q, d', a\}$. We show that $\mathrm{N}_{\mathcal{G}}(t) \setminus \{s, t, q_{st}\}$ is connected in $\mathcal{G} \setminus \{s, t, q_{st}\}$. Clearly, q and d' are still connected. Thus, if $q_{st} = a$ there is nothing to show. So assume $q_{st} \neq a$. Consider the following two pairwise internally vertex-disjoint paths in $\mathcal{G}$ that connect a and q: (1) acq; and (2) the path consisting of the edge ab and the path between a and q via z on the side of I_{sq} that does not contain d and t. None of these paths contain s or t and, thus, removing s, t, q_{st} can only destroy one of them. So, $\mathrm{N}_{\mathcal{G}}(t) \setminus \{s, t, q_{st}\}$ is connected in $\mathcal{G} \setminus \{s, t, q_{st}\}$, which yields a contradiction to Observation 30. Therefore, collapsing st does not create any separation pair, as claimed.

Altogether, this yields a contradiction to assumption (F4). ◁

Case 1.2: d and t are located on distinct sides of I_{sq}. The following steps are illustrated in Figure 43b. Note that b, s, q are separating in $\mathcal{G}$. Therefore, by 3-connectivity, the

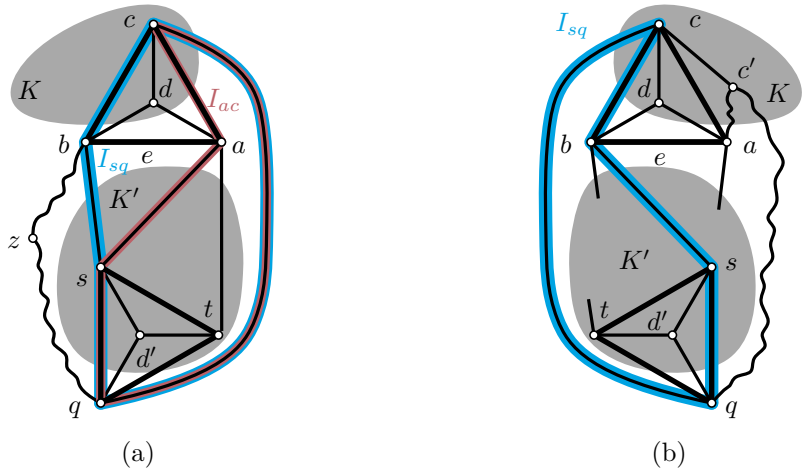

Figure 43: Illustrations for (a) Case 1.1 and (b) Case 1.2 in Lemma 38.

vertex b has some neighbor in K' on the side of I_{sq} that contains t. As a consequence, the degree bound of b, 3-connectivity, and planarity imply that

(F9) the side of I_{sq} that contains t does not contain any vertex of K

(such a vertex would have three disjoint paths to vertices of I_{sq} and none of them may end at s).

As in the previous case, our plan is to show that st is a candidate edge that can be collapsed without creating a separation pair, which is a contradiction to assumption (F4). As a first step, we will show that

(F10) $as \in \mathrm{E}(\mathcal{G})$ and c has a neighbor c' on the side of I_{sq} that contains d such that there exist two internally vertex-disjoint paths from c' to a and q, respectively, whose interior vertices belong to $K \setminus \{c\}$, as depicted in Figure 43b.

By assumption (F4), the edge ac is either not a candidate edge or it is a candidate edge, but its collapse creates a separation pair. Consequently, the edge ac is constrained by a real 4-inhibitor or collapsing ac creates a separation pair. We will show that either of these properties already implies (F10).

So suppose first that ac is constrained by a real 4-inhibitor $I_{ac} = acvw$. We distinguish two cases regarding the identity of v. First, assume that v is some neighbor $c' \in K$ of c (by (F9) on the side of I_{sq} that contains d). We have that $w \neq q$, as $w = q$ would imply that collapsing $e = ab$ creates an adjacent separation pair, contradicting our choice of e. Therefore, we have $w \in K$. However, since I_{ac} is separating and by 3-connectivity, three vertices of I_{ac} have a neighbor on the side of I_{ac} that does not contain t. This contradicts the degree bounds of a or c (recall that a has a neighbor in K'). Altogether, we obtain that v is not a neighbor of c in K. By Observation 27, the only other option is $v = q$.

Let us study the identity of w. By planarity, we have $w \notin \{t, d'\}$. We claim that $w = s$. Towards a contradiction, assume that $w \neq s$, i.e., $w \in \mathrm{N}_{\mathcal{G}}(q) \setminus \{s, t, d', c\}$. By planarity, w belongs to the side of I_{sq} that contains d. Since I_{ac} is separating, there exists a vertex z on the side of I_{ac} that does not contain t. By 3-connectivity, there exists three paths from z to I_{ac} that are pairwise vertex-disjoint except for the common endpoint z. We consider two subcases. First, suppose that $w \in K$. Then one of the three paths leads to a or to q, contradicting the degree bound of a or q, respectively. Secondly, suppose that $w \in K'$. In this case, two of the three paths lead to a and q, respectively (otherwise, there would be a path between $c \in K$ and $s \in K'$ via z that does not pass through any of a, b, q). However, this contradicts the degree bound of q. Altogether, this shows that $w = s$. Hence, $as \in \mathrm{E}(\mathcal{G})$, which establishes the first statement of (F10). Moreover, since I_{ac} is separating, there is a vertex c' on the side of I_{ac} that does not contain t. By 3-connectivity, there exist three paths from c' to I_{ac} that are pairwise vertex-disjoint except for the common endpoint c'. Due to the degree bound of s, these paths lead to c, a, q, respectively. Since the path to c cannot pass through any of a, s, q, we have that c' belongs to K, and without loss of generality we may assume that c' is a neighbor of c. This establishes the second part of (F10) and, hence, shows that (F10) holds if ac is constrained by a real 4-inhibitor.

It remains to prove (F10) for the case that collapsing ac creates a separation pair p_{ac}, q_{ac}. By Observation 31, we may assume without loss of generality that p_{ac} is the new vertex resulting from collapsing ac in $\mathcal{G}$. As b, s, q are separating and by 3-connectivity, the vertex t has three paths to the vertices b, s, q that are pairwise vertex-disjoint except for the common endpoint t. Together with the path bsq, we obtain the existence of two internally vertex-disjoint paths between b and q that do not pass through ac. As a consequence, the neighbors b, d, q (in case $q_{ac} \neq q$) of c remain connected in $\mathcal{G} \setminus \{a, c, q_{ac}\}$. Thus, in order for a, c, q_{ac} to be separating, c needs to have some neighbor $c' \in K$ (by (F9) on the side of I_{sq} that does contain not contain t). By 3-connectivity and since a, c, q are separating, the vertex c' has three paths to a, c, q that are pairwise vertex-disjoint except for the common endpoint c'. This saturates the degree bounds of a, c, q. Note that for example with $q_{ac} = q$ the vertices a, c, q_{ac} are actually separating. The 3-connectivity of $\mathcal{G}$ implies that the side of I_{sq} that does not contain t does not contain any vertex of K', as the degree bound of b or q would be violated due to 3-connectivity. Thus, the neighbor of a in K' is s. This establishes (F10) for the case that collapsing ac creates a separation pair. Moreover, together with the discussion in the previous paragraph, it follows that (F10) holds, regardless of whether ac is a candidate edge (whose collapse creates a separation pair) or not.

We proceed by showing that the edge st can neither be constrained by a real 4-inhibitor nor can its collapse create a separation pair. First, assume towards a contradiction that st is constrained by a 4-inhibitor $I_{st} = stgh$. If $h = a$, we have $g = c$ by Lemma 28 applied to I_{sq}. However, his implies that $c \in K$ and $t \in K'$

are adjacent; a contradiction. Otherwise, we have $h = b$ due to Observation 27. By planarity, the vertex g must be located on the side of I_{sq} that contains t or on I_{sq} itself. If $g = c$, we have an edge between $c \in K$ and $t \in K'$; a contradiction. If $g = q$, collapsing $e = ab$ creates an adjacent separation pair. Finally, assume that g is some new (that is, not named so far) vertex of K'. Since I_{st} is separating and by 3-connectivity, three vertices of I_{st} have a neighbor on the side of I_{st} that does not contain a. This violates the degree bound of b or s. Hence, we conclude that the edge st is not constrained by a real 4-inhibitor.

Finally, assume that collapsing st creates a separation pair p_{st}, q_{st}. By Observation 31, we may assume without loss of generality that p_{st} is the new vertex resulting from collapsing st in $\mathcal{G}$ and that $q_{st} \neq d', q$. We will show that the vertices $\mathrm{N}_{\mathcal{G}}(s) \setminus \{s, t, q_{st}\}$ are connected in $\mathcal{G} \setminus \{s, t, q_{st}\}$, which yields a contradiction to Observation 30. The vertices q and d' are adjacent, as are a and b. There are two internally vertex-disjoint paths from q to b that do not pass through s or t: (1) qcb; and (2) the edge ab together with the path from a to q via c'. Similarly, there are two internally vertex-disjoint paths from q to a that do not pass through s or t: (1) qca; and (2) the path from a to q via c'. Removing s, t, q_{st} destroys at most one of the paths between q and a and, similarly, at most one of the paths between q and b. Altogether, this yields the claimed contradiction to Observation 30. Consequently, collapsing st does not create a separation pair. Together with the fact that st is not constrained by a real 4-inhibitor, we obtain a contradiction to assumption (F4). ◁

Overall, we obtain a contradiction to the assumption that I_{sq} has a vertex in K. Hence, (F7) is established, that is, the vertices of $I_{sq} = sqxy$ belong to $K' \cup \{a, b, q\}$.

We can turn to our main objective to resolve Case 1: show that

- I_{sq} appears as a 4-cycle of real edges in the triconnected component $R \in \mathcal{R}$ of $\mathcal{G}'$ that contains sqt *and*
- that this 4-cycle is hyperseparating in R, or collapsing sq in R creates an *adjacent* separation pair.

Note that $x \notin \{a, b\}$, as otherwise collapsing $e = ab$ would create an adjacent separation pair. Let $I'_{sq} = sqxy'$ denote the 4-cycle of $\mathcal{G}'$ corresponding to I_{sq}, i.e., if $y \notin \{a, b\}$ we have $y' = y$; and otherwise $y' = p$, where p is the vertex created by collapsing ab. Let $R \in \mathcal{R}$ denote the rigid triconnected component of $\mathcal{G}'$ that contains sqt.

We proceed by showing that $I'_{sq} = sqxy'$ is a 4-*cycle* of R. It suffices to show that for each separation pair of $\mathcal{G}'$ the vertices s, q, t, d', x, y' are incident to edges of a common separation class. In fact, we show something stronger: Since I_{sq} is hyperseparating in $\mathcal{G}$, there is some vertex z that is not the interior vertex of some real separating triangle and that is located on the side of I_{sq} that does not contain the vertices of K; see Figure 44 for illustration. Further, by 3-connectivity, there exist three simple paths P_1, P_2, P_3 that connect z with I_{sq} that are pairwise vertex-disjoint

except for their common endpoint z. Note that in $\mathcal{G}'$ the vertex z and P_1, P_2, P_3 are also on the side of I'_{sq} that does not contain the vertices of K. We show that for each separation pair of $\mathcal{G}'$ the vertices s, q, t, d', x, y', and, in addition, z are incident to edges of a common separation class. Recall that, by Observation 31, each separation pair of $\mathcal{G}'$ involves p. So let pq' be a separation pair of $\mathcal{G}'$. Note that we do not necessarily have $q' = q$.

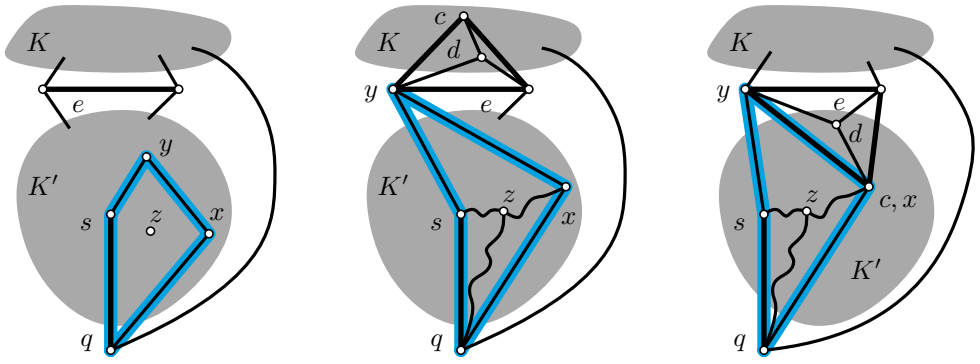

(a) $y' \neq p$ (b) $y' = p$ and $x \neq c$ (c) $y' = p$ and $x = c$

Figure 44: The possible ways the inhibitor I_{sq} can interact with the triangle abc.

The vertices s, q, t, d', x, y' induce a 2-connected subgraph of $\mathcal{G}'$. Adding the paths P_1, P_2, and P_3 to this subgraph yields another 2-connected subgraph H' of $\mathcal{G}'$. If $y' \neq p$, we have $p \notin H'$ and, hence, the vertices $\{s, q, t, d', x, y', z\} \setminus \{p, q'\}$ belong to the same connected component of $\mathcal{G}' \setminus \{p, q'\}$; for an illustration see Figure 44a. In fact, this statement is true even for $y' = p$: Recall that in this case $y \in \{a, b\}$ in $\mathcal{G}$ and, therefore, y belongs to the real separating triangle abc and has at least one neighbor in K, see Figure 44b and 44c. Thus, even if the edge yc of abc belongs to I_{sq}, y cannot have a neighbor on the side of I_{sq} that contains z since the degree of y is bounded by five. Hence, the endpoints of P_1, P_2, P_3 on I_{sq} are s, q, x and the degree of y' in H' is 2. Consequently, removing $p = y'$ from H results in yet another 2-connected subgraph of $\mathcal{G}'$ and, thus, the vertices $\{s, q, t, d', x, y', z\} \setminus \{p, q'\}$ belong to the same connected component of $\mathcal{G}' \setminus \{p, q'\}$. This implies that, regardless of whether $y' \neq p$ or $y' = p$, the vertices s, q, t, d', x, y', z belong to edges of a common separation class of the separation pair p, q'. Consequently, I'_{sq} is a 4-cycle in R, as claimed. Moreover, the vertex z also appears in R.

It remains to study whether I'_{sq} is hyperseparating in R. This is clearly the case if $y' \neq p$ since p is an endpoint of every $\mathcal{G}'$-virtual edge (see Figure 45a); or if the vertex t of the triangle sqt is located on the side of I'_{sq} that does not contain z (see Figure 45b). Thus, in these cases, Property (J3) of R ensures that $sq \notin \mathcal{K}_R$. So assume that $y' = p$ and that t, d' are located on the same side of I'_{sq} as z (see Figure 45c). Recall that the former assumption implies that y' has no neighbor on the side of I'_{sq} that contains z, t, d'. This implies that the triple s, q, x separates t

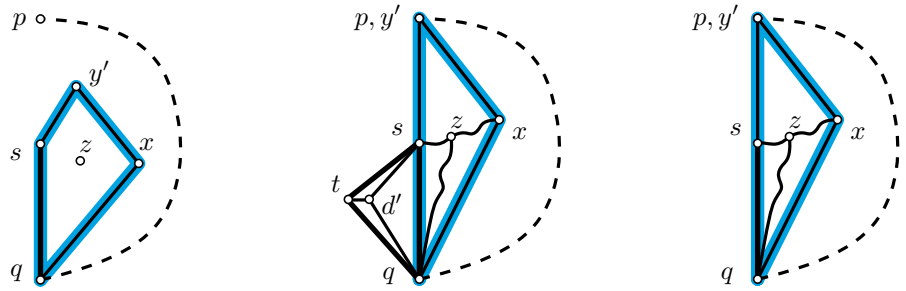

(a) $y' \neq p$ (b) $y' = p$ (c) $y' = p$

Figure 45: The inhibitor I_{sq} appears as a 4-cycle I'_{sq} in the rigid triconnected component R of $\mathcal{G}'$ that contains the triangle stq. Depending on how I_{sq} interacts with p and the triangle stq, it may be hyperseparating or not. Regardless, Properties (J1)–(J4) for $\mathcal{K}_R$ ensure that sq is not collapsed.

from $y' = p$ in R and, hence, collapsing sq in R creates an adjacent separation pair. Hence, Property (J4) for R ensures that $sq \notin \mathcal{K}_R$. Altogether, we have shown that if sq is constrained by a real 4-inhibitor in $\mathcal{G}$, then the edge sq is not in any set $\mathcal{K}_R$, for any rigid triconnected component R of $\mathcal{G}'$, that satisfies Properties (J1)–(J4). ◁

We have shown that if the edge sq of the separating triangle sqt is constrained by a real 4-inhibitor, the path $P = psq$ is an obstruction that is already taken into account implicitly when selecting the set $\mathcal{K}_R$ of the rigid triconnected component $R \in \mathcal{R}$ that contains sqt. So from now on, we may assume that:

Case 2: Collapsing sq in $\mathcal{G}$ creates a separation pair p_{sq}, q_{sq}. (This follows from the assumptions that sq is not constrained by a real 4-inhibitor and (F4), as explained in the beginning of Section 4.5.7.2.) By Observation 31, we may assume that p_{sq} is the vertex resulting from collapsing sq and that $q \notin \{t, d'\}$. Our plan is to show that this is only possible in two very constrained special cases for which we will describe separate treatments.

We begin by showing that the vertex t of the triangle sqt is positioned in a specific way. Recall that, in $\mathcal{G}'$, the new vertex p that is created by the collapse of $e = ab$ is adjacent to s and, hence, s is adjacent to a or b in $\mathcal{G}$. Without loss of generality, we assume that s is adjacent to b. By 3-connectivity, there exists a simple path from b to q whose inner vertices belong to K. Together with the edges bs and sq, this path forms a simple cycle C, see Figure 46a. We will show that the vertex t belongs to the side of C that also contains a. Towards a contradiction, assume otherwise, that is:

Case 2.1: t and a belong to distinct sides of C. For an illustration, refer to Figure 46a. By planarity and since a, b, q separate the vertices in K' (in particular, t) from the vertices in K, the vertices b, s, q form a separating triple. Thus, by 3-connectivity, the vertex s has a neighbor s' on the side of C that contains a (possibly $s' = a$). Further, the vertex b has some neighbor $b' \in K'$ on the side of C that contains t

(possibly $b' = t$). By the degree constraints for s, we have $\mathrm{N}_{\mathcal{G}}(s) = \{b, t, q, d', s'\}$, where d' is the interior vertex of stq.

Let us study the location of the vertices c, d. By Property (F3), we have $c \neq s$. Further, c cannot belong to the side of C that contains t, since otherwise the edge ac would cross C. Recall that the degree of b is bounded by five and, hence, $\mathrm{N}_{\mathcal{G}}(b) = \{a, c, d, s, b'\}$. Since b has at least one neighbor in K it follows that c, d belong to K, see Figure 46b.

Our plan is to show that at least one of the edges st and tq is not constrained by a real 4-inhibitor and can be collapsed without creating a separating pair, which yields a contradiction to assumption (F4).

First, assume that st is constrained by a real 4-inhibitor $I_{st} = stxy$. If $y = s'$, we have $x = b$ or $x = q$ since b, s, q is separating. In both cases, we obtain a contradiction to Lemma 28. By Observation 27, it follows that $y = b$. Since b, s, q is separating and by Observation 27, we obtain that $x = b' \neq t$. As I_{st} is separating, three of its vertices have a neighbor on the side of I_{st} that does not contain a. This is a contradiction to the degree bound of b or s.

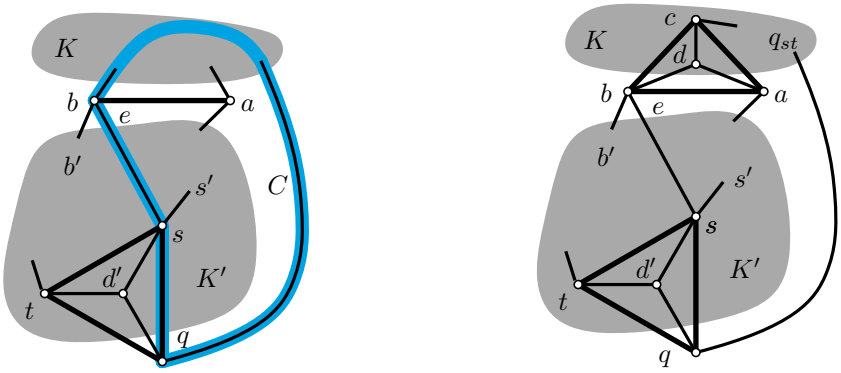

Figure 46: Illustrations for Case 2.1 in Lemma 38.

It remains to consider the case that collapsing st creates a separation pair p_{st}, q_{st}, where by Observation 31, we may assume that p_{st} is the vertex resulting from collapsing st and that $q_{st} \notin \{q, d'\}$. We will show that $|K \cap \mathrm{N}_{\mathcal{G}}(q)| = 1$ and that q_{st} is the (only) vertex in this set. As a first step, we show that each vertex z in $\mathcal{G} \setminus \{s, t, q_{st}\}$ is still connected to at least one of b or q, regardless of the identity of q_{st} (in particular, if $q_{st} = b$, every vertex is connected to q). For the case $z = d'$ the statement is clearly true. Next, assume that z belongs to the connected component of $\mathcal{G} \setminus \{b, s, q\}$ that does not contain t. By 3-connectivity, there exist three paths from z to b, s, q that are pairwise vertex-disjoint except for the common endpoint z. Only up to two of these paths can be destroyed by removing s, t, and q_{st}, as t belongs to the connected component of $\mathcal{G} \setminus \{b, s, q\}$ that does not contain z. As removing s destroys the path between z and s, we obtain that the path between z and q or the path between z

and b remains after the removal of s, t, q_{st}. It remains to consider the case that z ($\neq d'$) belongs to the connected component of $\mathcal{G} \setminus \{b, s, q\}$ that contains t. Again, there exist three paths from z to b, s, q that are pairwise vertex-disjoint except for the endpoint z. Due to the degree bound of s, the path to s must contain t. Thus, the removal of s, t, and q_{st} can only destroy the path from z to b or the path from z to q, but not both. Altogether, this shows that

(F11) each vertex is connected to b or q in $\mathcal{G} \setminus \{s, t, q_{st}\}$.

We now show that, unless $q_{st} \in K \cap \mathrm{N}_{\mathcal{G}}(q)$ and $|K \cap \mathrm{N}_{\mathcal{G}}(q)| = 1$, every vertex of $\mathcal{G} \setminus \{s, t, q_{st}\}$ belongs to the same connected component, which yields a contradiction to Observation 30. If $q_{st} = b$, all vertices are connected to q by (F11). So assume that $q_{st} \neq b$. If b and q are connected in $\mathcal{G} \setminus \{s, t, q_{st}\}$, then all vertices are connected. So assume that b and q are not connected. This implies that $q_{st} \in K$, as otherwise there clearly is some path between b and q whose inner vertices belong to K. Assume that there is a vertex $k \neq q_{st}$ that belongs to $K \cap \mathrm{N}_{\mathcal{G}}(q)$. From k there exist three paths to a, b, q in $\mathcal{G}$ that are pairwise vertex-disjoint except for the common endpoint k and whose inner vertices belong to K. We may assume the path to q to be the edge qk. Thus, removal of q_{st} may only destroy the path to a or to b but not both. Since a and b are adjacent, together with (F11) we obtain that all vertices in $\mathcal{G} \setminus \{s, t, q_{st}\}$ are connected, a contradiction. In conclusion, s, t, q_{st} is a separating triple where q_{st} is the singleton vertex in $K \cap \mathrm{N}_{\mathcal{G}}(q)$, as depicted in Figure 46b.

In order to prepare for the final step, in which we consider the remaining edge tq, we will show that $\mathrm{N}_{\mathcal{G}}(q) = \{s, t, d', q_{st}\}$. Our choice of e implies that $a, b \notin \mathrm{N}_{\mathcal{G}}(q)$. So assume that there is a neighbor $z \notin \{s, t, d', q_{st}\}$ of q. We have $z \in K'$, since $K \cap \mathrm{N}_{\mathcal{G}}(q) = \{q_{st}\}$, as established in the previous paragraph. Assume that z belongs to the side of C that contains t. By 3-connectivity and since b, s, q are separating, there exist three paths from z to b, s, q that are pairwise vertex-disjoint except for the common endpoint z. The path to s must contain t due to the degree bound for s. Therefore, in $\mathcal{G} \setminus \{s, t, q_{st}\}$ the vertices q and b are still connected via z since $q_{st} \in K$ cannot belong to any of these paths. However, by (F11), this is a contradiction to the fact that s, t, q_{st} are separating. Thus, it remains to consider the case that z belongs to the side of C that does not contain t. By 3-connectivity and since a, b, q are separating, there exist three paths from z to a, b, q that are pairwise vertex-disjoint except for the common endpoint z. We may assume that the path to q is the edge qz. By planarity, any of the paths to a, b can pass through t only after going through s; and none of the two paths passes through $q_{st} \in K$. Hence, due to these paths and the edge ab, the vertex z is connected to both b and q in $\mathcal{G} \setminus \{s, t, q_{st}\}$. Again this is a contradiction to (F11) and the fact that s, t, q_{st} is separating. Altogether this shows that $\mathrm{N}_{\mathcal{G}}(q) = \{s, t, d', q_{st}\}$, as claimed.

Let us now study the remaining edge tq. First, assume that tq is constrained by a real 4-inhibitor $I_{tq} = tqvw$. We study the identity of v. By Observation 27 we have $v \notin \{s, d'\}$. Together with $\mathrm{N}_{\mathcal{G}}(q) = \{s, t, d', q_{st}\}$ it follows that $v = q_{st}$.

Since a, b, q is separating, we have $w \in \{a, b\}$. As a and t are on different sides of the cycle C, by planarity we have $w = b$ and so $I_{tq} = tqq_{st}b$. The degree bound of b enforces $q_{st} = c$. The vertices s, d', a belong to the same side of I_{tq}. Since I_{tq} is separating, there is some vertex z on the side of I_{tq} that does not contain a. By 3-connectivity, there exist three paths from z to vertices of I_{tq} that are pairwise vertex-disjoint except for the common endpoint z. We obtain a contradiction to the degree bound of b or the fact that $\mathrm{N}_{\mathcal{G}}(q) = \{s, t, d', q_{st}\}$. Altogether, this shows that tq is not constrained by a real 4-inhibitor.

It remains to consider the case that collapsing tq creates a separation pair $p_{tq}q_{tq}$ where by Observation 31, we may assume that p_{tq} is the vertex resulting from collapsing tq and that $q_{tq} \notin \{s, d'\}$. We claim that $\mathrm{N}_{\mathcal{G}}(q) \setminus \{t, q, q_{tq}\}$ is connected in $\mathcal{G} \setminus \{t, q, q_{tq}\}$, which yields a contradiction to Observation 30. Recall that $\mathrm{N}_{\mathcal{G}}(q) = \{s, t, d', q_{st}\}$. Hence, if $q_{tq} = q_{st}$, our claim obviously holds. Otherwise, it suffices to show that q_{st} is still connected to s. There exist three paths from q_{st} to a, b, q that are pairwise vertex-disjoint except for the common endpoint q_{st} and whose inner vertices belong to K. We may assume that the path to q is the edge qq_{st}. Since $t \in K'$, the removal of t, q, q_{tq} may only destroy the path from q_{st} to a or the path to b, but not both. Thus, if $q_{tq} \neq b$, the edges bs and ab imply the claim. So assume that $q_{tq} = b$. There exist three paths from s to a, b, q that are pairwise vertex-disjoint except for the common endpoint s. The path to a cannot contain t and, thus, in $\mathcal{G} \setminus \{t, q, q_{tq}\}$ there exists a path from s to a. Since $\mathcal{G} \setminus \{t, q, q_{tq}\}$ also contains a path from a to q_{st} (through K), we conclude that s and q_{st} are connected, which yields the desired contradiction. Altogether this shows that collapsing tq does not create any separation pair. Since tq is also not constrained by a real 4-inhibitor, we obtain a contradiction to assumption (F4). Hence, the vertices t and a cannot belong to distinct sides of the cycle C, that is, Case 2.1 does not occur, as claimed. ◁

Case 2.2: t and a belong to the same side of C. See Figure 47a for illustration. Our plan is to establish that $q_{sq} \in K' \cup \{a\}$ and then to study the constraints implied by assumption (F4) for the remaining two edges st and tq of the triangle stq. As a preparation, we show that

(F12) the side of C that does not contain a and t does not contain any vertices of K' and, hence,

- the cyclic order of neighbors around s is t, d', q, b and then, potentially some other vertex s';
- the cyclic order of neighbors around b is s, then some vertices of K, a, and then potentially some vertices in K'; and
- the cyclic order of neighbors around q is s, d', t, potentially some vertex in K', and then at least one vertex in K.

Towards a contradiction, assume that (F12) does not hold. This implies that b, s, q are separating and by 3-connectivity each of them has a neighbor in K' on the side of C that does not contain a and t, see Figure 47b. Let s' denote the neighbor of s. We show that st is not constrained by a real 4-inhibitor and, further, that collapsing st cannot create a separation pair, which contradicts assumption (F4).

First, assume that st is constrained by a real 4-inhibitor $I_{st} = stxy$. We study the identity of y. Assume that $y = s'$. In this case, since b, s, q is separating, we have $x \in \{b, q\}$. In both cases we obtain a contradiction to Lemma 28. So assume that $y = b$. Since a, b, q are separating, we have $x \notin K$. Since b, s, q are separating, x cannot be the neighbor of b on the side of C that does not contain t. Finally, since ab is not constrained by a real 4-inhibitor (because $e = ab$ is a candidate edge), it follows that $x \neq a$ and, hence, $x \in K'$ and it belongs to the side of C that contains t. Three of the vertices of I_{st} have a neighbor on the side of I_{st} that does not contain a. This is a contradiction to the degree bound of b or s. Hence, the edge st is not constrained by a real 4-inhibitor.

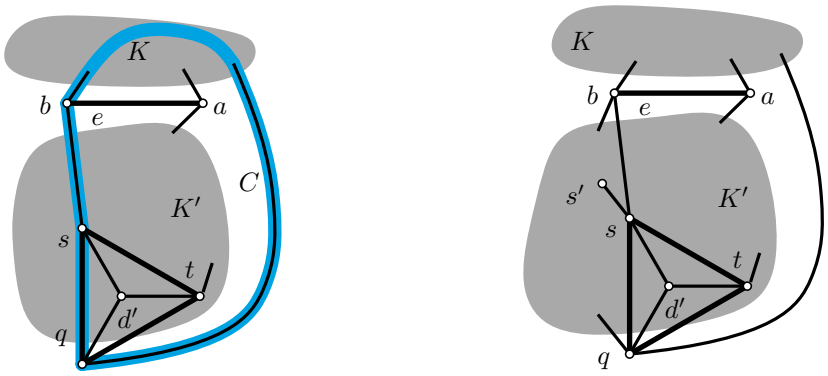

Figure 47: Illustrations for Case 2.2 in Lemma 38.

Next, assume that collapsing st creates a separation pair p_{st}, q_{st} where by Observation 31, we may assume that p_{st} is the vertex resulting from collapsing st and that $q_{st} \notin \{q, d'\}$. We will show that the vertices in $\mathrm{N}_{\mathcal{G}}(s) \setminus \{s, t, q_{st}\} \subseteq \{b, q, d', s'\}$ are connected in $\mathcal{G} \setminus \{s, t, q_{st}\}$, which yields a contradiction to Observation 30. The vertices q and d' remain connected regardless of the choice of q_{st}. If $q_{st} = b$, the vertex s' is still connected to q in $\mathcal{G} \setminus \{s, t, q_{st}\}$ by 3-connectivity and since b, s, q are separating. If $q_{st} = s'$, the vertex b is still connected to q via some path through K in $\mathcal{G} \setminus \{s, t, q_{st}\}$ by 3-connectivity and since a, b, q are separating. If $q_{st} \in K' \setminus \{s'\}$, the vertex s' remains connected to at least one of b or q in $\mathcal{G} \setminus \{s, t, q_{st}\}$ since b, s, q is separating and t and s' are in distinct components of $\mathcal{G} \setminus \{s, t, q_{st}\}$ and by 3-connectivity. Further, b and q are connected via some path through K as a, b, q is separating and by 3-connectivity. Finally, if $q_{st} \in K \cup \{a\}$ then, by 3-connectivity and since b, s, q are separating, there exists a path in $\mathcal{G} \setminus \{s, t, q_{st}\}$ between b and s' and a path between s'

and q. Thus, regardless of the choice of q_{st}, we obtain a contradiction and, hence, collapsing st does not create a separation pair. Since st is also not constrained by a real 4-inhibitor, we obtain a contradiction to assumption (F4). Altogether, this establishes (F12).

The identity of q_{sq}. We claim that

(F13) $q_{sq} \in K' \cup \{a\}$.

First we will establish $q_{sq} \neq b$ by showing that the graph $\mathcal{G} \setminus \{s, q, b\}$ is connected. Clearly every vertex of K remains connected to a. It suffices to show that every $z \in K' \setminus \{s\}$ is also connected to a. By 3-connectivity, there exist three paths P_a, P_b, P_q between z and a, b, q, respectively, that are pairwise vertex-disjoint except for the common endpoint z and whose interior vertices belong to K'. The vertices b, q cannot belong to the path P_a. Assume that $s \in P_a$ and, hence, $s \notin P_b, P_q$. Then, by planarity, fact (F12), and 3-connectivity, the subpath of P_a that connects s with a intersects $P_b - z$ or $P_q - z$, contradicting the assumption that these paths are pairwise vertex-disjoint. It follows that $s \notin P_a$ and, hence, the vertex z is connected to a in $\mathcal{G} \setminus \{s, q, b\}$. Altogether, this shows that $\mathcal{G} \setminus \{s, q, b\}$ is connected and, consequently, $q_{sq} \neq b$.

Next, assume that that $q_{sq} \in K$. We show that every vertex of $\mathcal{G} \setminus \{s, q, q_{sq}\}$ is connected to b and, thus, obtain a contradiction. The vertex a is connected to b via the edge ab. For each vertex $v' \in K'$ there exist three paths to a, b, q in $\mathcal{G}$ that are pairwise vertex-disjoint except for the common endpoint v'. As q_{sq} is located in K, removing s, q, q_{sq} from $\mathcal{G}$ leaves the path from v' to a or the path from v' to b intact. In any case, this implies that v' is connected to b in $\mathcal{G} \setminus \{s, q, q_{sq}\}$ due to the edge ab. Finally, each vertex v of K has three paths to a, b, q that are pairwise vertex-disjoint except for the common endpoint v. As s is located in K', removing s, q, q_{sq} from $\mathcal{G}$ leaves the path from v to a or the path from v to b intact. Again, this implies that v remains connected to b due to the edge ab. Altogether, this shows that $q_{sq} \notin K$ and establishes (F13).

Next we show that (F13) implies that

(F14) t and b are not connected in $\mathcal{G} \setminus \{s, q, q_{sq}\}$.

Assume otherwise. Each of s, q, q_{sq} has a neighbor in the component of $\mathcal{G} \setminus \{s, q, q_{sq}\}$ that does not contain b, t, d'. Let $\tilde{s}$ and $\tilde{q}$ denote these neighbors of s and q, respectively. A path from b to t in $\mathcal{G} \setminus \{s, q, q_{sq}\}$ together with the edges st and bs form a cycle C' in $\mathcal{G}$, for an illustration see Figure 48. By (F12), the vertices $\tilde{s}$ and $\tilde{q}$ are located on distinct sides of C' (none of these vertices belongs to C' since they are not is connected to t and b in $\mathcal{G} \setminus \{s, q, q_{sq}\}$). By planarity, the path from $\tilde{q}$ to $\tilde{s}$ in $\mathcal{G} \setminus \{s, q, q_{sq}\}$ intersects the path from b to t in some vertex. However, this implies that $\mathrm{N}_{\mathcal{G}}(s) \setminus \{s, q, q_{sq}\}$ is connected in $\mathcal{G} \setminus \{s, q, q_{sq}\}$, since the two intersecting paths are in $\mathcal{G} \setminus \{s, q, q_{sq}\}$. We obtain a contradiction to Observation 30. Altogether, this establishes (F14).

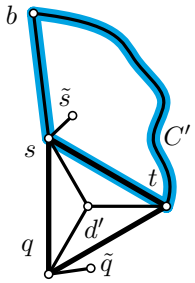

Figure 48: Illustrations for Case 2.2 in Lemma 38.

We proceed by studying the remaining edges of stq. Recall that, by assumption (F4), each of st and tq is constrained by a real 4-inhibitor or its collapse creates a separation pair. Let us start by studying the edge st.

Case 2.2.1: st is constrained by a real 4-inhibitor $I_{st} = stxy$. We will show that this case cannot occur. We begin by studying all possible identities of the vertices x and y. We will see that there are two principal options for the identity of I_{st}, namely $I_{st} = stq_{sq}b$ as depicted in Figure 49a, or $I_{st} = stt's'$ as depicted in Figure 49b. Having these two options in mind, we proceed to show that qt is a candidate edge. Using assumption (F5), we show that the situation looks as depicted in Figure 50b. To obtain the desired contradiction, we then show that the edge $\hat{s}t'$ is a candidate edge that can be collapsed without creating a separation pair, which contradicts (F4). Let us proceed to carry out this plan.

To establish the first possible identity of I_{st}, assume that $y = b$. By (F14), the vertices b and t are in distinct components of $\mathcal{G} \setminus \{s, q, q_{sq}\}$ (in particular, they are nonadjacent). It follows that $x \in \{q, q_{sq}\}$. If $x = q$, then collapsing $e = ab$ creates an adjacent separation pair, in contradiction to the choice of e. So we obtain $I_{st} = stq_{sq}b$, see Figure 49a. Note that in this case $q_{sq} \neq a$ because $e = ab$ is a candidate edge.

For the other identity of I_{st}, assume that $y \neq b$. By Observation 27, we have $q, d' \notin I_{st}$. Further, $x \neq b$ since b and t are in distinct components of $\mathcal{G} \setminus \{s, q, q_{sq}\}$. Thus, $x, y \in K' \cup \{a\}$. By 3-connectivity three vertices of I_{st} need a neighbor on the side of I_{st} that does not contain q and b. The degree bound of s implies that these three vertices are x, y, t. Now the degree bound of a implies that none of x, y can be a. In particular, even if $c \in K'$ and I_{st} shares the edge ac with abc the degree bound of a is violated since it has at least one neighbor in K.

Hence, we either have $I_{st} = stq_{sq}b$; or $I_{st} = stt's'$ where s' and t' are neighbors of s and t, respectively, in K'; see Figure 49. In both cases, 3-connectivity implies that t has a neighbor $\hat{t}$ on the side of $I_{st} = stxy$ that does not contain q: if $I_{st} = stt's'$, this follows from the degree bound of s as explained above. If $I_{st} = stq_{sq}b$, it follows from the degree bound of b. In particular, even if $c \in K'$ and I_{st} shares the edge bc with abc the degree bound of b is violated since it has at least one neighbor in K.

Let us proceed by studying the last edge tq of sqt. For the sake of contradiction,

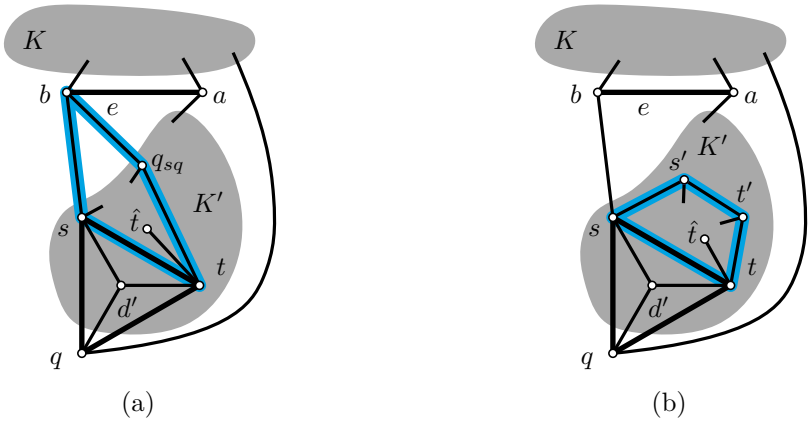

Figure 49: Illustrations for Case 2.2.1 in Lemma 38.

assume that qt is constrained by a real 4-inhibitor $I_{tq} = tqvw$. If $w = \hat{t}$, then by Lemma 28, v is the vertex y of $I_{st} = stxy$. If $I_{st} = stq_{sq}b$ and, thus, $y = b$, we obtain a contradiction to the degree bound of b. If $I_{st} = stt's'$ and, thus, $y = s'$, the triple sqs' forms a real separating triangle that is not vertex-disjoint from sqt, which is a contradiction to (F3). By Observation 27, the only other option is $w = x$. However, since $I_{tq} = tqvw$ is separating and by 3-connectivity, three vertices of I_{tq} have a neighbor on the side of I_{tq} that does not contain $\hat{t}$. This contradicts the degree bound of t or q regardless of whether $I_{st} = stq_{sq}b$ or $I_{st} = stt's'$. Hence, qt is not constrained by a real 4-inhibitor.

Next suppose that collapsing tq creates a separation pair $p_{tq}q_{tq}$ where by Observation 31, we may assume that p_{tq} is the vertex resulting from collapsing tq and that $q_{tq} \notin \{s, d'\}$. We claim that

(F15) $I_{st} = stt's'$, $q_{tq} = s'$, and t' is not connected to a in $\mathcal{G} \setminus \{t, q, s'\}$.

To prove the claim we show that if any of the three conditions is violated, we obtain a contradiction to the fact that, by Observation 30, the vertices $\mathrm{N}_{\mathcal{G}}(t) \setminus \{t, q, q_{tq}\}$ are not connected in $\mathcal{G} \setminus \{t, q, q_{tq}\}$. Obviously, d' and s remain connected $\mathcal{G} \setminus \{t, q, q_{tq}\}$.

Assume towards a contradiction that $I_{st} = stq_{sq}b$. By 3-connectivity, there exist paths from $\hat{t}$ to s, t, q_{sq}. These paths are pairwise vertex-disjoint, except for the common endpoint $\hat{t}$. Thus, if $q_{tq} = \hat{t}$ or $q_{tq} = q_{sq}$, we obtain the claimed contradiction to Observation 30 (in the former case, due to the edges sb, bq_{sq}), so assume otherwise. As $\hat{t}$ is on the side of I_{st} that does not contain q, removing t, q, q_{tq} from $\mathcal{G}$ leaves the path from $\hat{t}$ to s or the path from $\hat{t}$ to q_{sq} intact. If the former path is intact, we obtain the claimed contradiction to Observation 30 (q_{sq} can reach s via the path to $\hat{t}$ or the path $q_{sq}bs$). So assume that this is not the case and, hence, that the path to q_{sq} is intact, which together with the edges bq_{sq} and bs yields the claimed contradiction to Observation 30.

This establishes the first statement of Claim (F15), namely, $I_{st} = stt's'$. By

3-connectivity, there exist paths from $\hat{t}$ to t, s', t'. These paths are pairwise vertex-disjoint, except for the common endpoint $\hat{t}$. Thus, if $q_{tq} = \hat{t}$ or $q_{tq} = t'$, we obtain the claimed contradiction to Observation 30 due to the edges ss' and $s't'$, so assume otherwise. Removing t, q, q_{tq} from $\mathcal{G}$ leaves the path from $\hat{t}$ to s' or the path from $\hat{t}$ to t' intact. If the path to s' is intact, the contradiction to Observation 30 follows due to the edges ss' and $s't'$. So assume that the path to s' is destroyed by removing q_{tq} and, thus, the path to t' remains. If $q_{tq} \neq s'$, we again obtain a contradiction to Observation 30 due to the edges $s't'$ and ss'. It follows that $q_{tq} = s'$, which establishes the second statement of Claim (F15). The final statement, t' is not connected to a in $\mathcal{G} \setminus \{t, q, s'\}$, is also satisfied, since otherwise the edges ab and bs yet again yield the claimed contradiction to Observation 30. Altogether this proves Claim (F15).

By 3-connectivity, there is a path $P_{s'a}$ from s' to a that does not pass through any of b, q. Let z the first vertex of I_{st} that is encountered when traversing $P_{s'a}$ from a towards s'. We have $z \notin \{s, t\}$ due to the degree bounds of s and t and because $P_{s'a}$ does not pass through q. Further, $z \neq t'$ since t' is not connected to a in $\mathcal{G} \setminus \{t, q, s'\}$ by assumption. It follows that $z = s'$ and that s' has an edge to a neighbor $\check{s}$ on the side of I_{st} that does not contain $\hat{t}$ such that $\check{s} \in P_{s'a}$ (it is possible that $\check{s} = a$), see Figure 50a.

We remark that this situation is actually plausible, e.g. consider the case that $\check{s} = a$ and that t' has an edge to q. Recall that qt is not constrained by a real 4-inhibitor and that the only separation pair created by collapsing qt is p_{tq}, s', which cannot be adjacent as this would imply an edge between s' and t, which violates the degree bound of t, or an edge between s' and q, which implies that sqs' is a separating triangle (real or not) that is not vertex-disjoint from the real separating triangle stq and, hence, yields a contradiction to (F3). Thus, qt is a candidate edge.

We now make use of the fact that, by assumption (F6), s' belongs to a real separating triangle $s'vw$ with inner vertex d''. Recall that s' has edges to s, t', $\check{s}$, and a vertex, which we denote by $\hat{s}$, on the side of I_{st} that does not contain $\check{s}$. Due to the degree constraints for s' (implied by (G2)), at least two of these edges lead to v, w, d''. The vertex s cannot be part of real separating triangle other than sqt by (F3). Moreover, due to the degree of s, we have $s \neq d''$. By planarity, we have $\{\hat{s}, \check{s}\} \not\subseteq \{v, w, d''\}$, as otherwise there would be an edge between $\hat{s}$ and $\check{s}$. Hence, $t' \in \{v, w, d''\}$. Recall that t' and a belong to distinct connected components of $\mathcal{G} \setminus \{q, t, s'\}$. Moreover, the vertex $\check{s}$ belongs to the path $P_{s'a}$ that connects s' with a without visiting any of s, q, t, t', b. It follows that $\check{s} \notin \{v, w, d''\}$, as otherwise $\check{s}$ would be adjacent to t' and, thus, t and a would be connected in $\mathcal{G} \setminus \{q, t, s'\}$, which is not the case by assumption. Hence, $\hat{s} \in \{v, w, d''\}$.

So far we have shown that $\{s', \hat{s}, t'\} \subseteq \{s', v, w, d''\}$. By planarity, the fourth vertex $\{s', v, w, d''\} \setminus \{s', \hat{s}, t'\}$ cannot belong to the side of I_{st} that does not contain $\hat{s}$. It cannot belong to I_{st} itself, either, as this would imply a chord of I_{st}, contradicting Lemma 28. By planarity, we have $t' \neq d''$. Hence, we may assume without loss of generality that $v = t'$ and $\hat{s} = w$, see Figure 50b.

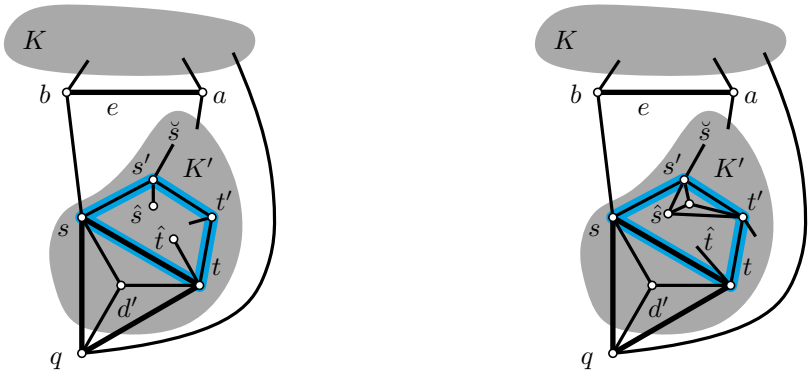

Figure 50: Illustrations for Case 2.2.1 in Lemma 38.

By 3-connectivity, the vertex t' has three chordless paths to t, q, s' that are pairwise vertex-disjoint except for the common endpoint t'. We may assume that the paths to t and s' are the edges $t't$ and $t's'$, respectively. Hence, due to the degree bound of s, the path to q implies an edge incident to t' on the side of I_{st} that does not contain $\hat{t}$, see Figure 50b.

The degree bounds of s, s', and t' (implied by (G2)) and 3-connectivity now imply that the side of I_{st} that contains $\hat{s}$ contains no other vertex aside from $\hat{s}$ and d'', since otherwise $\hat{s}$ and t would be a separation pair in $\mathcal{G}$. This implies that $\hat{s} = \hat{t}$ and, hence, $\mathrm{N}_{\mathcal{G}}(\hat{s}) = \{s', t, t', d''\}$. It follows, by Observation 27 and Lemma 28, that the edge $\hat{s}t'$ cannot be constrained by a real 4-inhibitor (since tt' would be a chord of such an inhibitor). Moreover, it follows by Observation 30 that collapsing $\hat{s}t'$ cannot create a separation pair: in $\mathcal{G} \setminus \{\hat{s}, t', q_{\hat{s}t'}\}$ with $q_{\hat{s}t'} \in \mathrm{V}(\mathcal{G}) \setminus \{s', d''\}$ the vertices $\mathrm{N}_{\mathcal{G}}(\hat{s}) \setminus \{\hat{s}, t', q_{\hat{s}t'}\}$ are connected regardless of the identity of $q_{\hat{s}t'}$: if $q_{\hat{s}t'} \neq s$, this follows due to the edges ss' and st. Otherwise, it follows due to the edge qt, path between q and a via K and the path $P_{s'a}$. Hence, $\hat{s}t'$ is a candidate edge that can be collapsed without creating a separation pair, in contradiction to (F4). In summary, we have shown that st is not constrained by a real 4-inhibitor. ◁

Having obtained a contradiction in Case 2.2.1, it follows that we are in

Case 2.2.2: Collapsing st creates a separation pair p_{st}, q_{st}. (This follows from the assumption that st is not constrained by a real 4-inhibitor and (F4), as explained in the beginning of Section 4.5.7.2.) By Observation 31, we may assume that p_{st} is the vertex resulting from collapsing st and that $q_{st} \notin \{q, d'\}$. Recall that by (F13) we have $q_{sq} \in K' \cup \{a\}$. We distinguish several cases regarding the position of q_{st}. Let us assume that:

Case 2.2.2.1: $q_{st} \in K$. We show that this is possible in a constrained special case only and describe a separate treatment for this case. For illustration refer to Figure 51.

Preparation. We start by establishing some statements which will be useful for the upcoming considerations regarding the last edge qt of stq. First, we show that

(F16) $K \cap \mathrm{N}_{\mathcal{G}}(q) = \{q_{st}\}$.

Assume otherwise and let $k \in K \cap \mathrm{N}_{\mathcal{G}}(q)$ with $k \neq q_{st}$. We claim that the vertices in $\mathrm{N}_{\mathcal{G}}(s) \setminus \{s, t, q_{st}\}$ are connected in $\mathcal{G} \setminus \{s, t, q_{st}\}$, which yields a contradiction to Observation 30. We proceed by proving this claim. By 3-connectivity and since a, b, q is a separator, there exist three paths from k to a, b, q that are pairwise vertex-disjoint except for the common endpoint k. Moreover, the interior vertices of these paths belong to K. We may assume that the path to q is the edge kq. Hence, since $s, t \in K'$, removing s, t, q_{st} can only destroy the path to a or to b, but not both. Due to the edge ab, this implies that q and d' belong to the same component as a and b in $\mathcal{G} \setminus \{s, t, q_{st}\}$. If $|\mathrm{N}_{\mathcal{G}}(s)| = 4$ of if s is adjacent to a, we obtain the claimed contradiction to Observation 30. So assume that s has another neighbor $s' \in K' \setminus \{b, t, q, d'\}$. By 3-connectivity there exist three paths from s' to s, t, q_{st} that are pairwise vertex-disjoint except for the common endpoint s'. The path to q_{st} passes through a, b, or q. This path shows that s' is also connected to b, q, d' in $\mathcal{G} \setminus \{s, t, q_{st}\}$ and, hence, we obtain the claimed contradiction to Observation 30. Altogether, this establishes (F16).

We observe that in $\mathcal{G} \setminus \{s, t, q_{st}\}$ all vertices of $K \setminus \{q_{st}\}$ are connected to a and b: this follows from the fact that each vertex in $K \setminus \{q_{st}\}$ has three disjoint path to a, b, q and only the path to q is destroyed by removing s, t, q_{st} by (F16) and since s and t belong to K'. In particular, this holds for the neighbors of q_{st} in K. Thus,

(F17) in $\mathcal{G} \setminus \{s, t, q_{st}\}$, the vertices a and b belong to a different component than q,

since by Observation 30, the vertices $\mathrm{N}_{\mathcal{G}}(q_{st)\setminus\{s,t,q_{st}\}}$ are not connected in the graph $\mathcal{G} \setminus \{s, t, q_{st}\}$. We proceed by showing that

(F18) $K' \cap \mathrm{N}_{\mathcal{G}}(q) = \{s, t, d'\}$.

Assume otherwise and let $k \in K' \setminus \{s, t, d'\}$ be a neighbor of q. Recall that by (F12) the side of C that does not contain a and t does not contain any vertices of K'. By 3-connectivity there exist three paths from k to a, b, q that are pairwise vertex-disjoint except for the common endpoint k. We may assume that the path to q is the edge kq. Assume that one of the two remaining paths passes through s. Then, by planarity, the other remaining path cannot pass through t. Thus, in $\mathcal{G} \setminus \{s, t, q_{st}\}$, there is a path from k to a or to b since q_{st} belongs to K. We obtain a contradiction to (F17) due to the edge kq. Altogether, this establishes (F18).

Algorithmic considerations. Before we continue with the argument, let us briefly note how to efficiently recognize this case ($q_{st} \in K$), for the given vertices a, b, c, d, s, t, q, d'. We have seen that (1) q_{st} is the only neighbor of q in K and (2) the only neighbors of q in K' are s, d', t. Therefore, $\deg_{\mathcal{G}}(q) = 4$. Conversely, if

$\deg_{\mathcal{G}}(q) = 4$, then the unique neighbor of q outside of s, d', t together with s, t forms a separating triple. Hence, we are in this case if and only if $\deg_{\mathcal{G}}(q) = 4$, which can be tested in constant time.

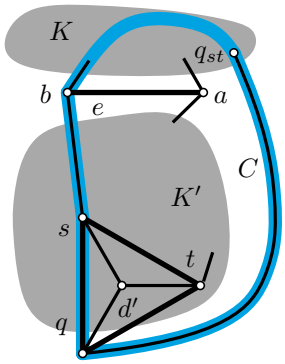

Figure 51: Illustration for Case 2.2.2.1 in Lemma 38.

The last edge qt. Let us continue our argument by considering the third edge qt of stq. We will start by showing that qt cannot be constrained by a real 4-inhibitor. Assume otherwise and let $I_{qt} = qtvw$ be a real 4-inhibitor constraining qt. By (F16), (F18), and Observation 27, we have $w = q_{st}$. As a, b, q is separating, we obtain that $v = a$ or $v = b$. The latter case would imply an edge between b and t, a contradiction to (F14). So $v = a$ and it follows that a and t are adjacent. Since $I_{qt} = qtvw = qtaq_{st}$ is separating, there is a vertex z on the side of I_{qt} that does not contain b. By 3-connectivity there exist three paths from z to three distinct vertices of I_{qt} that are pairwise vertex-disjoint except for the common endpoint z. As $t \in K'$ and $q_{st} \in K$, two of these paths have their endpoints at q and a, respectively. This is a contradiction to (F16) and (F18). Altogether, this shows that qt cannot be constrained by a real 4-inhibitor.

It remains to consider the case that collapsing qt creates a separation pair p_{qt}, q_{qt} where by Observation 31, we may assume that p_{qt} is the vertex resulting from collapsing qt and that $q_{qt} \notin \{s, d'\}$. We will show that the only possibility is $q_{qt} = b$. Towards a contradiction, assume otherwise:

Case 2.2.2.1.1: $q_{qt} \neq b$. We show that $\mathrm{N}_{\mathcal{G}}(q) \setminus \{q, t, q_{qt}\} = \{s, d', q_{st}\} \setminus \{q_{qt}\}$ is connected in $\mathcal{G} \setminus \{q, t, q_{qt}\}$, which yields a contradiction to Observation 30. It suffices to show that each vertex of the set is connected to b. This is clear for s and d' due to the edge bs. By (F16) and (F18) the only other vertex that may be contained in the set is q_{st}; this is the case only if $q_{qt} \neq q_{st}$. The vertex q_{st} (assuming that $q_{qt} \neq q_{st}$) has three paths to the vertices a, b, q that are pairwise vertex-disjoint except for the common endpoint q_{st} and whose interior vertices belong to K. Removing q, t, q_{qt} can destroy at most two of these paths, as $t \in K'$. Moreover, one of the destroyed paths is the one to q. Hence, the path to to a or the path to b remains intact. In the latter

case, there is nothing to show; in the former case the edge ab implies the claim, which yields the claimed contradiction to Observation 30. ◁

Having obtained a contradiction in the previous case, it follows that collapsing qt creates a *unique* separation pair, namely $p_{qt}, q_{qt} = b$.

Case 2.2.2.1.2: $q_{qt} = b$. We proceed in three steps. First, we show that in $\mathcal{G} \setminus \{q, t, b\}$ there is no path from s to a. From this we can conclude that s cannot have any neighbor in K' except for the vertices t, d'. Finally, we will derive that the rigid triconnected component $R \in \mathcal{R}$ of $\mathcal{G}'$ that contains sqt is a specific graph of constant size. We will use this fact and choose the edge set $\mathcal{K}_R$ explicitly, rather than inductively. Let us proceed by carrying out this plan.

First, assume, towards a contradiction, that in $\mathcal{G} \setminus \{q, t, b\}$ there is a path P_{sa} from s to a. We show that this path implies that $\mathrm{N}_{\mathcal{G}}(b) \setminus \{q, t, b\}$ is connected in $\mathcal{G} \setminus \{q, t, b\}$, which yields a contradiction to Observation 30. This is obvious for the vertices s, a, c, d. So let b' denote a fifth neighbor of b. If $b' \in K$, it is still connected to a, by 3-connectivity and since a, b, q is separating and $t \in K'$. So assume that $b' \in K'$. The edges bs and ab together with P_{sa} form a cycle C_{sa}. If b' is on P_{sa}, then it is connected to a via this path. So assume otherwise. By planarity b' is located on the side of C_{sa} that does not contain t. By 3-connectivity there exist three paths from b' to three distinct vertices of C_{sa}. Since q, t belong to the side of the C_{sa} that does not contain b', two of these paths remain in $\mathcal{G} \setminus \{q, t, b\}$. Hence, b' is connected to a, b, c, d, s and we obtain the claimed contradiction to Observation 30. Altogether, this shows that

(F19) in $\mathcal{G} \setminus \{q, t, b\}$ there is no path from s to a.

Next, we show that

(F20) $\mathrm{N}_{\mathcal{G}}(s) = \{q, d', t, b\}$.

Assume otherwise and let $s' \in (K' \cup \{a\}) \setminus \{t, d'\}$ be a neighbor of s. By (F19), we have $s' \neq a$. Recall that, by (F12), the side of C that does not contain a and t does not contain any vertices of K'. Observe that the path $abstq$ separates s' from K since a, b, q is separating. Thus, by 3-connectivity there exist three paths connecting s' to distinct vertices of $abstq$. By (F18) ($K' \cap \mathrm{N}_{\mathcal{G}}(q) = \{s, t, d'\}$) and (F19), the paths from s' lead to s, t, and b. This implies the existence of a path P_{bt} between t and b via s' that does not pass through s, a, or q. By (F14), the vertices t and b are not connected in $\mathcal{G} \setminus \{s, q, q_{sq}\}$. As a consequence, q_{sq} belongs to P_{bt}.

We will show that there is *another* path between b and t that does not pass through s and q and, in addition, it does not pass through q_{sq}, either; altogether, this yields a contradiction to (F14). By (F19), the vertices a and s belong to distinct components of $\mathcal{G} \setminus \{q, t, b\}$. Moreover, the vertex q_{sq} belongs to the component of $\mathcal{G} \setminus \{q, t, b\}$ that contains s since s' belongs to P_{bt} and s and s' are adjacent.

This implies that there exist three paths from $a \neq q_{sq}$ to q, t, b that are pairwise vertex-disjoint except for that common endpoint a and that do not contain any of s, q, q_{sq} as interior vertices. But then b and t belong to the same connected component of $\mathcal{G} \setminus \{s, q, q_{sq}\}$, which is the desired contradiction to (F14). Altogether this shows (F20).

We observe that t has some neighbor t' other than s, q, d', as otherwise s, q would be a separation pair in $\mathcal{G}$. By (F14), this neighbor cannot be b. It could be that there is an edge between t and a, as shown in Figure 52a. Alternatively, t' could be some vertex of $K' \setminus \{s, d'\}$, see Figure 52b. As the path $abstq$ separates t' from the vertices in K, we obtain that there are three paths from t' to three distinct vertices of $abstq$. By (F18) and (F20), the endpoints of these paths are a, b, t. Therefore, in this case, a, b, t is a separating triple and, thus, collapsing $e = ab$ also creates the separation pair p, t (in addition to p, q). Note our choice of e implies that these two possibilities are mutually exclusive, that is, *either* there is an edge between t and a *or* there is the neighbor $t' \in K' \setminus \{s, d'\}$ as otherwise collapsing e would create an *adjacent* separation pair. Further, note that these two cases can be distinguished in constant time (by checking for the presence of the edge at).

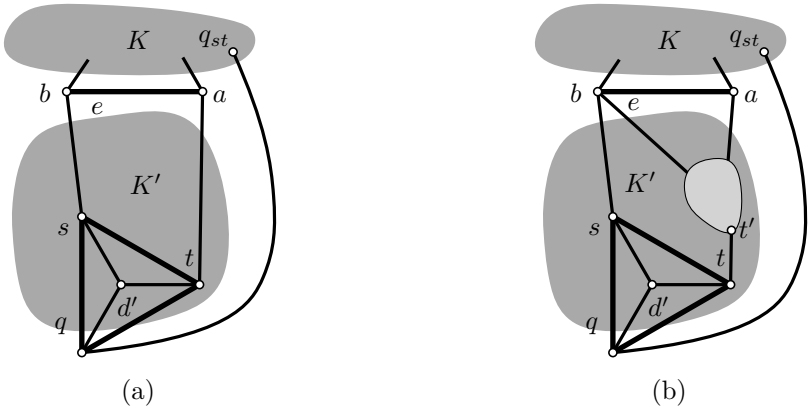

Figure 52: Illustrations for Case 2.2.2.1.2 in Lemma 38.

It also follows that in both cases the triconnected component $R \in \mathcal{R}$ of $\mathcal{G}'$ that contains sqt has an explicit form, see Figure 53: it contains sqt together with its inner vertex d', the edge ps, the $\mathcal{G}'$-virtual edge pq and the edge pt which is either $\mathcal{G}'$-virtual or not.

Recall that currently, in Section 4.5.7.2, we are only concerned with ensuring that the separation pairs of $\mathcal{G}'$ do not become adjacent when collapsing the edges in $\mathcal{K}'$ if the sets $\mathcal{K}_R$ are chosen, inductively or explicitly, such that (J1)–(J4) hold for the respective rigid triconnected component R of $\mathcal{G}'$ (the *remaining* separation pairs of $\mathcal{G}''$ are dealt with in Section 4.5.7.3).

If pt is not $\mathcal{G}'$-virtual in R (and thus p, t is not a separation pair in $\mathcal{G}'$), we set $\mathcal{K}_R = \{st\}$. This choice of $\mathcal{K}_R$ clearly satisfies the Properties (J1)–(J4) for the

Figure 53: Illustrations for Case 2.2.2.1.2 in Lemma 38.

graph R (note that collapsing st in R results in a K_3). Moreover, this choice does not create any new adjacencies for the vertices of $\mathcal{G}'$. In particular, the separation pair p, q remains nonadjacent.

If pt is $\mathcal{G}'$-virtual in R (and so p, t is a separation pair in $\mathcal{G}'$), we set $\mathcal{K}_R = \{qt\}$. Again, this choice of $\mathcal{K}_R$ clearly satisfies the Properties (J1)–(J4) for the graph R (note that collapsing qt in R results in a K_3). Recall that, by Observation 31, every separation pair of $\mathcal{G}'$ involves p. The collapse of qt in $\mathcal{G}'$ does not create any new adjacencies for p since our choice of e guarantees that $pq \notin \mathrm{E}(\mathcal{G}')$ and $pt \notin \mathrm{E}(\mathcal{G}')$. In particular, the separation pairs p, q and p, t remain nonadjacent. ◁

Altogether, we have shown how to explicitly deal with the obstruction presented by the path psq for the case that collapsing st creates a separation pair p_{st}, q_{st} where $q_{st} \in K$ (Case 2.2.2.1). Thus, from now on, we may assume that none of the separation pairs p_{st}, q_{st} has $q_{st} \in K$. ◁

Recall that we are considering the case that collapsing sq creates a separation pair p_{sq}, q_{sq} where p_{sq} is the new vertex created by the contraction and $q_{sq} \in (K' \cup \{a\}) \setminus \{t, d'\}$ due to (F13). Moreover, collapsing st creates a separation pair p_{st}, q_{st} where p_{st} is the vertex resulting from collapsing st and $q_{st} \notin \{q, d'\}$. We are in the process of determining the possible identities of q_{st}.

Note that collapsing st may create *multiple* separation pairs (all involving p_{st}), i.e., our choice of q_{st} might not be unique. So far, we have dealt with the case that there exists a choice for q_{st} such that $q_{st} \in K$. It remains to consider the case that $q_{st} \in K' \cup \{a, b\}$. First, assume that:

Case 2.2.2.2: $q_{st} = b$. (And collapsing st does not create a separation pair p_{st}, q^2_{st} with $q^2_{st} \in K$.) We will show that this case cannot occur. To this end, we claim that $\mathrm{N}_{\mathcal{G}}(s) \setminus \{s, t, b\}$ is connected in $\mathcal{G} \setminus \{s, t, b\}$, which yields a contradiction to Observation 30. If s has no neighbor except b, t, q, d', our claim clearly holds. So assume that s has some neighbor s' in $(K' \cup \{a\}) \setminus \{t, d'\}$. By 3-connectivity and since a, b, q is separating in $\mathcal{G}$, there is a path between a and q whose interior vertices belong to K. Consequently, we have that

(F21) a, q, d' belong to the same connected component of $\mathcal{G} \setminus \{s, t, b\}$.

This implies our claim for the case that s' and a are connected in $\mathcal{G} \setminus \{s, t, b\}$. It follows that it is sufficient to deal with the case that

(F22) s' and a belong to distinct connected components of $\mathcal{G} \setminus \{s, t, b\}$.

In particular, this implies that $s' \neq a$, i.e., $s' \in K' \setminus \{t, d'\}$. There exist three paths from s' to b, s, t, respectively, that are pairwise vertex-disjoint except for the common endpoint s'; see Figure 54 for illustration. This implies the existence of a path P_{bt} between b and t via s' such that, by (F21) and (F22),

(F23) the interior vertices of P_{bt} belong to the connected component of $\mathcal{G} \setminus \{s, t, b\}$ that does not contain a, q, d'.

In other words, P_{bt} does not pass through any of s, q, d', a. By (F14), the path P_{bt} passes through q_{sq} and, in particular, $q_{sq} \neq a$.

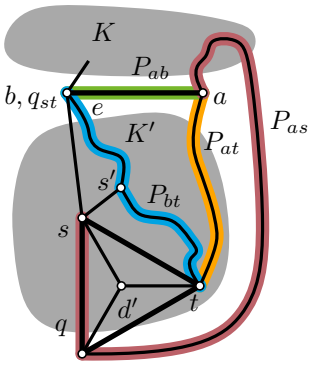

Figure 54: Illustration for Case 2.2.2.2 in Lemma 38.

We will show that there is *another* path between b and t that does not pass through s nor q and, in addition, it does not pass through q_{sq}, either; altogether, this yields a contradiction to (F14). By 3-connectivity, there exist three paths P_{as}, P_{at}, P_{ab} from a ($\neq q_{sq}$) to s, t, b, respectively, that are pairwise vertex-disjoint except for the common endpoint a. We may assume that $P_{ab} = ab$. By (F23), none of P_{as} and P_{at} contain an interior vertex of P_{bt} (because the interior vertices of P_{as} and P_{at} belong to the connected component of $\mathcal{G} \setminus \{s, t, b\}$ that contains a). In particular, P_{at} does not go through q_{sq}, and P_{as} does not go through s'. Therefore, P_{as} must pass through q, because all other neighbors of s (except d') are excluded. It follows that P_{at} cannot pass through q. Thus, P_{at} together with the edge ab constitutes a path between b and t in $\mathcal{G} \setminus \{s, q, q_{sq}\}$, which is the desired contradiction to (F14). ◁

It remains to treat the case that:

Case 2.2.2.3: $q_{st} \in K' \cup \{a\}$. (And collapsing st does not create a separation pair p_{st}, q^2_{st} where $q^2_{st} \in K \cup \{b\}$.) We show that this is possible in a constrained

special case only, for which we will describe a separate treatment. We start by establishing some basic properties implied by the case assumption and then study the constraints implied by (F4) for the final edge qt of sqt.

Preparation. The main goal of this preparatory step is to show that

(F24) there exists a path P_{sa} between s and a (as depicted in Figure 55) whose length is at least two and whose interior vertices belong to $K' \setminus \{t, d'\}$ (in particular, it does not pass through any of a, b, q, t, d').

Since s, t, q_{st} belong to $K' \cup \{a\}$, there is a path P_{bq} between b and q in $\mathcal{G} \setminus \{s, t, q_{st}\}$ (via K). The vertices $\mathrm{N}_{\mathcal{G}}(s) \setminus \{s, t, q_{st}\}$ are not connected in $\mathcal{G} \setminus \{s, t, q_{st}\}$ due to Observation 30. The vertices b, q, d' belong to this set and are connected in $\mathcal{G} \setminus \{s, t, q_{st}\}$ due to P_{bq}. Consequently, s has a neighbor $s' \in (K' \cup \{a\}) \setminus \{d'\}$ that belongs to the connected component of $\mathcal{G} \setminus \{s, t, q_{st}\}$ that does not contain b, q, d'. The vertex a also does not belong to said component due to the edge ab and consequently, we have $s' \neq a$. Hence, $s' \in K' \setminus \{d'\}$, for an illustration see Figure 55. Recall that, by (F12), the side of C that does not contain a and t does not contain any vertices of K'. In particular, this implies that the cyclic order of edges around s is as depicted.

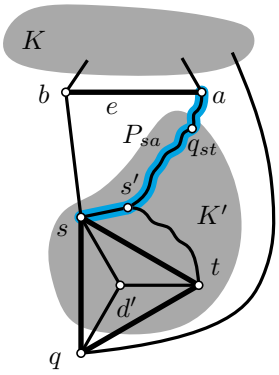

Figure 55: Illustrations for Case 2.2.2.3 in Lemma 38.

By 3-connectivity, there exist three paths connecting s' with s, t, q_{st}, respectively, that are pairwise vertex-disjoint except for the common endpoint s and that, by the discussion in the previous paragraph, do not contain $a, b, q, d', s, t, q_{st}$ as interior vertices.

We claim that there is a simple path between q_{st} and a, that does not pass through any of b, s, t, q. If $q_{st} = a$, the claim holds. So assume that $q_{st} \neq a$. By 3-connectivity, there exist three simple paths $P_s, P_t, P_{q_{st}}$ that connect a with s, t, q_{st}, respectively, that are pairwise vertex-disjoint except for the common endpoint a. Note that, by the definition of s', these paths belong to the connected component of $\mathcal{G} \setminus \{s, t, q_{st}\}$ that does not contain s'. We will show that $P_{q_{st}}$ has the desired properties. Let ℓ

denote the last vertex of $\{a, b, q\}$ visited by $P_{q_{st}}$ when traversed from a towards q_{st}. If $\ell = a$, there is nothing to show.

So assume that $\ell = q$, for an illustration see Figure 56a. Consider a cycle formed by the edge sq, the subpath of $P_{q_{st}}$ that connects q with q_{st}, and a simple path $P_{sq_{st}}$ between q_{st} and s via s' that does not pass through t (which exists by 3-connectivity and since s, t, q_{st} is a separating triple). By planarity (recall that, by (F12), the cyclic order of edges around s is as depicted), this cycle separates t from a. Moreover, it is vertex-disjoint from P_t (by definition of P_t, $P_{q_{st}}$, $P_{sq_{st}}$). Hence, we obtain a contradiction to planarity and, thus, $\ell \neq q$.

It remains to consider the case that $\ell = b$, for an illustration see Figure 56b. We will show that there is a path between t and b in $\mathcal{G} \setminus \{s, q, q_{sq}\}$ and, thereby, obtain a contradiction to (F14). Consider a cycle C'' formed by the subpath of $P_{q_{st}}$ that connects b with q_{st}, the edge tq, the path $C - s$, and a simple path between q_{st} and t via s' that does not pass through s (which exists by 3-connectivity and since s, t, q_{st} is a separating triple). This cycle separates s from a (by the cyclic order of edges around b implied by (F12)) and, hence, it must be intersected by P_s. We claim that P_s passes C'' via q. To this end, we analyze the different parts of C'': the paths P_s and $P_{q_{st}}$ are disjoint by definition. Hence, P_s does not intersect the part of C'' formed by the subpath of $P_{q_{st}}$. By definition of P_s, it does not intersect the part of C'' formed by the simple path between q_{st} and t via s' (that does not pass through s). Hence, P_s intersects C'' in the part formed by $(C - s - b)$. Towards a contradiction, assume that P_s does not intersect q. It follows that all vertices in the intersection of C'' with P_s belong to K. However, this implies that P_s contains a subpath from $s \in K'$ to some vertex of K that does not pass through any of a, b, q, contradicting the 3-connectivity. Consequently, P_s intersects C'' in q. This implies that P_t does not pass through any of s, q, b. By (F14), we have $q_{sq} \in (P_t \cup ab)$. However, there is another path from b to t that does not pass through s and q and that is internally vertex-disjoint from $P_t \cup ab$: the path formed by the subpath of $P_{q_{st}}$ that connects b with q_{st} and the path between q_{st} and t via s'. Hence, b and t are connected in $\mathcal{G} \setminus \{s, q, q_{sq}\}$ and we obtain a contradiction to (F14).

Altogether, this shows that there is a simple path $P_{q_{st}}$ between q_{st} and a, that does not pass through any of b, s, t, q (to unify notation, we set $P_{q_{st}} = a$ in case $q_{st} = a$).

The path formed by $P_{q_{st}}$ together with some simple path between q_{st} and s via s' that does not pass trough t is the desired path P_{sa} from (F24), see Figure 55. We conclude the preparatory step by showing that

(F25) all neighbors of b in K' belong to the connected component of a in $\mathcal{G} \setminus \{q, t, b\}$.

This is clear for s due to existence of the path P_{sa}. So suppose that b has a neighbor $b' \in K'$ with $b' \neq s$. (F25) holds for b' if it belongs to P_{sa}. Hence, assume that this is not the case. By 3-connectivity and (F12), there exist three paths from b' to the cycle $P_{sa} \cup sba$, which separates b' from q. These paths are pairwise vertex-disjoint except for the common endpoint b'. We may assume that one of the paths

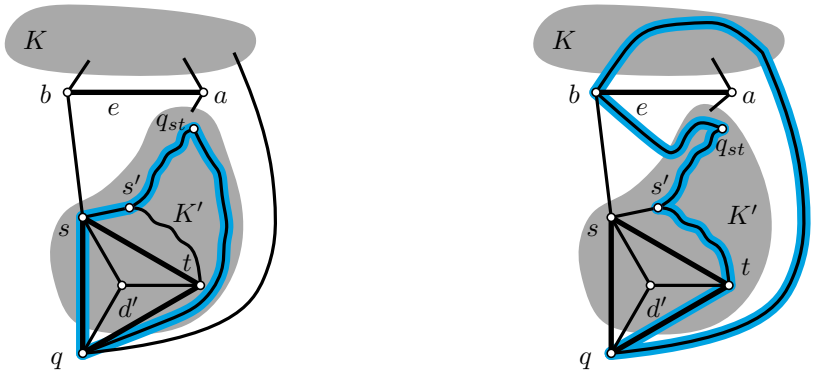

(a) $\ell = q$ (b) $\ell = b$

Figure 56: Illustrations for Case 2.2.2.3 in Lemma 38.

is the edge bb'. Since s, t, q_{st} is separating b' from s' and due to the degree bound of s, the remaining two paths lead to $P_{q_{st}}$ (which implies that $q_{st} \neq a$ in case a vertex such as b' exists) and, therefore, b' is connected to a in $\mathcal{G} \setminus \{q, t, b\}$, which establishes claim (F25).

The final edge qt. Let us proceed by considering the final edge qt of stq. We distinguish two cases.

Case 2.2.2.3.1: qt is constrained by a real 4-inhibitor $I_{qt} = qtvw$. We will show that this cannot actually be the case. To this end, we study all possible identities of v and w.

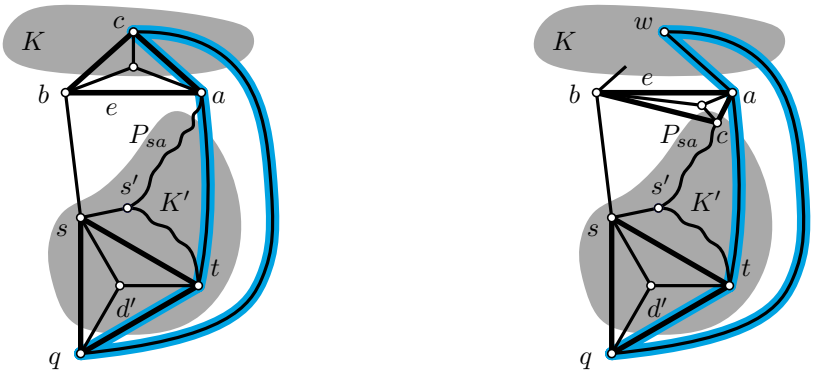

(a) $\ell = q$ (b) $\ell = b$

Figure 57: Illustrations for Case 2.2.2.3.1 in Lemma 38.

The identity of I_{qt}. By Observation 27, we have $\{v, w\} \cap \{s, d'\} = \emptyset$. Since e is a candidate edge, we have $w \notin \{a, b\}$ (otherwise, collapsing e creates an adjacent

separation pair). Assume that $w \in K$. Since a, b, q is separating, this implies that $v = a$ or $v = b$. The latter would imply an edge between b and t, which violates planarity since b and t are separated (due to the cyclic order of edges around s implied by (F12)) by a cycle formed by the path $P_{sa} \cup sq$ and a path between between q and a via K. Thus, $v \neq b$. So assume that $v = a$ and, thus, there is an edge between a and t. Recall that a belongs to the separating triangle abc with inner vertex d. Due to the degree bound of a we obtain that $c = w$ (Figure 57a); or c belongs to P_{sa} (Figure 57b). Since I_{qt} is separating, there is a vertex on the side of I_{qt} that does not contain s. Moreover, this vertex has three paths to distinct vertices of I_{qt}, which violates the degree bound of a or t. Altogether, this shows that $w \notin K$.

It remains to consider the case that $w \in K'$. Assume that the vertex v does not belong to the component of $\mathcal{G} \setminus \{s, t, q_{st}\}$ that does not contain q (in particular, this is the case if $w \neq q_{st}$). This brings the degree of t up to five. Since I_{qt} is separating, there is a vertex on the side of I_{qt} that does not contain s. By 3-connectivity, this violates the degree bound of t or q.

So assume that v belongs to the component of $\mathcal{G} \setminus \{s, t, q_{st}\}$ that does not contain q, which implies that $w = q_{st}$ and, moreover, $q_{st} \neq a$ by our choice of e. Since I_{qt} is separating, there is vertex z on the side of I_{qt} that does not contain s. By 3-connectivity, there exist three paths from z to I_{qt} that are pairwise vertex-disjoint except for the common endpoint z. Due to the degree bound of q, these paths lead to v, q_{st}, t, see Figure 58. Again, this brings the degree of t up to five. W.l.o.g. we may assume that z is a neighbor of t.

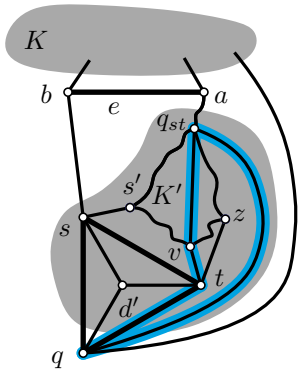

Figure 58: Illustrations for Case 2.2.2.3.1 in Lemma 38.

Obtaining the contradiction. Now that we have established that the situation looks as depicted in Figure 58, our plan is as follows: first, we show that collapsing st creates precisely one separation pair, i.e, the vertex q_{st} is unique. It easily follows that st is a candidate edge. We then use the fact that, by (F6), the vertex q_{st} belongs to a real separating triangle, which yields the desired contradiction due to the degree bound implied by (G2).

So first, let us establish that p_{st}, q_{st} is the unique separation pair created by collapsing st. Towards a contradiction, assume that this is not the case and let p^2_{st}, q^2_{st} be another separation pair created by collapsing st. By Observation 31, we may assume that $p^2_{st} = p_{st}$ and $q^2_{st} \neq q, d'$. Assume that q^2_{st} belongs to the component of $\mathcal{G} \setminus \{s, t, q_{st}\}$ that contains q. Then, by 3-connectivity, all vertices in the component of $\mathcal{G} \setminus \{s, t, q_{st}\}$ that does *not* contain q, in particular the two neighbors v and z of t, remain connected to q_{st} in $\mathcal{G} \setminus \{s, t, q^2_{st}\}$. However, due to the path $q_{st}qd'$, this implies that the vertices $\mathrm{N}_{\mathcal{G}}(t) \setminus \{s, t, q^2_{st}\}$ are connected in $\mathcal{G} \setminus \{s, t, q^2_{st}\}$, which yields a contradiction to Observation 30. It remains to consider the case that q^2_{st} belongs to the component of $\mathcal{G} \setminus \{s, t, q_{st}\}$ that does not contain q. We claim that, once again, the vertices $\mathrm{N}_{\mathcal{G}}(t) \setminus \{s, t, q^2_{st}\}$ are connected in $\mathcal{G} \setminus \{s, t, q^2_{st}\}$, which yields a contradiction to Observation 30. If $q^2_{st} = z$, the claim follows due to the path $vq_{st}qd'$. So assume $q^2_{st} \neq z$. Recall that z has two vertex-disjoint paths to v and q_{st}, respectively. The vertex q^2_{st} can only belong to one of these paths and, so, z remains connected to v or to q_{st}. In any case, this implies that z remains connected to q_{st} (due to the edge vq_{st}) and, hence, q and d'. Finally, if $v \neq q^2_{st}$, the vertex v also remains connected to q_{st}, q, d' due to the edge vq_{st}. Altogether, this yields the claimed contradiction to Observation 30. Therefore, p_{st}, q_{st} is the unique separation pair created by collapsing st, as claimed.

Due to the degree bounds of s and t, neither s nor t is adjacent to q_{st}. Further, by the case assumption of Case 2.2.2, st is not constrained by a real 4-inhibitor and, hence, it is a candidate edge. By (F6), no candidate edge is safe. Thus, q_{st} belongs to a real separating triangle T and, hence, its degree is bounded by five by (G2). Let E_T denote the set of six edges that are incident to vertices of T and do not lead to the exterior of T. Due to the degree bound of q_{st}, at most two of its incident edges do not belong to E_T. By (F3), real separating triangles are pairwise vertex-disjoint and the edge qq_{st} cannot belong to E_T. Thus, at most one of the remaining edges incident to q_{st} does not belong to E_T. We distinguish two cases regarding the identity of v.

First, assume that $v = s'$, for an illustration see Figure 59a. In this case, the edge vq_{st} cannot belong to E_T: v cannot be the interior vertex of T since its degree is larger than three. Moreover, if vq_{st} belongs to T, the Precondition (G3) implies that vs is real. This, in turn, implies that $sqq_{st}v$ is a real 4-inhibitor constraining sq, which contradicts the assumption of Case 2. Thus, both vq_{st} and qq_{st} do not belong to E_T. But this implies that the vertices of $\mathrm{N}_{\mathcal{G}}(q_{st}) \setminus \{v, q\}$ are pairwise adjacent since the edges of E_T form a K_4. However, two of these vertices belong to distinct sides of I_{qt}, which yields the desired contradiction.

So assume that $v \neq s'$, for an illustration see Figure 59b. Recall that s' has three pairwise vertex-disjoint paths to s, t, q_{st}, respectively. The path to s is the edge ss', which together with the path to q_{st} forms a subpath of P_{sa}. The path to t must pass through v due to the degree bound of t and by planarity. Hence, P_{sa} does not pass through v. This implies that q_{st} has the following set of neighbors: q, v, two vertices along P_{sa} (recall that $q_{st} \neq a$), and a vertex on the side of I_{qt} that does not contain s. The last three of these neighbors are pairwise nonadjacent: the two

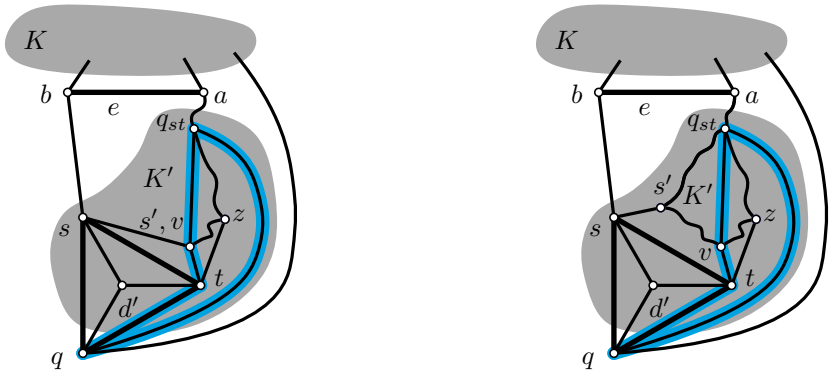

(a) $v = s'$ (b) $v \neq s'$

Figure 59: Illustrations for Case 2.2.2.3.1 in Lemma 38.

vertices along P_{sa} belong to the side of I_{qt} that contains s. Moreover, these two vertices belong to distinct components of $\mathcal{G} \setminus \{s, t, q_{st}\}$. Consequently, at most one of the edges between q_{st} and these three neighbors belongs to E_T. Together with the fact that qq_{st} does not belong to E_T, we obtain that E_T contains at most two edges incident to q_{st}, which yields the desired contradiction.

Altogether, this shows that qt is not constrained by a real 4-inhibitor. ◁

Having obtained a contradiction in the previous case, it follows that:

Case 2.2.2.3.2: Collapsing qt creates a separation pair p_{qt}, q_{qt}. (This follows from the assumption that qt is not constrained by a real 4-inhibitor and (F4), as explained in the beginning of Section 4.5.7.2). We will show that this is possible in a specific special case only and we will describe a separate treatment to deal with this situation.

We start by studying the identity of q_{qt} and the neighborhoods of q and t. The derived facts allow us to show that both st and qt are candidate edges (in particular, their collapses do not create adjacent separation pairs). We will then describe a replacement strategy ϱ to find, in at most a linear number of steps, a candidate edge for which the current Case 2.2.2.3.2 does not arise. More precisely, the first step replaces e with $\varrho(e) \in \{qt, st\}$. There is no guarantee that Case 2.2.2.3.2 does not arise for $\varrho(e)$. However, we will make sure that ϱ is acyclic, that is, after at most $i \in O(n)$ repetitions, we end up with a candidate edge $\varrho^i(e) \neq e$ for which Case 2.2.2.3.2 does not arise. So let us proceed by carrying out this plan.

The identity of q_{qt} and the neighborhood of q and t. By Observation 31, we may assume that p_{qt} is the vertex resulting from collapsing qt and that $q_{qt} \neq s, d'$. Let us study the identity of q_{qt}. It can not belong to K, as otherwise, by 3-connectivity and since a, b, q is separating, every vertex of $\mathcal{G} \setminus \{q, t, q_{qt}\}$ remains connected to the edge ab, as $q_{qt} \in K$ and $t \in K'$. So assume that $q_{qt} = b$. Note that the neighbors a, c, d, s of b belong to the same connected component of $\mathcal{G} \setminus \{q, t, b\}$ due

to the path P_{sa}. By Observation 30, this implies that b has a fifth neighbor b' such that b' and a belong to distinct components of $\mathcal{G} \setminus \{q, t, b\}$. By (F25), we have $b' \notin K'$. However, if $b' \in K$, the 3-connectivity also implies that b' remains connected to a in $\mathcal{G} \setminus \{q, t, b\}$ since $t \in K'$. Thus, we obtain a contradiction and it follows that $qqt \neq b$.

So far, we have established that $q_{qt} \in K' \cup \{a\}$ (and collapsing qt does not create a separation pair p_{qt}, q_{qt}^2 where $q_{qt}^2 \in K \cup \{b\}$). The vertex q_{qt} cannot belong to the component of $\mathcal{G} \setminus \{s, t, q_{st}\}$ that does not contain q since otherwise, the vertices in $\mathrm{N}_{\mathcal{G}}(q_{qt}) \setminus \{q, t, q_{qt}\}$ clearly belong to to the same connected component of $\mathcal{G} \setminus \{q, t, q_{qt}\}$, which yields a contradiction to Observation 30.

We claim that, as depicted in Figure 60, both q and t have some neighbor (denoted by t' and $\hat{q}$, respectively) in $K' \setminus \{d', s\}$ that does not belong to the connected component of s' in $\mathcal{G} \setminus \{s, t, q_{st}\}$ and that does not belong to the path P_{sa}.

Assume that a vertex with the properties of $\hat{q}$ does not exist. We will show that this implies that $\mathrm{N}_{\mathcal{G}}(q) \setminus \{q, t, q_{qt}\}$ is connected in $\mathcal{G} \setminus \{q, t, q_{qt}\}$, contradicting Observation 30. The neighbor(s) of q in K and the vertices d' and s are connected in $\mathcal{G} \setminus \{q, t, q_{qt}\}$ due to the path $d'sb$ and some path(s) through K (which exist by 3-connectivity and since none of q, t, q_{qt} belongs to $K \cup \{b\}$). Moreover, q cannot have a neighbor in the connected component of s' in $\mathcal{G} \setminus \{s, t, q_{st}\}$, and if it has a neighbor (distinct from q_{qt}) on P_{sa}, then this neighbor remains connected to b since it belongs to the cycle formed by P_{sa} and the path sba, which does not contain q or t. Altogether, this yields the claimed contradiction to Observation 30.

Similarly, we show that if no vertex with the properties of t' exists, then $\mathrm{N}_{\mathcal{G}}(t) \setminus \{q, t, q_{qt}\}$ is connected in $\mathcal{G} \setminus \{q, t, q_{qt}\}$, contradicting Observation 30. The neighbors d' and s of t are connected to b in $\mathcal{G} \setminus \{q, t, q_{qt}\}$. Each of the neighbors of t in the connected component of $\mathcal{G} \setminus \{s, t, q_{st}\}$ that contains s' remains connected to s (recall that q_{qt} does not belong to this component) and, hence, it also remains connected to b. Finally, if t has a neighbor on P_{sa}, then it belongs to the cycle formed by P_{sa} and sba, which does not contain q or t and, hence, this neighbor also remains connected to b. Altogether, this yields the claimed contradiction to Observation 30.

It follows that q has a unique neighbor in K, which we denote by k. Moreover, t has a unique neighbor in the connected component of s' in $\mathcal{G} \setminus \{s, t, q_{st}\}$, which is denoted by $\hat{t}$. Note that it is possible that $q_{qt} = q_{st}$ or $q_{qt} = a$ or both.

***st* and *qt* are candidate edges.** Let us proceed by showing that both st and qt are candidate edges. By the assumptions of Case 2.2.2 and 2.2.2.3.2, neither qt nor st is constrained by a real 4-inhibitor. To see that the collapse of st or qt cannot create an *adjacent* separation pair, we can make use of the fact that we are aware of the entire neighborhood of each of s, q, t. Hence, we can explicitly check all edges incident to p_{st} or p_{qt} in the graph created by collapsing st or qt, respectively. These edges are $p_{qt}k$, $p_{qt}\hat{q}$, $p_{qt}t'$, $p_{qt}\hat{t}$, $p_{st}t'$, $p_{st}\hat{t}$, $p_{st}s'$, and $p_{st}b$. Their endpoints are the potential adjacent separation pairs.

The pair p_{st}, b is nonseparating by the assumption of Case 2.2.2.3. The pairs p_{qt}, k

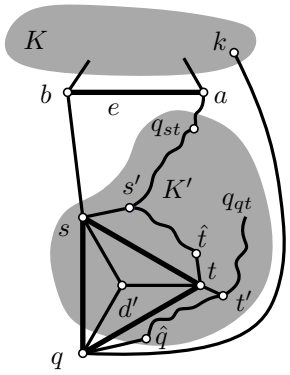

Figure 60: Illustrations for Case 2.2.2.3.2 in Lemma 38.

and $p_{qt}, \hat{t}$ are nonseparating since we have established above that $q_{qt} \in K' \cup \{a\}$ and that q_{qt} does not belong to the connected component of $\mathcal{G} \setminus \{s, t, q_{st}\}$ that does not contain q for any valid choice of q_{qt}. Moreover, by 3-connectivity, it is easy to see that the vertices

- $\mathrm{N}_{\mathcal{G}}(q) \setminus \{q, t, \hat{q}\} = \{s, d', k\}$ are connected in $\mathcal{G} \setminus \{q, t, \hat{q}\}$;
- $\mathrm{N}_{\mathcal{G}}(t) \setminus \{q, t, t'\} = \{s, d', \hat{t}\}$ are connected in $\mathcal{G} \setminus \{q, t, t'\}$;
- $\mathrm{N}_{\mathcal{G}}(t) \setminus \{s, t, t'\} = \{q, d', \hat{t}\}$ are connected in $\mathcal{G} \setminus \{s, t, t'\}$;
- $\mathrm{N}_{\mathcal{G}}(t) \setminus \{s, t, \hat{t}\} = \{q, d', t'\}$ are connected in $\mathcal{G} \setminus \{s, t, \hat{t}\}$; and
- $\mathrm{N}_{\mathcal{G}}(s) \setminus \{s, t, s'\} = \{q, d', b\}$ are connected in $\mathcal{G} \setminus \{s, t, s'\}$;

respectively. Altogether, by Observation 30, this shows that collapsing st or qt cannot create an adjacent separation pair and, thus, st and qt are candidate edges.

Implications of (F5). By (F5), no candidate edge is safe. Hence, we may assume without loss of generality (recall that q_{qt} is not necessarily unique) that there is a path $\tilde{P} = \tilde{p}\tilde{s}q_{qt}$ where $\tilde{p} = t$ or $\tilde{p} = q$, and $\tilde{s}q_{qt}$ is an edge that belongs to a real separating triangle.

We have established that $q_{qt} \in K' \cup \{a\}$. We will now show that $q_{qt} \neq a$. So, towards a contradiction, assume that $q_{qt} = a$. To obtain the desired contradiction, we proceed in two steps. First, we show that $c = k$. From this we derive that the edge bc violates (F4).

So let us start be showing that $c = k$. Since $\tilde{s}q_{qt}$ belongs to a separating triangle, we have $\tilde{s} = b$ or $\tilde{s} = c$. Assume $\tilde{s} = b$. This implies that there is an edge between q and b; or between t and b. The former is not possible by our choice of e. The latter is not possible due to the degree bound of t. Hence, in both cases we obtain a contradiction and so $\tilde{s} = c$.

We show that this implies that $\tilde{p} \neq t$. Assume otherwise. We have $c \notin K$, as otherwise there would be an edge between $c \in K$ and $t \in K'$. So assume that $c \in K'$ and that there is an edge between c and t. Due to the degree bound of t it follows that $c \in \{t', \hat{t}\}$. In both cases we obtain a contradiction due to the edge bc: the case $c = t'$ is not possible since t' and b belong to distinct components of $\mathcal{G} \setminus \{q, t, q_{qt}\}$ (due to the path bss' the vertex b belongs to the same component as s', which does not belong to the component of t' by definition of t'). The case $c = \hat{t}$ is also not possible since $\hat{t}$ and b belong to distinct components of $\mathcal{G} \setminus \{s, t, q_{st}\}$ (by definition of $\hat{t}$, it belongs to the component of s', which does not belong to the component of b by definition of s'). Altogether, this shows that $\tilde{p} \neq t$ and, thus, $\tilde{p} = q$.

This implies an edge between q and c and, thus, due to the degree bound of q we have $c \in \{k, \hat{q}\}$. Yet again, the case $c = \hat{q}$ is not possible due to the edge bc since, as discussed above, b belongs to the connected component of $\mathcal{G} \setminus \{q, t, q_{qt}\}$ that does not contain t', which, by definition, is in the same component as $\hat{q}$. Therefore, $c = k$ as claimed, for an illustration see Figure 61.

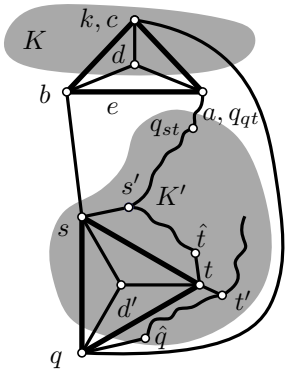

Figure 61: Illustrations for Case 2.2.2.3.2 in Lemma 38.

To prepare for the next step, we show that

(F26) the cycle $cbsq$ is nonseparating.

Indeed, towards a contradiction, assume that $cbsq$ is separating and let z be a vertex on the side of $cbsq$ that does not contain t. By 3-connectivity, there exist three paths connecting z with $cbsq$ that are pairwise vertex-disjoint except for the common endpoint z. Since $c \in K$ and $s \in K'$, two of these paths lead to b and q, respectively. However, this contradicts the degree bound of q. Altogether, this shows (F26).

It remains to show that bc violates (F4). To this end, we show that bc is a candidate edge that can be collapsed without creating a separation pair. So ,towards a contradiction, assume that bc is constrained by a real 4-inhibitor $I_{bc} = bcxy$. Assume that $y \in K'$. Since a, b, q is separating, we have $x = a$ or $x = q$. The former can be excluded by Observation 27. In the latter case, planarity and the degree bound of q imply that $y = s$, $y = t$, or that $y = \hat{q}$. The first option contradicts (F26). The

case $y = t$ violates the degree bound of t. Finally, the case $y = \hat{q}$ can also be excluded since, as discussed above, $\hat{q}$ and b are nonadjacent. Thus, in all cases we obtain a contradiction and, therefore, $y \notin K'$.

Next assume that $y \in K$. This implies that $bcqs$ is separating, which contradicts (F26). Finally, the case $y = a$ can be excluded by Observation 27. Altogether, this shows that bc is not constrained by a real 4-inhibitor.

It remains to sow that bc can be collapsed without creating a separation pair. So, towards a contradiction, assume that collapsing bc creates a separation pair $p_{bc}q_{bc}$ where by Observation 31, we may assume that p_{bc} is the vertex resulting from collapsing bc and that $q_{bc} \neq a, d$. By (F26), there is no vertex of K on the side of $bcqs$ that does not contain a and t. The other side of $bcqs$ also cannot contain a vertex of $K \setminus \{d\}$: by 3-connectivity, such a vertex z would have three vertex path to a, b, q that are disjoint except for the common endpoint z. Moreover, the interior vertices of these paths belong to K. By (F26), the only neighbors of b in K are c and d. Hence, the path from z to b has to pass through c. However, due to the degree bound of q, the vertex c is the only neighbor of q in K and, therefore, the path from z to q also passes through q; a contradiction. Thus, there cannot be a vertex in K other than c, d, which implies that $\mathrm{N}_{\mathcal{G}}(c) = \{a, b, d, q\}$. Further, since $q_{bc} \neq a, b, c, d$, this implies that $q_{bc} \in K' \cup \{q\}$.

If $q_{bc} = q$, then clearly the vertices $\mathrm{N}_{\mathcal{G}}(c) \setminus \{b, c, q_{bc}\} = \{a, d\}$ are connected in $\mathcal{G} \setminus \{b, c, q_{bc}\}$, which yields a contradiction to Observation 30. So assume that $q_{bc} \neq q$ and, hence $q_{bc} \in K'$. Again, we show that the vertices $\mathrm{N}_{\mathcal{G}}(c) \setminus \{b, c, q_{bc}\} = \{a, d, q\}$ are connected in $\mathcal{G} \setminus \{b, c, q_{bc}\}$, which yields a contradiction to Observation 30. It suffices to show that there is a path between a and q in $\mathcal{G} \setminus \{b, c, q_{bc}\}$. This is true, since there are two interior vertex-disjoint paths between a and q, none of which passes through b or c: the path P_{sa} together with the edge sq; and the path through the component of $\mathcal{G} \setminus \{q, t, q_{qt}\}$ (recall that we are assuming $q_{qt} = a$) that contains t'. This yields the desired contradiction and so, altogether, establishes that collapsing bc cannot create a separation pair. Since bc is also not constrained by a real 4-inhibitor, as shown above, we obtain a contradiction to (F4). Altogether, this establishes that $q_{qt} \neq a$ and, consequently, $q_{qt} \in K'$.

Replacement strategy. Recall that $\tilde{P} = \tilde{p}\tilde{s}q_{qt}$ is a path where $\tilde{p} = t$ or $\tilde{p} = q$, and $\tilde{s}q_{qt}$ is an edge that belongs to a real separating triangle. By (F3), the triangle sqt is vertex-disjoint from all other separating triangles. It follows that $\tilde{s} \notin \{s, t, q, d'\}$ (since $q_{qt} \notin \{s, t, q, d'\}$). Together with the degree bounds of q and t, this lets us narrow down the possible identities of $\tilde{P}$: if $\tilde{p} = q$, then $\tilde{P} = q\hat{q}q_{qt}$ since there cannot be an edge between $k \in K$ and $q_{qt} \in K'$. If $\tilde{p} = t$, then $\tilde{P} = tt'q_{qt}$ or $\tilde{P} = t\hat{t}q_{qt}$, where the latter case is only possible if $q_{qt} = q_{st}$ (since q_{qt} does not belong to the component of $\mathcal{G} \setminus \{s, t, q_{st}\}$ that contains s' and $\hat{t}$).

We choose the replacement $\varrho(e)$ for the edge e according to the following rules: note that q_{st} and q_{qt} are not necessarily unique. Moreover, $\tilde{P}$ is not necessarily

unique even for fixed choices of q_{st} and q_{qt}. If there exist valid choices of q_{st}, q_{qt}, and $\tilde{P} = \tilde{p}\tilde{s}q_{qt}$ (according to all the restrictions established for this case so far), such that $q_{st} = q_{qt}$ and $\tilde{P} = t\hat{t}q_{st}$ $(= t\hat{t}q_{qt})$, then we choose $\varrho(e) = st$. Otherwise, we choose $\varrho(e) = qt$. We will argue below that ϱ is acyclic, that is, after at most $i \in O(n)$ repetitions, we end up with a candidate edge $\varrho^i(e) \neq e$ for which the current Case 2.2.2.3.2 does not arise.

Algorithmic considerations. But before we dive into the proof of acyclicity, let us discuss the algorithmic complexity of computing $\varrho(e)$ (given the vertices a, b, c, d, s, t, q, d'), that is, a single iteration of the replacement strategy. We will show below that if $\varrho(e) = st$, then p_{st}, q_{st} is the unique separation pair created by collapsing st. Therefore, if collapsing st creates more than one separation pair, we choose $\varrho(e) = qt$. With the data structures that were set up in the very beginning of the proof (in particular, the list $\mathcal{P}(st)$ of separating pairs created by collapsing st), this decision can be made in constant time. Otherwise, we search for the vertex q_{st} in the list $\mathcal{P}(qt)$ of separation pairs created by collapsing qt. By (F9), this test can be done in constant time. If q_{st} does not appear in said list, then we set $\varrho(e) = qt$. Otherwise, we find $\hat{t}$ as the neighbor of t that follows d', s in the cyclic order of neighbors around t, and test whether $\hat{t}$ is adjacent to q_{st} and the edge $\hat{t}q_{st}$ is part of a real separating triangle. Due to the degree bounds for vertices of real separating triangles, this check can also be carried out in constant time. In case of a positive answer, we set $\varrho(e) = st$. Otherwise, we choose $\varrho(e) = qt$. It follows that $\varrho(e)$ can be computed in constant time overall.

Acyclicity. Let us now prove that ϱ is acyclic. There is no guarantee that Case 2.2.2.3.2 does not arise for the new choice $\varrho(e) = \bar{a}\bar{b}$ of e as well. However, we claim that for any separating triple $\bar{a}, \bar{b}, \bar{q}$ and for any path $\bar{P} = \bar{p}\bar{s}\bar{q}$ where $\bar{p} = \bar{a}$ or $\bar{p} = \bar{b}$, where $\bar{p}\bar{s}$ is a real edge, and where $\bar{s}\bar{q}$ belongs to a real separating triangle, our choice of $\varrho(e)$ guarantees that the vertices of the component $\bar{K}'$ of $\mathcal{G} \setminus \{\bar{a}, \bar{b}, \bar{q}\}$ that contains the vertex $\bar{s}$ form a proper subset of K'. This claim implies that the replacement procedure may be iterated without cycling, since the original edge e is not incident to a vertex of K', but the replacement edge $\varrho(e)$ is incident to $t \in K'$. Thus, after at most a linear number of steps, we find an edge for which Case 2.2.2.3.2 does not arise.

If $\varrho(e) = qt$, the rules according to which $\varrho(e)$ is chosen directly imply our claim since, as discussed above, for any choice of q_{qt} and $\tilde{P}$, we have $\tilde{P} = q\hat{q}q_{qt}$ or $\tilde{P} = tt'q_{qt}$ and, hence, $\tilde{s} \in \{\hat{q}, t'\}$ belongs to the connected component of $\mathcal{G} \setminus \{q, t, q_{qt}\}$ that does not contain s. So assume our choice was $\varrho(e) = st$ and, hence, we may assume that $q_{st} = q_{qt}$ and $\tilde{P} = t\hat{t}q_{qt}$ $(= t\hat{t}q_{st})$. We claim that this implies that p_{st}, q_{st} is the *unique* separation pair created by collapsing st. Assume otherwise and let p^2_{st}, q^2_{st} be another separation pair created by collapsing st. By Observation 31, we may

assume that $p^2_{st} = p_{st}$ and $q^2_{st} \neq q, d'$. We show that the vertices $\mathrm{N}_{\mathcal{G}}(t) \setminus \{s, t, q^2_{st}\}$ are connected in $\mathcal{G} \setminus \{s, t, q^2_{st}\}$, which yields a contradiction to Observation 30.

For an illustration, see Figure 62. Consider the cycle composed of the subpath $P_{q_{st}}$ of P_{sa} that connects $q_{st} = q_{qt}$ with a, some path between a and q whose interior vertices belong to K, and, finally, some path between q and q_{qt} $(= q_{st})$ via t' through the connected component of $\mathcal{G} \setminus \{q, t, q_{qt}\}$ that contains t'. This cycle does not contain s or t as a vertex. Hence, even if it contains q^2_{st} as a vertex, its remaining vertices are connected in $\mathcal{G} \setminus \{s, t, q^2_{st}\}$. Moreover, we have $q^2_{st} \notin \{q, q_{st}\}$ and, hence, due to the edges $\hat{t}q_{st}$ and qd', the vertices $\{s, d', q, t', \hat{t}\} \setminus \{s, t, q^2_{st}\}$ are connected in $\mathcal{G} \setminus \{s, t, q^2_{st}\}$. As discussed above, this yields the desired contradiction.

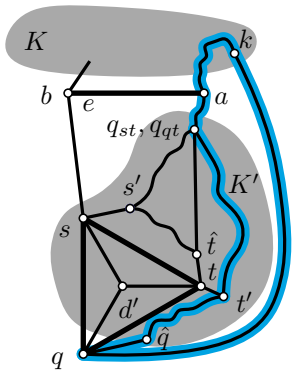

Figure 62: Illustrations for Case 2.2.2.3.2 in Lemma 38.

Altogether, this establishes that p_{st}, q_{st} is the unique separation pair created by collapsing st. Thus, it suffices to show that for any path $\bar{P} = \bar{p}\bar{s}q_{st}$ where $\bar{p} = s$ or $\bar{p} = t$, where $\bar{p}\bar{s}$ is real, and where $\bar{s}q_{st}$ belongs to a real separating triangle, the vertices of the component $\bar{K}'$ of $\mathcal{G} \setminus \{s, t, q_{st}\}$ that contains the vertex $\bar{s}$ form a proper subset of K'. This is clear, since q_{st} can belong to only one real separating triangle and its remaining two vertices (one of which is $\hat{t}$ since $\varrho(e) = st$), which are the only possible identities of $\bar{s}$, belong to the connected component of $\mathcal{G} \setminus \{s, t, q_{st}\}$ that does not contain q. Altogether, this shows that the replacement strategy is acyclic as claimed and, so, after a linear number of steps, we end up with a candidate edge of $\mathcal{G}$ for which Case 2.2.2.3.2 does not arise.

This finally concludes Case 2.2.2.3.2 and, by extension, the Cases 2.2.2.3, 2.2.2, 2.2, and 2. ◁

This concludes the first step of our proof that $\mathcal{G}''$ contains no adjacent separation pairs: we have established Property (J4) for the separation pairs that are already present in $\mathcal{G}'$, that is, simultaneously collapsing all edges in $\mathcal{K}'$ results in a biconnected graph $\mathcal{G}''$ in which no two vertices that correspond to a separation pair of $\mathcal{G}'$ are adjacent. It remains handle the second step, where we have to show that $\mathcal{G}''$ does not contain any *new* adjacent separation pairs.

4.5.7.3 $\mathcal{G}''$ contains no new adjacent separation pairs

To show that $\mathcal{G}''$ does not contain any adjacent separation pair whose vertices do not correspond to a separation pair in $\mathcal{G}'$, we recall several facts from Section 4.5.7.1:

- each separation pair of $\mathcal{G}'$ corresponds to a separation pair of $\mathcal{G}''$;
- to obtain the graph $\mathcal{G}''$, instead of collapsing the edges of $\mathcal{K}'$ directly in $\mathcal{G}'$, we may collapse the edges of $\mathcal{K}'$ in the decomposition of $\mathcal{G}'$ into its triconnected components, and then merge all pairs of identical $\mathcal{G}'$-virtual edges;
- $\mathcal{G}'$ has no parallel triconnected components, and its series components are 3-cycles; and
- each of the triconnected components of $\mathcal{G}'$ remains biconnected when collapsing its edges of $\mathcal{K}'$.

Towards a contradiction, assume that $\mathcal{G}'$ contains an edge $p'q'$ such that the vertices $\hat{p}, \hat{q}$ of $\mathcal{G}''$ that correspond to p' and q' form an adjacent separation pair. Let N denote the triconnected component of $\mathcal{G}'$ that contains the edge $p'q'$, which is not $\mathcal{G}'$-virtual since it belongs to $\mathcal{G}'$. We distinguish two cases depending on whether N is series or rigid.

Series components. Assume that N is series, i.e., N is a 3-cycle $p'q'x'$. At least one of the edges of N has to be $\mathcal{G}'$-virtual, say $p'x'$. Since p' and x' form a separation pair in $\mathcal{G}'$ that separates q' from the subgraph represented by $p'x'$, it follows that the vertex $\hat{p}$ and the vertex $\hat{x}$ corresponding to x' form a separation pair of $\mathcal{G}''$ that separates $\hat{q}$ from the subgraph of $\mathcal{G}''$ corresponding to the collapsed version of the subgraph of $\mathcal{G}'$ corresponding to $p'x'$. Since $\mathcal{G}''$ is biconnected, all vertices contained in this subgraph of $\mathcal{G}''$ remain connected to $\hat{x}$ in $\mathcal{G}'' \setminus \{\hat{p}, \hat{q}\}$. In case $q'x'$ is also $\mathcal{G}'$-virtual, symmetric arguments apply to the collapsed version of the subgraph represented by $q'x'$. Therefore, all vertices of $\mathcal{G}'' \setminus \{\hat{p}, \hat{q}\}$ are connected to $\hat{x}$; a contradiction to the assumption that $\hat{p}, \hat{q}$ are separating.

Rigid components. It remains to consider the case that N is rigid, i.e., $N \in \mathcal{R}$. By Property (J4) for $\mathcal{K}_N$, collapsing all edges of $\mathcal{K}_N$ in N results in a biconnected graph N' without adjacent separation pairs. Thus, since $\hat{p}, \hat{q}$ are adjacent, $N' \setminus \{\hat{p}, \hat{q}\}$ is connected. Moreover, since no separation pair of $\mathcal{G}'$ becomes adjacent when collapsing all edges of $\mathcal{K}' \supseteq \mathcal{K}_N$, the edge $\hat{p}\hat{q}$ cannot correspond to a $\mathcal{G}'$-virtual edge (as such an edge corresponds to a separation pair of $\mathcal{G}'$). Therefore, for each $\mathcal{G}'$-virtual edge e' in N', one of its endpoints remains in $N' \setminus \{\hat{p}, \hat{q}\}$. Since only one vertex is removed from the biconnected graph represented by e', it is still connected. In particular, all remaining vertices are still connected to the remaining endpoint of e'. Thus, they are also connected to all the remaining vertices in $N' \setminus \{\hat{p}, \hat{q}\}$ and, so, all vertices of $\mathcal{G}'' \setminus \{\hat{p}, \hat{q}\}$ are connected; a contradiction to the assumption that $\hat{p}, \hat{q}$ are separating.

Altogether, we have established Property (J4) for $\mathcal{K}'$.

4.5.8 Ensuring Property (J3)

It remains to show that $\mathcal{K}'$ satisfies Property (J3).

For the first statement of Property (J3), suppose that there is a k-inhibitor I in $\mathcal{G}'$ with respect to the set of real separating triangles of $\mathcal{G}'$ such that $\mathcal{K}'$ contains *more* than $k-3$ edges of I. Since the real separating triangles of $\mathcal{G}'$ are pairwise vertex-disjoint, we have by (J1) and (J2) that at most $\lfloor \frac{k}{2} \rfloor$ edges of I belong to $\mathcal{K}'$ and this number can be larger than $k-3$ only if $k=4$ and $\mathcal{K}'$ contains exactly two edges of I. However, this implies that the graph $\mathcal{G}''$ obtained by the simultaneous collapse of the edges in $\mathcal{K}'$ contains an adjacent separation pair formed by the two vertices u, v that correspond to the two collapsed edges of I , which yields a contradiction to (J4) of $\mathcal{K}'$. (The vertices u and v are indeed separating in $\mathcal{G}''$ because I is hyperseparating in $\mathcal{G}'$ and no edge between a vertex of $\mathrm{V}(\mathcal{G}') \setminus \mathrm{V}(I)$ and a vertex of $\mathrm{V}(I)$ can be in $\mathcal{K}'$ by (J1) and (J2) and since the separating triangles of $\mathcal{G}'$ are pairwise vertex-disjoint.)

For the second statement of Property (J3), suppose that there is some real k-inhibitor I in $\mathcal{G}'$ such that $\mathcal{K}'$ contains exactly $k-3$ edges of I and the triangle in $\mathcal{G}''$ corresponding to I is separating. We distinguish two cases regarding whether the edges of I belong to a single triconnected component of $\mathcal{G}'$ or not.

First, assume that all edges of I belong to a single (rigid) triconnected component R of $\mathcal{G}'$. By Property (J3) for R, the cycle I is not hyperseparating in R (with respect to the set of real separating triangles of R), but as a real inhibitor it is hyperseparating in $\mathcal{G}'$ (with respect to the set of real separating triangles of $\mathcal{G}'$). In particular, there is at least one vertex that is not interior to a separating triangle of $\mathcal{G}'$ on either side of I. This situation can arise only if all such vertices on one side of I belong to other triconnected components. Hence, two vertices p, q of I form a separation pair in $\mathcal{G}'$, which is nonadjacent by Property (J4) for $\mathcal{K}'$. Hence, R contains a chord pq of I that is $\mathcal{G}'$-virtual. However, after collapsing the edges of $\mathcal{K}'$, the vertices p and q become adjacent; a contradiction to Property (J4) for $\mathcal{K}'$.

It remains to consider the case that the edges of I belong to at least two distinct triconnected components of $\mathcal{G}'$. Then I passes through both vertices of a separation pair. By Property (J4), no separation pair of $\mathcal{G}'$ becomes adjacent and, hence, the length of I after collapsing all edges of $\mathcal{K}'$ is at least four (at least two edges in each triconnected component that I passes through); a contradiction.

We have shown that $\mathcal{K}'$ satisfies the Properties (J1)–(J4) for $\mathcal{G}$. By Lemma 39, this establishes the Properties (J1)–(J4) of the set $\mathcal{K} = \mathcal{K}' \cup \{e\}$ for $\mathcal{G}$, which concludes the proof. □

4.6 Stellation

Recall that our general assumption is that G is a 3-connected simple plane graph such that every vertex of a separating triangle has degree at most five. Moreover, all separating triangles of G are trivial, G is not a double kite and, hence, the separating triangles of G are pairwise vertex-disjoint (by Lemma 23 and Observation 25). Finally, we are assuming that G does not contain a subgraph isomorphic to a $\mathcal{G}_1$ or $\mathcal{G}_2$ as in Lemma 29. (We deal with these special cases in Section 4.8.)

Let $K \subset \mathrm{E}(G)$ be a set of edges to collapse as described in Theorem 33, and let G' denote the graph that results from simultaneously collapsing the edges from K in G. Then by Property (I3) of Theorem 33 the graph G' does not contain any separating triangle. Let G'' denote the graph that results from stellating all faces in G', that is, for every nontriangular face f of G' we insert a new vertex v_f into f and we add an edge between v_f and each vertex on ∂f. As discussed in Section 4.5.2, the following lemma is an easy consequence of Property (I4) of Theorem 33.

Lemma 42. *The graph G'' does not contain any separating triangle.*

Proof. Suppose for the sake of a contradiction that there exists a separating triangle T in G''. Given that G' does not contain any separating triangle, at least one vertex v_f of T is a stellation vertex, which has been inserted into a nontriangular face f of G'. As all neighbors of v_f are on ∂f, so are the remaining two vertices u and w of T. As T is separating, u and w are not consecutive along ∂f. Therefore, the edge uw of T is a chord of the cycle ∂f (so that uw is drawn in the exterior of f) in G', and so $\{u, w\}$ forms a separation pair in G'. However, according to Property (I4) of Theorem 33 the graph G' does not contain any adjacent separation pair, which leads to a contradiction. □

Therefore, we can apply Theorem 21 to G'' to obtain a Hamiltonian cycle C'' for G''. It remains to address the case that one or two edges of the outer face $T_\circ$ of G are prescribed.

Observation 43. *If any edge of $T_\circ$ is prescribed, then $T_\circ$ is also the outer face of G''.*

Proof. By Property (P3) we know that $T_\circ$ is also the outer face of G'. By Property (P2) we know that $T_\circ$ is a triangle and, therefore, it is not subdivided when going from G' to G''. □

By Observation 43 we can pass any possibly prescribed edge of $T_\circ$ to Theorem 21 so that the obtained Hamiltonian cycle C'' for G'' passes through the(se) prescribed edge(s).

4.7 Reconstructing collapses

As a final step to prove Theorem 22, it remains to handle the edge collapses, that is, to go back from the modified graph G'', from Section 4.6, to the original graph G. We proceed in three phases. In a first phase, we deal with some easy cases where we can reconstruct collapsed triangles, one by one, while maintaining a plane subhamiltonian cycle in the current graph. At the end of the first phase, we move on to the graph G, that is, we reconstruct all remaining triangles in one step. The previously plane subhamiltonian cycle then becomes a plane nonspanning cycle that visits all vertices of G but a pair of vertices from each of the remaining triangles (that have not been reconstructed during the first phase). We can classify these remaining triangles into five different types according to how they interact with the current cycle. During the second phase we then maintain this classification: although a triangle may change its type, it always remains one of these five types. As in the first phase, we process the triangles one by one. Processing a triangle amounts to extending the current cycle to visit the two missing vertices. At the end of Phase 2 we either have a plane subhamiltonian cycle for G (and are done), or we are in a situation where all remaining triangles to be handled are of a very specific type with respect to the current cycle. These remaining triangles are then processed during a third phase, at the end of which we obtain the desired plane subhamiltonian cycle for G.

During the whole reconstruction process, we modify the current cycle in specific ways only. In particular, we only modify edges of the cycle that share at least one endpoint with a separating triangle of G (including vertices that result from the collapse of an edge in K). By Observation 43 and Property (P3) this assertion suffices to ensure that prescribed edges on the outer face of G (if any) are part of the cycle that is constructed.

Notation. Next we define some notation for of our reconstruction algorithm. Let $\mathcal{T}$ with $|\mathcal{T}| = s$ denote the set of separating triangles of G. We will incorporate the vertices of the triangles of $\mathcal{T}$ in some order $T_1, \dots, T_s$. This linearly ordered sequence is not fixed in advance. Instead, we will iteratively specify the order as we go and describe how to select the next triangle to be processed.

Specifically for Phase 1, for $i \in \{1, \dots, s\}$, let $e_i = u_i v_i$ denote the edge of T_i in K, let t_i denote the interior vertex of $G^-_{T_i}$, let $\overline{u}_i$ denote the vertex of G' that corresponds to the collapsed edge e_i, and let w_i denote the vertex in $\mathrm{V}(T_i) \setminus \mathrm{V}(e_i)$.

Recall that the collapse operation merges parallel edges in order to ensure that its output is always a simple graph, cf. Section 4.4. This merging step is justified since the operation is only applied to collapsible edges, which ensures that whenever parallel edges are created during a collapse, these edges form a nonseparating 2-cycle. We define a sequence of plane graphs G_i, for $i = 0, \dots, s$, where the graph G_i is obtained by collapsing the edges $e_{i+1}, \dots, e_s$ in G, but *without* merging the created

parallel edges. Note that $G_s = G$. The graph G_0 corresponds to G', but potentially with some parallel edges that form nonseparating bigons. The advantage of not merging edges is that for each pair G_i, G_j, the faces of G_i and G_j that are noninterior to separating triangles are in one-to-one correspondence, which allows us to develop a consistent labeling scheme.

In Phase 1, we start from G_0 and maintain a subhamiltonian plane cycle C_i in G_i, for increasing $i = 0, \ldots$. To incrementally maintain these cycles, we impose an additional restriction on the cycles C_i that we call unichordality, which is defined in the next paragraph.

The vertices in $\{\overline{u}_j : j > i\} \cup \{w_j : j > i\}$ as well as the triangles T_j, for $i < j \leq s$, are called *relevant* for G_i and for C_i. A subhamiltonian plane cycle C_i for G_i consists of two possible types of edges: *proper* edges of G_i and *chords* of faces of G_i. Here, a *chord* of a face f is a chord of the cycle ∂f that is drawn inside f. Denote by $G_i^{\otimes}$ the *chordal completion* of G_i, that is, the graph (drawn in general with crossings) obtained from G_i by adding as edges all chords of the faces of G_i. Note that if G_i is 2-connected only, adding the chords in order to obtain $G_i^{\otimes}$ might create additional multiple edges, i.e., $G_i^{\otimes}$ might not be simple even if G_i is simple. We are looking for a *plane* Hamiltonian cycle in $G_i^{\otimes}$, that is, a Hamiltonian cycle that does not use two chords of $G_i^{\otimes}$ that cross each other. Specifically, we are interested in plane Hamiltonian cycles with the following property:

Definition 44. A cycle in $G_i^{\otimes}$ is *unichordal* if it is a plane Hamiltonian cycle that uses at most one chord of every face of G_i.

Lemma 45. *There exists a unichordal subhamiltonian cycle C_0 for G_0.*

Proof. Take the Hamiltonian cycle C'' for G'' obtained as described in Section 4.6. Removing all stellation vertices from C'' we obtain a subhamiltonian plane cycle C' for G'. The cycle C' may be interpreted as a plane subhamiltonian cycle C_0 of G_0. To this end, we simply replace each proper edge e of C' of which there exists multiple proper copies in G_0 with any one of these copies. We claim that C_0 is unichordal for G_0.

Every stellation vertex v_f has degree two in C'', and both neighbors a and b of v_f lie on ∂f. Thus, there are two options: Either a and b are adjacent in ∂f, in which case they are connected by a proper edge of G_0 in C_0; or a and b are not adjacent in ∂f, in which case they are connected by a chord of the face corresponding to f in C_0. Therefore, for every face f of G', the corresponding face of G_0 has at most one chord in C_0 since f has a unique stellation vertex v_f. □

4.7.1 Reconstruction Phase 1: the easy cases

We start with G_0 and its subhamiltonian unichordal cycle C_0 in $G_0^{\otimes}$ that is guaranteed by Lemma 45. In each step of Phase 1, we have a current subhamiltonian cycle C_{i-1} for $G_{i-1}^{\otimes}$ and we need to select a relevant triangle to be reconstructed next so that

we can extend the current subhamiltonian cycle to also visit the vertices of the reconstructed triangle. For notational convenience we denote the candidate triangle to be reconstructed by T_i, although in principle we may choose any of the not yet reconstructed triangles. Recall that since each of u_i, v_i, and w_i have degree three in $G^-_{T_i}$, we have $\deg_{G'}(w_i) \leq 3$ and $\deg_{G'}(\overline{u}_i) \leq 5$.

If the current cycle C_{i-1} is unichordal, then, for *any* choice of T_i, it is quite easy to find a subhamiltonian cycle C_i that also visits the vertices of T_i. However, to iterate the procedure, it is also necessary to maintain unichordality of the current cycle, which does not always seem easy to accomplish. The goal of Phase 1 is to handle some easy cases where the next triangle T_i can be chosen such that a subhamiltonian unichordal cycle C_i is easily obtained. At the end of Phase 1, we have chosen and reconstructed the triangles $T_1, T_2, \ldots, T_k$ and a subhamiltonian unichordal cycle C_k. The remaining cases are more challenging and will be the subject of the Phases 2 and 3.

Each case in Phase 1 is phrased as a lemma in the following that lists a condition on the (any) triangle T_i in relation to the current cycle C_{i-1} under which T_i can be reconstructed and C_{i-1} can be locally extended to a subhamiltonian unichordal cycle C_i in $G_i^{\otimes}$. Note that in order to show unichordality of C_i we need to consider *new chords* only, that is, chords in C_i that are not in C_{i-1} already.

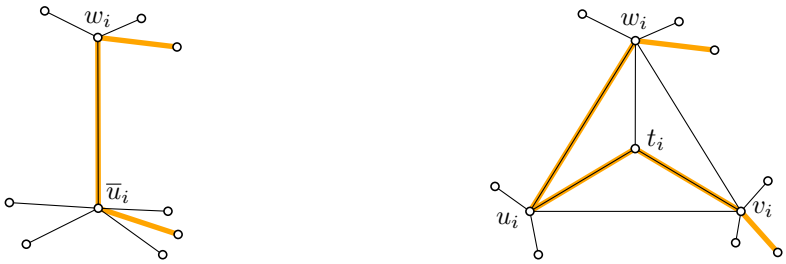

(a) C_{i-1} in $G^{\otimes}_{i-1}$ (b) C_i in $G_i^{\otimes}$

Figure 63: Reconstruction of T_i if C_{i-1} uses the edge $\overline{u}_i w_i$.

Lemma 46. *If $\overline{u}_i w_i \in \mathrm{E}(C_{i-1})$, then there is a subhamiltonian unichordal cycle C_i in $G_i^{\otimes}$.*

Proof. We obtain C_i from C_{i-1} by replacing the edge $\overline{u}_i w_i$ by one of the paths $u_i t_i v_i w_i$ or $v_i t_i u_i w_i$, depending on the other neighbor of $\overline{u}_i$ in C_{i-1} (which is adjacent to at least one of u_i or v_i in $G_i^{\otimes}$); see Figure 63. In particular, we prefer a proper edge over a chord, that is, if $\overline{u}_i$ is incident to a proper edge $\overline{u}_i z, z \neq w_i$ along C_{i-1}, we replace $\overline{u}_i w_i$ such that C_i uses the proper edge corresponding to $\overline{u}_i z$. Thus, no new chord is introduced and therefore C_i is unichordal in $G_i^{\otimes}$. □

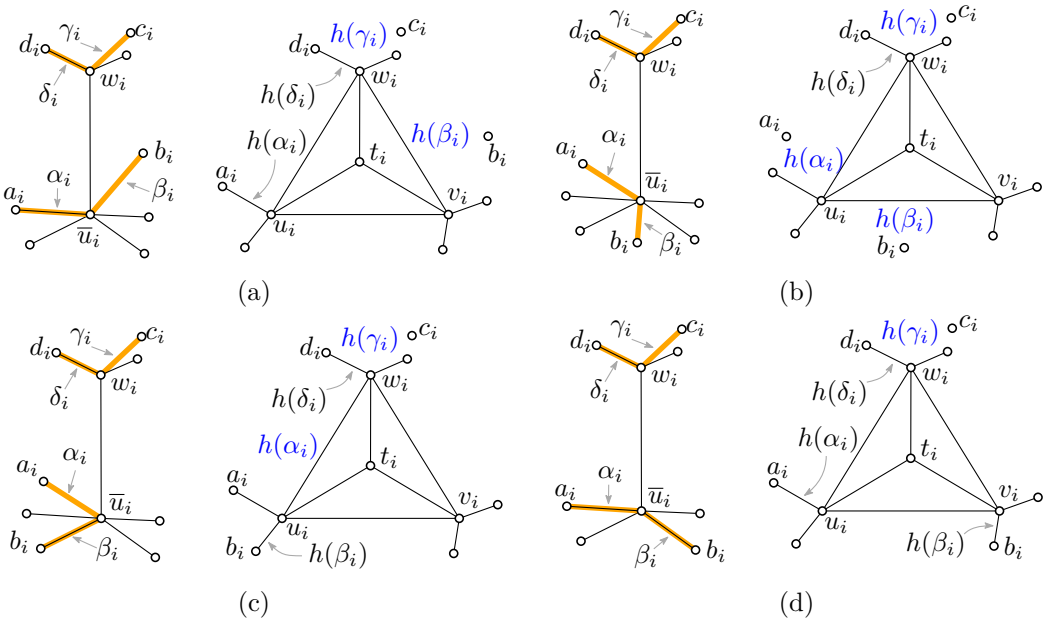

Figure 64: Illustrations for the normalization in Phase 1. In (a), (b), and (c), the orientation and labeling is uniquely defined. In (d), both u_i and v_i are incident to the same number of elements in $\{h(\alpha_i, h(\beta_i)\}$, and both these elements are edges. Hence, the roles of u_i and v_i could be exchanged.

Normalization. Using Lemma 46 we may assume that $\overline{u}_j w_j \notin \mathrm{E}(C_{i-1})$, for all $j \geq i$. In other words, C_{i-1} visits $\overline{u}_j$ and w_j independently. Let us apply some normalization in order to name the vertices of relevant triangles consistently, for illustrations refer to Figure 64.

Let a_i and b_i denote the neighbors of $\overline{u}_i$ in C_{i-1}. We denote the edges $\overline{u}_i a_i$ and $\overline{u}_i b_i$ in C_{i-1} by α_i and β_i, respectively. If $\alpha_i \in \mathrm{E}(G_{i-1})$, we define $\alpha'_i = \alpha_i$. Otherwise, we define α'_i to be the face in which the edge α_i is drawn. Analogously, we define β'_i. Finally, denote by $h(\alpha_i)$ and $h(\beta_i)$ the element (edge or face) of G_i that corresponds to α'_i and β'_i, respectively. For instance, if α_i is a chord in a face f of G_{i-1}, then $h(\alpha_i)$ denotes the corresponding face of G_i.

Exchange the roles of u_i and v_i if necessary so that u_i is incident to at least as many elements of $\{h(\alpha_i), h(\beta_i)\}$ as v_i (see Figures 64b and 64c), and if both belong to the same number of elements, then u_i is incident to at least as many edges in $\{h(\alpha_i), h(\beta_i)\}$ as v_i (see Figure 64a). Reflect the drawing if necessary to ensure that the triple $u_i v_i w_i$ is oriented anticlockwise. Exchange the roles of a_i and b_i if necessary so that the edges $\overline{u}_i w_i$, α_i, and β_i appear in this order anticlockwise around $\overline{u}_i$. As a consequence of this normalization, u_i is a vertex of $h(\alpha_i)$. Moreover, if $h(\alpha_i) \notin \mathrm{E}(G_i)$ and $h(\beta_i) \in \mathrm{E}(G_i)$, then $h(\beta_i) = u_i b_i$ (since $h(\beta_i) = v_i b_i$ would imply that u_i is incident to at most one element of $\{h(\alpha_i), h(\beta_i)\}$, namely the face $h(\alpha_i)$, while v_i is incident to at least one element of $\{h(\alpha_i), h(\beta_i)\}$, namely the edge $h(\beta_i)$),

see Figure 64c.

(Note that this normalization may prescribe different orientations for different triangles. This is no problem since we consider the triangles individually.)

Finally, let c_i and d_i denote the two (distinct) neighbors of w_i in C_{i-1}, denote the edges c_iw_i and d_iw_i by γ_i and δ_i, respectively, and exchange the roles of c_i and d_i if necessary so that $\overline{u}_iw_i$, γ_i, and δ_i appear in this order anticlockwise around w_i. The elements $h(\gamma_i)$ and $h(\delta_i)$ are defined analogously to $h(\alpha_i)$.

The vertices u_i, v_i, and w_i in G_i are surrounded by altogether up to six neighbors and up to six faces. Starting from the face f_ℓ to the left of the edge u_iw_i and then in clockwise order we denote the faces around T_i in G_i by f_ℓ, f_t, f_r, f_{br}, f_b, and $f_{b\ell}$; see Figure 65b. The faces f_ℓ, f_r, and f_b always exist. On the other hand, $f_{b\ell}$, f_t, and f_{br} exist only if the degree of the corresponding vertex of T_i is five in G_i. Further, note that some of these three faces (if they exist) might be bigons (if their bounding edges where merged due to the collapse of some $e_j \in K, j > i$. We use g_ℓ, g_t, g_r, g_{br}, g_b, and $g_{b\ell}$ to denote the face of G_{i-1} that corresponds to f_ℓ, f_t, f_r, f_{br}, f_b, and $f_{b\ell}$, respectively (if it exists); see Figure 65a. Note that g_b is a bigon if f_b is a triangle.

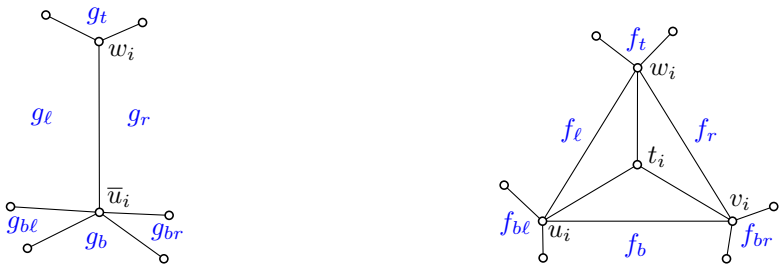

Figure 65: Labeling the faces around T_i in (a) G_{i-1} and (b) G_i.

Let $k \in \{\ell, t, r, br, b, b\ell\}$. We define $E_k^i = \mathrm{E}(\partial f_k)$ and we use X_k^i to denote the set of chords of f_k. Moreover, we define $\bigotimes_k^i = E_k^i \cup X_k^i$. The sets E_k^{i-1}, X_k^{i-1}, and $\bigotimes_k^{i-1}$; and E_k, X_k, and $\bigotimes_k$ are defined analogously except that the role of f_k is taken over by the face corresponding to f_k in G_{i-1} and G, respectively (in the former case the corresponding face is g_k). If a face does not exist, the corresponding edge set is a singleton edge and its set of chords is empty. For instance, if $\deg_{G_i}(w_i) = 4$, then $\bigotimes_t^i = E_t^i \subseteq E_\ell^i \cap E_r^i$ and it consists of the single edge of G_i that connects w_i to a vertex outside of T_i.

Let us proceed by reconstructing further triangles where the current cycle may easily be extended.

Lemma 47. *If $\beta_i \in \bigotimes_r^{i-1} \cup \bigotimes_{br}^{i-1} \cup X_b^{i-1}$, then there exists a subhamiltonian unichordal cycle C_i for $G_i^{\otimes}$.*

Proof. We obtain C_i from C_{i-1} by replacing the path $a_i\overline{u}_ib_i$ by the path $a_iu_it_iv_ib_i$; see Figure 66. If α_i is a chord, we draw the edge a_iu_i in $h(\alpha_i)$; and if α_i is an edge,

we use the corresponding proper edge $h(\alpha_i)$ of G_i. We proceed analogously, with β_i and v_ib_i.

If $\beta_i \in \mathrm{E}(G_{i-1})$, then no new chord is introduced and, hence, C_i is unichordal in $G_i^{\otimes}$. So suppose that $\beta_i \in X_r^{i-1} \cup X_{br}^{i-1} \cup X_b^{i-1}$. Then the edge β_i is the only chord of its face of G_{i-1} in C_{i-1} (because C_{i-1} is unichordal). Therefore, the edge b_iv_i retains that role for the corresponding face $h(\beta_i)$ of G_i in C_i. It follows that the cycle C_i is unichordal in $G_i^{\otimes}$. □

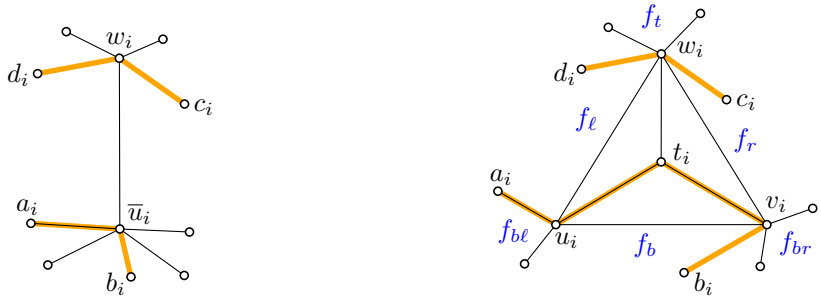

(a) C_{i-1} in $G_{i-1}^{\otimes}$ (b) C_i in $G_i^{\otimes}$

Figure 66: Vertex b_i is adjacent to v_i in $G_i^{\otimes}$ but not to u_i in G_i.

By Lemma 47 we may assume that $\alpha_i, \beta_i \in \bigotimes_{b\ell}^{i-1} \cup X_\ell^{i-1}$, for every possible choice of T_i from $T_i, \ldots, T_s$.

Lemma 48. *If $\gamma_i \in X_r^{i-1}$, then there exists a subhamiltonian unichordal cycle C_i for $G_i^{\otimes}$.*

Proof. We obtain C_i from C_{i-1} by replacing the path $c_iw_id_i$ by the path $c_iv_it_iw_id_i$; see Figure 67. If δ_i is a proper edge, we use the corresponding edge $h(\delta_i)$. In this case, no new chord is created and, hence, C_i is unichordal for $G_i^{\otimes}$. Otherwise, we draw w_id_i in $h(\delta_i)$, so that the role of the chord δ_i is taken over by w_id_i and, hence, C_i is unichordal for $G_i^{\otimes}$. Note that c_iv_i might be a proper edge of G_i even tough γ_i is a chord by assumption. In this case, we choose the corresponding edge on the boundary of f_r and, thus, no new chord is created. Otherwise, if $c_iv_i \notin \mathrm{E}(G_i)$, we draw it in $h(\gamma_i)$, so that the role of γ_i is taken over by c_iv_i. Therefore, C_i is unichordal in $G_i^{\otimes}$. □

Lemma 49. *If (1) $\alpha_i \in X_\ell^{i-1}$, (2) $\{\gamma_i, \delta_i\} \subset \bigotimes_t^{i-1}$, and (3) $\{\gamma_i, \delta_i\} \cap X_t^{i-1} \neq \emptyset$, then there exists a subhamiltonian unichordal cycle C_i for $G_i^{\otimes}$.*

Proof. We obtain C_i from C_{i-1} by (1) replacing the path $c_iw_id_i$ by the edge c_id_i and (2) replacing the edge $a_i\overline{u}_i$ by the path $a_iw_iv_it_iu_i$; see Figure 68. If $\beta_i \in \mathrm{E}(G_{i-1})$, we use the corresponding edge $h(\beta_i)$. Otherwise, we draw b_iu_i in $h(\beta_i)$ $(= f_{b\ell}$, by unichordality of $C_{i-1})$, where it takes over the role of the chord β_i as the only chord

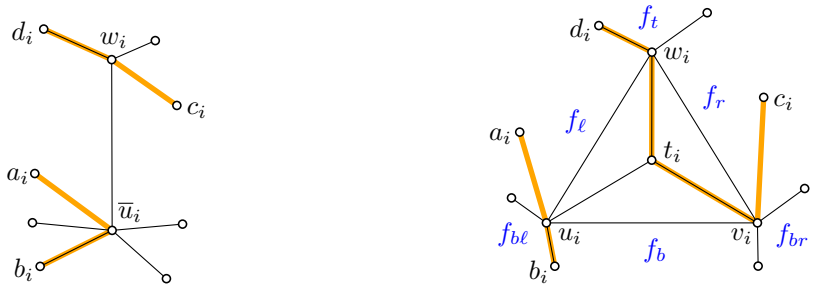

(a) C_{i-1} in $G_{i-1}^{\otimes}$ (b) C_i in $G_i^{\otimes}$

Figure 67: If c_iw_i is a chord in f_r, then we reroute the cycle via v_i.

of its face. By assumption, α_i is a chord. So if $a_iw_i \notin \mathrm{E}(G_i)$, we draw it in $h(\alpha_i)$ $(= f_\ell)$, where it takes over the role of α_i as the only chord of its face. Otherwise, we simply use the proper edge a_iw_i on the boundary of f_ℓ, so that no new chord is created. Finally, by assumption and unichordality of C_{i-1}, exactly one of γ_i, δ_i is a chord of g_t and the other is a proper edge of ∂g_t. By drawing c_id_i in $h(\gamma_i)$ or $h(\delta_i)$, depending on which edge is the chord of g_t, the edge c_id_i takes over the role of the previous chord as the only chord of its face (or it is a proper edge of the boundary of f_t). Altogether, no new chord is created and, thus, C_i is unichordal in $G_i^{\otimes}$. □

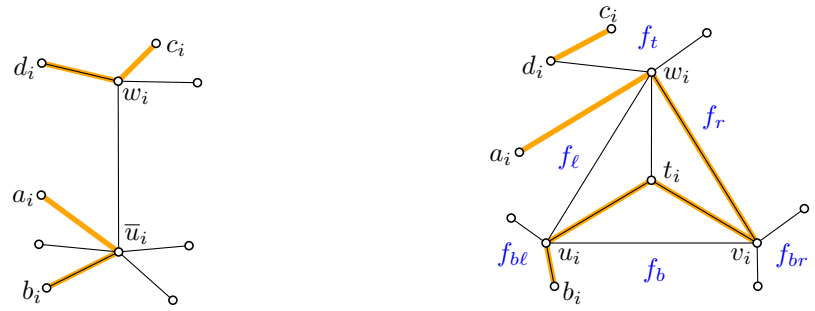

(a) C_{i-1} in $G_{i-1}^{\otimes}$ (b) C_i in $G_i^{\otimes}$

Figure 68: If a_iu_i is a chord in f_ℓ and we can shortcut $c_iw_id_i$, then we reroute a_iu_i via w_i.

In the same way we obtain the following symmetric statement.

Lemma 50. *If (1) $\delta_i \in X_\ell^{i-1}$, (2) $\{\alpha_i, \beta_i\} \subset \bigotimes_{b\ell}^{i-1}$, and (3) $\{\alpha_i, \beta_i\} \cap X_{b\ell}^{i-1} \neq \emptyset$, then there exists a subhamiltonian unichordal cycle C_i for $G_i^{\otimes}$.* □

Classification. The five statements in Lemma 46–50 summarize the cases that we handle in Phase 1 of our algorithm. All triangles that satisfy any of the listed conditions are selected and reconstructed in an arbitrary order. Hence for Phase 2 of the algorithm we may suppose that none of the remaining relevant triangles (if any)

satisfy any of the conditions listed in Lemma 46–50. As a result we can classify how the relevant triangles are traversed by the current subhamiltonian unichordal cycle C_{i-1} at the end of Phase 1.

Lemma 51. *If Phase 1 of the algorithm ends with a subhamiltonian unichordal cycle C_{i-1} in $G^{\otimes}_{i-1}$, for some $i \leq s$, then every remaining relevant triangle T_i is of one of the following types (see Figure 69):*

(S1) $\alpha_i \in X^{i-1}_{\ell}$, $\beta_i \in \bigotimes^{i-1}_{b\ell}$, *and* $\{\gamma_i, \delta_i\} \subseteq E^{i-1}_{t}$,

(S2) $\delta_i \in X^{i-1}_{\ell}$, $\gamma_i \in \bigotimes^{i-1}_{t}$, *and* $\{\alpha_i, \beta_i\} \subseteq E^{i-1}_{b\ell}$, *or*

(S3) $\{\alpha_i, \beta_i\} \subseteq \bigotimes^{i-1}_{b\ell}$ *with at least one of* α_i *or* β_i *in* $E^{i-1}_{b\ell}$; *and* $\{\gamma_i, \delta_i\} \subseteq \bigotimes^{i-1}_{t}$ *with at least one of* γ_i *or* δ_i *in* E^{i-1}_{t}.

Proof. By Lemma 47 we know that $\{\alpha_i, \beta_i\} \subseteq \bigotimes^{i-1}_{b\ell} \cup X^{i-1}_{\ell}$, and by Lemma 48 we know that $\{\gamma_i, \delta_i\} \subseteq \bigotimes^{i-1}_{t} \cup X^{i-1}_{\ell}$. We distinguish three cases.

Case 1: $\alpha_i \in X^{i-1}_{\ell}$. Then the unichordality of C_{i-1} implies that it uses no other chord of f_ℓ and, thus, $\beta_i \in \bigotimes^{i-1}_{b\ell}$ and $\{\gamma_i, \delta_i\} \subseteq \bigotimes^{i-1}_{t}$. By Lemma 49 we have $\{\gamma_i, \delta_i\} \subseteq E^{i-1}_{t}$ and, hence, the triangle T_i is of type (S1). ◁

Case 2: $\delta_i \in X_\ell$. Symmetrically to Case 1 we conclude that T_i is of type (S2) for C_{i-1}. ◁

If neither of the two cases above applies, the combination of all conditions yields $\{\alpha_i, \beta_i\} \subseteq \bigotimes^{i-1}_{b\ell}$ and $\{\gamma_i, \delta_i\} \subseteq \bigotimes^{i-1}_{t}$. The unichordality of C_{i-1} implies that T_i is of type (S3). □

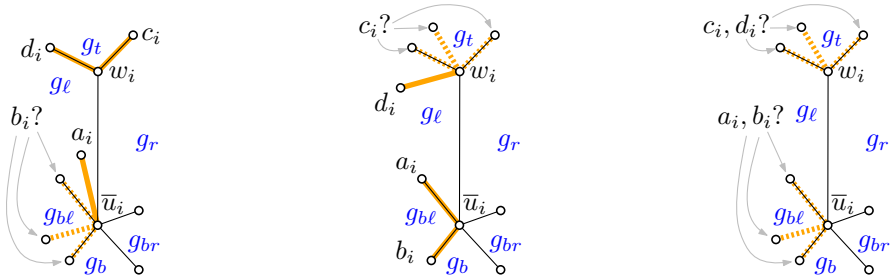

(a) Type (S1) (b) Type (S2) (c) Type (S3)

Figure 69: The three different types of relevant triangles in G_i at the end of Phase 1.

Transition to G. For the remainder of the algorithm, we will work in G and $G^{\otimes}$, rather than in the collapsed graphs G_i and $G^{\otimes}_i$. This simplifies our notation scheme: by 3-connectivity of G, the graph $G^{\otimes}$ is simple. Hence, given a vertex u and its two neighbors a, b in a cycle of $G^{\otimes}$, the edges au and ub are uniquely defined.

The classification in Lemma 51 allows us to interpret a final cycle $C_k, k < s$ of Phase 1 as an *almost spanning cycle* in $G^{\otimes}$, meaning that it passes through all vertices but the pairs t_j and v_j, for $k+1 \leq j \leq s$. To this end, we simply replace every vertex $\overline{u}_j$ along C_k with the corresponding vertex u_j.

In the following lemma we restate the classification of Lemma 51 in terms of this interpretation of C_k in $G^{\otimes}$. Since we are no longer transitioning between graphs G_{i-1} and G_i in every step, our earlier notation is no longer suited for this task and, hence, we will slightly repurpose it (intuitively, the meaning remains the same): when we refer to a relevant triangle T_i of $G^{\otimes}$ as being of a certain type with respect to some cycle C_j, then, in the description of the type, the vertices u_i, v_i, w_i of the triangle T_i are assumed to appear in this counterclockwise order; and the cycle C_j is visiting u_i and w_i, but not v_i. The neighbors of u_i along C_j are denoted by a_i and b_i and the (unique) edges u_it_i, u_ia_i, and u_ib_i appear in this counterclockwise order around u_i. Further, the neighbors of w_i along C_j are denote by c_i and d_i and the edges w_it_i, w_ic_i, and w_id_i appear in this counterclockwise order around w_i. We also reuse f_k for $k \in \{\ell, t, r, br, b, b\ell\}$ to denote the faces of G that surround T_i as in the earlier definition of these faces in G_i.

Lemma 52. *If Phase 1 of the algorithm ends with an almost spanning unichordal cycle C_k in $G^{\otimes}$ for some $k < s$, then every remaining relevant triangle T_i is of one of the following types with respect to C_k (see Figure 70):*

(T1) $a_iu_i \in X_\ell$, $b_iu_i \in \bigotimes_{b\ell}$, *and* $\{c_iw_i, d_iw_i\} \subseteq E_t$,

(T2) $d_iw_i \in X_\ell$, $c_iw_i \in \bigotimes_t$, *and* $\{a_iu_i, b_iu_i\} \subseteq E_{b\ell}$, *or*

(T3) $\{a_iu_i, b_iu_i\} \subseteq \bigotimes_{b\ell}$ *with at least one of* a_iu_i *or* b_iu_i *in* $E_{b\ell}$ *and* $\{c_iw_i, d_iw_i\} \subseteq \bigotimes_t$ *with at least one of* c_iw_i *or* d_iw_i *in* E_t. □

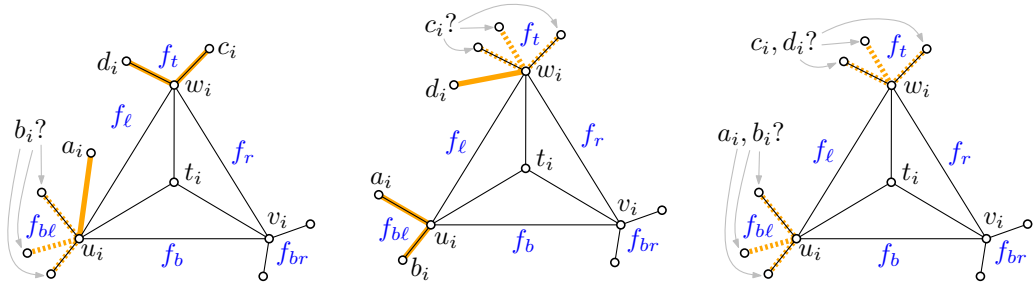

(a) Type (T1) (b) Type (T2) (c) Type (T3)

Figure 70: The three different types of relevant triangles in $G^{\otimes}$ at the end of Phase 1.

We remark that in the upcoming phases of the algorithm, we will treat symmetric cases explicitly. In particular, (T1) and (T2) are symmetric, and (T3) is selfsymmetric. Hence, the orientation of the drawing is globally consistent (which was not the case in Phase 1).

4.7.2 Reconstruction Phase 2: the hard cases

In this phase, we handle the remaining relevant triangles that do not classify as easy cases, as discussed in the previous section. In analogy to Phase 1, we use the term *relevant* with respect to a cycle C_i to denote those triangles $\mathcal{T} \setminus \{T_1, \ldots, T_i\}$, for which not all vertices are visited by C_i. Also the vertices of a relevant triangle are called relevant, regardless of whether or not C_i visits them.

The main challenge in this phase is not so much to handle any single triangle but rather to handle it *and* at the same time maintain the property that all relevant triangles interact with the current cycle according to the classification in Lemma 52. In fact, we do not know how to always achieve both of these goals. Instead, we introduce two relaxations of our classification: First, we allow two more types of triangles, (T4) and (T5) defined below, which may arise after reconstructing some relevant triangles (which, initially, are of the types (T1), (T2), and (T3) only). Second, we allow (up to) one exceptional triangle to leave the classification, in a controlled way. Whenever such an exceptional triangle appears, we handle it next, thereby making it irrelevant for the correspondingly updated cycle. For such a scheme to work, we need to show that every step in Phase 2 creates at most one exceptional triangle.

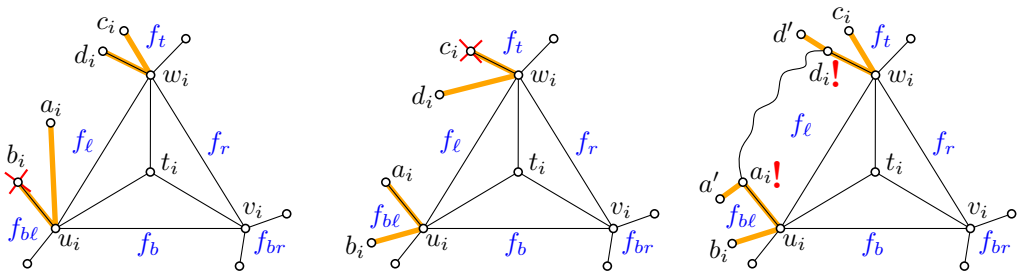

(a) Type (T4) (b) Type (T5) (c) Type (T6)

Figure 71: Three more types of relevant triangles in $G^{\otimes}$ in Phase 2.

Additional types. As in the first phase, we aim to maintain and iteratively extend a cycle C, but this time in $G^{\otimes}$. Also, in contrast to the first phase we do not insist that the cycle remains unichordal. Instead, we demand that (almost) every triangle T_i that is relevant with respect to C is of one of the types (T1)–(T3) (described in Lemma 52), or of one of the following two types (T4) and (T5) with respect to C; see also Figure 71. Note that the types (T4) and (T5) are symmetric to each other.

(T4) $a_iu_i \in X_\ell$, $c_iw_i \in X_t$, $b_iu_i, d_iw_i \in E_\ell$, and b_i is irrelevant for C.

(T5) $d_iw_i \in X_\ell$, $b_iu_i \in X_{b\ell}$, $a_iu_i, c_iw_i \in E_\ell$, and c_i is irrelevant for C.

Phase 2 will end with a cycle C for which every remaining relevant triangle T_i

has the following type (T6) with respect to C. Note that type (T6) is selfsymmetric, and a specialization of (T3).

(T6) $a_iu_i, d_iw_i \in E_\ell$, $b_iu_i \in X_{b\ell}$, $c_iw_i \in X_t$, and both a_i and d_i are relevant for C, with $a_ia' \in X_{b\ell}$ and $d_id' \in X_t$, where a' and d' is the other (not in T_i) neighbor of a_i and d_i, respectively, on C (in particular, $a_i \neq d_i$).

Slide operations. Whenever we modify the cycle, the type of a triangle may change. However, we change the cycle in a very controlled way only. For once, going from C_{i-1} to C_i, thereby making some triangle T_i part of the cycle, only edges in the neighborhood of T_i may change. Moreover, we allow edges to change by the following three *slide* operations only; see Figure 72 for examples.

- In a *chord slide* at x, a chord xy of some face f is replaced by either a proper edge $xy' \in E$ of G along ∂f (*chord-to-edge slide*) or by another chord xy' of f (*chord-to-chord slide*);
- in an *edge slide* at x, a proper edge xy is replaced by a chord xy' of a face with xy on its boundary (*edge-to-chord slide*) or by another proper edge xy' so that xy and xy' are consecutive in the circular order of edges around x in G (*edge-to-edge slide*);
- finally, a *cone slide* at x combines an edge slide at x that replaces a proper edge xy and a chord-to-edge slide at x that replaces a chord with xy.

A triangle T_i is *normal* for a cycle C in $G^\otimes$ if it is of one the five types (T1)–(T5) with respect to C. Recall that our plan is to maintain the types (T1)–(T5) for almost all the remaining relevant triangles—we only allow one triangle that is not normal with respect to the current cycle, and if such an exceptional triangle exists, it is handled next. When the next triangle T_i is normal for the current cycle C_{i-1}, then its type precisely describes the interaction of C_{i-1} with T_i. If T_i is not normal, then the interaction of C_{i-1} with T_i may be derived from its type in the previous cycle C_{i-2}, given that only a limited number of slide operations is used to modify the neighborhood of T_i when going from C_{i-2} to C_{i-1}. Hence, in Phase 2, we will not only keep track of the current cycle C_{i-1}, but also its predecessor.

Lemma 53. *Let C_k be the almost spanning cycle obtained during Phase 1, cf. Lemma 52. Let $\mathcal{T}_k$ denote the triangles of G that are still relevant for C_k. We can extend C_k to an almost spanning cycle C_ℓ for which all of its remaining relevant triangles $\mathcal{T}_\ell \subseteq \mathcal{T}_k$ (if any) are of type (T6).*

Proof. We inductively generate a sequence of almost spanning cycles $C_k, C_{k+1}, \ldots, C_\ell$ with $\ell \leq s$, and a sequences of triangles $T_{k+1}, T_{k+2}, \ldots, T_{\ell'}$ from $\mathcal{T}_k$ where $\ell' = \min\{\ell+1, s\}$. Denote $V_i = \mathrm{V}(G^-_{T_i}) = \{u_i, v_i, w_i, t_i\}$ and $V_i^- = \{t_i, v_i\}$. We ensure that for each index $i \in \{k, \ldots, \ell\}$ the following invariants are satisfied.

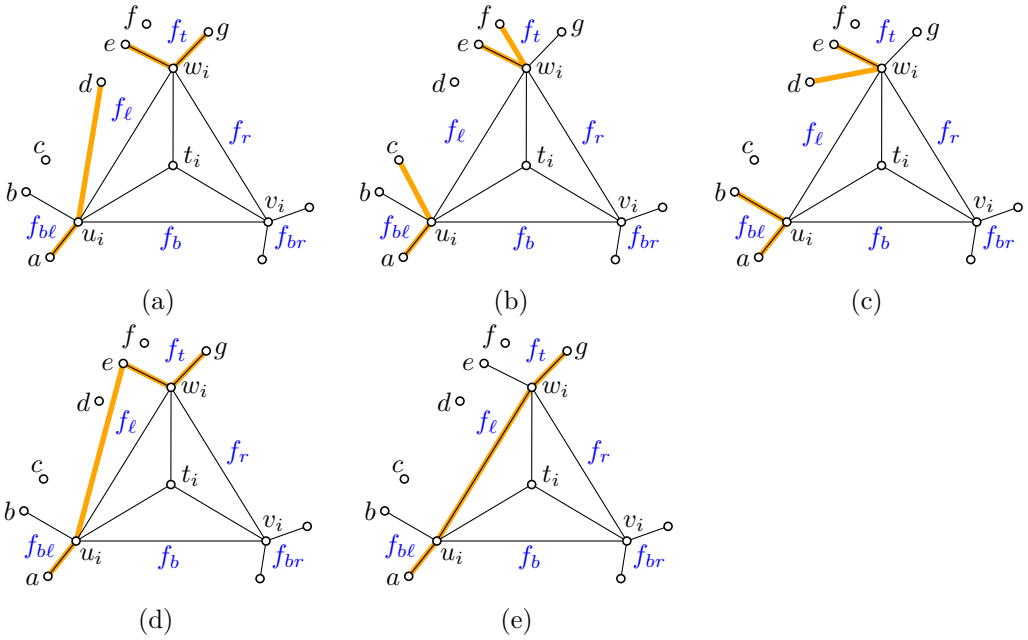

Figure 72: Examples of slides: From (a) to (b) there is a chord-to-chord slide ($u_id \to u_ic$) at u_i and an edge-to-chord slide ($w_ig \to w_if$) at w_i. From (b) to (c) there is a chord-to-edge slide ($u_ic \to u_ib$) at u_i and a cone slide (around the edge w_ie) at w_i. From (d) to (e) there is an edge-to-edge slide ($w_ie \to w_iu_i$) at w_i and a chord-to-edge slide from ($u_ie \to u_iw_i$). Note that the resulting undirected edge is the same for both slides.

(N1) The set of triangles that are still relevant for C_i is $\mathcal{T}_i$, where $\mathcal{T}_i = \mathcal{T}_{i-1} \setminus \{T_i\}$ if $k > 1$.

(N2) All triangles in $\mathcal{T}_i$ are normal for C_i, with the possible exception of T_{i+1}.

(N3) If $i > k$, to create C_i from C_{i-1}, some edges of C_{i-1} are deleted and some new edges are added. Each endpoint of an added cycle edge $e^+ \in \mathrm{E}(C_i) \setminus \mathrm{E}(C_{i-1})$ belongs to V_i or some neighbor of a vertex from V_i along the cycle C_{i-1}, that is, $e^+ \in G^{\otimes}[V_i \cup \mathrm{N}_{C_{i-1}}(V_i)]$. For the deleted edges $e^- \in \mathrm{E}(C_{i-1}) \setminus \mathrm{E}(C_i)$, we distinguish two cases: if $e^- \in \mathrm{E}(G)$, then we demand that e^- has at least one endpoint in V_i. If $e^- \notin \mathrm{E}(G)$, then e^- has at least one endpoint that belongs to a relevant triangle of $\mathcal{T}_{i-1}$ and e^- passes through a face with a vertex of V_i on its boundary.

(N4) If $i > k$, orient the cycles C_{i-1} and C_i counterclockwise and consider a directed edge $e = xy$ of C_{i-1} where $x \notin V_i$. Then the directed edge $e' = xy'$ of C_i is either unchanged (that is $y' = y$), or it is obtained from e by an edge slide or chord slide (possibly as part of a cone slide) at x. Analogously, the incoming edge zx of x in C_{i-1} is either unchanged in C_i or replaced by an edge $z'x$

obtained from zx by an edge slide or chord slide (possibly as part of a cone slide) at x. For examples, consider Figure 73 or any of the other figures in the remainder of Section 4.7.2.

(N5) If $i < s$, either T_{i+1} is normal with respect to C_i, or $i > k$ and the edges of C_i that are incident to a vertex of T_{i+1} are obtained from the edges of C_{i-1} that are incident to a vertex of T_{i+1} by performing a *single* edge slide or cone slide at a vertex of T_{i+1}, plus optionally a *single* chord slide at *another* vertex of T_{i+1}, for an example see Figure 73. Note that $\mathrm{E}(C_i) \triangle \mathrm{E}(C_{i-1})$ may also contain edges that are not incident to vertices of T_{i+1}. Such edges are not affected by (N5). (See for example Figure 74, where the transition from C_{i-2} to C_{i-1} affects not only T_i, but also w_j of T_j. Property (N5) holds, since only one vertex of T_i is affected by the transition, by a single edge slide.)

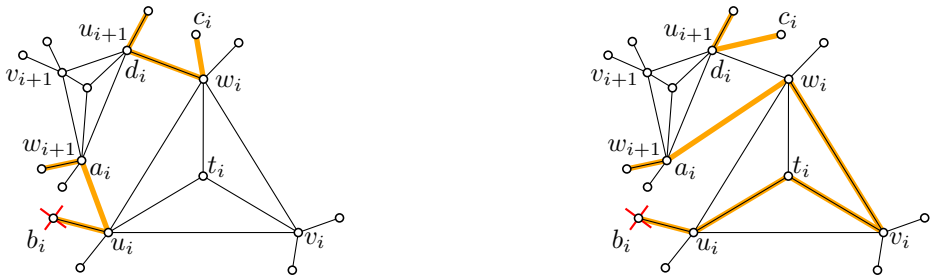

(a) C_{i-1} (b) C_i

Figure 73: For the cycle C_{i-1}, the triangle T_i is of type (T4), and the triangle T_{i+1} is of type (T2). The transition to C_i makes T_{i+1} special. Property (N5) is satisfied since the transition is perceived as a single edge slide at u_{i+1} and a single chord slide at w_{i+1} of T_{i+1}.

We start with the cycle C_k that is obtained during Phase 1, and some arbitrary triangle T_{k+1} of $\mathcal{T}_k$. By Lemma 52, all triangles of $\mathcal{T}_k$ are normal with respect to C_k. Hence, the Conditions (N1)–(N5) are satisfied. To compute, for some $i > k$, the cycle C_i and, if $i < s$, the triangle T_{i+1}, we only need the following inductively obtained objects: the cycle C_{i-1}, the triangle T_i, and, if $i > k+1$, the cycle C_{i-2}.

If the remaining relevant triangles $\mathcal{T}_{i-1}$ of the given cycle C_{i-1} are all of type (T6) with respect to C_{i-1} (in particular, this is the case if $i = s+1$), we set $\ell = i-1$ and abort the process; the cycle C_ℓ has the desired properties.

So assume that not all triangles in $\mathcal{T}_{i-1}$ are of type (T6) with respect to C_{i-1}. To establish (N1) for C_i we need to make sure that C_i visits the vertices of V_i^-, in addition to the vertices $\mathrm{V}(C_{i-1})$. The remaining properties (N2)–(N5) for C_i and T_{i+1} follow, for the most part, from the properties (N2)–(N5) for C_{i-1} and T_i, except that we have to separately consider the (few) edges that change when going from C_{i-1} to C_i —if these edges are incident to a vertex of a triangle of $\mathcal{T}_i$. Note that by (N2)

all triangles of $\mathcal{T}_i$ are normal for C_{i-1}. By property (N5), either T_i is also normal for C_{i-1} or it is normal for C_{i-2} (by (N2)) and changed by a single edge slide or cone slide at a vertex of T_i, plus optionally a single chord slide at another vertex of T_i, between C_{i-2} and C_{i-1}. If T_i is not normal for C_{i-1}, then we call T_i *special* for C_{i-1}; the proper edge of C_{i-2} that was affected the edge slide (either directly or as part of a cone slide) that made T_i special for C_{i-1} is called the *special edge* of T_i, see Figure 74. Accordingly, we will distinguish ten cases depending on how C_{i-1} and C_{i-2} interact with T_i. Some of these cases are symmetric to each other.

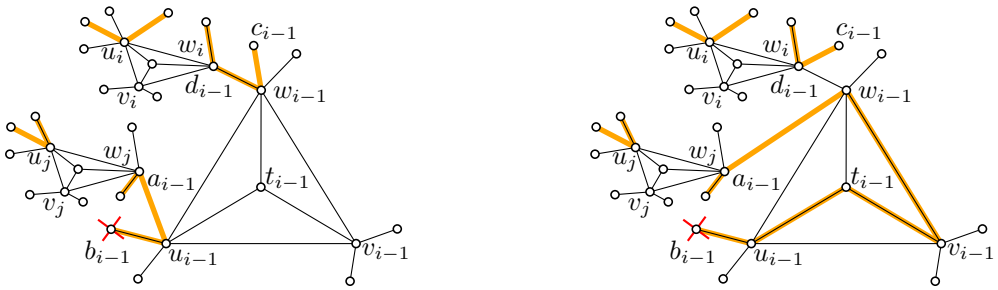

(a) C_{i-2} (b) C_{i-1}

Figure 74: For the cycle C_{i-2}, the triangle T_{i-1} is of type (T4), the triangle T_i is of type (T1), and the triangle T_j is of type (T3). The transition to C_{i-1} is perceived as an edge slide at w_i of T_i, which makes T_i special for C_{i-1}. The edge $d_{i-1}w_{i-1}$ of C_{i-2} is the special edge of T_i. Note that the vertex w_j of the triangle T_j is also affected by a slide operation, but its type (T3) is maintained.

Regarding (N2), we need to ensure that going from C_{i-1} to C_i creates at most one special triangle (which is then chosen to be T_{i+1}). Using chord slides at a vertex of T_j, a triangle T_j cannot leave the classes (T1)–(T5): a single chord-to-chord slide cannot change the type of a given triangle. Moreover, under a single chord-to-edge slide, a (T3)-triangle remains (T3), a (T1)– or (T2)-triangle remains what it is or it becomes (T3), and a (T4)– or (T5)-triangle remains what it is or it becomes (T1) or (T2), respectively. Multiple simultaneous chord slides at vertices of T_j may be thought of as a *sequence* of chord slides, given that, regardless of the type of T_j, at most two of the four cycle edges incident to some vertex of T_j are chords, and if there are two such chords, then they belong to distinct faces. Hence, by the above argument, even when the change from C_{i-1} to C_i results in multiple chord slides at T_j, it remains of one of the types (T1)–(T5).

Therefore, in order to check whether a triangle $T_j \in \mathcal{T}_i$, that is normal for C_{i-1} remains normal for C_i, it suffices to consider triangles where at least one of its vertices is affected by an edge slide or a cone slide when going from C_{i-1} to C_i. Whenever the change from C_{i-1} to C_i does not create any special triangle, we choose some arbitrary $T_{i+1} \in \mathcal{T}_i$ as the next triangle.

Case 1: $b_iv_i \in \bigotimes_b$ (regardless of whether T_i is normal or special for C_{i-1}). Then we obtain C_i from C_{i-1} by replacing the edge b_iu_i by the path $b_iv_it_iu_i$; see Figure 75. Compared to C_{i-1}, the only new edge in C_i outside of T_i is b_iv_i. If b_i is irrelevant, then there is nothing to show. Hence suppose that b_i is relevant, and let T_p denote the corresponding triangle that is relevant for C_i.

If b_iv_i is obtained from b_iu_i by a chord slide (note that it might be that $b_iu_i \in X_b$ if T_i is special), then T_p remains normal. However, if b_iv_i is obtained from b_iu_i by an edge slide, then T_p may not be normal for C_i. In that case, we set $T_{i+1} = T_p$. So in any case (N1)–(N5) hold. ◁

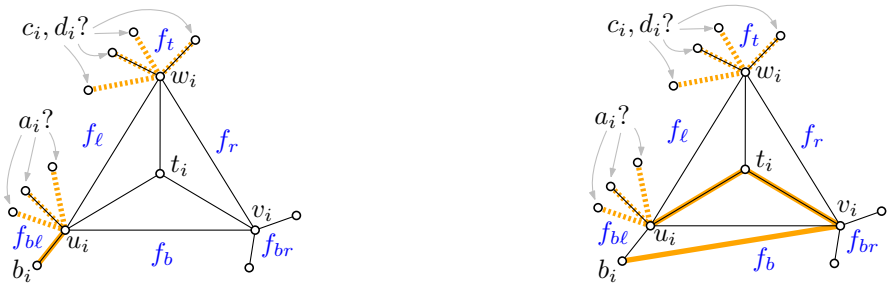

(a) before (b) after

Figure 75: Case 1 ($b_iv_i \in \mathrm{E}(G^{\otimes})$) in Lemma 53. We remark that the figure does not depict all possibilities. Since T_i could be special, it is possible that $a_i = w_i$ and $d_i = u_i$. Moreover, it could be that b_i is some other vertex of $f_b \setminus \{u_i, v_i\}$, that is, $u_ib_i \in X_b$. Analogously, we could have $w_ic_i \in X_r$.

Case 2: $c_iv_i \in \bigotimes_r$. Symmetric to Case 1. ◁

Hence we may suppose that $b_iv_i \notin \bigotimes_b$ and $c_iv_i \notin \bigotimes_r$. In particular, it follows that $b_iu_i \notin \bigotimes_b$ and $c_iw_i \notin \bigotimes_r$, and that T_i is neither of type (T1) nor of type (T2) for C_{i-1}.

Case 3: T_i is special for C_{i-1} and of type (T1) for C_{i-2}. Then T_i differs by a single edge slide or cone slide at a vertex of T_i, plus optionally a single chord slide at another vertex of T_i, from a triangle of type (T1) for C_{i-2}.

Since $c_iw_i \notin \bigotimes_r$, it follows that the special edge of T_i is c_iw_i. Since w_i is not incident to a chord in C_{i-2}, the special edge c_iw_i can only be affected by an edge-to-chord slide that is not part of a cone slide when going from C_{i-2} to C_{i-1}, that is, $c_iw_i \in X_t$ in C_{i-1}, see Figure 76a. If there is an optional chord slide, it has to affect u_i. It suffices to consider the case that it is a chord-to-edge slide. If it affects b_iu_i, the situation does not change , i.e., we are still applying an edge-to-chord slide to the special edge c_iw_i in a type (T1) triangle. Therefore we may assume that the chord-to-edge slide affects a_iu_i. It cannot become u_iw_i as this would imply a vertex of degree three in the cycle C_{i-1}. However, the only remaining option would

imply that T_i is of type (T3), which contradicts the case assumption. Hence, we may assume that there is no optional chord slide.

We obtain C_i from C_{i-1} by replacing the edge a_iu_i with the path $a_iw_iv_it_iu_i$ and shortcutting $c_iw_id_i$ to c_id_i; see Figure 76. Regarding (N4), there are three directed edges of C_{i-1} that are affected by this change. The changes from a_iu_i to a_iw_i and from c_iw_i to c_id_i are perceived as chord slides at a_i and c_i, respectively; and the change from d_iw_i to d_ic_i is an edge slide. Regarding (N5), if d_i is not relevant, there is nothing to show. So assume that d_i is relevant and let T_p denote the corresponding relevant triangle. If T_p is not normal for C_i, we set $T'_{i+1} = T_p$. The vertex d_i perceives the change from C_{i-1} to C_i as an edge slide or a cone slide (the latter if $a_i = d_i$). It remains to consider the slide operations affecting other vertices ($\neq d_i$) of T_p. The only slide operations that may affect T_p at a vertex $z \neq d_i$ are chord slides, namely if $z \in \{c_i, a_i\}$. Note that $a_i \neq c_i$ since otherwise w_i and a_i form a separation pair, which contradicts the 3-connectivity. Hence, z may be affected by *at most* one chord slide, and so (N5) holds. In any case, (N1)–(N5) hold. ◁

(a) before (b) after

Figure 76: Case 3 (T_i is special for C_{i-1} and of type (T1) for C_{i-2}) in Lemma 53.

Case 4: T_i is special for C_{i-1} and of type (T2) for C_{i-2}. Symmetric to Case 3. ◁

Case 5: T_i is of type (T4) for C_{i-1}. We obtain C_i from C_{i-1} by replacing the edge a_iu_i by the path $a_iw_iv_it_iu_i$ and by replacing the path $c_iw_id_i$ by the edge c_id_i; see Figure 77. Analogous to Case 3, we obtain that (N1)–(N5) hold. ◁

Case 6: T_i is of type (T5) for C_{i-1}. Symmetric to Case 5. ◁

Case 7: T_i is special for C_{i-1} and of type (T4) for C_{i-2} (but not for C_{i-1}). Then T_i differs by a single edge slide or cone slide at a vertex of T_i, plus optionally a single chord slide at another vertex of T_i, from a triangle of type (T4) for C_{i-2}. As b_i is irrelevant for C_{i-2} by definition of (T4), the incident cycle edge b_iu_i does not change between C_{i-2} and C_{i-1} by (N3). It follows that the special edge is d_iw_i.

Hence, if there is an optional chord slide when going from C_{i-2} to C_{i-1}, then it affects a_iu_i. It suffices to consider the case that this edge is affected by a chord-to-edge

(a) before (b) after

Figure 77: Case 5 (T_i is of type (T4) for C_{i-1}) in Lemma 53.

slide. The only way to turn a_iu_i into a proper edge is a chord slide to u_iw_i, which is only possible if $a_i = d_i$ and the special edge d_iw_i is affected by an edge slide to w_iu_i (otherwise, there is a vertex of degree three along C_{i-1}). In this case, we may obtain C_i from C_{i-1} by simply replacing the edge u_iw_i by the path $u_iv_it_iw_i$. Obviously, all triangles in $\mathcal{T}_i$ are normal for C_i and (N1)–(N5) hold.

Thus, we may assume that there is no optional chord slide and, moreover, that d_iw_i is affected by an edge-to-*chord* slide (possibly as part of a cone slide). An edge-to-chord slide that transforms d_iw_i to a chord of f_t allows us to proceed in the same way as in Case 5. In fact, here the argument is even easier because the shortcut of the path $c_iw_id_i$ to c_id_i encompasses chord slides only.

It remains to consider an edge-to-chord slide that makes d_iw_i a chord of f_ℓ, possibly as part of a cone slide that turns c_iw_i into a proper edge. Regardless, we obtain C_i from C_{i-1} by replacing the edge d_iw_i by the path $d_iu_it_iv_iw_i$ and by shortcutting $a_iu_ib_i$ to a_ib_i; see Figure 78. Regarding (N4), there are three directed edges of C_{i-1} that are affected by this change. The changes from d_iw_i to d_iu_i and from a_iu_i to a_ib_i are perceived as chord slides at d_i and a_i, respectively. The change from b_iu_i to b_ia_i is an edge slide at b_i. Since b_i is irrelevant by the condition of (T4), it follows that no relevant triangle becomes special. In any case, (N1)–(N5) hold. ◁

(a) before (b) after

Figure 78: Case 7 (T_i is special for C_{i-1} and of type (T4) for C_{i-2}) in Lemma 53.

Case 8: T_i is special for C_{i-1} and of type (T5) for C_{i-2} (but not for C_{i-1}). Symmetric to Case 7. ◁

Case 9: T_i is of type (T3) for C_{i-1}. Given that $c_iv_i \notin \bigotimes_r$, we have $c_iw_i \in X_t$ and $d_iw_i \in E_t$ and, symmetrically, $b_iu_i \in X_{b\ell}$ and $a_iu_i \in E_{b\ell}$; see Figure 79a. As T_i being of type (T3) for C_{i-1} is the last of the nonspecial cases, we may assume without loss of generality that all remaining relevant triangles in $\mathcal{T}_{i-1}$, are of type (T3) and look as in Figure 79a with respect to C_{i-1}. (Otherwise, since T_i is not special, we could pick a different triangle in the role of T_i and apply one of the other Cases 1, 2, 5, or 6.) We distinguish two subcases.

Case 9.1: At least one of a_i or d_i is irrelevant for C_{i-1}. Suppose without loss of generality that d_i is irrelevant for C_{i-1}. We obtain C_i from C_{i-1} by shortcutting $a_iu_ib_i$ to a_ib_i and replacing the edge d_iw_i by the path $d_iu_it_iv_iw_i$; see Figure 79b. The change from b_iu_i to b_ia_i is a chord slide at b_i. The changes from a_iu_i to a_ib_i and from d_iw_i to d_iu_i are edge slides. Since d_i is irrelevant by assumption, the only way to create a special triangle is if a_i is relevant (in this case $a_i \neq d_i$). In this case, if the relevant triangle T_p that contains a_i is not normal for C_i, we set $T_{i+1} = T_p$. The vertex a_i perceives the change from C_{i-1} to C_i as a single edge slide (since $a_i \notin \{b_i, d_i\}$). Hence, even if b_i also belongs to T_p, Property (N5) holds. ◁

(a) before (b) after

Figure 79: Case 9.1 (T_i is of type (T3) and d_i is irrelevant for C_{i-1}) in Lemma 53.

Case 9.2: Both a_i and d_i are relevant for C_{i-1}. We may suppose without loss of generality that for every triangle $T_x \in \mathcal{T}_{i-1}$ the vertices a_x and d_x are both relevant, that is, each of them belongs to some other triangle from $\mathcal{T}_{i-1} \setminus \{T_x\}$. Otherwise, we could select T_x to take the role of T_i and proceed as in Case 9.1.

Since not all triangles in $\mathcal{T}_{i-1}$ are of type (T6) (otherwise, we are done), we may suppose that T_i is not of type (T6). Consider the relevant triangle T_j that contains d_i. There are two edges of C_{i-1} incident to d_i: one is the proper edge w_id_i, the other edge e is in one of the two faces X_t and X_ℓ incident to w_id_i because T_j is of type (T3). In other words, either $d_i = w_j$, as depicted in Figure 81, or $d_i = u_j$ as depicted in Figure 80a. We claim that we may suppose without loss of generality that $d_i = w_j$

and $w_jc_j \in X_\ell$, as shown in Figure 81. Note that both statements ("$d_i = w_j$" and "$e \in X_\ell$") are equivalent, given that T_j is of type (T3).

To prove the claim, suppose to the contrary that $d_i = u_j$; see Figure 80a. Let T_h denote the relevant triangle that contains a_i as a vertex. Then, as T_i is not of type (T6), not both edges of C_{i-1} incident to a_i are in $\bigotimes_{b\ell}$ (and so the situation depicted in Figure 80a does *not* occur). Instead, the other ($\neq a_iu_i$) edge of C_{i-1} incident to a_i is in X_ℓ; see Figure 80b. In particular, it follows that $a_i \neq w_j$, i.e., a_i belongs to a relevant triangle $T_h \neq T_j$. Now the claim follows by symmetry: Reflecting $\mathbb{R}^2$ exchanges the roles of u_i and w_i, as well as the roles of T_h and T_j.

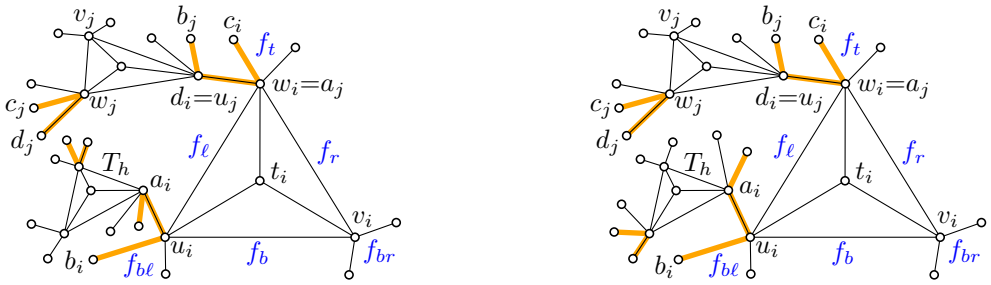

(a) T_i is (T6) (b) T_i is not (T6)

Figure 80: Using symmetry in Case 9.2 of Lemma 53.

So we face a situation as depicted in Figure 81, with $d_i = w_j$ and $w_jc_j \in X_\ell$. We observe that $a_i \neq u_j$ due to 3-connectivity (otherwise, u_j and w_j would form a separation pair in G). In particular, the relevant triangles of d_i and a_i, respectively, are distinct.

We obtain C_i from C_{i-1} by shortcutting $a_iu_ib_i$ to a_ib_i and replacing the path $c_iw_iw_jc_j$ by the path $c_iw_jw_iv_it_iu_ic_j$; see Figure 81b. Regarding (N4), the changes from c_iw_i to c_iw_j, from b_iu_i to b_ia_i, and from c_jw_j to c_ju_i are perceived as chord slides. The changes from a_iu_i to a_ib_i and from w_jw_i to w_jc_i are edge slides (the latter as part of a cone slide and the former possibly, if $a_i = c_j$, also as part of a cone slide).

Regarding (N5), it remains to show that at most one special triangle is created and that all other remaining relevant triangles belong to one of the types (T1)–(T6). The edge slide from w_jw_i to w_jc_i together with the chord-to-edge slide from c_jw_j to w_iw_j is perceived as a cone slide at w_j, which belongs to the relevant triangle T_j. If the cycle edges incident to u_j do not change, then T_j is of type (T5) with respect to C_i (because w_i is irrelevant for C_i). Given that $u_j \neq a_i$, the cycle edges incident to u_j change only if $u_jb_j = c_iw_i$ or if $u_jb_j = b_iu_i$. The former case would imply that u_j is a cut vertex, which contradicts the 3-connectivity. So assume that $u_jb_j = b_iu_i$. In this case, the slide b_iu_i to b_ia_i amounts to a slide u_ju_i to u_ja_i. If this change is perceived as a chord-to-chord slide, then T_j is of type (T5) with respect to C_i, as above. However, if the slide u_ju_i to u_ja_i is a chord-to-edge slide, that is, if $u_ja_i \in \mathrm{E}(G)$, then T_j is of type (T2) with respect to C_i. So, in any case, T_j is normal with respect to C_i.

It remains to consider the edge slide a_iu_i to a_ib_i. Note that $a_i \neq c_i$, as this would imply that a_i and w_i form a separation pair. Hence, the edge slide from a_iu_i to a_ib_i is either perceived as an edge slide or a cone slide (the latter if $c_j = a_i$) at the vertex a_i and its relevant triangle T_p ($\neq T_j$). If T_p is not normal for C_i, we set $T_{i+1} = T_p$. It remains to consider the slide operations affecting other vertices ($\neq a_i$) of T_p. Since w_j and w_i cannot belong to T_p, the only slide operations that may affect T_p at a vertex $\neq a_i$ are chord slides, namely if one of b_i, c_i, c_j belongs to T_p. Given that T_p is of type (T3), each vertex of T_p may only be affected by a *single* chord slide. Hence, (N5) is satisfied. In any case (N1)–(N5) are satisfied. ◁

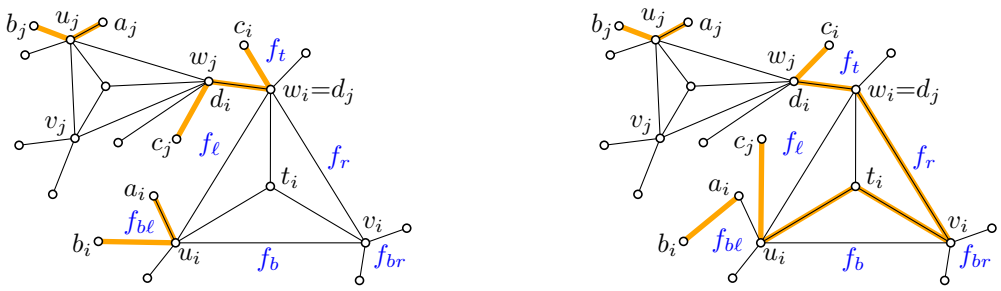

(a) before (b) after

Figure 81: Case 9.2 (Both a_i or d_i are relevant for C_{i-1} and $d_i = w_j$) in Lemma 53.

Case 10: T_i is special for C_{i-1} and of type (T3) with respect to C_{i-2} (but not for C_{i-1}). Given that $c_iv_i \notin \bigotimes_r$, we have $c_iw_i \notin \bigotimes_r$ and, symmetrically, $b_iu_i \notin \bigotimes_b$. As T_i is not of type (T3) with respect to C_{i-1}, by (N5) its normality is destroyed by an edge slide or cone slide between C_{i-2} and C_{i-1} at either w_i or u_i. Assume without loss of generality that this slide affects w_i. Then, if T_i is also affected by an optional chord slide when going from C_{i-2} to C_{i-1}, it affects u_i. However, note that such a slide does not change the situation: we are still applying an edge or cone slide to the vertex w_i of a type (T3) triangle. Hence, we may suppose that there is no optional chord slide.

We have already established that $c_iw_i \notin \bigotimes_r$. Since T_i is not of type (T3) it follows that $d_iw_i \notin E_t$. Consequently, we have $d_iw_i \in X_\ell$ or $d_iw_i \in X_t$.

Case 10.1: $d_iw_i \in X_\ell$. (in this case d_iw_i was affected by the edge or cone slide and, hence, c_iw_i remains a chord in X_t, or it assumes the former role of d_iw_i and becomes the proper edge in $E_t \cap E_\ell$.) We obtain C_i from C_{i-1} by shortcutting the path $a_iu_ib_i$ to the edge a_ib_i and by replacing the path $c_iw_id_i$ by $c_iw_iv_it_iu_id_i$; see Figure 82. The changes from d_iw_i to d_iu_i and from b_iu_i to b_ia_i are chord slides. The change from a_iu_i to a_ib_i is an edge slide (recall that $b_iu_i \notin \bigotimes_b$). Hence, regarding (N5), there is nothing to show if a_i is not relevant. So assume that T_p is a relevant triangle containing a_i. If T_p is not normal for C_i, we set $T_{i+1} = T_p$. The vertex a_i perceives the change from C_{i-1} to C_i is an edge slide or a cone slide (the latter if $a_i = d_i$).

It remains to consider the slide operations affecting other vertices ($\neq a_i$) of T_p. The only slide operations that may affect T_p at a vertex $z \neq a_i$ are chord slides, namely if $z \in \{b_i, d_i\}$. Note that $b_i \neq d_i$ since otherwise u_i and b_i form a separation pair, which contradicts the 3-connectivity. Hence, z may be affected by *at most* one chord slide, and so (N5) holds. In any case (N1)–(N5) are satisfied. ◁

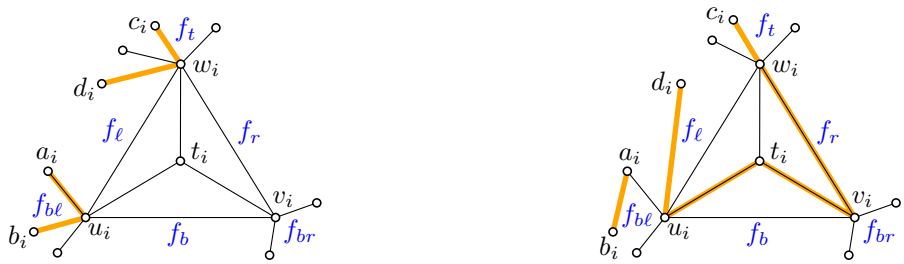

(a) before (b) after

Figure 82: Case 10.1 (T_i is special for C_{i-1} and of type (T3) for C_{i-2} and $d_iw_i \in X_\ell$) in Lemma 53.

Case 10.2: $d_iw_i \in X_t$. Then we obtain C_i from C_{i-1} by shortcutting $c_iw_id_i$ to c_id_i and replacing the edge a_iu_i by the path $a_iw_iv_it_iu_i$; see Figure 83. The change from d_iw_i to d_ic_i and from c_iw_i to c_id_i are chord slides. The change from a_iu_i to to a_iw_i is an edge slide (recall that $b_iu_i \notin \bigotimes_b$).

Hence, regarding (N5), there is nothing to show if a_i is not relevant. So assume that T_p is a relevant triangle that uses a_i. If T_p is not normal for C_i, we set $T_{i+1} = T_p$. The vertex a_i perceives the change from C_{i-1} to C_i is an edge slide (not a cone slide, since this would imply that $a_i \in \{c_i, d_i\}$, which implies that a_i and w_i form a separation pair and, hence, contradicts the 3-connectivity).

It remains to consider the slide operations affecting other vertices ($\neq a_i$) of T_p. The only slide operations that may affect T_p at a vertex $z \neq a_i$ are chord slides, namely if $z \in \{c_i, d_i\}$. Since $c_i \neq d_i$, z may be affected by *at most* one chord slide, and so (N5) holds. In any case (N1)–(N5) are satisfied. ◁

This concludes our case analysis and also the proof of Lemma 53. □

4.7.3 Reconstruction Phase 3: wrapup

In this section we discuss the third phase that we use to finish the (re)construction of our cycle in case the cycle C_ℓ obtained in Phase 2 is not yet subhamiltonian, that is, $\ell < s$. In this case, all remaining relevant triangles $\mathcal{T}_\ell$ interact with C_ℓ in a very specific way only, which makes it easy to adapt C_ℓ so as to incorporate the missing vertices.

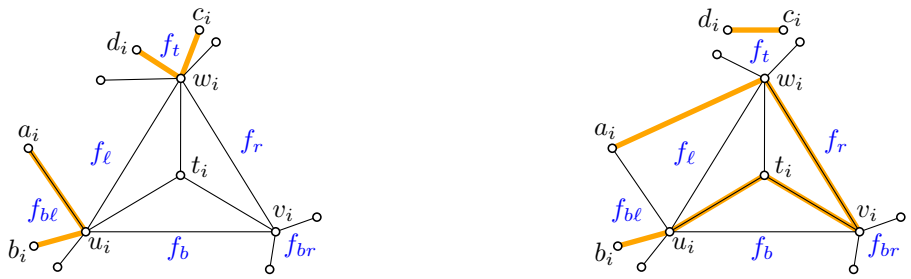

(a) before (b) after

Figure 83: Case 10.2 (T_i is special for C_{i-1} and of type (T3) for C_{i-2} and $d_iw_i \in X_t$) in Lemma 53.

Lemma 54. *Let $C = C_\ell$ be the almost spanning cycle obtained in Phase 2, cf. Lemma 53. Assume that $\ell < s$ and, hence, the set $\mathcal{T}_\ell$ of remaining relevant triangles is nonempty, and all these triangles are of type (T6) with respect to C. Then there exists a Hamiltonian plane cycle in $G^\otimes$.*

Proof. Consider a triangle T_i of type (T6) for C. By definition both edges u_ia_i and w_id_i are proper edges of G along ∂f_ℓ and both a_i and d_i belong to other relevant triangles T_k and T_j, respectively, with possibly $j = k$. Moreover, all edges of C adjacent to the edges u_ia_i and w_id_i in C are chords of faces other than f_ℓ.

As the situation is symmetric among T_i, T_j, T_k, we conclude that there exists a cycle $T_j, T_i, T_k, \ldots$ of (at least two) relevant triangles so that consecutive triangles are connected by an edge of $\partial f_\ell \cap \mathrm{E}(C)$ and $X_\ell \cap \mathrm{E}(C) = \emptyset$; see Figure 84a. We say that f_ℓ is the *central face* of the cycle $T_j, T_i, T_k, \ldots$.

We obtain a new cycle C' from C by replacing the edge d_iw_i by the path $\partial f_\ell \setminus \{dw_i\}$ and removing all other occurrences of the vertices on ∂f_ℓ from the cycle; see Figure 84b. For every edge of ∂f_ℓ in C the two adjacent edges belong to the same face and so the corresponding shortcuts are possible without crossing any edge of C or G. Moreover, all these "shortcut faces" are pairwise distinct (because G is 3-connected). In other words, the cycle C' is plane. Now every triangle in our cycle $T_j, T_i, T_k, \ldots$ has an edge in common with C' and we can modify C' locally to visit all vertices of those triangles as in Lemma 46; see Figure 84c.

This procedure may be repeated to handle all remaining central faces and their relevant triangles. Resolving one of the central faces may change the edges of the current almost spanning cycle that pass through the shortcut faces of other central faces. However, each of these edges will be replaced by a new edge that still belongs to the original shortcut face, so that the shortcuts may be carried out as described above. Intuitively, this strategy produces no crossings since, in each shortcut face, performing all shortcuts simultaneously corresponds to a homotopic deformation of the maximal paths of cycle edges along which the chords of the shortcut face and proper edges of the adjacent central faces alternate.

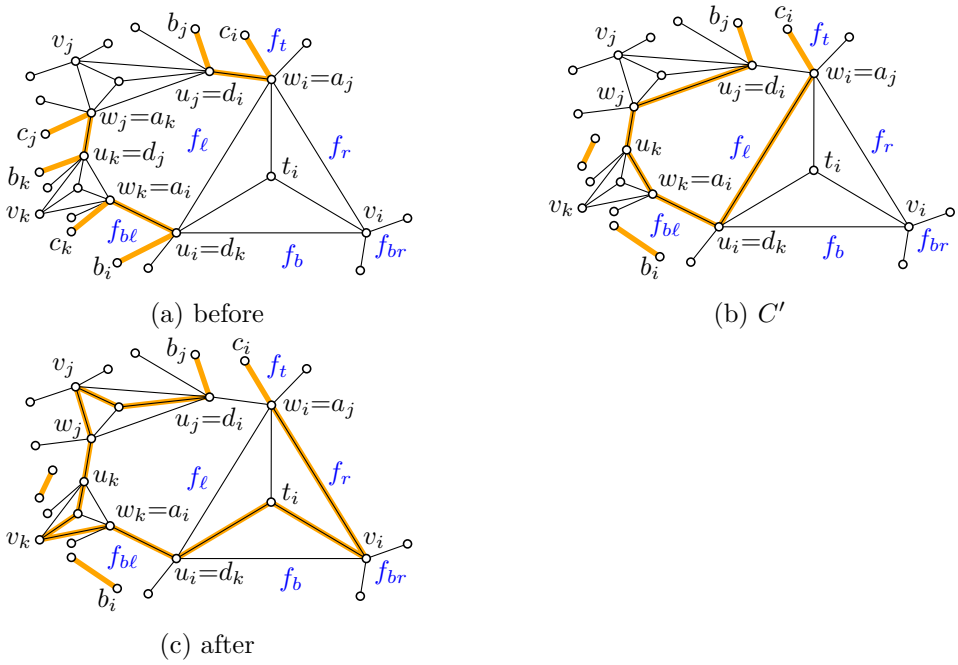

(a) before (b) C'

(c) after

Figure 84: A step of the third reconstruction phase in Lemma 54.

More formally, let us consider four maps: Every relevant vertex v on the current cycle (initially C) is incident to exactly two cycle edges; one of these edges, the *central edge* $c(v)$, is a proper edge of G that belongs to the central face $\zeta(v)$ of the corresponding relevant triangle of v; the other edge is the *shortcut edge* $s(v)$. Initially, $s(v)$ is a chord of some face $\sigma(v) \neq \zeta(v)$. While the shortcut edge incident to a relevant vertex v may change throughout our algorithm in Phase 3, we claim that the following properties are invariants of our algorithm.

(O1) For every relevant vertex v on the current cycle, all of $c(v)$, $\zeta(v)$, and $\sigma(v)$ remain constant and

(O2) $s(v)$ is either a chord or a proper edge along the boundary of $\sigma(v)$; and

(O3) for every shortcut edge $uv \in \mathrm{im}(s)$ with relevant vertex u, we have $s(u) = uv$; moreover, if v is also relevant, we have $s(u) = s(v) = uv$ and $\sigma(u) = \sigma(v)$.

It is easy to see that (O1)–(O3) hold at the beginning of Phase 3. We prove by induction on the number of steps that they continue to hold throughout Phase 3.

Suppose that, after performing some number of steps, we obtain an almost spanning plane cycle C_k and let f_ℓ be a central face whose vertices are still relevant for C_k. By (O1), the original central edges of f_ℓ are still on C_k. Moreover, by (O1) and (O2), the shortcut edges are chords or proper edges of the original shortcut faces. Hence, we can make all vertices of relevant triangles around f_ℓ irrelevant by

replacing one of the central edges d_iw_i by an appropriate path and by shortcutting the remaining central edges, as described above. It remains to argue that the resulting almost spanning cycle C_{k+1} is plane and that the invariants hold for those vertices that are relevant for C_{k+1} and incident to edges along C_{k+1} that were affected by the shortcuts.

Consider a shortcut that replaces a path $uvwx$ by the edge ux, where $c(v) = c(w) = vw$, $s(v) = uv$ and $s(w) = wx$. By (O1) we know that $vw = c(v) = c(w)$ is a central edge of $\zeta(v) = \zeta(w)$ and $\sigma(v) = \sigma(w)$. Hence, by (O2) both $s(v)$ and $s(w)$ are either a chord or a proper edge along the boundary of $\sigma(v)$, and so the shortcut can be performed without introducing edge crossings. The resulting edge ux is either a chord or a proper edge along the boundary of $\sigma(v)$. By (O3) we have $\sigma(v) = \sigma(u)$ if u is relevant; and $\sigma(w) = \sigma(x)$ if x is relevant. Therefore, the new shortcut edge ux incident to u and x satisfies (O2) and (O3) holds. Moreover, (O1) holds for σ.

To see that (O1) holds for c and ζ, observe that every shortcut involves a central edge vw and its two adjacent edges in the current cycle. Initially all central edges are pairwise nonadjacent (both in the graph G and in the cycle C). As the maps c and ζ remained constant up to this step, neither uv nor wx are a central edge of any relevant vertex, and so (O1) holds for c and ζ.

This concludes the proof that properties (O1)–(O3) are invariants of our algorithm. Consequently, we can perform all steps of the algorithm to eventually obtain a Hamiltonian plane cycle in $G^{\otimes}$. □

It remains to argue that the prescribed edges (if any) are still on the cycle.

Lemma 55. *Let $H \in \{C_k, C_\ell, C_s\}$ be the (almost) spanning cycle obtained in Phase 1, Phase 2, or Phase 3, respectively, cf. Lemmata 52–54. Then H uses all of the prescribed edges in F.*

Proof. In each step of the reconstruction algorithm, we obtain a cycle C_i from some cycle C_{i-1} by deleting some edges of C_{i-1} and adding some new edges. Each of the deleted edges $\mathrm{E}(C_{i-1}) \setminus \mathrm{E}(C_i)$ is incident to at least one vertex of a separating triangle of G (including vertices that result from the collapse of an edge in K): this is easily seen for the Phases 1 and 3, and for Phase 2 this follows by Property (N3). Hence, proper cycle edges that are not incident to a vertex of a separating triangle are never deleted.

By Observation 43, the (proper) edges of F belong to the initial cycle C_0 for the collapsed graph G_0 in the beginning of Phase 1. By Property (P3), the edges of F have no endpoint that belongs to a separating triangle of G and, hence, they also do not have an endpoint that belongs to an edge of G_0 that corresponds to a collapsed separating triangle of G. By the above argument, this implies that the edges of F never leave the cycle. □

From an algorithmic perspective, the reconstruction algorithm is straightforward.

Lemma 56. *The reconstruction algorithm can be implemented to run in linear time.*

Proof. The reconstruction step starts with a subhamiltonian cycle C'' for the stellation G'' of the graph G' in which all separating triangles of our general assumption G have been collapsed. To initiate Phase 1, we turn C'' into a subhamiltonian cycle C_0 for the graph G_0, which corresponds to G', but potentially with some duplicate edges. This is easily done in linear time.

In the beginning of Phase 1 and of Phase 2, we classify all (relevant) separating triangles, which takes constant time per triangle. We maintain the classification with respect to the evolving cycle throughout Phase 1 and 2. To this end, for each of the lemmata and cases of the respective phase, we maintain a list of all relevant triangles that satisfy the respective precondition. In Phase 2, the list of type (T3) triangles is split into a list of type (T6) triangles and a list of type (T3) triangles that are not of type (T6).

In Phase 1 and 2, the next triangle to handle is either determined by the previous step, or we pick an arbitrary triangle from one of our lists according to the priorities described in the respective phase. In every step, a constant number of cycle edges are modified. These changes affect constantly many remaining triangles, whose classification we need to check and possibly update, which can be done in constant time per affected triangle. Hence, every step of Phase 1 and Phase 2 takes constant time.

During Phase 3, we do not care for the classification anymore. A single step may affect many remaining triangles, and its complexity is linear in the number of these triangles (constant per triangle). As the number of separating triangles is linear, so is the overall runtime of Phase 3.

Since each phase is carried out in linear time, the entire reconstruction algorithm takes linear time as claimed. □

4.8 Special cases and proof summary

So far, we have considered the case that our general assumption G is not a trivial double kite and that it contains no subgraph isomorphic to a $\mathcal{G}_1$ or $\mathcal{G}_2$. In this section, we deal with these remaining special cases and we formally summarize the proof of Theorem 22, the generalization of our main result Theorem 18.

Theorem 22. *Let $\mathcal{G} = (V, E)$ be a 3-connected simple plane graph on n vertices where every vertex that belongs to a separating triangle has degree at most five. Then there is a plane augmentation of $\mathcal{G}$ that contains a Hamiltonian cycle C, which can be computed in $O(n^2)$ time.*

Moreover, for certain graphs, we may prescribe two edges to be part of the cycle C. Let $F \subset E$ be a set of up to two edges such that if $F \neq \emptyset$, the following conditions are satisfied:

(P1) the edges of F belong to the outer face $T_\circ$ of $\mathcal{G}$;

(P2) $T_\circ$ is a triangle;

(P3) no vertex of $T_\circ$ belongs to a separating triangle of $\mathcal{G}$; and

(P4) either at least one vertex of $T_\circ$ has degree three in $\mathcal{G}$, or all vertices of $T_\circ$ have degree four in $\mathcal{G}$.

The cycle C uses all edges in F.

Proof. The proof is by induction on the number n of vertices. For the base of the induction, note that the smallest plane simple 3-connected graph is K_4, where the statement holds. Similarly, there is a unique maximal planar graph on 5 vertices (the triangular bipyramid). In this graph, every face shares a vertex with a separating triangle and, hence, Property (P3), we have $F = \emptyset$. The claim follows since the graph is obviously Hamiltonian.

Hence suppose $n \geq 6$ and that the statement holds for all graphs on at most $n-1$ vertices. If $\mathcal{G}$ does not contain a separating triangle, then we are done by Theorem 21. So let $\mathcal{S}$ denote the nonempty set of separating triangles of $\mathcal{G}$. By Lemma 23, we may assume that each triangle in $\mathcal{S}$ is trivial.

If $\mathcal{G}$ is not a trivial double kite or contains a subgraph isomorphic to a $\mathcal{G}_1$ or $\mathcal{G}_2$, we obtain a set $\mathcal{K} \subset \mathrm{E}(\mathcal{G})$ of edges to collapse as described in Theorem 33 and proceed as in Sections 4.6 and 4.7. So it remains to consider the cases where $\mathcal{G}$ is a trivial double kite or contains a subgraph isomorphic to a $\mathcal{G}_1$ or $\mathcal{G}_2$. Accordingly, we distinguish three cases.

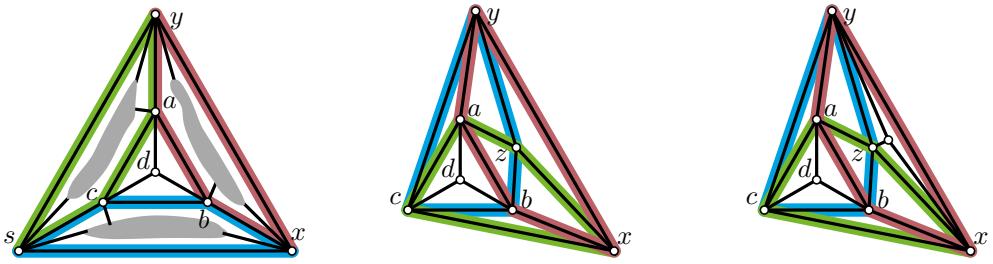

(a) $\mathcal{G}_1$ (b) $\mathcal{G}_2$ (c) $\mathcal{G}_3$

Figure 85: If all three edges of a separating triangle in the graph $\mathcal{G}$ are constrained by a 4-inhibitor, there is a subgraph $\mathcal{G}'$ isomorphic to a $\mathcal{G}_1$ or $\mathcal{G}_2$. In fact, if all separating triangles of $\mathcal{G}$ are trivial, then $\mathcal{G}$ itself is isomorphic to a $\mathcal{G}_1$ or one of $\mathcal{G}_2, \mathcal{G}_3$. The gray parts in (a) represent arbitrary subgraphs.

Case 1: $\mathcal{G}$ is a trivial double kite. Every face of $\mathcal{G}$ shares a vertex with a separating triangle. Hence, by Property (P3), we have $F = \emptyset$. The claim follows since the trivial double kite is obviously Hamiltonian. ◁

Case 2: $\mathcal{G}$ contains a subgraph isomorphic to $\mathcal{G}_2$. We use a, b, c, d, x, y, z to denote the vertices of the $\mathcal{G}_2$ subgraph as depicted in Figure 85b. By Observation 25, the triangles in $\mathcal{S}$ are pairwise vertex-disjoint. Hence, the only face of the $\mathcal{G}_2$ subgraph that is not necessarily a face of $\mathcal{G}$ is xyz. Thus, if $\mathcal{G} \not\simeq \mathcal{G}_2$, then xyz is a separating triangle of $\mathcal{S}$ that, by assumption, is trivial. Therefore, $\mathcal{G}$ is isomorphic to the graph $\mathcal{G}_3$ depicted in Figure 85c.

If $\mathcal{G} \simeq \mathcal{G}_3$, then every face shares a vertex with a separating triangle. Hence, $F = \emptyset$ by Property (P3) and the claim follows since $\mathcal{G}_3$ is easily seen to be Hamiltonian. On the other hand, if $\mathcal{G} \simeq \mathcal{G}_2$, then it might be that $F \neq \emptyset$ if $T_\circ = xyz$. It is easy to verify that $\mathcal{G}_2$ admits a Hamiltonian cycle that passes through any given pair of edges of xyz. ◁

Case 3: $\mathcal{G}$ contains a subgraph isomorphic to a $\mathcal{G}_1$. We use a, b, c, d, x, y, s to denote the vertices of the $\mathcal{G}_1$ subgraph as depicted in Figure 85a. The triangle xys has to be facial in $\mathcal{G}$ since, otherwise, it is separating and its vertices violate the degree bound. If $T_\circ = xys$, then $F = \emptyset$ by Property (P4). So it suffices to consider the case where $T_\circ \neq xys$. We may assume without loss of generality that $T_\circ$ is located on the side of $C_{ab} = abxy$ that does not contain d.

Our general plan is to proceed similarly as in the proof of Lemma 23: inductively, we will obtain a path that uses all edges in F and visits all vertices exterior to C_{ab} and some of the vertices on C_{ab}. Note that, contrary to what is depicted in Figure 85a, the exterior of C_{ab} is the side that does not contain d, due to the assumption about the position of the outer face. Depending on how the vertices of C_{ab} are visited, we perform two more inductive calls to obtain two paths that visit the vertices interior to $C_{bc} = bcsx$ and $C_{ca} = cays$, respectively. We make use of the fact that we may prescribe edges in order to be able to glue the three obtained paths appropriately.

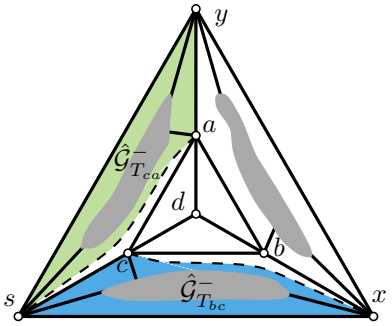

Figure 86: The graph $\hat{\mathcal{G}}$. The vertices interior to C_{bc} and C_{ca} are visited by making inductive calls to the highlighted subgraphs.

Covering the interior of C_{bc} and C_{ca}. Let us first describe how we visit the vertices interior to C_{bc} and C_{ca}. Add the edges sa and cx in a planar fashion to $\mathcal{G}$ and denote the resulting graph by $\hat{\mathcal{G}}$, for an illustration refer to Figure 86. In this graph,

consider the triangles $T_{bc} = csx$ and $T_{ca} = ays$. We now introduce notation for a total of six plane subhamiltonian cycles for the graphs $T_{bc}^- = \hat{\mathcal{G}}_{T_{bc}}^-$ and $T_{ca}^- = \hat{\mathcal{G}}_{T_{ca}}^-$. Each of these cycles uses two of the edges of the outer face of the respective graph, which allows us to combine these cycles in order to ultimately obtain a plane subhamiltonian cycle for $\mathcal{G}$. We will show that these cycles always exist and how to find them.

- $\mathcal{H}_{cx}$ is a plane subhamiltonian cycle in T_{bc}^- that uses the edges sc and sx.
- $\mathcal{H}_{cs}$ is a plane subhamiltonian cycle in T_{bc}^- that uses the edges xs and xc.
- $\mathcal{H}_{sx}$ is a plane subhamiltonian cycle in T_{bc}^- that uses the edges cs and cx.
- $\mathcal{H}_{ay}$ is a plane subhamiltonian cycle in T_{ca}^- that uses the edges sa and sy.
- $\mathcal{H}_{as}$ is a plane subhamiltonian cycle in T_{ca}^- that uses the edges ya and ys.
- $\mathcal{H}_{sy}$ is a plane subhamiltonian cycle in T_{ca}^- that uses the edges as and ay.

We show how to obtain the three cycles for T_{bc}^-; the cycles for T_{ca}^- are obtained analogously. It suffices to show that T_{bc}^- satisfies the preconditions of the theorem for any two prescribed outer edges. The graph T_{bc}^- is clearly 3-connected. The degree of each vertex of T_{bc}^- is at most as large as in $\mathcal{G}$. Hence, if T_{bc}^- does not satisfy the degree bounds, then it contains a separating triangle that is not part of $\mathcal{G}$. Such a triangle must use the added edge cx. However, the degree bounds of $\mathcal{G}$ imply that the degree of c in T_{bc}^- is exactly three. If follows that c and, consequently, cx cannot belong to a separating triangle in T_{bc}^- since in a 3-connected graph every vertex of a separating triangle has degree at least four. The Conditions (P1), (P2), and (P4) obviously hold. Regarding Condition (P3), we have already established that neither c nor cx belong to a separating triangle. Consequently, if one of the vertices x or s belongs to a separating triangle in T_{bc}^-, then it also belongs to this triangle in $\mathcal{G}$. However, if x or s belongs to a separating triangle in T_{bc}^-, its degree in this graph is at least four and, consequently, its degree in $\mathcal{G}$ is at least six, which contradicts the degree bounds for $\mathcal{G}$. Altogether, this show that T_{bc}^- satisfies the preconditions of the theorem and, hence, the desired cycles may be obtained by induction.

Covering the exterior of C_{ab}. Let us discuss how to obtain the path that covers the vertices exterior to C_{ab}. We remove, from $\mathcal{G}$, all vertices interior to C_{bc}. The resulting graph $\tilde{\mathcal{G}}$ clearly satisfies all preconditions of the theorem. Moreover, the outer face $T_\circ$ of $\mathcal{G}$ is also the outer face of $\tilde{\mathcal{G}}$. Hence, we may inductively obtain a plane subhamiltonian cycle $\mathcal{H}$ for $\tilde{\mathcal{G}}$ that uses the edges of F.

Covering $\mathcal{G}$. Removing the vertices interior to C_{ab} from $\mathcal{H}$, yields a set of paths. The endpoints of these paths belong to the set $\{a, b, x, y\}$. Consequently, there is either exactly one such path P that passes through all vertices in the exterior of C_{ab},

or exactly two such paths P', P'' where the disjoint union of their interior vertices coincides with the set of vertices in the exterior of C_{ab}.

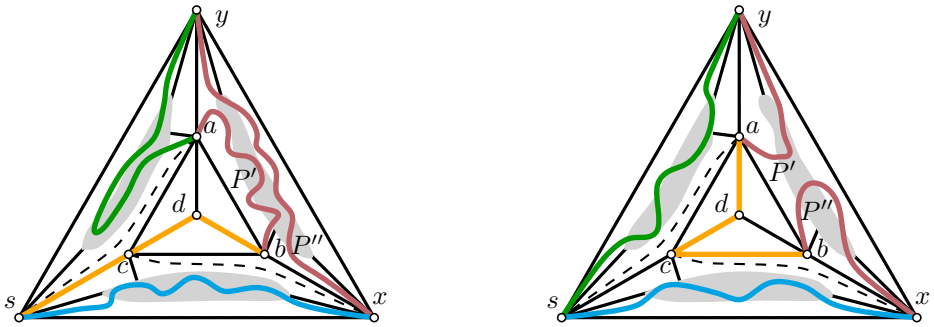

(a) $P' = a \dots b$ and $P'' = x \dots y$ (b) $P' = a \dots y$ and $P'' = b \dots x$

Figure 87: Obtaining a plane subhamiltonian cycle in the case that the exterior of C_{ab} is covered by two paths.

If $\mathcal{H}$ covers the exterior of C_{ab} by two paths P' and P'', then, by planarity, either, say, P' ends at a and b and P'' ends at x and y; or, say, P' ends at a and y and P'' ends at x and b. In the former case, the following cycle is plane subhamiltonian for $\mathcal{G}$, for illustration see Figure 87:

$$P' \cup bdcs \cup (\mathcal{H}_{sx} \setminus \{c\}) \cup P'' \cup (H_{ay} \setminus \{s\})$$

In the latter case, the desired cycle is:

$$P' \cup (\mathcal{H}_{sy} \setminus \{a\}) \cup (\mathcal{H}_{sx} \setminus \{c\}) \cup P'' \cup bcda$$

It remains to consider the case that the exterior of C_{ab} is covered by a single subpath P of $\mathcal{H}$. We distinguish several cases regarding the endpoints u, v of P.

Case 3.1: $\{u, v\} = \{y, a\}$. For an illustration refer to Figure 88. The desired cycle is

$$P \cup (\mathcal{H}_{sy} \setminus \{a\}) \cup P_1 \cup P_2$$

where P_1 is a path whose endpoints are s and c such that

$$P_1 = \begin{cases} \mathcal{H}_{cs} \setminus \{x\}, & \text{if } P \text{ visits } x \\ sx \cup (\mathcal{H}_{cx} \setminus \{s\}), & \text{if } P \text{ does not visit } x \end{cases}$$

and $P_2 = cda$ or $P_2 = cbda$ depending on whether P visits b or not. ◁

Case 3.2: $\{u, v\} = \{y, b\}$. We proceed as in Case 3.1 except that we replace P_2 with cdb or $cadb$ depending on whether P visits a or not. ◁

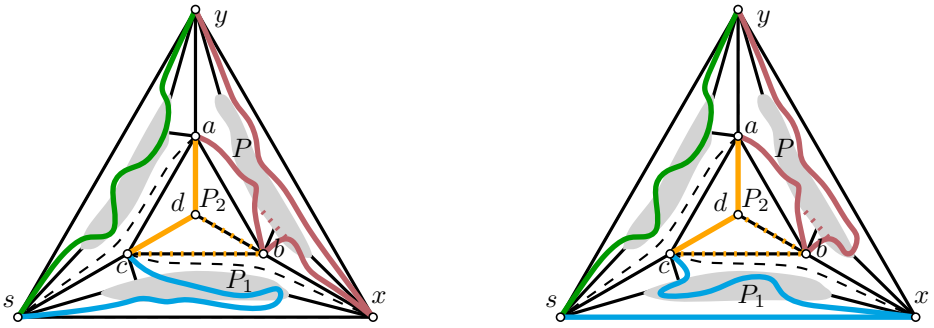

(a) P visits x. (b) P does not visit x.

Figure 88: Case 3.1.

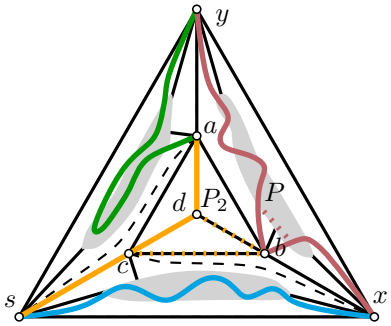

Figure 89: Case 3.3.

Case 3.3: $\{u, v\} = \{y, x\}$. For an illustration refer to Figure 89. Without loss of generality we may assume that P does not visit a, since otherwise we could simply shortcut P at a, that is, replace the subpath $a'aa''$ of P with the edge $a'a''$, given that a has no neighbor exterior to C_{ab}. Note that this shortcut cannot affect the edges of F: if $T_\circ$ is triangular, the vertex a cannot belong to $T_\circ$ as this would imply the existence of the edge by, which by 3-connectivity and the degree bound of b implies that C_{ab} is not separating, contradictory to the definition of $\mathcal{G}_1$.

The desired cycle is

$$P \cup (\mathcal{H}_{ay} \setminus \{s\}) \cup P_2 \cup cs \cup (H_{sx} \setminus \{c\})$$

where P_2 is defined as in Case 3.1. ◁

Case 3.4: $\{u, v\} = \{a, b\}$. For an illustration refer to Figure 90. The desired cycle is

$$P \cup P_3 \cup P_1 \cup cdb$$

where P_3 is a path whose endpoints are a and s such that

$$P_3 = \begin{cases} \mathcal{H}_{as} \setminus \{y\}, & \text{if } P \text{ visits } y \\ (H_{ay} \setminus \{s\}) \cup ys, & \text{if } P \text{ does not visit } y \end{cases}$$

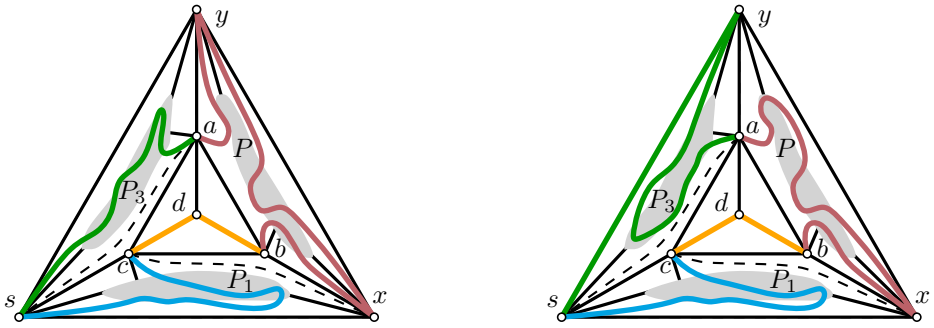

(a) P visits y. (b) P does not visit y.

Figure 90: Case 3.4.

and P_1 is defined as in Case 3.1. ◁

Case 3.5: $\{u, v\} = \{a, x\}$. For an illustration refer to Figure 91. We distinguish three subcases regarding the incidences of y and b with P.

Case 3.5.1: P does not visit y. The desired cycle is

$$P \cup xy \cup (\mathcal{H}_{sy} \setminus \{a\}) \cup (\mathcal{H}_{cs} \setminus \{x\}) \cup P_2$$

where P_2 is defined as in Case 3.1. ◁

Case 3.5.2: P visits y, but does not pass through b. The following

$$P \cup xbdc \cup (\mathcal{H}_{cs} \setminus \{x\}) \cup (\mathcal{H}_{as} \setminus \{y\})$$

is the desired plane subhamiltonian cycle for $\mathcal{G}$. ◁

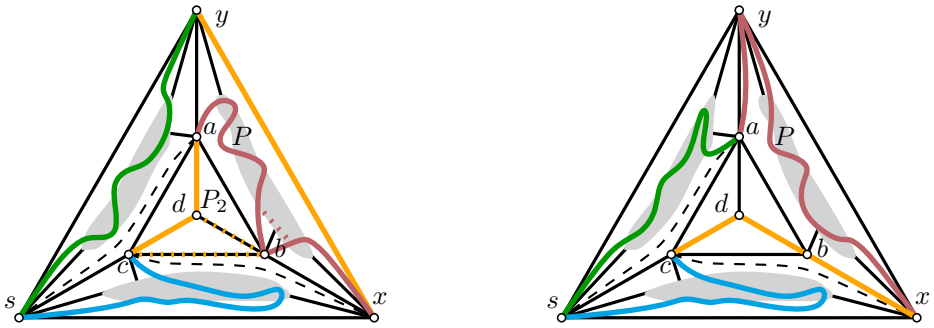

(a) Case 3.5.1. (b) Case 3.5.2.

Figure 91: Case 3.5.

Case 3.5.3: P visits both y and b. For an illustration refer to Figure 92a. We make use of the way how $\mathcal{H}$ traverses $\tilde{\mathcal{G}}$. Given that P has b as an interior vertex and a as

one of its endpoints, the cycle $\mathcal{H}$ cannot contain the edge bd. Hence, the path adc is a subpath of $\mathcal{H}$. This implies that s has to be a neighbor of x along $\mathcal{H}$. It follows that

$$\mathcal{H} = P \cup adc \cup P_4 \cup sx$$

where P_4 is a path whose endpoints are c and s and whose interior vertices are exactly all vertices in the interior of C_{ca}. Thus, we may obtain the desired cycle by replacing the edge sx of $\mathcal{H}$ with the path $\mathcal{H}_{sx} \setminus \{c\}$. ◁

Case 3.6: $\{u, v\} = \{b, x\}$. For an illustration refer to Figure 92b. As in Case 3.3, we may assume without loss of generality that P does not visit a. The desired path is

$$P \cup (\mathcal{H}_{sx} \setminus \{c\}) \cup P_3 \cup adcb$$

where P_3 is defined as in Case 3.4. ◁

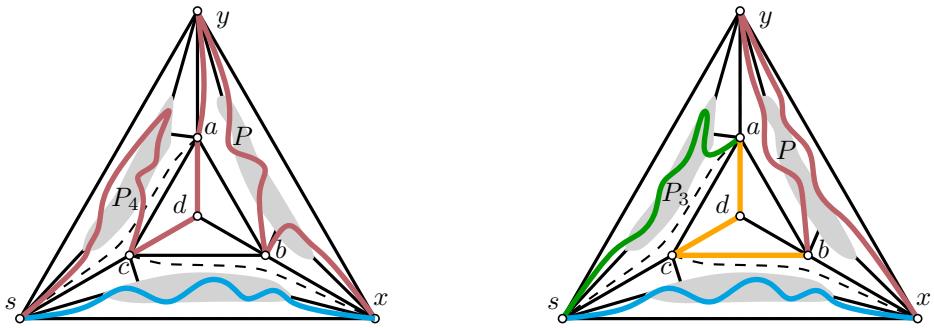

(a) Case 3.5.3. (b) Case 3.6.

Figure 92: Cases 3.5.3 and 3.6.

Altogether, we have shown how to obtain a plane subhamiltonian cycle for $\mathcal{G}$ that uses the edges of F for the cases that $\mathcal{G}$ is a trivial double kite or contains a subgraph isomorphic to a $\mathcal{G}_1$ or $\mathcal{G}_2$. This concludes the proof. □

Theorem 22 implies our main result:

Theorem 18. *Let G be a 3-connected simple planar graph on n vertices where every vertex that belongs to a separating 3-cycle has degree at most five. Then G is subhamiltonian planar. Moreover, a subhamiltonian plane cycle for G can be computed in $O(n^2)$ time.* □

Corollary 19. *Every 3-connected simple planar graph with maximum vertex degree five can be embedded on two pages, and such an embedding can be computed in quadratic time.* □

4.9 Summary of the algorithm

The details of the algorithm are scattered throughout the proofs of the different lemmata that contribute to the proof of Theorem 22. In this section, we collect these pieces into a high-level summary of the algorithm.

If $n = 4$, the given graph is a K_4, and if $n = 5$, we face a triangular bipyramid. In these cases, we simply return a Hamiltonian cycle for the respective graph, as described in Section 4.8. Otherwise, as described in Lemma 23, we begin by determining the maximal separating triangles in linear time.

If there is no separating triangle, we proceed as in Theorem 21: by using the algorithm of Biedl, Kant, and Kaufmann [23], we augment, if possible, the graph to be 4-connected, and then use the algorithm by Chiba and Nishizeki [30] to find a Hamiltonian cycle using the prescribed edges in the augmented graph. If it is not possible to find a 4-connected plane augmentation, the graph can be augmented to a wheel with at least four spokes, which makes it easy to directly determine the cycle. Each of the two cases takes linear time. For the case $n \geq 6$, this is the first way the algorithm can terminate.

In case there exist separating triangles, we check if all of them are trivial. If not, we split the graph at a nontrivial separating triangle as described in Lemma 23, recursively find subhamiltonian plane cycles for the two parts, and glue them together in constant time to find the desired cycle for $\mathcal{G}$.

If all separating triangles are trivial, we test if we face one of the following constant sized cases: the trivial double kite (Figure 24b), the graph $\mathcal{G}_2$ (Figure 85b), or the graph $\mathcal{G}_3$ (Figure 85c). If so, it is easy to determine the desired cycle in constant time as described in the Cases 1 and 2 in Section 4.8, which is another way for the algorithm to terminate.

If $\mathcal{G}$ is not a trivial double kite, or $\mathcal{G}_2$ or $\mathcal{G}_3$, we proceed by setting up the data structures described in Section 4.5.4. Using these data structures, it is easy to determine whether $\mathcal{G}$ belongs to the family $\mathcal{G}_1$ of graphs (Figure 85a). If so, we may obtain the desired cycle by means of three recursive calls as described in Case 3 in Section 4.8.

If $\mathcal{G}$ is not a $\mathcal{G}_1$, we now face the subproblem of finding a set K of edges to collapse as described by Theorem 33. As long as possible, we pick edges that are not constrained by 4-inhibitors and whose collapse maintains 3-connectivity and we perform the collapse. The next edge to pick can be efficiently determined by means of our data structures. Maintaining these data structures may take linear time per step and is one of the bottlenecks of our algorithm. If at some point it is not possible to choose an edge whose collapse preserves 3-connectivity, we check, using our data structures, if there is a safe candidate edge. If so, we collapse this edge. If not, we pick an arbitrary unsafe candidate edge e. For each of the constantly many (by (F2) and the degree bounds) paths psq that make e unsafe, there are three possibilities, namely, the two easy Cases 1 and 2.2.2.3.1 and the difficult Case 2.2.2.3.2 in Lemma 38. If

the difficult Case 2.2.2.3.2 arises for at least one of the paths psq, we need to choose a new candidate edge to replace e. Testing for Case 1 amounts to checking whether the edge sq is constrained by a 4-inhibitor, which can be done in constant time due to our data structures. Testing whether we are in Case 2.2.2.3.1, can be done by checking whether the degree of q is four, which is trivially possible in constant time. Hence, checking whether the Case 2.2.2.3.2 arises for at least one of the paths can be done in constant time. If this is not the case, we collapse e. Otherwise, we use our replacement strategy ϱ to find a replacement for e. Finding the edge $\varrho(e)$ takes constant time. However, a linear number $i \in O(n)$ of iterations may be required to find an edge $\varrho^i(e)$ for which Case 2.2.2.3.2 does not arise. This is another bottleneck of our algorithm.

In any case, we end up collapsing a candidate edge. The resulting graph is 2-connected only and we recursively find an edge set K_R for each rigid triconnected component R. Here, we face the generalized version of the subproblem, as described by Lemma 38: some of the edges of R are virtual, and these edges may be part of separating triangles that do not belong to the original graph. These triangles are not necessarily trivial and they may have vertices of degree larger than five. In particular, an original separating triangle T may now be part of a nontrivial double kite. We deal with this case next. As shown in Lemma 38, we can safely collapse any of the two edges of T that belong to only one separating triangle (namely T). Other than that, the algorithm operates as in the initial phase, except that when collapsing an edge, we need to worry about creating a 4-inhibitor only if its edges are real. The recursion has two base cases. One of them is a rigid triconnected component that does not have one of the original separating triangles (in this case, we are done), the other is described in Case 2.2.2.3.1, where the rigid triconnected component has constant size.

Once we have dealt with the subproblem and obtained the set K, we return to our original problem of finding a subhamiltonian cycle that uses the prescribed edges (if any). We collapse the edges in K and stellate all nontriangular faces of the resulting graph, which can be done in linear time. We then apply the algorithm for the case of graphs without separating triangles as described above. Finally, we apply our reconstruction algorithm, described in Section 4.7. The algorithm has three phases. In each step of one of the first two phases, we pick a triangle and incorporate its vertices into the current (almost) spanning cycle. Choosing the next triangle can be done in constant time by means of the data structures described in Lemma 56. To make its vertices part of the cycle amounts to changing a constant number of cycle edges in the neighborhood of the triangle and updating the data structures, both of which is easily done in constant time. In each step of Phase 3, multiple triangles are made part of the cycle by modifying a constant number of cycle edges in the neighborhood of each affected triangle. Overall, the reconstruction takes linear time and it terminates with a subhamiltonian plane cycle that uses the prescribed edges. This concludes the description of our algorithm.

4.10 Planar graphs that are not subhamiltonian

A maximal planar graph is subhamiltonian if and only if it is Hamiltonian. In order to establish upper bounds on the vertex degrees, we are looking for graphs that are maximal planar and nonhamiltonian to begin with. A large class of these graphs is known as *Kleetopes* [68]: a Kleetope G' is obtained from a maximal planar graph G by stellating each face, that is, by adding a new vertex v_f for each face f of G and connecting v_f to each of the three vertices on the boundary of f. If G has n vertices, then it has $2n-4$ faces and so G' has $3n-4$ vertices. Every vertex in $G' \setminus G$ has only three neighbors, all of which are vertices of G. In other words, the set of vertices in $G' \setminus G$ is independent in G'. For $n \geq 5$, we have $2n-4 > n$ and so this independent set contains more than half of the vertices of G'. No graph with such a large independent set can be Hamiltonian because every vertex needs two neighbors in a Hamiltonian cycle. Therefore, every Kleetope that is obtained from a maximal planar graph other than K_3 or K_4 is nonhamiltonian—the smallest such Kleetope is obtained from a triangular bipyramid and has 11 vertices [51, 68], and is known under the name Goldner-Harary graph [51], see Figure 93.

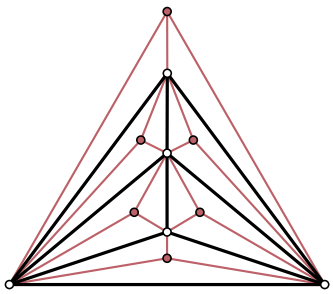

Figure 93: The Goldner-Harary graph contains an independent set (red) of size 6, but only 11 vertices in total and, hence, is not subhamiltonian.

Observation 57 ([51, 68])**.** *There exists a 3-connected planar graph of vertex degree at most eight that is not subhamiltonian.* □

The argument above generalizes to planar 3-connected graphs: start with any 3-connected planar graph G, add a new vertex v_f for each face f and connect v_f to some of the vertices on the boundary of f such that the resulting planar graph G' is 3-connected and, hence, has a unique combinatorial embedding (up to its orientation). Thus, in any plane augmentation of G' the neighborhood of each of the new vertices v_f consists exclusively of vertices of G, i.e., $G' \setminus G$ is an independent set. Hence, if G' admits a plane subhamiltonian cycle, then G has at most as many faces as vertices. This idea was also used in the $\mathcal{NP}$-hardness proof for the recognition of subhamiltonian planar graphs by Bauernöppel [20] (in fact, his construction is only 2-connected—Bauernöppel connects some of the vertices v_f to two nonadjacent

vertices of a face f only). We will use this idea to prove the following theorem.

Theorem 20. *There exists an infinite family of 3-connected simple planar graphs that are not subhamiltonian planar and where every vertex of a separating 3-cycle has degree at most six.*

Proof. For an integer $k \geq 3$, we build a graph G_k on $8k + 25$ vertices as follows. Start from the Cartesian product $C_4 \square P_k$ of a fourcycle with a path on k vertices. (The Cartesian product of two graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ is the graph $G_1 \square G_2 = (V, E)$ on the vertex set $V = V_1 \times V_2$ with $(u_1, u_2)(v_1, v_2) \in E$ if and only if $u_1 = v_1$ and $u_2v_2 \in E_2$ or $u_1v_1 \in E_1$ and $u_2 = v_2$.) This product is a planar quadrilateralization on $4k$ vertices with $8k - 4$ edges and $4k - 2$ faces; see Figure 94a. Next pick any set of three pairwise nonadjacent faces; to each picked face f, attach a cube: Denote $\partial f = u_1u_2u_3u_4$ and add a fourcycle $v_1v_2v_3v_4$ of four new vertices to the graph so that u_i and v_i are connected by an edge, for $i \in \{1, 2, 3, 4\}$; moreover, add one chord v_1v_3 to the new fourcycle. We call the resulting graph the *frame* $F_k \subset G_k$; it has $4k + 12$ vertices, $8k + 23$ edges and $4k + 13$ faces; see Figure 94b.

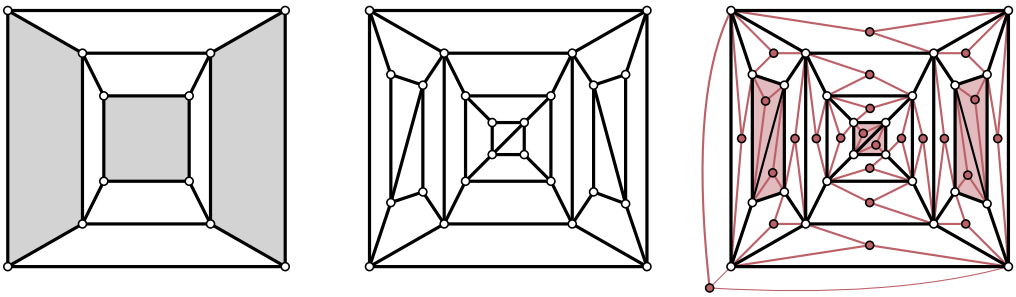

(a) $C_4 \square P_3$ (b) F_3 (c) G_3

Figure 94: The construction of G_3, a 3-connected planar graph that is not subhamiltonian and where every vertex of a separating triangle has degree at most six. We start from the Cartesian product $C_4 \square P_3$, where we pick three pairwise nonadjacent faces (shaded in (a)). Then we plant a cube on each picked face, obtaining the frame F_3 (b). Finally, to obtain G_3 we add a new vertex in every face of F_3 and connect it to three vertices on the boundary (c). The separating triangles of G_3 are shaded red; their vertices have degree six. The red vertices form an independent set, and no edge between any two red vertices can be added while maintaining planarity. As there are 25 red vertices and 24 black vertices, no plane augmentation of G_3 is Hamiltonian.

Note that F_k has exactly six triangular faces—created by the chords added to the three rectangular prisms—while all other faces are bounded by fourcycles; let us call the latter faces *quadrilaterals*. Also note that F_k has one more face than it has vertices and, therefore, adding vertices to F_k in a stellating fashion yields a nonhamiltonian maximal planar graph. In fact, for every quadrilateral face f of F_k,

it suffices to connect the added vertex v_f only to three out of the four vertices on ∂f in order to make the resulting graph 3-connected. So this is exactly what we do to obtain G_k from F_k: In every face f, add a new vertex v_f and connect v_f to three selected vertices on ∂f; see Figure 94c. Although G_k is not maximal planar, by being 3-connected its combinatorial embedding is unique (up to its orientation), and by the Kleetope condition (an independent set of $4k + 13$ out of $8k + 25$ vertices, no two of which bound a common face) G_k is not subhamiltonian.

It remains to argue about the degree restriction, for which we will use the choice of which three vertices on the boundary of a face the added vertex is connected to. For a triangular face f, there is no choice: We have to add edges from v_f to all vertices on ∂f. But in every quadrilateral face f, we may select one vertex on ∂f to not receive an edge from v_f. In many faces, we can make that choice arbitrarily. But in quadrilaterals that are adjacent to a triangular face, we will make a specific choice, so as to ensure the degree restriction for vertices of separating triangles in G_k.

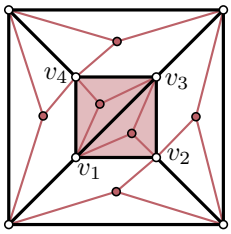

Figure 95: Distributing the edges of the added vertices v_f to achieve a degree of at most six at v_1, v_2, v_3, v_4.

For every vertex v we have $\deg_{G_k}(v) \leq 2 \deg_{F_k}(v)$ because a vertex can get at most one additional edge from each of its faces. The triangles in F_k appear along the added fourcycles in the three selected faces only. Let $v_1v_2v_3v_4$ denote such a fourcycle and let v_1v_3 denote the added chord, see Figure 95. Then $\deg_{F_k}(v_1) = \deg_{F_k}(v_3) = 4$ and $\deg_{F_k}(v_2) = \deg_{F_k}(v_4) = 3$. It immediately follows that $\deg_{G_k}(v_2) \leq 6$ and $\deg_{G_k}(v_4) \leq 6$. The vertices v_1 and v_3 get two more edges from their two triangles of F_k, bringing their degree up to six. But we can ensure that they do not get any more edge from the remaining two (quadrilateral) faces of F_k that contain them: since there is no quadrilateral face whose boundary passes through both v_1 and v_3, for every quadrilateral face f that contains v_1 or v_3, we select v_1 or v_3, respectively, to be the vertex on ∂f that does not get an additional edge from v_f. (The three faces to which we attached the cubes are pairwise nonadjacent in $C_4 \,\square\, P_3$ and, thus, every quadrilateral of F_k is adjacent to at most one triangle.) Therefore, we have $\deg_{G_k}(v_1) = \deg_{G_k}(v_3) = 6$. It follows that all vertices of separating triangles in G_k have degree at most six, as claimed. □

4.11 Conclusion

We have shown that $k = 5$ is the largest $k \in \mathbb{N}$ such that each 3-connected planar graph where the degree of vertices that belong to separating 3-cycles is bounded by k is subhamiltonian planar. It remains to investigate what happens when the connectivity condition is dropped. It was observed already by Bernhart and Kainen [22] that the book thickness of a graph is dominated by the maximum book thickness of its maximal 2-connected subgraphs (since it is easy to glue two book embeddings at a cut vertex). Therefore it suffices to consider the case of 2-connected graphs. Our main open question becomes:

1. Is every 2-connected planar graph in which every vertex of a separating 3-cycle has degree at most five subhamiltonian planar?

As a corollary of our main result, we also obtain that 3-connected planar graphs of maximum degree five are always subhamiltonian planar. In terms of 3-connected graphs, this improves the previous result of Bauernöppel [20], and, independently, Bekos, Gronemann, and Raftopoulou [21] who showed that (2-connected) planar graphs of maximum degree four are subhamiltonian planar. We are not aware of examples that show either of these two statements to be tight.

Bauernöppel's $\mathcal{NP}$-hardness proof [20] for the problem of recognizing subhamiltonian planar graphs generates 2-connected instances with maximum degree seven.

2. Is every 2-connected planar graph with maximum degree five subhamiltonian planar?

3. Is every 2-connected planar graph with maximum degree six subhamiltonian planar?

The Goldner-Harary graph (Figure 93) is a 3-connected planar graph with maximum degree eight that is not subhamiltonian planar.

4. Is every 3-connected planar graph with maximum degree six subhamiltonian planar?

5. Is every 3-connected planar graph with maximum degree seven subhamiltonian planar?

Chapter 5

Convexity-increasing morphs

5.1 Introduction

Broadly speaking, morphing refers to the act of continuously transforming a given graphical object or shape into another [65]. Gomes, Darsa, Costa, and Velho [65] provide an extensive survey and list several applications from the field of computer graphics such as animation, modeling, and geometric correction and matching. Here, we focus on 2-dimensional morphs between objects whose shapes can be represented by graph drawings. More specifically, a *morph* between two straight-line drawings Γ_0 and Γ_1 of a graph G is a continuous deformation that transforms Γ_0 into Γ_1 while preserving straight-line edges throughout the motion [3, 4, 7, 12]. In other words, each vertex moves from its position in Γ_0 to its position in Γ_1 and the edges follow along as line segments. A morph is *planar* if Γ_0 and Γ_1 are planar and the morph preserves planarity throughout the entire deformation.

It was established long ago [28, 120] that there *exists* a planar morph between any two planar straight-line drawings of the same plane graph. More recent results [3, 4, 7, 12] aim to *efficiently* compute morphs and decrease the "complexity" of the deformation. Instead of morphing between two given drawings, our focus is somewhat different, and more aligned with graph drawing goals: we are interested in transforming an initial drawing Γ into one that is *strictly convex*, that is, the boundary of each face is described by a strictly convex polygon. This drawing convention is motivated by visual qualities [109]. We say that such a morph *convexifies* Γ.

It is easy, using known results, to find a planar morph that convexifies a given planar straight-line drawing Γ of a plane graph G: we can just create a strictly convex drawing Γ_C of G (assuming such a drawing exists), and morph between Γ and Γ_C using the known planar morphing algorithms (for a discussion of the techniques used, see the Section 5.1.1). However, it has been established by empirical investigations [107] that, for the purposes of visualization, it is important to maintain the viewer's "mental map", which means changing as little as possible while making observable progress towards a goal. Hence, we would like our morphs to be *convexity-increasing*,

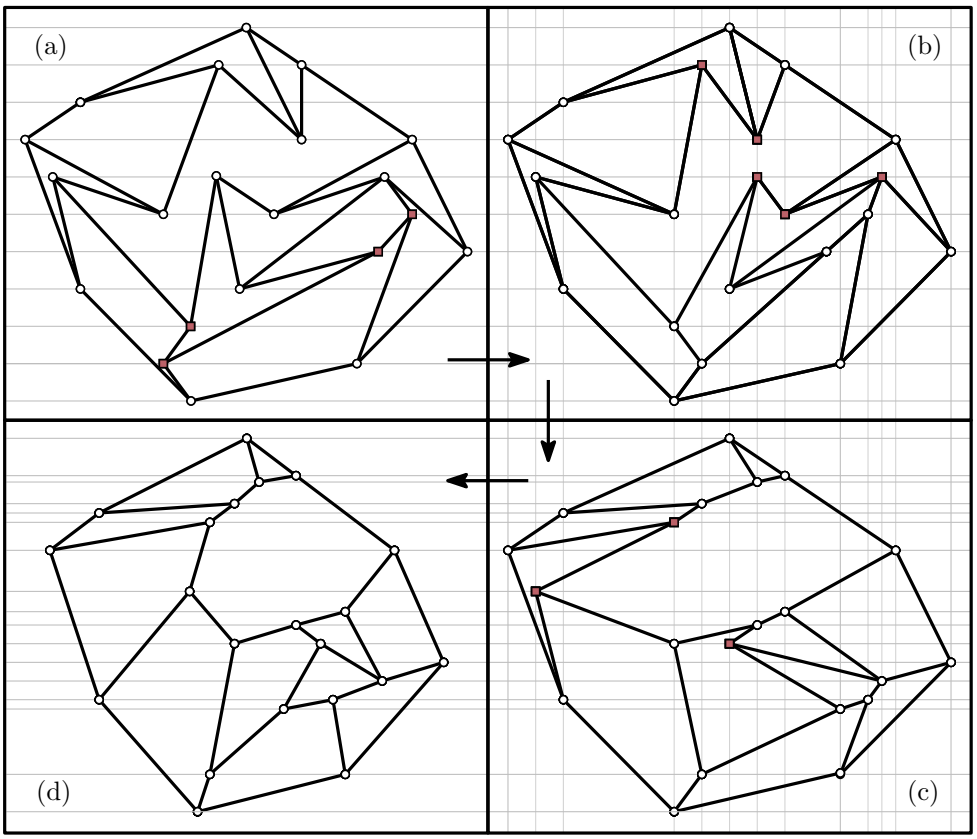

Figure 96: A sequence of convexity-increasing morphs (horizontal, vertical, horizontal) that morph a straight-line drawing of a graph G (a) into a strictly convex drawing of G (d). The marked vertices have reflex angles that are eliminated in the next step. All convex angles are preserved.

meaning that once an internal angle is strictly convex, it remains strictly convex, and once an outer angle is reflex, it remains reflex. We illustrate a convexity-increasing morph in Figure 96. Most previous morphing algorithms fail to provide convexity-increasing morphs even if the target is a strictly convex drawing because they start by triangulating the drawing. Therefore, an original strictly convex angle may be subdivided by new triangulation edges, so there is no constraint that keeps it strictly convex.

5.1.1 Related work

To the best of our knowledge, previous work on convexity-increasing morphs only considers the case when the input graph is a simple cycle: Connelly, Demaine, and Rote [37] and Streinu [117, 118] showed how to morph a simple polygon (i.e., a drawing of a simple cycle) to a convex polygon while preserving the length of each edge. Moreover, their motions are *expansive*, meaning that the pairwise vertex

distances do not decrease. The combination of these two properties implies that once an angle is 180°, it remains 180°. Hence, their morphs are convexity-increasing in a slightly different sense, namely, once an angle is convex, it remains convex (but not strictly convex). Other authors [2] showed how to convexify a given simple polygon while preserving vertex visibilities. This condition also implies the condition of being convexity-increasing.

There is an algorithm to morph a strictly convex drawing to another strictly convex drawing of the same graph while preserving planarity and strict convexity [10]. Such morphs are convexity-increasing by default, but they are not applicable to solve our problem since our initial drawing is not strictly convex.

Many previous morphing algorithms find "piece-wise linear" morphs [3, 4, 7, 12], which are composed of some number of linear morphing steps that can be efficiently encoded: a morph is called *linear* if each vertex moves along a straight-line segment at constant speed (depending on the length of the line segment) such that it arrives at its final position at the end of the morph. Note that vertices are allowed to move at different speeds, and some vertices may remain stationary. A linear morph is completely specified by the initial and final drawings. Hence, a morph that is a composition of linear morphs can be encoded as a sequence of drawings. A special case of linear morphs are *unidirectional* morphs. In these morphs, all the lines along which vertices move are parallel.

Alamdari et al. [3] gave an algorithm with runtime $O(n^3)$ that takes as input two planar straight-line drawings of a plane graph on n vertices. It then constructs a planar morph between the two drawings that consists of a sequence of $O(n)$ unidirectional morphs, which they show to be optimal in the worst case.

5.1.2 Main results and organization

We present an algorithm that convexifies a given straight-line planar drawing of a plane graph G via a planar convexity-increasing morph. The only requirement is that G admits a strictly convex drawing. This is the case if and only if G is internally 3-connected, see Section 5.2.1 for the definition and related discussions.

Our morphs are composed of a linear number of *horizontal* and *vertical* morphs. A *horizontal* morph is a unidirectional morph that moves vertices along horizontal lines. *Vertical* morphs are defined analogously. Figure 96 illustrates a sequence of horizontal and vertical morphs. Orthogonality is a well-studied graph drawing convention [50] with empirically proven visual qualities [13, 101, 108]. Similarly, it seems natural that orthogonal *motion* should be easier to comprehend. By using a shearing transformation, it is not difficult to turn a single linear morph into two horizontal / verticals morphs, but this will of course double the number of total steps. We count the number of morphing steps of our algorithm explicitly, and show that orthogonality can be achieved with a less significant increase in the number of steps.

Our algorithm has another advantage in terms of visualization: previous morphing

algorithms, such as the one by Alamdari et al. [3] or by Aichholzer et al. [2], move vertices very close together such that they almost coincide, effectively "contracting" the vertices, which destroys the user's "mental map" [107] of the graph. We do not use such contractions, which might result in more visually pleasing morphs in practice.

Our main result is summarized in the following theorem.

Theorem 58. *Let Γ be a planar straight-line drawing of an internally 3-connected graph G on n vertices. Then Γ can be morphed to a strictly convex drawing via a sequence of at most $3.5n+2$ convexity-increasing planar morphing steps each of which is either horizontal or vertical.*

In the special case that Γ has a convex outer face, the upper bound on the number of morphing steps is $\max\{2, r+1\}$, where r denotes the number of internal reflex angles of Γ.

Furthermore, there is a $O(n^{1+\omega/2} + n^2 \log n)$ time algorithm to find the sequence of morphs, assuming that two $n \times n$ matrices can be multiplied with $O(n^\omega)$ arithmetic operations.

The term ω appears in the run time due to the fact that one of the steps of our algorithm uses a variant of Tutte's graph drawing algorithm [124], for which we need to solve a linear system of equations. The run time is $O(n^{2.5})$ with Gaussian elimination and $O(n^{2.1865})$ when using the current fastest matrix multiplication method with $\omega \approx 2.3728639$ [95].

In Section 5.3, we will first prove a weaker version (Theorem 68) of Theorem 58, where the morphs are only required to maintain strict convexity of internal angles—outer reflex angles are allowed to switch their convexity status arbitrarily often. We call such a morph *weakly* convexity-increasing. We then extend the algorithm to also maintain outer reflex angles in Section 5.5.

Theorem 58 guarantees the existence of a convexity-increasing morph to a strictly convex drawing where the morph is composed of $O(n)$ horizontal/vertical morphs. In Section 5.6, we show that this bound is optimal in the worst case. In fact, we show something stronger:

Theorem 59. *For any $n \geqslant 3$, there exists a drawing of an internally 3-connected graph on n vertices for which any convexifying planar morph composed of a sequence of linear morphing steps requires $\Omega(n)$ steps.*

Our model of computation is the real-RAM. In particular, we do not have a polynomial bound on the bit-complexity of the coordinates of the vertices in the sequence of drawings that specify the morph. This is a common problem of most previous morphing algorithms. An exception is a morphing approach by Barrera-Cruz, Haxell, and Lubiw [19], which is based on Schnyder drawings and maintains a grid size of $O(n) \times O(n)$. Since this method only applies to special types of triangulations (so-called weighted Schnyder drawings), it does not seem applicable to solve our

problem. A main ingredient of most other morphing algorithms [3, 10, 40], including the one presented in this thesis, is a procedure for redrawing a given planar drawing to have convex faces while preserving the y-coordinate of each vertex (we formally define this concept in Section 5.2.4). In Section 5.7, we show that a single application of this procedure can blow up the width of a given drawing from $W \in \Theta(n)$ to $2^{W-2}(W-2)!$. Hence, it seems not possible to achieve a good resolution with our current approach for proving Theorem 58 and, more generally, with any approach that naively uses this redrawing procedure.

In the following Section 5.2, we introduce several techniques and concepts used throughout this chapter.

5.2 Preliminaries

We define (strictly) convex drawings and internally 3-connected graphs in Section 5.2.1. In Section 5.2.2, we discuss drawings with y-monotone faces, which play an important role for our morphing procedure. We proceed by stating several useful properties of unidirectional morphs in Section 5.2.3. Finally, we address the concept of finding (strictly) convex redrawings while preserving the y-coordinate of each vertex in Section 5.2.4.

Throughout this chapter, we use $\mathrm{x}(v)$ and $\mathrm{y}(v)$ to denote the x-coordinate and y-coordinate of a vertex v, respectively. All graphs in this chapter are simple.

5.2.1 Convex drawings and internal 3-connectivity

Given a planar straight-line drawing Γ of a graph, its *angles* are formed by pairs of consecutive edges around a face, with the angle measured inside the face. An *internal angle* is an angle of an inner face. We say an angle is *reflex* if it exceeds π, *convex* if it is at most π, and *strictly convex* if it is less than π. A drawing Γ is *convex* if the boundary of every face is a convex polygon, i.e., angles of the inner faces are convex and angles of the outer face are reflex or of size π. The drawing is *strictly* convex if the boundary of every face is a strictly convex polygon.

Conditions for the existence of convex drawings. Throughout, we assume that our input is a drawing of a plane graph that admits a strictly convex drawing. Necessary and sufficient conditions for the existence of a strictly convex drawing were given by Tutte [123], Thomassen [121], and Hong and Nagamochi [79]. These conditions can be tested in linear time using an algorithm presented by Chiba, Onogushi, and Nishizeki [31].

Such conditions are usually stated for a fixed (strictly) convex drawing of the outer face, but the conditions become simpler when, as in our case, the drawing of the outer face may be freely chosen. In particular, when it may be chosen to have no three consecutive collinear vertices. Internal vertices of degree 2 play a special role in

these conditions: in a convex drawing, 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. This has two implications: (1) a graph with an internal vertex of degree 2 has no strictly convex drawing, and (2) for a convex drawing we may eliminate every internal degree 2 vertex by repeatedly replacing a degree 2 vertex and its two incident edges by a single edge. Whenever this produces multiple edges, then there exists no convex drawing.

With these observations, the necessary and sufficient conditions for the existence of a (strictly) convex drawing become quite simple to state. A plane graph G is *internally 3-connected* if the graph is 2-connected and any separation pair u, v is *external*, meaning that u and v lie on the outer face and that every connected component of $G \setminus \{u, v\}$ contains a vertex of the outer face of G. Observe that the two neighbours of an internal vertex of degree 2 form a separation pair that is not external. The results of Tutte [123], Thomassen [121], and Hong and Nagamochi [79] become:

Lemma 60 ([79, 121, 123])**.** *Let G be a plane graph. Then*

1. *G has a strictly convex drawing if and only if G is internally 3-connected.*
2. *G has a convex drawing if and only if repeatedly eliminating internal vertices of degree 2 produces a graph that has no multiple edges and is internally 3-connected.*

Note that a separation pair that is not external *can* have both of its vertices on the outer face, see Figures 97(b) and 97(c). For this reason, we refer to a separation pair that is not external as *nonexternal*, rather than using the more straightforward, but misleading term *internal*.

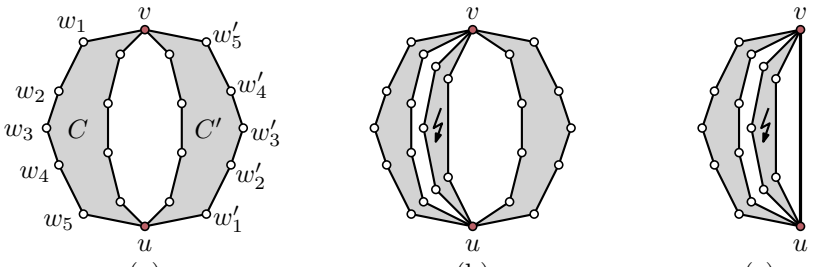

Figure 97: (a) An external separation pair u, v and its two components C and C'. (b,c) In these cases, u, v form a nonexternal separation pair since the component marked with a jagged arrow has no vertex on the outer face, which implies that there is no convex drawing of G.

Structure of internally 3-connected graphs. There are multiple well-known equivalent definitions of internal 3-connectivity. Each of them provides a different perspective on the concept and it will be convenient to refer to all of them. Hence, we state the following well-known characterization:

Lemma 61. *Let G be a plane 2-connected graph and let f_o denote its outer face. The following statements are equivalent:*

(K1) G is internally 3-connected.

(K2) Inserting a new vertex v in f_o and adding edges between v and all vertices of f_o results in a 3-connected graph.

(K3) From each internal vertex w of G there exist three paths to f_o that are pairwise disjoint except for the common vertex w.

Proof. (K1) obviously implies (K2), which in turn implies (K3) by 3-connectivity. It remains to show that (K3) implies (K1). So suppose that from each internal vertex w there exist three paths to f_o that are disjoint except for the common vertex w. Consider a pair of vertices p, q and an arbitrary vertex w of $G \setminus \{p, q\}$. Then the connected component of w in $G \setminus \{p, q\}$ has at least one vertex that belongs to f_o: this is clear if w belongs to f_o. Otherwise, w is an internal vertex of G and, hence, there are three disjoint path from w to f_o, and only two of these paths can be destroyed by removing p and q.

Moreover, we claim that if p, q is a separation pair of G, then both p and q belong to f_o. Towards a contradiction, assume that the conclusion of this claim does not hold. Since G is 2-connected, f_o is a *simple* cycle and, hence, all vertices of $f_o \setminus \{p, q\}$ belong to the same connected component of $G \setminus \{p, q\}$. However, by assumption from each internal vertex w there exist at least one path to f_o in $G \setminus \{p, q\}$ and, thus, $G \setminus \{p, q\}$ is connected, which yields the desired contradiction. □

The following statement gives a characterization of external separation pairs and describes the structure of internally 3-connected graphs. For an illustration see Figure 97(a).

Observation 62. *Let H be a plane 2-connected graph and let u, v be a separation pair of H. Then, the separation pair u, v is external if and only if all of the following conditions hold:*

(E1) Vertices u and v belong to the outer face of H.

(E2) The outer face of H decomposes into two internally disjoint paths $(u, w_1, \ldots, w_j, v)$ and $(v, w'_1, \ldots, w'_\ell, u)$ each with at least 3 vertices, i.e. $j \geq 1$ and $l \geq 1$.

(E3) Vertices $w_1, \ldots, w_j$ belong to a connected component C of $H \setminus \{u, v\}$.

(E4) Vertices $w'_1, \dots, w'_j$ belong to a connected component C' of $H \setminus \{u, v\}$.

(E5) The graph $H \setminus \{u, v\}$ has no connected component other than C, C'.

(E6) The components C and C' are distinct.

Proof. If the six conditions hold, then clearly u, v form an external separation pair. On other hand, if u, v form an external separation pair, then u and v belong to the outer face of H (E1). Since H is 2-connected, its outer face is a *simple* cycle. Further, the removal of u and v splits the graph into at least two connected components each of which has a vertex that belongs to the outer face of H. Hence, the removal of u and v decomposes the outer face into two internally disjoint paths $(u, w_1, \dots, w_j, v)$ and $(v, w'_1, \dots, w'_\ell, u)$ each with at least 3 vertices (E2). Since $(w_1, \dots, w_j)$ is a path in $H \setminus \{u, v\}$, its vertices belong to a connected component C (E3). Analogously, $(w'_1, \dots, w'_j)$ is a path in $H \setminus \{u, v\}$, its vertices belong to a connected component C' (E4). Since these two paths together with u and v cover the entire outer face of H, there can not be any more components (E5). Finally, since there are at least two components, C and C' have to be distinct (E6). □

5.2.2 Drawings with y-monotone faces

Let G be a 2-connected planar graph and consider a planar drawing Γ of G. A face of Γ is *y-monotone* if its boundary consists of two y-monotone paths. A path is *y-monotone* if the y-coordinates of the points along the curve realizing the path are strictly increasing. These definitions apply to general planar graph drawings, not just straight-line drawings. We remark that if all faces of Γ are y-monotone and, additionally, the y-coordinates of vertices are pairwise distinct, then Γ is y-monotone (cf. Section 3.1.4). The converse is not true in general. *Directed graphs* that have drawings with y-monotone faces are the so-called st-planar graphs, which are well-studied [47].

Assume that each edge of Γ is a y-monotone curve. We say a vertex v is a *local minimum* (*local maximum*) of a face f in Γ if both neighbors of v in f lie above (below) v. A *local extremum* refers to a local minimum or a local maximum. Note that a face f of Γ is y-monotone if and only if it has exactly one local maximum and exactly one local minimum. Equivalently, in case the edges of Γ are polygonal chains, an interior (exterior) face f is y-monotone if and only if it has no *reflex* (*convex*) local extremum, that is, a local extremum whose angle in f is reflex (convex).

5.2.3 Linear and unidirectional morphs

A linear morph is completely specified by the initial and the final drawing. To denote the unique linear morph from a drawing Γ_1 to a drawing Γ_2, we use the notation $\langle \Gamma_1, \Gamma_2 \rangle$. Restricting to linear morphs is a sensible way to discretize morphs. Essentially, it amounts to requiring piece-wise linear vertex trajectories. At first

glance, further restriction to unidirectional morphs seems arbitrary and too restrictive. However, unidirectional morphs are conceptually much simpler than linear morphs and they have several useful properties, some of which are described in this section. Moreover, we show that the necessary number of *linear* morphing steps for convexifying a given drawing (Theorem 59) is always sufficient even when restricted to *unidirectional* morphs (Theorem 58).

Suppose we perform a horizontal morph. Then every vertex keeps its y-coordinate. Alamdari et al. [3] gave conditions on the initial and final drawing that guarantee that the horizontal morph between them is planar:

Lemma 63 ([3])**.** *If Γ and Γ' are two planar straight-line drawings of the same graph such that every line parallel to the x-axis crosses the same ordered sequence of edges and vertices in both drawings, then the linear morph $\langle \Gamma, \Gamma' \rangle$ is planar.*

The statement of Lemma 63 is formulated as in intermediate claim in the proof of [3, Lemma 13], but the authors do not state it as an independent theorem. It is proven in the proof of [3, Lemma 13]. The idea is that it suffices to show that if two points p and q move along a horizontal line at constant speeds for a unit of time and q is to the right of p in the initial and final position, then q is to the right of p at every time instant throughout the movement. The latter claim is formulated as [3, Lemma 14] and follows directly by linearity.

Observe that the conditions of Lemma 63 imply that every vertex is at the same y-coordinate in Γ and Γ' and, hence, the linear morph between them is horizontal. Also note that the lemma generalizes in the obvious way to any direction, not just the direction of the x-axis. We note several useful consequences:

Lemma 64. *Let $\Gamma_1, \Gamma_2, \Gamma_3$ be three planar straight-line drawings of the same graph where the linear morphs $\langle \Gamma_1, \Gamma_2 \rangle$ and $\langle \Gamma_2, \Gamma_3 \rangle$ are horizontal and planar. Then the linear morph $\langle \Gamma_1, \Gamma_3 \rangle$ is a horizontal planar morph.*

Proof. The morphs $\langle \Gamma_1, \Gamma_2 \rangle$ and $\langle \Gamma_2, \Gamma_3 \rangle$ are horizontal and planar, so every line parallel to the x-axis crosses the same ordered sequence of edges and vertices in Γ_1 and Γ_3. Then by Lemma 63 the morph $\langle \Gamma_1, \Gamma_3 \rangle$ is horizontal and planar. □

Lemma 64 allows us to merge two morphs $\langle \Gamma_1, \Gamma_2 \rangle$ and $\langle \Gamma_2, \Gamma_3 \rangle$ as in the statement into a single morph, which we will use to decrease the overall number of morphing steps.

Lemma 65. *Let Γ_1, Γ_2 be two planar straight-line drawings of the same graph such that $\langle \Gamma_1, \Gamma_2 \rangle$ is a horizontal planar morph. Then the convexity status of an angle β formed by two edges $\{a, b\}, \{b, c\}$ at a vertex b can change at most once during the morph $\langle \Gamma_1, \Gamma_2 \rangle$, i.e., β cannot change from $< 180°$ to $\geq 180°$ and then back to $< 180°$, or from $\geq 180°$ to $< 180°$ and then back to $\geq 180°$.*

Moreover, if every strictly convex internal angle of Γ_1 is also strictly convex in Γ_2, then the morph $\langle \Gamma_1, \Gamma_2 \rangle$ is weakly convexity-increasing. If additionally, every outer

reflex angle of Γ_1 *is also reflex in* Γ_2*, then the morph* $\langle \Gamma_1, \Gamma_2 \rangle$ *is convexity-increasing.*

Lemma 65 is a generalization of [10, Lemma 7], which states that a horizontal morph between two strictly convex drawings that satisfy the preconditions of Lemma 63 preserves strict convexity. The proof of [10, Lemma 7] uses basic properties of unidirectional morphs from [3]. In contrast, we prove the generalization Lemma 65 using the orientation test, a standard operation in computational geometry [42].

Proof of Lemma 65. The orientation test [42] determines the orientation of a sequence (p, q, r) of points $p = (p_x, p_y)$, $q = (q_x, q_y)$, $r = (r_x, r_y)$ in the plane by checking the sign of the determinant

$$\begin{vmatrix} p_x & q_x & r_x \\ p_y & q_y & r_y \\ 1 & 1 & 1 \end{vmatrix} = p_x q_y + q_x r_y + p_y r_x - q_y r_x - p_y q_x - p_x r_y =: \mathrm{ori}(p, q, r).$$

Specifically, $\mathrm{ori}(p, q, r) > 0$ if (p, q, r) appear in counterclockwise order, $\mathrm{ori}(p, q, r) < 0$ if (p, q, r) appear in clockwise order, and $\mathrm{ori}(p, q, r) = 0$ if p, q, r are colinear. Hence, the convexity status of an angle β formed by two edges $\{a, b\}$ and $\{b, c\}$ can be checked by considering the sign of $\mathrm{ori}(a, b, c)$.

The morph $\langle \Gamma_1, \Gamma_2 \rangle$ is horizontal. Hence, the y-coordinates a_y, b_y, c_y of a, b, c are maintained, while their x-coordinates $a_x(t), b_x(t), c_y(t)$ vary linearly with the time t. Plugging $(a_x(t), a_y)$, $(b_x(t), b_y)$, $(c_x(t), c_y)$ into ori, we obtain the function

$$a_x(t)b_y + b_x(t)c_y + a_y c_x(t) - b_y c_x(t) - a_y b_x(t) - a_x(t)c_y =: f(t),$$

which is linear in t, and whose sign describes the convexity status of the angle β at time t. By linearity, $f(t)$ is either constantly 0, or it is 0 at most once during any time interval. Consequently, the convexity status changes at most once during $\langle \Gamma_1, \Gamma_2 \rangle$, which proves the first statement of the lemma.

We prove the second statement of the lemma by contrapositive. Suppose the horizontal morph $\langle \Gamma_1, \Gamma_2 \rangle$ is not (weakly) convexity-increasing. Then some angle α looses its desired status during the morph. The first statement of the lemma implies that α also does not have its desired status in Γ_2. □

Alamdari et al. gave the following further condition that implies the precondition of Lemma 63. We emphasize that the statement applies to planar graph drawings in general, that is, edges are not required to be straight-line.

Observation 66 ([3])**.** *Let* G *be a plane graph and let* Γ *and* Γ' *be two planar graph drawings of* G*. Further, assume that all faces in both drawings are* y*-monotone and each vertex of* G *has the same* y*-coordinate in both drawings. Then every line parallel to the* x*-axis crosses the same ordered sequence of edges and vertices in both drawings.*

The statement of Observation 66 is formulated as an intermediate claim in the proof of [3, Lemma 13] for the case of straight-line drawings without horizontal edges. It is not stated as an independent theorem. We observe that the proof of Alamdari et al., also discussed in the proof of [3, Lemma 13], does not actually make use of the fact that the edges are straight-line segments. Instead, it suffices that they are y-monotone. For the sake of self-containment, we summarize the proof by Alamdari et al. in terms of graph drawings with y-monotone edges:

Proof of Observation 66. Consider a horizontal line ℓ that intersects a face f of Γ. Since f is y-monotone, ℓ crosses the boundary of f in up to 2 points. Since the y-coordinates of vertices are the same in both drawings, the crossings belong to the same edges/vertices of f in both drawings. Moreover, since Γ and Γ' are drawings of the same plane graph, the order in which the edges/vertices of f are crossed by ℓ is identical. Hence, it remains to argue that ℓ crosses the same sequence of faces. Again this follows from the fact that Γ and Γ' are drawings of the same plane graph and each vertex has the same y-coordinate in both drawings. This implies that if ℓ crosses the interior of two faces f and f' consecutively in Γ, then it also crosses the interiors of f and f' consecutively in Γ', even in the case that the intersection of f, f', and ℓ is a vertex. □

To make use of the above lemmata, we will again follow Alamdari et al. [3] and use a procedure for redrawing a given planar drawing to have convex faces while preserving the y-coordinate of each vertex, which is described in the next section.

5.2.4 Redrawing with convex faces while preserving y-coordinates

A main ingredient of many morphing algorithms [3, 10, 40], including the one presented in this thesis, is a procedure for redrawing a given planar (straight-line) drawing to have convex faces while preserving the y-coordinate of each vertex. An idea for obtaining such a procedure was described already in 1996 by Chrobak, Goodrich, and Tamassia in a paper [33] that was only published as a "preliminary version". Their approach is based on using Tutte's graph drawing algorithm [124] and it has been reused and extended, for instance, in [116] and [111]. However, the idea does not appear to be widely known. Instead, modern graph morphing papers [3, 10, 40] build on a well-known convex redrawing technique described by Hong and Nagamochi [79] in 2012. The drawings generated by the algorithm of Hong and Nagamochi [79] are in general not strictly convex since the corresponding recursive algorithm draws certain paths of the graph by placing all its vertices on a common line. Angelini, Da Lozzo, Frati, Lubiw, Patrignani, and Roselli [10] extended the result to strictly convex faces by perturbing vertices to avoid angles of 180°. Since the idea by Chrobak et al. is based on Tutte's algorithm, it automatically produces strictly convex drawings. The algorithm of Hong and Nagamochi has a runtime of $O(n^2)$. Angelini et al. did not analyze the runtime of their extension. Chrobak et al. pointed out [33] that their

approach can be implemented in time $O(n^{\omega/2} + n \log n)$ by using the generalized nested dissection method by Lipton, Rose, and Tarjan [98, 99].

For the morphing algorithm presented in this thesis, we rekindle the idea of Chrobak et al. [33]. Using an improved version of generalized nested dissection due to Alon and Yuster [5] from 2013, we give a very simple and short proof, based on the ideas of Chrobak et al., of the strengthened version of Hong and Nagamochi's well-known result while achieving the improved runtime of $O(n^{\omega/2} + n \log n)$.

Tutte's algorithm is a very flexible framework that can be used to achieve additional properties, which is another advantage of using the method by Chrobak et al. In Section 5.5, we use the freedom that this framework provides to extend our algorithm for the weakly convexity-increasing case to also maintain outer reflex angles.

In the following, we give the precise formulation of the redrawing technique. The proof is discussed in Section 5.4.

Lemma 67 ([10, 33, 79])**.** *Let Γ be a planar drawing of an internally 3-connected graph G such that every face is y-monotone. Let C be a strictly convex straight-line drawing of the outer face of G such that every vertex of C has the same y-coordinate as in Γ. Then there is a strictly convex straight-line drawing Γ' of G such that the subdrawing of the outer face is C and every vertex of Γ' has the same y-coordinate as in Γ.*

Furthermore, the drawing Γ' can be found in time $O(n^{\omega/2} + n \log n)$, even if only the underlying abstract graph of G, the cycle corresponding to C, and the y-coordinates of vertices are given, assuming that two $n \times n$ matrices can be multiplied with $O(n^{\omega})$ arithmetic operations.

5.3 Computing weakly convexity-increasing morphs

In this section, we prove the following theorem, which is a weaker version of Theorem 58. It is mainly concerned with maintaining internal strictly convex angles (though, if the outer face of the given drawing is already convex, the produced morphs are convexity-increasing in the strong sense).

Theorem 68. *Let Γ be a planar straight-line drawing of an internally 3-connected graph G on n vertices. Then Γ can be morphed to a strictly convex drawing via a sequence of at most $3.5n + 2$ weakly convexity-increasing planar morphing steps each of which is either horizontal or vertical.*

In the special cases that G is 3-connected or Γ has a convex outer face, the upper bound on the number of morphing steps can be improved to $1.5n + 2$ or $\max\{2, r + 1\}$, respectively, where r denotes the number of internal reflex angles of Γ. Moreover, if Γ has a convex outer face, the morph is convexity-increasing.

Furthermore, there is an $O(n^{1+\omega/2} + n^2 \log n)$ time algorithm to find the sequence of morphs, assuming that two $n \times n$ matrices can be multiplied with $O(n^{\omega})$ arithmetic operations.

In fact, we show multiple variants of Theorem 68, starting with a highly specialized version and proceeding to more and more general ones, which use the more specialized cases as building blocks. So let Γ be a planar straight-line drawing of an internally 3-connected graph G on n vertices. Our goal is to convexify Γ by means of a weakly convexity-increasing morph.

Section overview. In Section 5.3.1, we start with the simple case that all faces of Γ are y-monotone. The purpose of this section is to demonstrate how all the lemmata and observations from Sections 5.2.3 and 5.2.4 fit together. Indeed, by simply combining the tools from these sections, we obtain a horizontal convexity-increasing morph to a strictly convex drawing of G.

The first more challenging case is discussed in Section 5.3.2, where we prove Theorem 68 for the case that the outer face of Γ is already drawn as a convex polygon. The idea is to augment Γ to be y-monotone and apply the algorithm for the y-monotone case to convexify some angles. After that, we repeat the same strategy, but we augment the drawing to be x-monotone and perform a vertical morph to convexify further angles. We continue to alternate between horizontal and vertical morphs until the drawing is convexified.

In Section 5.3.3, we allow the outer face to be realized as a nonconvex polygon, but we restrict our attention to the case that G is 3-connected (rather than internally 3-connected). The idea is to augment the outer face by adding its convex hull edges and then apply the algorithm for the convex case. After that, we iteratively remove the additional edges while performing a constant number of morphing steps after each removal.

Finally, in Section 5.3.4, we prove the general case of Theorem 68, that is, we allow G to be internally 3-connected. The reason why the algorithm for the 3-connected case does not apply to internally 3-connected graphs is that adding the convex hull edges may create nonexternal separation pairs. The idea to overcome this problem is to augment the outer face of Γ by creating a "buffer" layer of new vertices and edges that ensure that adding the convex hull edges preservers internal 3-connectivity. After that, we apply the algorithm for the 3-connected case and then iteratively remove the vertices that do not belong to G while performing a constant number of morphing steps after each removal.

5.3.1 A simple case: morphing y-monotone drawings

To give some intuition about our general proof strategy, we first consider an easy case where all faces (including the outer face) of Γ are y-monotone. We choose some strictly convex polygon C as a redrawing of the outer face of Γ that preservers the y-coordinate of each vertex. We can immediately apply Lemma 67 with C as the prescribed drawing of the outer face to obtain a new straight-line strictly convex drawing Γ' of G in which each vertex has the same y-coordinate as in Γ. By

Observation 66, every line parallel to the x-axis crosses the same ordered sequence of edges and vertices in Γ and in Γ'. Thus, by Lemma 63 the morph $\langle \Gamma, \Gamma' \rangle$ is planar. Moreover, it is a horizontal morph. Hence, we have a morph from Γ to a strictly convex drawing Γ' by way of a single horizontal morph. Furthermore, this morph is weakly convexity-increasing by Lemma 65 since every internal angle is strictly convex in the final drawing Γ' (in fact, the morph is even convexity-increase in the strong sense since *every* angle has its desired status in Γ').

5.3.2 Morphing drawings with a convex outer face

In this section, we consider the case that the subdrawing C of the outer face in Γ is convex (but not necessarily strictly convex) and the inner faces of Γ are not necessarily y-monotone. Additionally, we assume that Γ has no horizontal edge; we will later show how to ensure this.

Overview. As in Section 5.3.1, we begin by performing a horizontal morph. Observe that such a morph preserves the local extrema and does not change their convexity status. Thus, the only inner angles that can be convexified via a horizontal morph are the *h-reflex* angles, that is, those reflex or 180° angles of an inner face f that are not local extrema of f. We will show that a single horizontal morph suffices to convexify all these angles. The plan is to then repeat the procedure after conceptually "turning the paper" by 90°, that is, we proceed by performing a vertical morph to make every *v-reflex* angle strictly convex, where a v-reflex angle is defined as an inner angle that becomes h-reflex after rotating the drawing by 90°. By continuing to alternate between the horizontal and the vertical direction, we eventually end up with a strictly convex drawing and, thus, prove Theorem 68 for the case of a convex outer face.

A horizontal morphing step. To find the desired horizontal morph, we will apply Lemma 67. Therefore, we must first augment Γ to a drawing with y-monotone faces by inserting y-monotone edges (that are not necessarily straight-line). For an example see Figure 98. This is a standard operation in the context of drawing styles that use y-monotone edges, see for example [47, Lemma 4.1] [104, Lemma 3.1], and it is also used in an algorithm for triangulating a simple polygon [42, Section 3.2]. However, the approaches in [47, 104] are not directly applicable in our case since we need our augmented drawing to satisfy an additional constraint: the new edges are only allowed to be incident to local extrema since otherwise we would relinquish control of convexity at an incident vertex.

Lemma 69. *Let Γ be straight-line planar drawing of an internally 3-connected graph G without horizontal edges. Then Γ can be augmented by adding edges into the inner faces to obtain a planar drawing Γ' such that each additional edge is a*

y-monotone curve joining two local extrema of some face in Γ, every inner face is y-monotone, and the augmented plane graph G' is internally 3-connected. Furthermore, the plane graph G' can be computed in $O(n \log n)$ time.

Proof. Recall that a face of a drawing with y-monotone polygonal edges is y-monotone if and only if it has no reflex local extremum. Our proof is by induction on the number of reflex local extrema in all inner faces of the drawing. If there are none, then all inner faces are y-monotone. Otherwise, consider an inner face f that has a local reflex extremum u. Without loss of generality, u is a local minimum. For an illustration consider Figure 98.

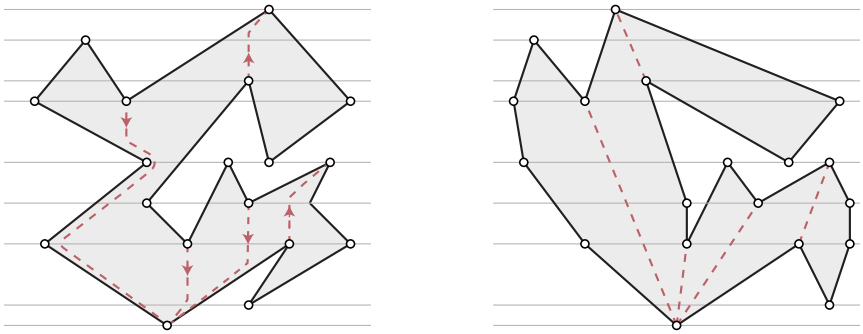

Figure 98: (a) A face that is not y-monotone. The dashed edges inside the face are added by Lemma 69. (b) The face after application of Lemma 70.

We want to find a local extremum v below u such that we can insert a y-monotone curve from v to u within face f. To construct the curve go vertically downwards from u to the first point p_u on the boundary of f, and then follow a y-monotone path along the boundary of f downwards from p_u to some local minimum v. Adding the edge $\{u, v\}$ splits f into two faces, and decreases the total number of local reflex extrema.

Note that the procedure cannot create multiple edges since all edges remain y-monotone and the new edge connects a vertex u with a vertex v such that all neighbors of u have larger y-coordinates than u, but the vertex v has a smaller y-coordinate than u. Since we only insert inner edges, the graph remains internally 3-connected.

We remark that Pach and Tóth [104] used a similar idea to triangulate with monotone curves, although their curves stopped at the first vertex on f's boundary and we must continue to the first local minimum.

To complete the proof we describe how the set of augmenting edges can be found in time $O(n \log n)$. We deal with the local reflex minima; the maxima can be dealt with in a second phase. Find a trapezoidal decomposition [42] of the drawing in $O(n \log n)$ time (in such a decomposition, each face is partitioned into trapezoids by shooting two rays r_t and r_b to the top and bottom at each vertex v of the face, and

adding segments between v and the first edge intersected by r_t and r_b, respectively). This gives the point p_u for each local reflex minimum u in face f. By traversing each face once, we can preprocess the graph in total time $O(n)$ to find, for each edge e in face f, the local minimum v that is reached by following a y-monotone curve downward from e in f. This gives the set of augmenting edges and, hence, the underlying *abstract* graph of the desired *plane* graph G'. To determine G' (i.e., the combinatorial embedding), for each vertex v, we sort in each incident face f the new edges incident to v according to the order in which their endpoints appear along f, which is easily done in total time $O(n \log n)$. (The embedding can even be determined in $O(n)$ time by traversing each face and adding the edges at each vertex in the order in which they are discovered.) □

This observation allows us to prove the following:

Lemma 70. *Let Γ be a straight-line planar drawing of an internally 3-connected graph G with a convex subdrawing C of the outer face and no horizontal edge. There exists a straight-line planar drawing Γ' of G such that Γ' has a strictly convex outer face and each internal angle of Γ' whose corresponding angle in Γ is not a local extremum is strictly convex in Γ'. Furthermore, the morph $\langle \Gamma, \Gamma' \rangle$ is convexity-increasing, horizontal, and planar and can be found in time $O(n^{\omega/2} + n \log n)$, assuming that two $n \times n$ matrices can be multiplied with $O(n^\omega)$ arithmetic operations.*

Proof. If C is not strictly convex, we choose some strictly convex redrawing C' of C that preserves the y-coordinate of each vertex. Otherwise, we set $C' = C$.

By Lemma 69 there exists a set A of edges that can be drawn as a set Γ_A of y-monotone curves such that $\Gamma \cup \Gamma_A$ is a planar drawing of $G \cup A$ in which all faces are y-monotone, and any edge of A is incident to two local extrema of some inner face. Its underlying plane graph G' can be computed in $O(n \log n)$ time.

Apply Lemma 67 to G' and C' with the y-coordinates as in Γ to obtain in time $O(n^{\omega/2} + n \log n)$ a new straight-line strictly convex drawing $\Gamma' \cup \Gamma'_A$ of G' with each vertex at the same y-coordinate as in Γ. Here, Γ'_A denotes the set of straight-line edges corresponding to A. By Observation 66 every line parallel to the x-axis crosses the same ordered sequence of edges and vertices in $\Gamma \cup \Gamma_A$ and in $\Gamma' \cup \Gamma'_A$. This property carries over to the two corresponding subdrawings of G, i.e, every line parallel to the x-axis crosses the same ordered sequence of edges and vertices in Γ and Γ'. Consequently, by Lemma 63 the horizontal morph $\langle \Gamma, \Gamma' \rangle$ is a planar.

Each internal angle of Γ that is not a local extremum has no edge of A incident to it, and thus becomes strictly convex in Γ'. Any internal angle of Γ that is a local extremum maintains its convexity status status in Γ' since $\langle \Gamma, \Gamma' \rangle$ is horizontal and planar. Moreover, since C' is strictly convex, every outer angle in Γ' is reflex. Thus by Lemma 65 the morph is convexity-increasing.

The time to find the morph (i.e., to find Γ') is dominated by the application of Lemma 67 and sums up to $O(n^{\omega/2} + n \log n)$. □

Making progress. Lemma 70 generalizes to directions other than the horizontal direction: for any direction d that is not parallel to an edge of Γ, there exists a convexity-increasing unidirectional (with respect to d) morph that convexifies all internal angles that are not extreme in the direction orthogonal to d. Thus, if we do not insist on a sequence of *horizontal and vertical* morphs, we immediately obtain a proof of Theorem 68 for the case of a convex outer face, since, for each inner angle that is not strictly convex, we can choose a direction d that convexifies it and, therefore, we need at most one morph for each of these angles (in the case that all inner angles are strictly convex, but the outer face is not strictly convex, we need a single morphing step).

In order to prove the stronger result that uses only two orthogonal directions, we need to alternate between the horizontal and the vertical direction. Note that after one application of Lemma 70, we do not necessarily obtain a drawing that contains a v-reflex vertex, see Figures 99(a) and 99(b). To ensure that our algorithm makes progress after every step, we prove a strengthened version of Lemma 70 that ensures that there is at least one h-reflex or v-reflex vertex available after each step.

Lemma 71. *Let Γ be a straight-line planar drawing of an internally 3-connected graph with a convex subdrawing of the outer face and no horizontal edge. There exists a convexity-increasing, horizontal, and planar morph to a straight-line drawing Γ'' such that*

(i) the outer face of Γ'' is strictly convex,

(ii) each internal angle of Γ that is not a local extremum is strictly convex in Γ'',

(iii) Γ'' has no vertical edge, and

(iv) if Γ'' is not strictly convex, then it has at least one v-reflex angle.

Furthermore, the morph can be found in time $O(n^{\omega/2} + n \log n)$, assuming that two $n \times n$ matrices can be multiplied with $O(n^\omega)$ arithmetic operations.

Proof. We first apply Lemma 70 to obtain a morph from Γ to a drawing Γ' that satisfies (i) and (ii). If Γ' satisfies all the requirements, we are done. Otherwise, we will achieve the properties (iii) and (iv) by shearing the drawing Γ'. Eliminating vertical edges via a horizontal shear is easy, so we concentrate on the requirement (iv) about v-reflex angles.

Suppose Γ' is not strictly convex and has no v-reflex angle. Consider any angle that is not strictly convex at some vertex u in an internal face f in Γ'. By property (ii), u must be a local extremum in Γ' and, hence, its angle in f is reflex (rather than 180°). We apply a shearing transformation along the x-axis to create a drawing Γ'' in which the angle at u in f becomes v-reflex, i.e., in which the x-coordinate of u is between the x-coordinates of its two neighbors in f, see Figure 99. Furthermore,

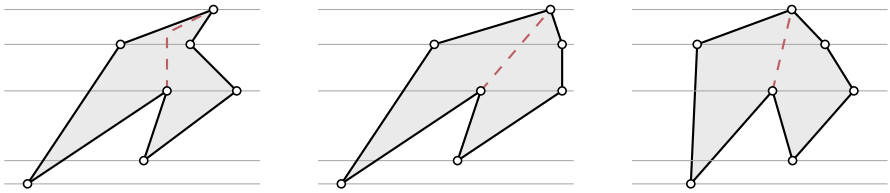

(a) (b) (c)

Figure 99: (a) A face that is not y-monotone. (b) The face after application of Lemma 70. There is a vertical edge and the unique reflex vertex is not v-reflex. (c) After applying a horizontal shear transformation, the reflex vertex is v-reflex and there are no vertical edges.

the shear should eliminate all vertical edges. It is easy to compute the corresponding shearing factor in linear time.

Since shearing is an affine transformation, it preserves the convexity status of each angle and, thus, Γ'' satisfies the properties (i)–(iv). The horizontal morph $\langle \Gamma', \Gamma'' \rangle$ is planar by Lemma 63 and it the preserves the convexity status of each angle by Lemma 65. By Lemma 64 the morph $\langle \Gamma, \Gamma'' \rangle$ is a horizontal planar morph. By Lemma 65 it is convexity-increasing.

The time to find the morph (i.e., to find Γ'') is dominated by the application of Lemma 70 and sums up to $O(n^{\omega/2} + n \log n)$. □

We are now ready to prove Theorem 68 for the case of a convex outer face.

Proof of Theorem 68 for the case of a convex outer face. If the given drawing Γ has a horizontal edge and/or if the set of internal reflex angles is nonempty and contains no h-reflex angle, then we use one vertical shear as in the proof of Lemma 71 to remedy this. Then, in the special case that there is no internal reflex angle, but there are angles of degree exactly 180°, we may apply Lemma 71 once to obtain the desired strictly convex drawing. Otherwise, there are internal reflex angles, and we apply Lemma 71 alternately in the horizontal and vertical directions until the drawing is strictly convex. In each step there is at least one h-reflex or v-reflex angle that becomes strictly convex. All internal angles of degree exactly 180° becomes strictly convex after the first application of Lemma 71. Thus, the number of morphing steps is at most $\max\{2, r+1\} \leq n$, where r is the number of internal reflex angles in Γ. The resulting total runtime is $O(n^{1+\omega/2} + n^2 \log n)$. □

5.3.3 Morphing drawings of 3-connected graphs

In this section, we no longer insist on the convexity of the subdrawing of the outer face of Γ. However, we are still not ready to prove Theorem 68 in its general form. Instead, we assume the given graph G to be 3-connected (rather than just *internally* 3-connected).

Overview. On a high level, our approach works as follows: we start by augmenting the outer face of Γ with edges from its convex hull to obtain a drawing of an augmented graph with a convex outer face. We then apply the results from Section 5.3.2 to morph to a strictly convex drawing and then remove the extra edges on the outer face one by one. After each removal of an edge, we morph to a strictly convex drawing of the reduced graph using at most three horizontal or vertical morphs. Let us proceed by discussing these steps in more detail.

Augmenting the outer face. We begin by computing the convex hull of Γ in $O(n)$ time [96]. Any segment of the convex hull that does not correspond to an edge of G becomes a new edge that we add to G. Let A denote the new edges and let $G \cup A$ denote the augmented graph with straight-line planar drawing $\Gamma \cup \Gamma_A$. Note that adding edges maintains 3-connectivity. Each edge $e \in A$ is part of the boundary of an inner face f_e of $\Gamma \cup \Gamma_A$. We call f_e the *pocket* of e.

We now apply the results of Section 5.3.2 to obtain a morph to a strictly convex drawing $\Gamma' \cup \Gamma'_A$ of $G \cup A$, see Figure 100(a). Recall that the techniques used in that section give us a drawing without horizontal or vertical edges. Note that this step makes the outer angles of all outer vertices of G that do not belong to the convex hull strictly convex, even though their desired convexity status is reflex. Hence, the morph is only weakly convexity-increasing in general.

We remark that this step is the reason why we limit ourselves to 3-connected graphs in this section: adding the convex hull edges in a drawing of an *internally* 3-connected graph may create nonexternal separations pairs, see Figure 101(a). This would prevent us from using the algorithm from Section 5.3.2 since it builds on Lemma 67, which requires the input graph to be internally 3-connected.

Popping a pocket outward. In this step, we describe a way to remove an edge of A and "pop" out the vertices of its pocket so that they become part of the convex hull. Lemma 67 serves once again as an important subroutine. We make ample use of the fact that we may freely specify the desired subdrawing of the outer face after each application of Lemma 67, as long as we maintain either the x-coordinates or the y-coordinates of all vertices.

We remark that the following lemma applies to *internally* 3-connected graphs, not only 3-connected graphs. We plan to re-use it in the upcoming Section 5.3.4, in which we prove Theorem 68 in its general form. In fact, the final algorithm will use the entire procedure described in the current section as a subroutine. We will augment the given internally 3-connected graph such that adding the convex hull edges does not create nonexternal separation pairs.

Lemma 72. *Let Γ be a strictly convex drawing of an internally 3-connected graph G without vertical edges and let e be an edge on the outer face. If $G \setminus \{e\}$ is internally 3-connected, then $\Gamma \setminus \{e\}$ can be morphed to a strictly convex drawing of $G \setminus \{e\}$*

without vertical edges via at most three weakly convexity-increasing planar morphs, each of which is horizontal or vertical. Furthermore, the morphs can be found in time $O(n^{\omega/2} + n \log n)$.

Proof. Our morph consists of up to three unidirectional morphs specified by a sequence of drawings Γ, Γ_1, Γ_2, and Γ_3. The first and the last morph are vertical. The second morph, which we can sometimes skip, is horizontal.

Let $e = \{u, v\}$. We show that there is a vertical planar morph $\langle \Gamma, \Gamma_1 \rangle$ such that Γ_1 is a strictly convex drawing of G in which the vertex u has either the largest or the smallest y-coordinate among all vertices. Moreover, Γ_1 does not contain vertical or horizontal edges.

Since Γ is strictly convex and has no vertical edges, all of its faces are x-monotone. Therefore, the desired drawing Γ_1 can be found by choosing some strictly convex subdrawing of the outer face in which u is top-most or bottom-most while maintaining the x-coordinates of all vertices, and then using one application of Lemma 67 (for vertical morphs). We may need to apply a vertical shearing transformation in order to get rid of horizontal edges. This is easily done while still guaranteeing that u is top-most or bottom-most. (In fact, we can also obtain the desired drawing Γ_1 directly by means a shearing transformation, by making some tangent of u that is not parallel to any edge horizontal.)

Analogous to Section 5.3.1, by combining Observation 66, Lemma 63, and Lemma 65 we conclude that the vertical morph $\langle \Gamma, \Gamma_1 \rangle$ is planar and weakly convexity-increasing (the term convexity-*preserving* is actually more fitting since Γ is already strictly convex). It follows that the vertical morph $\langle \Gamma \setminus \{e\}, \Gamma_1 \setminus \{e\} \rangle$ is also planar. Moreover, it is weakly convexity-increasing since the removal of e decreases the set of internal angles.

For the remainder of the proof, assume without loss of generality that u is the top-most vertex and that v lies to the right of u in Γ_1. The other cases are symmetric. Let p_{uv} denote the unique path from u to v in $f_e \setminus \{e\}$. We distinguish two cases depending on the shape of f_e in Γ_1.

Case 1: The path p_{uv} is x-monotone in Γ_1, see Figure 100(b). In this case we can skip the second step (to Γ_2) of our morphing sequence. We compute a vertical morph from $\Gamma_1 \setminus \{e\}$ to a strictly convex drawing Γ_3 of $G \setminus \{e\}$ without vertical edges. Once again, this can be done by combining Lemma 67 (for vertical morphs), Observation 66, Lemma 63, and Lemma 65 as long as we can specify a strictly convex drawing of the outer face of $\Gamma_1 \setminus \{e\}$ in which the x-coordinates match those of Γ_1. It suffices to compute a suitable new reflex chain for p_{uv}, which is easy to achieve, see Figure 100(b). ◁

Case 2: The path p_{uv} is not x-monotone. Our plan is to perform a reduction to Case 1, that is, we will morph to a drawing in which p_{uv} is x-monotone.

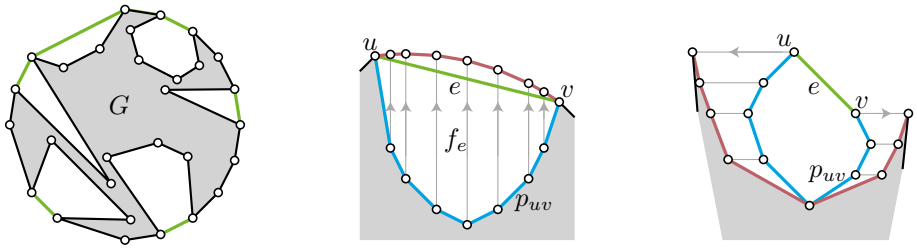

(a) (b) (c)

Figure 100: (a) Schematic of the convex drawing of $G \cup A$. Graph G is depicted in gray and the thick green edges belong to A. (b)–(c) Cases 1 and 2 for Lemma 72, where faint gray arrows indicate explicit placements on the convex hull.

To this end, we compute a horizontal, weakly convexity-increasing, and planar morph from Γ_1 to a strictly convex drawing Γ_2 in which u and v are the unique left-most and the unique right-most vertices, respectively. Moreover, we require that Γ_2 contains no vertical edge. Once again, this can be done by combining Lemma 67, a horizontal shearing transformation, Observation 66, Lemma 63, and Lemma 65 as long as we can specify a strictly convex subdrawing of the outer face of Γ_1 in which u and v are the left-most and right-most vertices and the y-coordinates match those of Γ_1. Given that u is top-most, such a drawing is easily obtained, see Figure 100(c).

The morph $\langle \Gamma_1, \Gamma_2 \rangle$ is horizontal, weakly convexity-increasing, and planar, and the morph $\langle \Gamma_1 \setminus \{e\}, \Gamma_2 \setminus \{e\} \rangle$ retains these properties since the removal of e decreases the set of internal angles. Moreover, in the drawing Γ_2 the pocket f_e is strictly convex with unique extreme points u and v in the x-direction. Therefore, the path p_{uv} is x-monotone in Γ_2 and, hence, analogous to Case 1, we obtain a vertical planar weakly convexity-increasing morph from $\Gamma_2 \setminus \{e\}$ to a strictly convex drawing Γ_3 of $G \setminus \{e\}$. □

To conclude the proof of Theorem 68 for the 3-connected case, we simply iterate the procedure described by Lemma 72. More formally, for each $e \in A$, Lemma 72 guarantees a sequences $\mathcal{M}$ of at most three vertical / horizontal weakly convexity-increasing planar morphs between $(\Gamma' \cup \Gamma'_A) \setminus \{e\}$ and some strictly convex drawing of $G \cup (A \setminus \{e\})$. The sequence $\mathcal{M}$ is also planar and weakly convexity-increasing if restricted to Γ' since the removal of A decreases the set of internal angles. Moreover, applying $\mathcal{M}$ to Γ' results in a straight-line planar drawing Γ'' of G such that the number of segments of the convex hull of Γ'' that do not correspond to edges of G is $|A| - 1$. Hence, by induction on this number, we obtain a sequence of vertical / horizontal weakly convexity-increasing planar morphs that transforms Γ' into a drawing with a strictly convex subdrawing of the outer face. Moreover, this drawing is strictly convex since the internal angles of Γ' are strictly convex.

It remains to analyze the number of morphing steps. Let ρ denote the number of pockets of $\Gamma' \cup \Gamma'_A$.We observe that $\rho \leq n/2$, since each pocket can be associated with

two private vertices of G: the clockwise first of the two convex hull vertices defining the pocket and its clockwise successor, which is private to the pocket.

Obtaining $\Gamma' \cup \Gamma'_A$ from $\Gamma \cup \Gamma_A$ requires at most $\max\{2, r+1\}$ morphing steps, where r is the number of internal reflex angles in the drawing $\Gamma \cup \Gamma_A$. We have $r \leq n-\rho$ since each pocket can be associated with a private vertex of the convex hull, that cannot have an internal reflex angle. Hence, the number of morphing steps for obtaining $\Gamma' \cup \Gamma'_A$ is bounded by $\max\{2, r+1\} \leq n - \rho + 1$.

We then use at most three horizontal and vertical morphs to pop out each pocket, which results in at most 3ρ additional morphing steps. This bound can be improved to at most $2\rho + 1$ since each application of Lemma 72 involves a vertical-horizontal-vertical morphing sequence, and Lemma 64 allows us to compress two consecutive vertical morphs into one.

Summing up, the total number of morphs is at most

$$n - \rho + 1 + 2\rho + 1 = n + \rho + 2 \leq 1.5n + 2,$$

where the last inequality uses the fact that $\rho \leq n/2$.

The run time of the algorithm is $O((1+\rho)n^{\omega/2} + n^2 \log n) \subseteq O(n^{1+\omega/2} + n^2 \log n)$.

5.3.4 Morphing drawings of internally 3-connected graphs

So far, if the given drawing does not have a convex outer face, we have restricted our attention to the class of 3-connected graphs. The only reason why the approach described in Section 5.3.3 is not able to handle the case of *internally* 3-connected graphs is that the addition of the convex hull edges may create nonexternal separation pairs, see Figure 101(a). As a result, the augmented drawing is no longer a valid input for Lemma 67. In this section, we extend our algorithm such that it is able to convexify drawings of internally 3-connected graphs and, thus, we conclude the proof of Theorem 68 in its general form.

Overview. Let Γ be a straight-line planar drawing of an internally 3-connected graph $G = (V, E)$. As a first step, we augment the outer face of Γ by adding new edges and vertices near each pocket as illustrated in Figure 102. The goal of this step is to ensure that we can add convex hull edges without introducing nonexternal separation pairs. We then apply the algorithm from Section 5.3.3, which results in a drawing in which all the new vertices appear on the strictly convex outer face. Finally, we remove the new vertices one by one, gradually turning the strictly convex drawing of the augmented graph into a strictly convex drawing of G.

In slightly more detail, our method involves the following steps:

Step 1: Augmenting the outer face. We will augment G, by adding vertices and edges, to obtain an internally 3-connected graph G'. We then extend Γ to obtain a drawing Γ' of G', which involves geometric arguments. The goal is to ensure that G'

has the additional property that adding convex hull edges to Γ' does not introduce nonexternal separation pairs, as, for example, in Figure 101(a). We will show that bad cases only arise when a pocket has a vertex that is part of a separation pair. Thus, our method will be to add an extra "buffer" layer of vertices to each pocket boundary as illustrated in Figure 102 while ensuring that no buffer vertex is part of a separation pair.

Step 2: Convexifying the augmented drawing. The goal of this step is to find a weakly convexity-increasing morph from the drawing Γ' of the augmented graph G' to a strictly convex drawing of G'. To this end, we apply the algorithm described in Section 5.3.3 to Γ'. Recall that this algorithms adds the missing edges of the convex hull and then iteratively removes them while performing morphing steps before and after each removal. The challenging aspect is that these morphs are computed via Lemma 67 and, hence, we need to ensure that the graph remains internally 3-connected throughout the entire procedure.

Step 3: Removing the additional vertices. At this point we have a strictly convex drawing of G'. We must now "reverse" the augmentation process, removing vertices of G' to get back to G. After each removal, we will perform a morph, involving Lemma 67, to obtain a strictly convex drawing. Hence, as in Step 2, we have to ensure that the graph remains internally 3-connected *at each step*. We will therefore interpret the augmentation process of Step 1 as an iterative process and we show that G' can be obtained from G by performing a sequence of operations that preserve internal 3-connectivity.

Let us proceed by discussing these steps in detail.

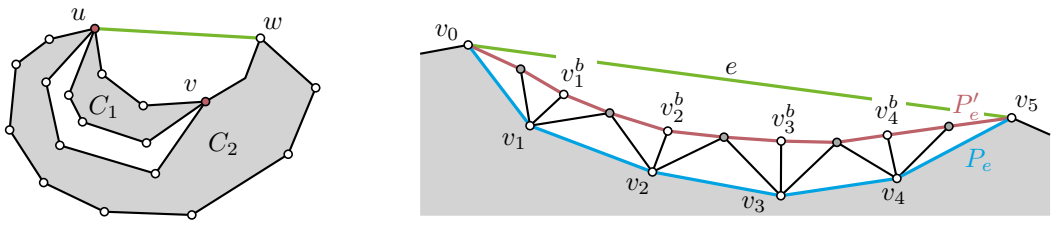

(a) (b)

Figure 101: (a) An internally 3-connected graph with an external separation pair u, v. Adding the convex hull edge uw turns u, v into a nonexternal separation pair both since v becomes an internal vertex and since C_1 no longer has a vertex on the outer face. (b) Schematic drawing of the path P'_e created for the pocket defined by the convex hull edge $e = \{v_0, v_5\}$.

Step 1: Augmenting the outer face. We compute the convex hull of Γ in $O(n)$ time [96]. Let e be a convex hull edge with $e \notin E$. The following steps are illustrated in Figure 101(b). Let $P_e = (v_0, v_1, \ldots, v_{k+1})$ be the unique path on the outer face of Γ such that $P_e \cup \{e\}$ is a cycle with $G \setminus P_e$ in its exterior. We introduce $2k+1$ new vertices that form a path P'_e connecting v_0 and v_{k+1}. Let

$$P'_e = (v_0,\ v_1^a, v_1^b, v_1^c\ =\ v_2^a, v_2^b, v_2^c\ =\ v_3^a, \ldots, v_{k-1}^c\ =\ v_k^a, v_k^b, v_k^c,\ v_{k+1}).$$

Every interior vertex v_i of P_e gets a "private copy" v_i^b in P'_e. Two consecutive private copies v_i^b and v_{i+1}^b are connected via another vertex which is equipped with two labels $v_i^c = v_{i+1}^a$, which will simplify the notation later on. Additionally, we add the edges $\{v_i, v_i^a\}$, $\{v_i, v_i^b\}$ and $\{v_i, v_i^c\}$ for $i = 1, \ldots, k$.

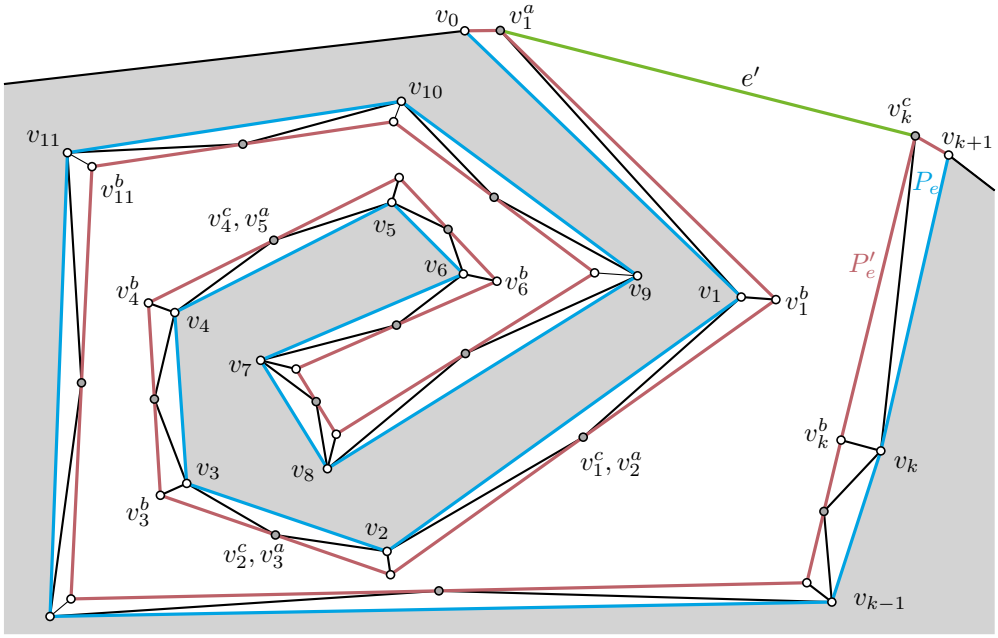

Figure 102: Geometrically, we embed P'_e very close to P_e. This is possible regardless of the shape of the pocket.

Geometrically, the new path P'_e is embedded in a planar fashion very close to P_e, see Figure 102. This can be accomplished regardless of the shape of P_e: let ε be the smallest distance between any edge of the cycle $P_e \cup \{e\}$ and any of its nonincident vertices along the cycle. For $i = 1, \ldots, k$, we place v_i^b on the angular bisector of the outer angle at v_i such that its distance to v_i is smaller than $\varepsilon/2$. For all $i > 1$, the vertex v_i^a is placed in the center of the line segment $v_{i-1}^b v_i^b$.

We make sure that during this procedure no vertex is placed in the exterior of the convex hull of Γ with the following exceptions: the vertices v_1^a and v_k^c (which have not been placed according to the above rules) play a special role and are placed close to v_0 and v_{k+1}, respectively, such that they appear between v_0 and v_{k+1} on the

convex hull of the augmented drawing, see Figure 102.

We repeat the process for all convex hull edges $e \notin E$ of Γ and use $G' = (V', E')$ and Γ' to denote the resulting plane graph and its straight-line planar drawing, respectively. The total runtime for this step is dominated by the time to compute the values ε. Given that the overall runtime of our algorithm is $\Omega(n^2)$, it suffices to compute these values by brute force in $O(n^2)$ total time. However, we remark that this step could be implemented more efficiently by means of more advanced tools [32] from computational geometry.

We will now prove that G' is internally 3-connected. Keeping in mind our plan for Step 3, we will add the vertices of G' one by one, showing that each addition preserves internal 3-connectivity. To this end, we define the following two operations.

Lemma 73. *Let H be an internally 3-connected graph with an edge $\{a, b\}$ on the outer face. Construct H' by adding a new vertex x in the outer face connected to a and b. Then H' is internally 3-connected.*

Proof. By internal 3-connectivity, Property (K3) of Lemma 61 holds for H. Thus, it is clear that H' also satisfies Property (K3) and, hence, H' is internally 3-connected. □

Lemma 74. *Let H be an internally 3-connected graph with two consecutive edges $\{a, b\}$ and $\{b, c\}$ on the outer face. Construct H' by adding a new vertex x in the outer face connected to a, b and c. Then H' is internally 3-connected.*

Proof. By internal 3-connectivity, Property (K3) of Lemma 61 holds for H. It suffices to show that Property (K3) also holds for H'. Let f_0 and f_0' denote the outer faces of H and H', respectively. Clearly, b has three paths to f_0' that are disjoint except for b. So let $v \neq b$ be some internal vertex of H' and note that v is also internal in H. Hence, by Property (K3) of H, v has three paths to f_0 that are disjoint except for v. At most one of these paths does not end at f_0', namely if its endpoint on f_0 is b. However, appending the edge $\{b, x\}$ to this path yields the desired three paths from v to f_0' that are disjoint except for v. Hence, Property (K3) holds for H'. □

With these operations in hand, we can show that G' is internally 3-connected. In fact, we can build G' by adding one vertex at a time, preserving internal 3-connectivity. Let $V' = V \cup V^b \cup V^{ac}$ where V^b is the set of all vertices whose upper index is b, and where V^{ac} is the set of the remaining vertices (whose upper index is a and/or c).

Lemma 75. *Starting with G and adding the vertices of V^{ac} one by one in any order and then the vertices of V^b one by one in any order produces an internally 3-connected graph at each step.*

Proof. The addition of each vertex of V^{ac} maintains internal 3-connectivity by Lemma 73. The addition of each vertex of V^b maintains internal 3-connectivity by Lemma 74. □

Step 2: Convexifying the augmented drawing. As mentioned above, the plan is to apply the algorithm described in Section 5.3.3 to Γ'. That algorithm adds the convex hull edges of Γ' and then iteratively removes these edges while performing some morphing steps before and after each removal. Each morphing step requires one application of Lemma 67. Therefore, in order to prove the correctness of Step 2, we need to ensure that before and after each removal of a convex hull edge, the input graph is internally 3-connected.

We begin by observing that the new vertices of G' are not part of separation pairs. Then we show that adding a convex hull edge is safe when none of the vertices of the pocket belong to a separation pairs.

Observation 76. *No vertex of $V' \setminus V$ belongs to a separation pair of G'.*

Proof. Let $v \in V' \setminus V$ and assume that it forms a separation pair together with a vertex u of G'. All vertices in $\mathrm{N}_{G'}(v)$ belong to the outer face of $G' \setminus \{v\}$, which, by construction, is a simple cycle. Hence, the removal of u from $G' \setminus \{v\}$ cannot disconnect $\mathrm{N}_{G'}(v) \setminus \{u, v\}$, which yields a contradiction to Observation 30. □

We require one more operation that preserves internal 3-connectivity.

Lemma 77. *Let H be an internally 3-connected graph, with two nonadjacent vertices a and b on the outer face. Let P be one of the paths from a to b along the outer face. Assume that no vertex of P is part of a separation pair in H. Let H' be the result of adding the edge $\{a, b\}$ embedded such that P becomes internal. Then H' is internally 3-connected.*

Proof. Suppose that H' has a separation pair u, v. This pair is also separating in H. Further, since H is internally 3-connected, u, v is an *external* separation pair of H. Since both u and v do not belong to P, the vertices a and b belong to the same component of $H \setminus \{u, v\}$. Thus, is easy to verify that adding the edge between a and b maintains the six conditions of Observation 62 (in particular, Property (E2) holds as a and b remain on the outer face). Therefore, u, v is an external separation pair of H'. □

Recall that the construction of Γ' ensures that all convex hull edges e' that do not correspond to edges of G' have the form $e' = \{v_1^a, v_k^c\}$, where v_1^a, v_k^c are the first and last internal vertex of one of the paths P'_e. By Observation 76 the interior vertices of P'_e form a path P on the outer face that does not contain any vertices that are part of a separation pair. Further, adding the edge e' encloses P in an internal face. Thus, we obtain:

Corollary 78. *Let A denote the set of convex hull edges of Γ' that do not correspond to edges of G'. Then, for any $S \subseteq A$ the plane graph $G' \cup S$ is internally 3-connected.* □

Step 3: Removing the additional vertices. At this point we have a strictly convex drawing of G' and want to convert it to a strictly convex drawing of G. We will remove the vertices of V^b iteratively, one by one. After each such removal, we will perform up to two morphing steps to recover a strictly convex drawing of the reduced graph. Recall that by Lemma 75, all the intermediary graphs are internally 3-connected and, thus, they are valid inputs for Lemma 67. The final strictly convex drawing of $G' \setminus V^b$ can then be turned into the desired drawing of G by simply removing all vertices of V^{ac}.

Let us describe a single step of the removal process for the vertices of V^b.

Lemma 79. *Let $B \subseteq V^b$ and let Γ'_B be a strictly convex drawing of $G' \setminus B$ without vertical or horizontal edges. Further, let $v_i^b \in V^b \setminus B$. Then there is a weakly convexity-increasing planar morph from $\Gamma'_B \setminus \{v_i^b\}$ to a strictly convex drawing Γ'' of $G' \setminus (B \cup \{v_i^b\})$ without vertical or horizontal edges. Moreover, there is such a morph that consists of a sequence of up to 2 horizontal / vertical morphs. The morphing sequence can be found in $O(n^{\omega/2} + n \log n)$ time, assuming that two $n \times n$ matrices can be multiplied with $O(n^\omega)$ arithmetic operations.*

Proof. Without loss of generality, we may assume that v_i^c is located to the bottom-right of v_i^a and that v_i^b is located to the right of the oriented line $\overrightarrow{v_i^c v_i^a}$, see Figure 103(a). We distinguish four cases regarding the position of the vertex v_i, for an illustration see Figure 103(b).

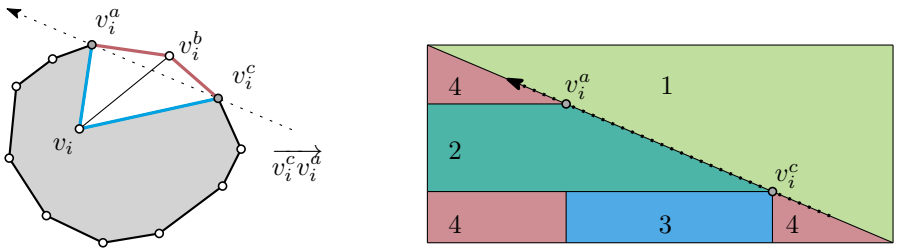

(a) (b)

Figure 103: (a) In Step 3, we iteratively remove the vertices v_i^b causing their counterparts v_i to become part of the outer face. (b) The regions corresponding to the Cases 1–4.

Case 1: v_i is located to the right of the oriented line $\overrightarrow{v_i^c v_i^a}$. We observe that the subdrawing $\Gamma'_B \setminus \{v_i^b\}$ is already a strictly convex drawing of $G' \setminus (B \cup \{v_i^b\})$, so there is nothing to show. ◁

Case 2: $\mathrm{y}(v_i^a) > \mathrm{y}(v_i) > \mathrm{y}(v_i^c)$ and we are not in Case 1. Let C be a strictly convex drawing of the outer face of $G' \setminus (B \cup \{v_i^b\})$ such that every vertex in C has the same y-coordinate as in $\Gamma'_B \setminus \{v_i^b\}$ (we can easily find such a drawing C by adding v_i to the convex hull of $\Gamma'_B \setminus \{v_i^b\}$ in a strictly convex fashion). Then Lemma 67 applied

to $\Gamma'_B \setminus \{v_i^b\}$ and C (potentially followed by a horizontal shearing transformation in order to remove vertical edges), yields the desired drawing Γ''. Analogous to Section 5.3.1, by combining Observation 66, Lemma 63, and Lemma 65, we conclude that the horizontal morph $\langle \Gamma'_B \setminus \{v_i^b\}, \Gamma'' \rangle$ is planar and weakly convexity-increasing. ◁

Case 3: $\mathrm{x}(v_i^a) < \mathrm{x}(v_i) < \mathrm{x}(v_i^c)$ and we are not in one of the Cases 1 or 2. This case can be handled analogous to Case 2 (by morphing vertically rather than horizontally). ◁

Case 4: We are not in one of the Cases 1, 2, or 3. We reduce to Case 2 or Case 3: using a shearing transformation along the x-axis (or along the y-axis), we obtain a drawing Γ''_B of $G' \setminus B$ satisfying the preconditions of Case 3 (or Case 2). Analogous to Section 5.3.1, by combining Observation 66, Lemma 63, and Lemma 65, we conclude that the linear morph $\langle \Gamma'_B, \Gamma''_B \rangle$ is planar and weakly convexity-increasing. □

Starting with the drawing Γ' of G' and iterating Lemma 79, we obtain a strictly convex drawing of $G' \setminus V^b$. By construction, we can simply remove all the vertices of V^{ac} to obtain a strictly convex drawing of G.

Observation 80. *Let Γ'_{ac} be a strictly convex drawing of $G' \setminus V^b$. Then, $\Gamma'_{ac} \setminus V^{ac}$ is a strictly convex drawing of G.* □

We summarize:

Proof of Theorem 68. We analyze the three steps of the algorithm individually.

Step 1: We begin by augmenting G and Γ to G' and Γ'. As discussed in the corresponding section, this can be done in $O(n^2)$ time.

Step 2: Next, we apply the algorithm from Section 5.3.3 to G' and Γ'. This algorithm was designed for 3-connected graphs. However, since Lemma 72 applies to internally 3-connected graphs as well, the algorithm also works for internally 3-connected graphs as long as adding convex hull edges and then successively removing them never creates a plane graph that is not internally 3-connected. This is the case by Corollary 78. Thereby, we obtain in

$$O((n')^{1+\omega/2} + n^2 \log n) \subseteq O(n^{1+\omega/2} + n^2 \log n)$$

time a weakly convexity-increasing morph from Γ' to a strictly convex drawing of G', where $n' \in O(n)$ is the number of vertices of the augmented graph G'. This morph is also weakly convexity-increasing with respect to the subdrawing of G as every internal angle of Γ is also internal in Γ'.

Let $\Gamma' \cup \Gamma'_A$ and $\Gamma \cup \Gamma_A$ denote Γ' and Γ' augmented by their convex hull edges, respectively. As discussed in the final paragraphs of Section 5.3.3, the upper bound on the number of morphing steps guaranteed by the algorithm is $\max\{2, r'+1\} + 2\rho' + 1$, which can be bounded by $1.5n' + 2$. Here, ρ' and r' denote the number of pockets

and internal reflex angles of $\Gamma' \cup \Gamma'_A$, respectively. In fact, the bound can be improved to $1.5n + 2$ ($< 1.5n' + 2$) by observing that $r = r'$ and $\rho = \rho'$, where ρ and r denote the number of pockets and internal reflex angles of the drawing $\Gamma \cup \Gamma_A$, respectively. The latter identity is obvious.

It remains to show that $r = r'$. Observe that each vertex $v_i^a \in V^{ac}, 2 \leq i \leq k$ has an angle of π and, hence, it has no reflex angle. The vertices $v_0, v_1^a, v_k^c, v_{k+1}$ of each pocket P'_e belong to the convex hull of $\Gamma' \cup \Gamma'_A$ and, thus, they do no have any internal reflex angles. By construction, an outer angle of Γ' at a vertex $v_i^b \in V^b$ is reflex if and only if the outer angle at the corresponding vertex v_i is reflex in Γ. Other angles at v_i^b cannot be reflex. Moreover, if the outer angle at v_i is reflex in Γ then v_i has no reflex angle in Γ'. Consequently, we can charge the reflex angles of the vertices v_i^b to their counterparts v_i. Finally, the convexity status of the remaining angles is untouched by the augmentation and, hence, $r = r'$ as claimed. Altogether, we obtain the improved bound

$$\max\{2, r' + 1\} + 2\rho' + 1 = \max\{2, r + 1\} + 2\rho + 1 \leq 1.5n + 2,$$

where the last inequality was already discussed in the last paragraph of Section 5.3.3.

Step 3: Finally, we iteratively apply Lemma 79 to the strictly convex drawing of G' that was obtained in the previous step. Each application increases the number of vertices of G on the convex hull. Thus, by induction we arrive at a strictly convex drawing of G. Each of the intermediary morphing steps is weakly convexity-increasing with respect to the respective augmented graph. Once again, since every internal angle of Γ remains internal in (all) the augmented graph(s), we have that the morphing sequence is weakly convexity-increasing for G as well. The number of morphing steps is bounded by $2n$ and the time required to obtain the entire sequence sums up to $O((n')^{1+\omega/2} + n^2 \log n) \subseteq O(n^{1+\omega/2} + n^2 \log n)$.

Summing up, we have performed a total of at most $3.5n + 2$ morphing steps. The algorithm has a total runtime of $O(n^{1+\omega/2} + n^2 \log n)$. □

5.4 Finding strictly convex redrawings while preserving y-coordinates

In this section, we follow the idea of Chrobak, Goodrich, and Tamassia [33] to prove Lemma 67 using Tutte's graph drawing algorithm (see Section 5.2.4 for a discussion about related results and literature). This reduces the problem to solving a linear system. Due to an improved version of generalized nested dissection by to Alon and Yuster [5] from 2013, this results in a very simple and short proof.

Tutte's theorem. In his paper, "How to Draw a Graph," [124] Tutte showed that every 3-connected plane graph $G = (V, E)$ admits a convex drawing such that the

convex subdrawing C of the outer face can be prescribed as part of the input. This method is sometimes referred to as Tutte's spring theorem since its idea is based on a physical analogy: the idea is to think of the vertices as small metal rings and of inner edges as springs that connect the two rings that correspond to its endpoints. The outer edges are thought of as a solid frame. The equilibrium state of this physical system corresponds to the desired drawing. Tutte computes his drawings by means of a linear system of equations: for each $v \in V$ let the variables (x_v, y_v) represent the coordinates of vertex v. Let V_I be the internal vertices of G and let V_B be the vertices of the outer face. For each vertex $v \in V_B$ let (x_v^b, y_v^b) be its (fixed) coordinates in C. Let d_v be the degree of vertex v. Consider the system of equations:

$$\forall u \in V_I \qquad (x_u, y_u) = \sum_{(u,v)\in E} \frac{1}{d_u}(x_v, y_v), \tag{5.1}$$

$$\forall u \in V_B \qquad (x_u, y_u) = (x_u^b, y_u^b). \tag{5.2}$$

Tutte proved that this system of equations has a unique solution and that the solution gives a convex drawing of G with C as the outer face. In fact, the drawing is even strictly convex if C is strictly convex.

Tutte's theorem can be generalized to more general "barycenter" weights other than $1/d_u$. Assign a weight $w_{u,v} > 0$ to each ordered pair (u, v) with $\{u, v\} \in E$ such that $\sum_v w_{u,v} = 1$ for each internal vertex u. We emphasize that $w_{u,v}$ may be different from $w_{v,u}$. Consider the system of equations:

$$\forall u \in V_I \qquad (x_u, y_u) = \sum_{(u,v)\in E} w_{u,v}(x_v, y_v), \tag{5.3}$$

$$\forall u \in V_B \qquad (x_u, y_u) = (x_u^b, y_u^b). \tag{5.4}$$

This system also has a unique solution that gives a convex drawing of G with outer face C and a strictly convex drawing of G if C is strictly convex. This generalization was first stated by Floater in 1997 [55] for triangulations and one year later [56] for general 3-connected planar graphs, though the result is not stated as a theorem in either case. Floater proved that the constraint matrix is non-singular and, for the rest, said that Tutte's proof[1] carries over. An explicit statement that the linear system (5.3–5.4) has a unique solution that gives a strictly convex drawing of G if C is strictly convex is due to Gortler, Gotsman, Thurston in 2006 [66, Theorem 4.1]. They give a new proof using "one-forms".

Tutte's theorem was originally stated for 3-connected graphs. However, it is well-known that Tutte's proof also applies to the more general class of internally 3-connected graphs since it only uses Property (K3) of Lemma 61. In fact, (Floater's version of) Tutte's theorem can be generalized even further: it is possible to allow

[1] Colin de Verdiere et al. [36] point out that Tutte's original proof is complicated because Tutte is also re-proving Kuratowski's theorem, and they recommend the simpler proof by Richter-Gebert [112].

weights $w_{u,v}$ of value $w_{u,v} = 0$, as long as the following relaxed version of Property (K3) of Lemma 61 is satisfied. Suppose we are given a non-negative weight $w_{u,v}$ for each ordered pair (u, v) with $\{u, v\} \in E$ such that $\sum_v w_{u,v} = 1$ for each internal vertex u. Then we say that an internal vertex x is *positively 3-connected to the boundary* if there exist three directed paths from x to the outer face C that are disjoint except for the common endpoint x and whose directed edges have positive weight. If every internal vertex is positively 3-connected to the boundary, we say that the graph itself is *positively 3-connected to the boundary* (with respect to the assignment of weights), which is sufficient to ensure that Tutte's method can be applied. This version of Tutte's theorem was formulated and proven by Haas et al. [73, Theorem 8]:

Theorem 81 ([73])**.** *Let G be an internally 3-connected graph and suppose we are given a non-negative weight $w_{u,v}$ for each ordered pair (u, v) with $\{u, v\} \in E$ such that $\sum_v w_{u,v} = 1$ for each internal vertex u. Assume that G is positively 3-connected to the boundary. Let C be a strictly convex drawing of the outer face of G. Then, the system (5.3–5.4) has a unique solution that corresponds to a strictly convex planar drawing of G with C as the outer face.*

We can now give an alternate proof of Hong and Nagamochi's well-known result based on the idea of Chrobak et al. We will only use positive weights—specifying weights of value 0 will be relevant for the upcoming Section 5.5, where we extend our algorithm for weakly convexity-increasing morphs to also maintain outer reflex angles.

Lemma 67 ([10, 33, 79])**.** *Let Γ be a planar drawing of an internally 3-connected graph G such that every face is y-monotone. Let C be a strictly convex straight-line drawing of the outer face of G such that every vertex of C has the same y-coordinate as in Γ. Then there is a strictly convex straight-line drawing Γ' of G such that the subdrawing of the outer face is C and every vertex of Γ' has the same y-coordinate as in Γ.*

Furthermore, the drawing Γ' can be found in time $O(n^{\omega/2} + n \log n)$, even if only the underlying abstract graph of G, the cycle corresponding to C, and the y-coordinates of vertices are given, assuming that two $n \times n$ matrices can be multiplied with $O(n^\omega)$ arithmetic operations.

Proof. We must show that there is a strictly convex drawing of G with C as the subdrawing of the outer face that preserves the y-coordinates of the vertices from the drawing Γ. The idea by Chrobak et al. [33] is to do this in two steps: first choose the barycenter weights to force the vertices to lie at the required y-coordinates, and then solve system (5.3–5.4) with these barycenter weights to determine the x-coordinates.

For the first step, we solve the following system separately for each $u \in V_I$:

$$y_u = \sum_{(u,v)\in E} w_{u,v} y_v, \qquad 1 = \sum_v w_{u,v} \tag{5.5}$$

Here the y's are the known values from Γ and the $w_{u,v}$'s are variables. There are two equations and $d_u > 2$ variables, so the system is under-determined and can easily be solved: since Γ has y-monotone faces, each vertex $u \in V_I$ has a neighbor with a larger y-coordinate and a neighbor with a smaller y-coordinate. Let N_u^+ be the neighbors of u that lie above u in Γ. Let $d_u^+ = |N_u^+|$. Similarly, let N_u^- be the neighbors of u that lie below u in Γ and let $d_u^- = |N_u^-|$. Compute the average y-coordinate of u's neighbors above and below:

$$y_u^+ = \sum_{v\in N_u^+} \frac{1}{d_u^+} y_v \qquad y_u^- = \sum_{v\in N_u^-} \frac{1}{d_u^-} y_v$$

Observe that y_u lies between y_u^+ and y_u^-. Thus, we can find a value t_u, $0 < t_u < 1$, such that

$$\begin{aligned} y_u &= t_u y_u^+ + (1-t_u) y_u^- \\ &= \sum_{v\in N_u^+} \frac{t_u}{d_u^+} y_v + \sum_{v\in N_u^-} \frac{1-t_u}{d_u^-} y_v \end{aligned}$$

Therefore, setting $w_{u,v} = \frac{t_u}{d_u^+}$ for $v \in N_u^+$ and $w_{u,v} = \frac{1-t_u}{d_u^-}$ for $v \in N_u^-$, yields a solution to (5.5). Observe that $w_{u,v} > 0$ for all (u, v) with $\{u, v\} \in E$.

Given values $w_{u,v} > 0$ satisfying the constraints (5.5) for all $u \in V_I$, we then solve the equations (5.3–5.4) to find values for the x-coordiantes x_u. By Theorem 81, this provides a strictly convex drawing of G with C as the subdrawing of the outer face while preserving y-coordinates.

It remains to discuss how to obtain the claimed runtime. Recall that we assume a real-RAM model of computation—in particular, each arithmetic operation takes unit time. Observe that solving the system (5.5) to find the appropriate weights $w_{u,v}$ based on the y-coordinates takes linear time. The significant aspect is solving Tutte's generalized system of equations (5.3–5.4).

Tutte's method gives rise to two linear systems $Ax = b$ and $Ay = b$ (for the x-coordinates and the y-coordinates, respectively) that have the the same constraint matrix A with a row and column for each vertex and where the nonzero entries in the matrix correspond to edges in the planar graph. In more detail, the equations (5.1) can be re-written as

$$\forall u \in V_I \qquad -d_u(x_u, y_u) + \sum_{(u,v)\in E} (x_v, y_v) = 0,$$

so the part of the constraint matrix that corresponds to the interior vertices V_I consists of entries $-d_u$ down the main diagonal, and $a_{u,v} = a_{v,u} = 1$ if $\{u, v\}$ is an

edge. The equations (5.2) for the vertices in V_B correspond to the remaining part of the constraint matrix, which is just an identity matrix. The crucial property is that the constraint matrix is symmetric. In fact, symmetry holds more generally if we choose weights or "stresses" $s_{u,v} = s_{v,u}$ and define $w_{u,v} = s_{u,v} / \sum_{(u,z)\in E} s_{u,z}$. Tutte's original theorem is the special case where $s_{u,v} = 1$ for all (u, v).

Chrobak et al. [33] point out that when the constraint matrix A is symmetric with the nonzero entries corresponding to the edges of a planar graph, the system $Ax = b$ can be solved in $O(n^{\omega/2} + n \log n)$ arithmetic operations using the generalized nested dissection method of Lipton, Rose, and Tarjan [98, 99]. However, this version of nested dissection does not apply when the matrix A is not symmetric. In particular, it does not apply to the linear system (5.3–5.4). As described in [33], it is also possible to compute symmetric weights that ensure the desired y-coordinates. However, our easily obtained system (5.3–5.4) with asymmetric weights can also be solved efficiently due to a recent generalization of nested dissection by Alon and Yuster [5]. They consider a linear system $Ax = b$ where A has a row and a column for each vertex of an *associated graph* G and there is an edge $\{u, v\}$ in G if and only if $a_{u,v} \neq 0$ or $a_{v,u} \neq 0$ (the diagonal entries of A play no role in the definition of G).

Theorem 82 (Theorem 1.1 in [5], specialized to $\mathbb{Q}$ and to planar graphs)**.** *Let $A \in \mathbb{Q}^{n\times n}$ be a nonsingular matrix and let $b \in \mathbb{Q}^n$. If the graph associated with A is planar, then $Ax = b$ can be solved in $O(n^{\omega/2} + n \log n)$ time, assuming that two $n \times n$ matrices can be multiplied with $O(n^\omega)$ arithmetic operations.*

This theorem also assumes that each arithmetic operation takes unit time and the matrix is assumed to be given in a sparse form.

Note that in our case the matrix A is non-singular because of Theorem 81. This completes the proof of Lemma 67. □

5.5 Maintaining outer reflex angles

So far, we have restricted our attention to weakly convexity-increasing morphs, that is, we have only insisted that the set of internal strictly convex angles is nondecreasing. However, outer reflex angles may become convex throughout the morphs generated by our current approach. In particular, the algorithm corresponding to the 3-connected case of Theorem 68 starts by convexifying each pocket, thereby making all reflex angles of interior vertices of the pocket strictly convex. This is not necessarily the case with the algorithm for the internally 3-connected case described in Section 5.3.4 (since it first adds a buffer layer of vertices in each pocket), but there is no guarantee yet that outer reflex angles are maintained. In this section, we modify the algorithm corresponding to Theorem 68 to ensure that the set of outer reflex angles is also nondecreasing, thereby obtaining the main result of this chapter:

Theorem 58. *Let Γ be a planar straight-line drawing of an internally 3-connected graph G on n vertices. Then Γ can be morphed to a strictly convex drawing via a sequence of at most $3.5n+2$ convexity-increasing planar morphing steps each of which is either horizontal or vertical.*

In the special case that Γ has a convex outer face, the upper bound on the number of morphing steps is $\max\{2, r+1\}$, where r denotes the number of internal reflex angles of Γ.

Furthermore, there is a $O(n^{1+\omega/2} + n^2 \log n)$ time algorithm to find the sequence of morphs, assuming that two $n \times n$ matrices can be multiplied with $O(n^\omega)$ arithmetic operations.

To prove Theorem 58, we only need to make a minor change to our current approach: it suffices to slightly modify the weights that are computed during the calls to Lemma 67. To this end, we will use the fact that Theorem 81 allows us to choose weights of value 0, as long as the graph remains positively 3-connected to the boundary.

Proof of Theorem 58. The statement for the case that Γ has a convex outer face follow directly from Theorem 68.

For the general case, we proceed as in the proof of Theorem 68, except that, for some vertices, we modify the weights that are computed when applying Lemma 67. Recall that the algorithm corresponding the proof of Theorem 68 has three steps: in Step 1, we augment the outer face of G and Γ to G' and Γ' by adding a path P'_e near each pocket P_e. Along this path, each interior vertex v_i of P_e gets a private copy v_i^b and two other new neighbors v_i^a and v_i^c, see Figure 104(a), which is a schematic version of Figure 102 (compared to Figure 101(b), the Figure 104(a) properly reflects the fact that the first and last interior vertex v_1^a and v_k^c, respectively, of P'_e are positioned such that they appear between v_0 and v_{k+1} on the convex hull). In Step 2, we add the missing convex hull edges and then iteratively remove them while performing morphing steps before and after each removal. Finally, in Step 3, we first iteratively remove all the private copies v_i^b while performing morphing steps before and after each removal to make the vertices v_i part of the outer face, and then remove the remaining vertices v_i^a and v_i^c that where added during Step 1.

The target drawing for each of the morphing steps during Step 2 and 3 is obtained by one of two means: a shearing transformation, or appealing to Lemma 67. As an affine transformation, the former preserves the convexity status of each angle (in particular those of the outer reflex vertices). It remains to ensure that drawings produced by Lemma 67 also preserve the outer reflex angles. To prove Lemma 67, we used Theorem 81. In particular, we determined, in a preprocessing step, weights w_{ij} that force the vertices to maintain their y-coordinates. This was done by interpolating, for each vertex v, between the average y-coordinate of the neighbors of v with a larger y-coordinate, and the average y-coordinate of the neighbors of v with a smaller

y-coordinate. For the current proof, we proceed in the same way, except that when determining the weights for an outer vertex of G sometimes not all of its neighbors are taken into account.

More specifically, let G'' and Γ'' denote the current graph and drawing, respectively, after performing some number (possibly zero) of morphs of Step 2 and Step 3. Let v_i be an outer vertex of G that is internal in G'' (and, hence, an interior vertex of a pocket P_e) and whose outer angle α (formed by the edges $\{v_{i-1}, v_i\}, \{v_i, v_{i+1}\}$) in the subdrawing of G in Γ'' is reflex. For each such v_i, we need to ensure that its angle α remains reflex. When computing the weights for v_i, we only take into account the vertices $S = \{v_{i-1}, v_i^a, v_i^b, v_i^c, v_{i+1}\}$ (the vertex v_i^b is still part of G'' since v_i is internal). Since α is reflex, the vertex v_i is located in the interior of the polygon $(v_{i-1}, v_i^a, v_i^b, v_i^c, v_{i+1})$ formed by vertices of S. Consequently, the position of v_i is in the interior of the convex hull of S. Therefore, at least one of the neighbors of v_i in S is located above and at least one them is located below v_i (moreover, at least one of them is to the left of and at least one of them is to the right of v_i). Hence, the weights for the directed edges from v_i to S can be computed with the method described in (the beginning of) the proof of Lemma 67 to ensure that the y-coordinate (or x-coordinate, depending on the current direction) of v_i is maintained. The remaining weights for edges outgoing from v_i are set to 0. For the remaining vertices (which are not outer vertices of G with a reflex outer angle in the subdrawing of G in Γ''), the weights are computed normally. For an illustration, see Figure 104(b).

Let us prove that this assignment of weights ensures that each internal vertex v of G'' is positively 3-connected to the boundary. We distinguish three cases. First assume that v belongs to one of the paths P'_e, which implies that the convex hull edge e' of the corresponding pocket P_e still belongs to G'' and that v is one of the vertices v_i^a, v_i^b, v_i^c of P'_e (but not first or last interior vertex v_1^a, v_k^c of P'_e). Two of the desired paths from v to the boundary are found on P'_e, and the remaining path is formed by (v, v_i) together with a part of P_e, see Figure 104(c). Next assume that v belongs to some of the pockets P_e and, hence, is one of the vertices v_i. Two of the desired paths are found on P_e, and the remaining path is formed by (v_i, v_i^b) and if e' belongs to G'' (so that v_i^b is internal), a part of P'_e, see Figure 104(d). Finally, assume that v is an internal vertex of G. By condition (K3) of Lemma 61 for G, there exist three disjoint paths from v to the outer face of G, and the weights of the directed edges along these paths are positive by construction. The only difficult case arises when the three endpoints of these paths all belong to the same pocket P_e. Denote by v_h, v_i, v_j the endpoints of the three paths in the order in which they appear along P_e, i.e., $h < i < j$. The first of the desired paths is obtained from the path to v_h by traversing P_e towards v_0, see Figure 104(e). The second path is obtained from the path to v_j by traversing P_e towards v_{k+1}. The final path is obtained from the path to v_i by appending the edge (v_i, v_i^b) and if e' belongs to G'' (so that v_i^b is internal), a part of P'_e.

Our new weight assignment ensures that for each outer reflex angle α formed by

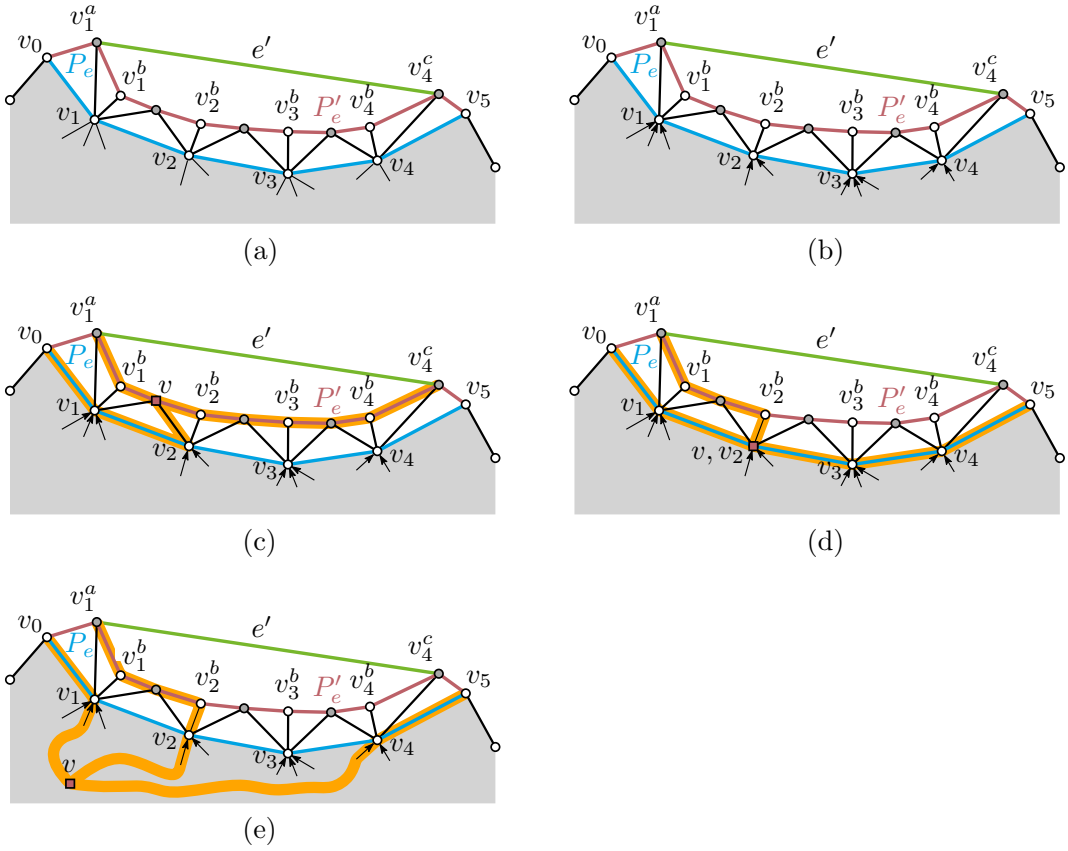

Figure 104: (a) An augmented pocket P_e. (b) The arrows correspond to edges that are possibly positively weighted in the indicated direction only. (c)–(e) Obtaining three disjoint paths of positively weighted edges to the outer face.

two edges $\{v_{i-1}, v_i\}, \{v_i, v_{i+1}\}$ at some vertex v_i of G that does not yet belong to the convex hull, the position of v_i is a strict convex combination of S in the generated drawing. Since v_{i-1} and v_{i+1} are consecutive in the cyclic order of the neighbors S of v_i, this implies that α remains reflex.

Overall, we have shown that once an outer angle is reflex, it remains reflex in all subsequent drawings of the morphing sequence. Hence, by Lemma 65, the set of outer reflex angles is nondecreasing.

The remaining parts of the statement follow analogous to the proof of Theorem 68. This concludes the proof. □

5.6 Lower bound on the number of morphing steps

In this section, we show a linear lower bound on the number of required morphing steps.

Theorem 59. *For any $n \geqslant 3$, there exists a drawing of an internally 3-connected graph on n vertices for which any convexifying planar morph composed of a sequence of linear morphing steps requires $\Omega(n)$ steps.*

Our proof of Theorem 59 builds on the following result by Alamdari et al. [3].

Theorem 83 ([3])**.** *There exist two straight-line planar drawings $\Gamma^{\Delta}(n')$ and $\Gamma^{-}(n')$ of a path with n' vertices such that any planar morph between them that is composed of a sequence of linear morphing steps requires $\Omega(n')$ steps.*

Alamdari et al. [3] describe the two drawings in the statement of Theorem 83 as follows. In the drawing $\Gamma^{-}(n')$ the n' vertices $a_1, \ldots, a_{n'}$ are placed on a horizontal line with a_i to the left of a_{i+1} for $1 \leq i < n'$, see Figure 105(a). In the drawing $\Gamma^{\Delta}(n')$ the path forms a spiral, see Figure 105(b).

More precisely, let e_i denote the edge $\{a_i, a_{i+1}\}$. Then for each i with $i \equiv 1 \pmod 3$, the edge e_i is horizontal and a_i is to the left of a_{i+1}. For each i with $i \equiv 2 \pmod 3$, the edge e_i is parallel to the line $y = \tan(2\pi/3)x$ and a_i is to the right of a_{i+1}. Finally, for each i with $i \equiv 0 \pmod 3$, the edge e_i is parallel to the line $y = \tan(-2\pi/3)x$ and a_i is to the right of a_{i+1}.

To prove Theorem 59, we present a drawing $\Gamma^{\underline{\Delta}}(n)$ of a cycle on n vertices that contains a subdrawing of $\Gamma^{\Delta}(n')$ for some $n' \in \Theta(n)$, see Figure 105(c). The existence of a convexifying planar morph for $\Gamma^{\underline{\Delta}}(n)$ with $o(n)$ linear morphing steps would imply the existence of a planar morph between $\Gamma^{\Delta}(n')$ and $\Gamma^{-}(n')$ with $o(n')$ linear morphing steps, contradicting Theorem 83.

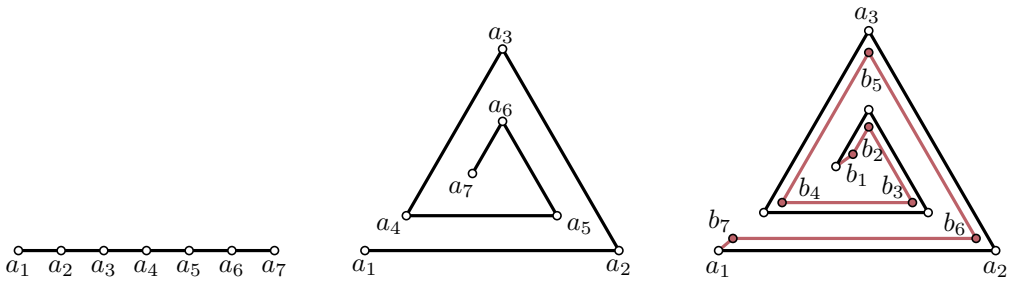

(a) (b) (c)

Figure 105: The drawings (a) $\Gamma^{-}(7)$, (b) $\Gamma^{\Delta}(7)$, and (c) $\Gamma^{\underline{\Delta}}(14)$.

Proof of Theorem 59. Let $n' = \lfloor n/2 \rfloor$. Let $\Gamma^{\underline{\Delta}}(n)$ be some planar straight-line drawing of the cycle $C = (a_1, \ldots, a_{n'}, b_1, \ldots, b_{n-n'})$ such that the subdrawing of the path $P = (a_1, \ldots, a_{n'})$ is $\Gamma^{\Delta}(n')$. The exact realization of the path $(b_1, \ldots b_{n-n'})$ is irrelevant for the purposes of this proof. We give an example in Figure 105(c).

Assume for a contradiction that there exists a morph $\mathcal{M}$ composed of sequence of $o(n)$ linear morphing steps that convexifies $\Gamma^{\underline{\Delta}}(n)$. Restricting the morph $\mathcal{M}$ to

the path P transforms $\Gamma^{\Delta}(n')$ into a reflex chain Γ_P on the boundary of a strictly convex polygon. It is easy to find $O(1)$ additional morphing steps that transform Γ_P into the drawing $\Gamma^{-}(n')$. For example, we can intermediately aim for coordinates of a_1 and a_n that are extreme in some direction as in the proof of Lemma 72. In fact, our situation here is much simpler, as we are not restricted to horizontal and vertical morphs anymore. Extending $\mathcal{M}$ by these additional morphs yields planar morph with $o(n) \subseteq o(n')$ linear morphing steps that transforms $\Gamma^{\Delta}(n')$ into $\Gamma^{-}(n')$. This is a contradiction to Theorem 83. □

5.7 Lower bound on grid size

Every 3-connected planar graph can be drawn with convex faces on a $O(n) \times O(n)$ grid [24, 34, 53] or with strictly convex faces on a $O(n^2) \times O(n^2)$ grid [18]. It would be desirable to find convexifying morphs in which the intermediate drawings also lie on a polynomial-sized grid. In this section, we show that this is not achievable with our current approach and, more generally, with any approach that naively uses a convex redrawing technique as described in Sections 5.2.4 and 5.4. To do so, we design a family of drawings to show that a single horizontal morph to a convex drawing may unavoidably blow up the width of the drawing from $W \in \Theta(n)$ to $2^{W-2}(W-2)!$.

A *grid-drawing* of a plane (or planar) graph G is a straight-line planar drawing of G in which all vertices are placed at integer coordinates. The *width* of a grid-drawing Γ is the length of the horizontal sides of the bounding box of Γ. The *height* is defined analogously.

We remark that each of our grid-drawings uses consecutive integer y-coordinates from an integer interval and, hence, may be thought of as a straight-line level planar drawing, cf. Section 3.1.1. Though, it does not seem sensible to express our result in terms of level graphs as this would unnecessarily clutter up the notation and require several new definitions.

We will prove the following lemma, which describes our family of drawings.

Lemma 84. *There exists a family of grid-drawings $(\Gamma_k)_k$ of internally 3-connected graphs G_k on $n_k = 6k+1$ vertices such that the width of Γ_k is $w(\Gamma_k) = 2k$, and the height of Γ_k is $4k-1$. Moreover, any grid-redrawing of Γ_k with convex inner faces that preservers the y-coordinate of each vertex has a width of at least*

$$w_c(\Gamma_k) := 4^{k-1}(2k-2)! = 2^{w(\Gamma_k)-2}(w(\Gamma_k)-2)!.$$

Our construction is inspired by a family of level graphs presented by Lin and Eades [97] to show that level planar straight-line drawings (that are not necessarily convex) may require large width. In particular, they used an elementary geometric observation very similar to the following Observation 85. The fundamental difference between our construction and the one by Lin and Eades is that our drawings Γ_k,

which are straight-line planar, have a small width of $O(n_k)$. Only the convexity requirement enforces a large width.

Observation 85. *Consider a grid-drawing of a path (a_1, b_1, c, b_2, a_2) such that $\mathrm{y}(c) = k$, b_1 is to the left of b_2 with $\mathrm{y}(b_1) = \mathrm{y}(b_2) = k + 1$, and a_1 is to the left of a_2 with $\mathrm{y}(a_1) = \mathrm{y}(a_2) = k + j$ where $j \geq 2$. Let D denote the distance of b_1 and b_2.*

If a_1 is not to the right of the oriented line $\overrightarrow{cb_1}$ and a_2 is not to the left of the oriented line $\overrightarrow{cb_2}$, then the distance between a_1 and a_2 is at least jD, for an illustration see Figure 106. □

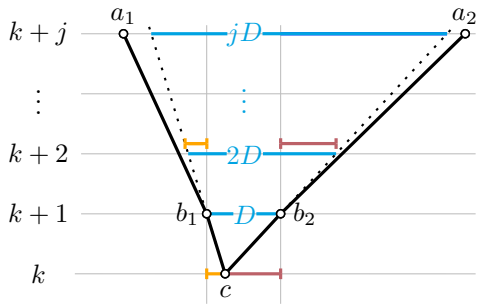

Figure 106: Illustration of Observation 85.

Proof of Lemma 84. The drawing Γ_1 of G_1 is depicted in Figure 107. The drawing Γ_{k+1} of G_{k+1} is obtained from Γ_k by introducing six new vertices on a total of four new y-coordinates. More precisely, we introduce a cycle on the six new vertices as the new outer face and four new internal edges as shown in Figure 107.

Note that G_k has exactly two vertices of degree two on the outer face. By construction, the graph G_k has $6k + 1$ vertices. Moreover, the drawing Γ_k uses $4k$ pairwise distinct y-coordinates and has width $w(\Gamma_k) = 2k$. We shift the coordinate system such that in each drawing Γ_k, the smallest y-coordinate is 1 and, consequently, the largest y-coordinate is $4k$.

The drawing Γ_1 is convex. Moreover, if Γ_{k-1} admits a grid-redrawing Γ'_{k-1} with convex inner faces that preserves the y-coordinates of Γ_{k-1}, then the analogous statement holds for Γ_k: we can simply augment Γ'_{k-1} to a redrawing of Γ_k and then shift the vertices with y-coordinates $(4k - 3), 2$, and $(4k - 1)$ far enough outwards, in this order, such that the four new inner faces become convex. This shows that every Γ_k admits a grid-redrawing with convex inner faces that preserves the y-coordinates. Moreover, we can turn this drawing into a convex drawing by further shifting the two vertices with y-coordinate 2, which are the only vertices that may have convex outer angles, while maintaining convex inner faces. Consequently, for each $k \geq 1$, the graph G_k has a convex grid-drawing. By Lemma 60, it follows that G_k is internally 3-connected. This concludes the construction of the family $(\Gamma_k)_k$. It remains to argue about the required width of the desired redrawings.

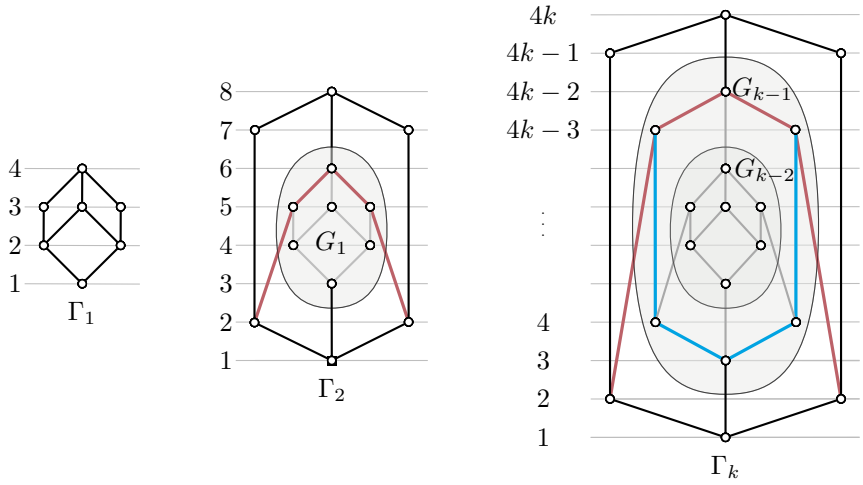

Figure 107: Illustration of the drawings Γ_k for small values of k. Any redrawing of Γ_k that preserves the y-coordinates and has convex inner faces needs an exponential increase in width.

Since Γ_1 contains three vertices with y-coordinate 3, every grid-redrawing that preserves the y-coordinate of each vertex has a width of at least $2 > \mathrm{w}_c(\Gamma_1) = 1$. For $k \geqslant 2$, we prove the claim by induction with the stronger induction hypothesis that in any grid-redrawing Γ'_k of Γ_k with convex inner faces that preserves the y-coordinate of each vertex, the two vertices with y-coordinate 2 have distance at least $\mathrm{w}_c(\Gamma_k)$. Note that this implies that the width of the redrawing is at least $\mathrm{w}_c(\Gamma_k)$ as in the statement of the lemma.

For the base case $k = 2$, consider any grid-redrawing Γ'_2 of Γ_2 with convex inner faces that preserves the y-coordinate of each vertex. Since there are three vertices with y-coordinate 5, the left-most and right-most of these vertices have distance at least 2. Consider the path induced by the outermost vertices with y-coordinate 6, 5, and 2, respectively, see Figure 107. By Observation 85 applied to this path, we obtain that the distance between the two vertices with y-coordinate 2 is at least $8 = \mathrm{w}_c(\Gamma_2)$.

For the induction step with $k \geqslant 3$, consider a grid-redrawing Γ'_k of Γ_k with convex inner faces that preserves the y-coordinate of each vertex. We apply Observation 85 twice in Γ'_k. By the induction hypothesis, the distance of the vertices with y-coordinate 4 in Γ'_k is at least $\mathrm{w}_c(\Gamma_{k-1})$ since Γ'_k contains a redrawing of Γ_{k-1} in which all inner faces are convex. Applying Observation 85 to the path induced by the outermost vertices with y-coordinates 3, 4, and $4k-3$ in Γ'_k (highlighted by the blue path in Figure 107) shows that the vertices with y-coordinate $4k-3$ in Γ'_k have a distance of at least $(4k-6)\mathrm{w}_c(\Gamma_{k-1})$.

Applying Observation 85 to the outermost vertices with y-coordinate 2, $4k-3$, and $4k-2$ (highlighted by the red path in Figure 107) shows that the vertices with

y-coordinate 2 in Γ'_k have distance at least

$$\begin{aligned}(4k-4)\cdot(4k-6)\cdot \mathrm{w}_c(\Gamma_{k-1}) &= 2(2k-2)\cdot 2(2k-3)\cdot \mathrm{w}_c(\Gamma_{k-1})\\ &= (2k-2)\cdot(2k-3)\cdot 4\cdot 4^{k-2}(2k-4)!\\ &= 4^{k-1}(2k-2)!\\ &= \mathrm{w}_c(\Gamma_k),\end{aligned}$$

which concludes the proof. □

5.8 Conclusion

We have shown how to morph any straight-line planar drawing of an internally 3-connected graph to a strictly convex drawing while preserving planarity and increasing convexity throughout the morph. Moreover, our morph is composed of a linear number of horizontal and vertical steps, which is asymptotically optimal. The following questions are open:

1. Our algorithm for finding a convexity-increasing morph to a strictly convex drawing can be executed in $O(n^{1+\omega/2} + n^2 \log n)$ time on a real-RAM. It would be interesting to find even more efficient algorithms, or to establish non-trivial lower bounds on the runtime.

2. A main open question is to design piece-wise linear morphs with a polynomial bound on the bit complexity of the intermediate drawings. This would be a step towards having intermediate drawings that lie on a polynomial-sized grid, i.e. with a logarithmic number of bits for each vertex's coordinates. This is open both for our problem of morphing to a (strictly) convex drawing and for the problem of morphing between two given planar straight-line drawings.

Chapter 6

Concluding remarks

"Graph drawing is the best possible field I can think of. It merges aesthetics, mathematical beauty and wonderful algorithms." This quote is attributed to Donald E. Knuth, on the occasion of the Symposium on Graph Drawing in 1996. In the author's opinion, one of the most beautiful aspects of Graph Drawing is that it is a very multifaceted discipline that borrows techniques from several research fields such as classical graph theory, algorithms and complexity, computational geometry, information visualization, the theory of topological and geometric graphs, knot theory, topology, order theory, physics, and algebra. Consequently, the individual problems can have very distinct flavor despite their common goal—drawing graphs. We hope that this aspect became apparent in this thesis. In this section, we briefly discuss the facets that played a role throughout Chapters 3–5, and then conclude by giving a summary of our results and discussing the most intriguing open problems.

In Chapter 4, we considered the problem of two-page book embeddability. We used a fact that allows for a reformulation as a graph augmentation problem: a graph can be embedded on two pages if and only if it is a subgraph of a Hamiltonian planar graph [22]. Consequently, our algorithms and the corresponding proofs make heavy use of many concepts of classical graph theory: connectivity, separators, edge contractions, decomposition trees, graph augmentation, and more. In particular, we devised an algorithm to solve the major subproblem of finding a set of edges that belong to separating triangles such that their simultaneous collapse results in a graph without separating triangles or adjacent separation pairs.

Very much in contrast are the techniques employed in Chapter 3, about Ordered Level Planarity. Here, one of our main goals was to reduce to Manhattan Geodesic Planarity such that the resulting instances are matchings. To achieve this goal, the reduction that shows the $\mathcal{NP}$-hardness of Ordered Level Planarity was designed to produce instances with maximum degree two. More precisely, the generated instances are disjoint unions of paths. As such, they do not exhibit much in terms of meaningful graph theoretic properties. Instead, we had to rely on direct planarity arguments to prove the correctness of the reduction, which is reminiscent

of problems related to topological graphs.

The objects of desire in Chapter 5 were straight-line drawings in which each face is described by a convex polygon. Naturally, we relied on techniques from computational geometry [42], such as convex hulls, trapezoidal decompositions, and the orientation test to accomplish our goals. In that chapter, we also used Tutte's graph drawing algorithm, which highlights two additional facets of Graph Drawing: Tutte's algorithm is based on a physical analogy, as explained in Section 5.4. The use of such analogies is no rarity in Graph Drawing; Brandes [26] provides a comprehensive survey. In addition, Tutte's algorithm involves solving a system of linear equations, which can be done efficiently by using advanced techniques from linear algebra.

Of course, given that we considered computational problems, we also made use of algorithmic techniques, data structures (in particular, for dynamic connectivity queries), and concepts from complexity theory.

Outlook

We conclude by summarizing our contributions and discuss the most intriguing open problems. Further open problems can be found in the conclusions (Sections 3.6, 4.11, and 5.8) of the three main Chapters 3, 4, and 5.

Ordered Level Planarity. Our research on Ordered Level Planarity was conclusive: we have established a complexity dichotomy with respect to both the maximum degree Δ and the level width λ. The most significant result is an $\mathcal{NP}$-hardness proof for the case $\Delta = \lambda = 2$.

We demonstrated the relevance of our result by stating reductions to several other graph drawing problems that have been studied in the literature: Clustered Level Planarity, T-Level Planarity, Constrained Level Planarity, Geodesic Planarity, and Bi-Monotonicity. Due to the simple structure of the instances for which we show Ordered Level Planarity to be $\mathcal{NP}$-hard, these reductions were readily obtained. Moreover, they produce very constrained instances of the targeted problems, which allowed us to solve multiple open questions that were posed by members of the Graph Drawing community.

Recently, Da Lozzo, Di Battista, and Frati [39] applied our result to show the $\mathcal{NP}$-hardness of another generalization of Ordered Level Planarity. We expect that the $\mathcal{NP}$-hardness of Ordered Level Planarity will serve as a useful tool for further reductions.

Only a minor problem regarding the complexity of the (Manhattan) Geodesic Planarity problem remains open: the case where the input graph is connected. We conjecture that $\mathcal{NP}$-hardness holds even in this case.

Convexity-increasing morphs. We presented an algorithm to convexify any given straight-line planar drawing of an internally 3-connected graph by means of a convexity-increasing morph. For graphs that are not internally 3-connected such a morph cannot exist since these graphs do not admit strictly convex drawings. Our morph is composed of a linear number of horizontal and vertical steps, which is asymptotically optimal.

The runtime of our current approach is $O(n^{1+\omega/2} + n^2 \log n)$. It would be interesting to find even more efficient algorithms. Of course, one could try to find an entirely new method to achieve convexity-increasing morphs. A much more intriguing question is whether the runtime of the procedure for finding convex redrawings while maintaining y-coordinates (Lemma 67) can be improved, as this is the bottleneck of our morphing algorithm, and would also speed up the runtime of the many other morphing algorithms that rely on a convex redrawing technique.

Hong and Nagamochi's [79] version of the redrawing method is a combinatorial algorithm that operates recursively. Its runtime is $O(n^2)$. In contrast, the version by Chrobak, Goodrich, and Tamassia [33] is based on using Tutte's graph drawing algorithm [124], which requires solving a linear system of equations and can be executed in $O(n^{\omega/2} + n \log n)$. To improve the runtime further, it would be interesting to explore entirely new methods for creating convex redrawings while maintaining y-coordinates.

Two-page book embeddings of triconnected planar graphs. We have shown that $k = 5$ is the largest $k \in \mathbb{N}$ such that each 3-connected planar graph where the degree of vertices that belong to separating 3-cycles is bounded by k is subhamiltonian. Equivalently, these graphs admit two-page book embeddings. Our results strengthen earlier work by Heath [75] and by Bauernöppel [20] and, independently, Bekos, Gronemann, and Raftopoulou [21], who showed that planar graphs of maximum degree three and four, respectively, can always be realized on two pages.

It is a very intriguing question, whether our result can be extended to the 2-connected case. It is not clear how one could design an algorithm for the 2-connected case that adopts the same high-level strategy as our current approach, as 3-connectivity is used throughout the entire proof. For instance, in a 2-connected graph, it is no longer true that separating triangles are pairwise vertex-disjoint. Moreover, it is also no longer true that a 4-inhibitor has at least three edges to each of its sides. An even more severe problem is that the graph can contain adjacent separation pairs. This prevents us from stellating the graph. Further, it implies that the edges of a separating triangle may belong to distinct triconnected components, so that it is no longer possible to recurse on the rigid triconnected components while ignoring separating triangles that use virtual edges.

However, maybe it is possible to use our result in a black box fashion to prove the result for the 2-connected case? One such approach would be to apply it directly to the rigid triconnected components. However, it is not clear how to glue the resulting

cycles together. Another option would be to augment the graph such that it becomes 3-connected. It does not always seem possible to do so without violating the degree bounds, but perhaps the augmentation can be carried out while increasing the degree of at most one vertex per separating 3-cycle? If so, it would be interesting whether our high-level strategy for the 3-connected case can be extended to this slightly relaxed degree bound.

Bibliography

[1] graphdrawing.org. http://www.graphdrawing.org. Accessed: 2019-10-15.

[2] 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 G. Aloupis and Bremner, editors, *Canadian Conference on Computational Geometry (CCCG)*, pages 229–234, 2011.

[3] 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. *SIAM J. Computing*, 46(2):29 pages, 2017.

[4] S. Alamdari, P. Angelini, T. M. Chan, G. Di Battista, F. Frati, A. Lubiw, M. Patrignani, V. Roselli, S. Singla, and B. T. Wilkinson. Morphing planar graph drawings with a polynomial number of steps. In S. Khanna, editor, *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6–8, 2013*, pages 1656–1667. SIAM, 2013.

[5] N. Alon and R. Yuster. Matrix sparsification and nested dissection over arbitrary fields. *Journal of the ACM*, 60(4):25, 2013.

[6] P. Angelini, E. Colasante, G. Di Battista, F. Frati, and M. Patrignani. Monotone drawings of graphs. *J. Graph Algorithms Appl.*, 16(1):5–35, 2012.

[7] P. Angelini, G. Da Lozzo, G. Di Battista, F. Frati, M. Patrignani, and V. Roselli. Morphing planar graph drawings optimally. In J. Esparza, P. Fraigniaud, T. Husfeldt, and E. Koutsoupias, editors, *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8–11, 2014, Proceedings, Part I*, volume 8572 of *Lecture Notes in Computer Science*, pages 126–137. Springer, 2014.

[8] P. Angelini, G. Da Lozzo, G. Di Battista, F. Frati, M. Patrignani, and I. Rutter. Beyond level planarity. In Y. Hu and M. Nöllenburg, editors, *Graph Drawing*

and Network Visualization - 24th International Symposium, GD 2016, Athens, Greece, September 19–21, 2016, Revised Selected Papers, volume 9801 of *Lecture Notes in Computer Science*, pages 482–495. Springer, 2016.

[9] P. Angelini, G. Da Lozzo, G. Di Battista, F. Frati, and V. Roselli. The importance of being proper: (in clustered-level planarity and T-level planarity). *Theor. Comput. Sci.*, 571:1–9, 2015.

[10] 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.

[11] P. Angelini, G. Di Battista, F. Frati, M. Patrignani, and I. Rutter. Testing the simultaneous embeddability of two graphs whose intersection is a biconnected or a connected graph. *J. Discrete Algorithms*, 14:150–172, 2012.

[12] P. Angelini, F. Frati, M. Patrignani, and V. Roselli. Morphing planar graph drawings efficiently. In S. K. Wismath and A. Wolff, editors, *Graph Drawing - 21st International Symposium, GD 2013, Bordeaux, France, September 23–25, 2013, Revised Selected Papers*, volume 8242 of *Lecture Notes in Computer Science*, pages 49–60. Springer, 2013.

[13] S. Appelle. Perception and discrimination as a function of stimulus orientation: the "oblique effect" in man and animals. *Psychological Bulletin*, 78(4):266, 1972.

[14] A. Asinowski, T. Miltzow, and G. Rote. Quasi-parallel segments and characterization of unique bichromatic matchings. *Journal of Computational Geometry*, 6:185–219, 2015.

[15] A. S. Asratian and N. K. Khachatrian. Some localization theorems on Hamiltonian circuits. *J. Combin. Theory Ser. B*, 49(2):287–294, 1990.

[16] C. Bachmaier, F. Brandenburg, and M. Forster. Radial level planarity testing and embedding in linear time. *J. Graph Algorithms Appl.*, 9(1):53–97, 2005.

[17] C. Bachmaier and W. Brunner. Linear time planarity testing and embedding of strongly connected cyclic level graphs. In D. Halperin and K. Mehlhorn, editors, *Algorithms - ESA 2008, 16th Annual European Symposium, Karlsruhe, Germany, September 15–17, 2008. Proceedings*, volume 5193 of *Lecture Notes in Computer Science*, pages 136–147. Springer, 2008.

[18] I. Bárány and G. Rote. Strictly convex drawings of planar graphs. *Documenta Math*, 11:369–391, 2006.

[19] F. Barrera-Cruz, P. Haxell, and A. Lubiw. Morphing Schnyder drawings of planar triangulations. *Discrete & Computational Geometry*, 61(1):161–184, Jan 2019.

[20] F. Bauernöppel. Degree bounds and subhamiltonian cycles in planar graphs. *J. Inf. Process. Cybern.*, 23(10–11):529–536, 1987.

[21] M. A. Bekos, M. Gronemann, and C. N. Raftopoulou. Two-page book embeddings of 4-planar graphs. *Algorithmica*, 75(1):158–185, 2016.

[22] F. Bernhart and P. C. Kainen. The book thickness of a graph. *J. Combin. Theory Ser. B*, 27(3):320–331, 1979.

[23] T. C. Biedl, G. Kant, and M. Kaufmann. On triangulating planar graphs under the four-connectivity constraint. *Algorithmica*, 19(4):427–446, 1997.

[24] N. Bonichon, S. Felsner, and M. Mosbah. Convex drawings of 3-connected plane graphs. *Algorithmica*, 47(4):399–420, 2007.

[25] K. S. Booth and G. S. Lueker. Testing for the consecutive ones property, interval graphs, and graph planarity using PQ-tree algorithms. *J. Comput. Syst. Sci.*, 13(3):335–379, 1976.

[26] U. Brandes. Drawing on physical analogies. In M. Kaufmann and D. Wagner, editors, *Drawing Graphs, Methods and Models*, volume 2025 of *Lecture Notes in Computer Science*, pages 71–86. Springer, 2001.

[27] G. Brückner and I. Rutter. Partial and constrained level planarity. In P. N. Klein, editor, *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16–19*, pages 2000–2011. SIAM, 2017.

[28] S. Cairns. Deformations of plane rectilinear complexes. *The American Mathematical Monthly*, 51(5):247–252, 1944.

[29] J. Cardinal, M. Hoffmann, V. Kusters, C. D. Tóth, and M. Wettstein. Arc diagrams, flip distances, and Hamiltonian triangulations. *Comput. Geom. Theory Appl.*, 68:206–225, 2018.

[30] N. Chiba and T. Nishizeki. The Hamiltonian cycle problem is linear-time solvable for 4-connected planar graphs. *J. Algorithms*, 10(2):187–211, 1989.

[31] N. Chiba, K. Onoguchi, and T. Nishizeki. Drawing plane graphs nicely. *Acta Informatica*, 22(2):187–201, 1985.

[32] F. Y. L. Chin, J. Snoeyink, and C. A. Wang. Finding the medial axis of a simple polygon in linear time. *Discrete & Computational Geometry*, 21(3):405–420, 1999.

[33] M. Chrobak, M. T. Goodrich, and R. Tamassia. Convex drawings of graphs in two and three dimensions (preliminary version). In S. Whitesides, editor, *Proceedings of the Twelfth Annual Symposium on Computational Geometry, Philadelphia, PA, USA, May 24–26, 1996*, pages 319–328. ACM, 1996.

[34] M. Chrobak and G. Kant. Convex grid drawings of 3-connected planar graphs. *Int. J. Comput. Geometry Appl.*, 7(3):211–223, 1997.

[35] F. R. K. Chung, F. T. Leighton, and A. L. Rosenberg. Embedding graphs in books: A layout problem with applications to VLSI design. *SIAM Journal on Algebraic Discrete Methods*, 8(1):33–58, 1987.

[36] É. Colin De Verdière, M. Pocchiola, and G. Vegter. Tutte's barycenter method applied to isotopies. *Computational Geometry*, 26(1):81–97, 2003.

[37] R. Connelly, E. D. Demaine, and G. Rote. Straightening polygonal arcs and convexifying polygonal cycles. *Discrete & Computational Geometry*, 30:205–239, 2003.

[38] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms, 3rd Edition.* MIT Press, 2009.

[39] G. Da Lozzo, G. Di Battista, and F. Frati. Extending upward planar graph drawings. In Z. Friggstad, J. Sack, and M. R. Salavatipour, editors, *Algorithms and Data Structures - 16th International Symposium, WADS 2019, Edmonton, AB, Canada, August 5–7, 2019, Proceedings*, volume 11646 of *Lecture Notes in Computer Science*, pages 339–352. Springer, 2019.

[40] G. Da Lozzo, G. Di Battista, F. Frati, M. Patrignani, and V. Roselli. Upward planar morphs. In T. C. Biedl and A. Kerren, editors, *Graph Drawing and Network Visualization - 26th International Symposium, GD 2018, Barcelona, Spain, September 26–28, 2018, Proceedings*, volume 11282 of *Lecture Notes in Computer Science*, pages 92–105. Springer, 2018.

[41] M. de Berg and A. Khosravi. Optimal binary space partitions for segments in the plane. *Int. J. Comput. Geometry Appl.*, 22(3):187–206, 2012.

[42] M. de Berg, M. Van Kreveld, M. Overmars, and O. Cheong. *Computational Geometry: Algorithms and Applications.* Springer, 3rd edition, 2008.

[43] N. G. de Bruijn. Problems 17 and 18. *Nieuw Archief voor Wiskunde*, 2:67, 1954. Answers in *Wiskundige Opgaven met de Oplossingen* 20:19–20, 1955.

[44] G. Di Battista, P. Eades, R. Tamassia, and I. G. Tollis. Algorithms for drawing graphs: an annotated bibliography. *Comput. Geom.*, 4:235–282, 1994.

[45] G. Di Battista, P. Eades, R. Tamassia, and I. G. Tollis. *Graph Drawing: Algorithms for the Visualization of Graphs*. Prentice-Hall, 1999.

[46] G. Di Battista and E. Nardelli. Hierarchies and planarity theory. *IEEE Trans. Systems, Man, and Cybernetics*, 18(6):1035–1046, 1988.

[47] G. Di Battista and R. Tamassia. Algorithms for plane representations of acyclic digraphs. *Theoretical Computer Science*, 61(2-3):175–198, 1988.

[48] R. Diestel. *Graph Theory, 4th Edition*, volume 173 of *Graduate texts in mathematics*. Springer, 2012.

[49] G. A. Dirac. Some theorems on abstract graphs. *Proc. London Math. Soc.*, s3-2(1):69–81, 1952.

[50] M. Eiglsperger, S. P. Fekete, and G. W. Klau. Orthogonal graph drawing. In M. Kaufmann and D. Wagner, editors, *Drawing Graphs, Methods and Models*, volume 2025 of *Lecture Notes in Computer Science*, pages 121–171. Springer, 2001.

[51] M. N. Ellingham, E. A. Marshall, K. Ozeki, and S. Tsuchiya. Hamiltonicity of planar graphs with a forbidden minor. *Journal of Graph Theory*, 90(4):459–483, 2019.

[52] G. Ewald. Hamiltonian circuits in simplicial complexes. *Geom. Dedicata*, 2:115–125, 1973.

[53] S. Felsner. Convex drawings of planar graphs and the order dimension of 3-polytopes. *Order*, 18(1):19–37, 2001.

[54] H. Fleischner. The uniquely embeddable planar graphs. *Discrete Mathematics*, 4(4):347–358, 1973.

[55] M. S. Floater. Parameterization and smooth approximation of surface triangulations. *Computer Aided Geometric Design*, 14(3):231–250, 1997.

[56] M. S. Floater. Parametric tilings and scattered data approximation. *International Journal of Shape Modeling*, 4(03n04):165–182, 1998.

[57] M. S. Floater and C. Gotsman. How to morph tilings injectively. *Journal of Computational and Applied Mathematics*, 101(1-2):117–129, 1999.

[58] M. Forster and C. Bachmaier. Clustered level planarity. In P. van Emde Boas, J. Pokorný, M. Bieliková, and J. Stuller, editors, *SOFSEM 2004: Theory and Practice of Computer Science, 30th Conference on Current Trends in Theory and Practice of Computer Science, Merin, Czech Republic, January 24–30, 2004*, volume 2932 of *Lecture Notes in Computer Science*, pages 218–228. Springer, 2004.

[59] R. Fulek, M. J. Pelsmajer, M. Schaefer, and D. Štefankovič. Hanani–Tutte, monotone drawings, and level-planarity. In J. Pach, editor, *Thirty Essays on Geometric Graph Theory*, pages 263–287. Springer, 2013.

[60] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.

[61] M. R. Garey, D. S. Johnson, and R. E. Tarjan. The planar Hamiltonian circuit problem is NP-complete. *SIAM J. Comput.*, 5(4):704–714, 1976.

[62] A. Garg and R. Tamassia. On the computational complexity of upward and rectilinear planarity testing. *SIAM J. Comput.*, 31(2):601–625, 2001.

[63] E. D. Giacomo, F. Frati, R. Fulek, L. Grilli, and M. Krug. Orthogeodesic point-set embedding of trees. *Comput. Geom.*, 46(8):929–944, 2013.

[64] E. D. Giacomo, L. Grilli, M. Krug, G. Liotta, and I. Rutter. Hamiltonian orthogeodesic alternating paths. *J. Discrete Algorithms*, 16:34–52, 2012.

[65] J. Gomes, L. Darsa, B. Costa, and L. Velho. *Warping & morphing of graphical objects*. Morgan Kaufmann, 1999.

[66] 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.

[67] C. Gotsman and V. Surazhsky. Guaranteed intersection-free polygon morphing. *Computers & Graphics*, 25(1):67–75, 2001.

[68] B. Grünbaum. *Convex Polytopes*, volume 221 of *Graduate Texts in Mathematics*. Springer, 2003.

[69] X. Guan and W. Yang. Embedding planar 5-graphs in three pages. *Discrete Applied Mathematics*, 2019.

[70] L. J. Guibas and F. F. Yao. On translating a set of rectangles. In *Proc. 12th Annual ACM Symposium Theory of Computing (STOC 1980)*, pages 154–160, 1980.

[71] L. J. Guibas and F. F. Yao. On translating a set of rectangles. In F. P. Preparata, editor, *Computational Geometry*, volume 1 of *Advances in Computing Research*, pages 61–77. JAI Press, Greenwich, Conn., 1983.

[72] C. Gutwenger and P. Mutzel. A linear time implementation of SPQR-trees. In *Graph Drawing, 8th International Symposium, GD 2000, Colonial Williamsburg, VA, USA, September 20–23, 2000, Proceedings*, volume 1984 of *LNCS*, pages 77–90. Springer, 2000.

[73] R. Haas, D. Orden, G. Rote, F. Santos, B. Servatius, H. Servatius, D. L. Souvaine, I. Streinu, and W. Whiteley. Planar minimally rigid graphs and pseudo-triangulations. *Comput. Geom.*, 31(1-2):31–61, 2005.

[74] C. Haslinger and P. F. Stadler. RNA structures with pseudo-knots: Graph-theoretical, combinatorial, and statistical properties. *Bulletin of Mathematical Biology*, 61(3):437–467, May 1999.

[75] L. S. Heath. *Algorithms for embedding graphs in books.* Phd thesis, University of North Carolina, Chapel Hill, 1985.

[76] L. S. Heath and S. V. Pemmaraju. Recognizing leveled-planar Dags in linear time. In F. Brandenburg, editor, *Graph Drawing, Symposium on Graph Drawing, GD '95, Passau, Germany, September 20–22, 1995, Proceedings*, volume 1027 of *Lecture Notes in Computer Science*, pages 300–311. Springer, 1995.

[77] M. Hoffmann and B. Klemz. Triconnected planar graphs of maximum degree five are subhamiltonian. In M. A. Bender, O. Svensson, and G. Herman, editors, *27th Annual European Symposium on Algorithms, ESA 2019, September 9-11, 2019, Munich/Garching, Germany*, volume 144 of *LIPIcs*, pages 58:1–58:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

[78] J. Holm, G. F. Italiano, A. Karczmarz, J. Łącki, and E. Rotenberg. Decremental SPQR-trees for planar graphs. In *Proc. 26th Annu. European Sympos. Algorithms (ESA 2018)*, volume 112 of *LIPIcs*, pages 46:1–46:16, 2018.

[79] 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.

[80] J. E. Hopcroft and R. E. Tarjan. Dividing a graph into triconnected components. *SIAM J. Comput.*, 2(3):135–158, 1973.

[81] J. E. Hopcroft and R. E. Tarjan. Efficient planarity testing. *J. ACM*, 21(4):549–568, 1974.

[82] M. Jünger, S. Leipert, and P. Mutzel. Pitfalls of using PQ-trees in automatic graph drawing. In G. D. Battista, editor, *Graph Drawing, 5th International Symposium, GD '97, Rome, Italy, September 18–20, 1997, Proceedings*, volume 1353 of *Lecture Notes in Computer Science*, pages 193–204. Springer, 1997.

[83] M. Jünger, S. Leipert, and P. Mutzel. Level planarity testing in linear time. In S. Whitesides, editor, *Graph Drawing, 6th International Symposium, GD'98, Montréal, Canada, August 1998, Proceedings*, volume 1547 of *Lecture Notes in Computer Science*, pages 224–237. Springer, 1998.

[84] P. C. Kainen and S. Overbay. Book embeddings of graphs and a theorem of Whitney. *Tech. Report GUGU-2/25/03 Georgetown University*, 13(12):8, 2003.

[85] P. C. Kainen and S. Overbay. Extension of a theorem of Whitney. *Appl. Math. Lett.*, 20(7):835–837, 2007.

[86] B. Katz, M. Krug, I. Rutter, and A. Wolff. Manhattan-geodesic embedding of planar graphs. In D. Eppstein and E. R. Gansner, editors, *Graph Drawing, 17th International Symposium, GD 2009, Chicago, IL, USA, September 22–25, 2009. Revised Papers*, volume 5849 of *Lecture Notes in Computer Science*, pages 207–218. Springer, 2009.

[87] M. Kaufmann and D. Wagner, editors. *Drawing Graphs, Methods and Models*, volume 2025 of *Lecture Notes in Computer Science*. Springer, 2001.

[88] L. Kleist, B. Klemz, A. Lubiw, L. Schlipf, F. Staals, and D. Strash. Convexity-increasing morphs of planar graphs. In A. Brandstädt, E. Köhler, and K. Meer, editors, *Graph-Theoretic Concepts in Computer Science - 44th International Workshop, WG 2018, Cottbus, Germany, June 27–29, 2018, Proceedings*, volume 11159 of *Lecture Notes in Computer Science*, pages 318–330. Springer, 2018.

[89] L. Kleist, B. Klemz, A. Lubiw, L. Schlipf, F. Staals, and D. Strash. Convexity-increasing morphs of planar graphs. In *Proceedings of the 34th European Workshop on Computational Geometry (EuroCG 2018)*, pages 65:1–65:6, Apr. 2018.

[90] L. Kleist, B. Klemz, A. Lubiw, L. Schlipf, F. Staals, and D. Strash. Convexity-increasing morphs of planar graphs. *Comput. Geom.*, 84:69–88, 2019.

[91] B. Klemz and G. Rote. Ordered level planarity and geodesic planarity. In *Proceedings of the 33rd European Workshop on Computational Geometry (EuroCG 2017)*, pages 269–272, Apr. 2017.

[92] B. Klemz and G. Rote. Ordered level planarity, geodesic planarity and bi-monotonicity. In F. Frati and K. Ma, editors, *Graph Drawing and Network Visualization - 25th International Symposium, GD 2017, Boston, MA, USA, September 25–27, 2017, Revised Selected Papers*, volume 10692 of *Lecture Notes in Computer Science*, pages 440–453. Springer, 2017.

[93] B. Klemz and G. Rote. Ordered level planarity and its relationship to geodesic planarity, bi-monotonicity, and variations of level planarity. *ACM Trans. Algorithms*, 15(4):53:1–53:25, 2019.

[94] D. E. Knuth. Computer-drawn flowcharts. *Commun. ACM*, 6(9):555–563, 1963.

[95] F. Le Gall. Powers of tensors and fast matrix multiplication. In *Proceedings of the 39th International Symposium on Symbolic and Algebraic Computation*, pages 296–303. ACM, 2014.

[96] D. T. Lee. On finding the convex hull of a simple polygon. *International Journal of Parallel Programming*, 12(2):87–98, 1983.

[97] X. Lin and P. Eades. Towards area requirements for drawing hierarchically planar graphs. *Theoretical Computer Science*, 292(3):679–695, 2003.

[98] R. J. Lipton, D. J. Rose, and R. E. Tarjan. Generalized nested dissection. *SIAM J. Numerical Analysis*, 16(2):346–358, 1979.

[99] R. J. Lipton and R. E. Tarjan. Applications of a planar separator theorem. *SIAM J. Comput.*, 9(3):615–627, 1980.

[100] S. Mac Lane. A structural characterization of planar combinatorial graphs. *Duke Math. J.*, 3(3):460–472, 1937.

[101] K. Marriott, H. Purchase, M. Wybrow, and C. Goncu. Memorability of visual features in network diagrams. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2477–2485, 2012.

[102] T. Nishizeki and M. S. Rahman. *Planar Graph Drawing*, volume 12 of *Lecture Notes Series on Computing*. World Scientific, 2004.

[103] Ø. Ore. Note on Hamilton circuits. *Amer. Math. Monthly*, 67(1):55, 1960.

[104] J. Pach and G. Tóth. Monotone drawings of planar graphs. *Journal of Graph Theory*, 46(1):39–47, 2004.

[105] H. C. Purchase. Effective information visualisation: a study of graph drawing aesthetics and algorithms. *Interacting with Computers*, 13(2):147–162, 2000.

[106] H. C. Purchase, D. A. Carrington, and J. Allder. Empirical evaluation of aesthetics-based graph layout. *Empirical Software Engineering*, 7(3):233–255, 2002.

[107] H. C. Purchase, E. E. Hoggan, and C. Görg. How important is the "mental map"? - an empirical investigation of a dynamic graph layout algorithm. In M. Kaufmann and D. Wagner, editors, *Graph Drawing, 14th International Symposium, GD 2006, Karlsruhe, Germany, September 18–20, 2006. Revised Papers*, volume 4372 of *Lecture Notes in Computer Science*, pages 184–195. Springer, 2006.

[108] H. C. Purchase, C. Pilcher, and B. Plimmer. Graph drawing aesthetics—created by users, not algorithms. *IEEE Transactions on Visualization and Computer Graphics*, 18(1):81–92, 2012.

[109] S. Rahman. Convex graph drawing. In M.-Y. Kao, editor, *Encyclopedia of Algorithms*, pages 1–7. Springer Berlin Heidelberg, Berlin, Heidelberg, 2015.

[110] R. C. Read. A new method for drawing a planar graph given the cyclic order of the edges at each vertex. *Congr. Numer.*, 56:31–44, 1987.

[111] A. Ribó Mor. *Realization and Counting Problems for Planar Structures: Trees and Linkages, Polytopes and Polyominoes*. PhD thesis, Freie Universität Berlin, 2006.

[112] J. Richter-Gebert. *Realization spaces of polytopes*, volume 1643 of *Lecture Notes in Mathematics*. Springer-Verlag, 1996.

[113] K. H. Rosen. *Discrete Mathematics and Its Applications*. McGraw-Hill Higher Education, 7th edition, 2011.

[114] D. P. Sanders. On paths in planar graphs. *J. Graph Theory*, 24(4):341–345, 1997.

[115] A. Schmid and J. M. Schmidt. Computing Tutte paths. In *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*, volume 107 of *LIPIcs*, pages 98:1–98:14, 2018.

[116] A. Schulz. Drawing 3-polytopes with good vertex resolution. *J. Graph Algorithms Appl.*, 15(1):33–52, 2011.

[117] I. Streinu. Acute triangulations of polygons. *Discrete & Computational Geometry*, 34(4):587–635, 2005.

[118] I. Streinu. Erratum to "pseudo-triangulations, rigidity and motion planning". *Discrete & Computational Geometry*, 35(2):358–358, 2006.

[119] R. Tamassia, editor. *Handbook on Graph Drawing and Visualization*. Chapman and Hall/CRC, 2013.

[120] C. Thomassen. Deformations of plane graphs. *Journal of Combinatorial Theory, Series B*, 34(3):244–257, 1983.

[121] 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.

[122] W. T. Tutte. A theorem on planar graphs. *Trans. Amer. Math. Soc.*, 82(1):99–116, 1956.

[123] W. T. Tutte. Convex representations of graphs. *Proceedings of the London Mathematical Society*, s3-10(1):304–320, 1960.

[124] W. T. Tutte. How to draw a graph. *Proceedings of the London Mathematical Society*, 3(1):743–767, 1963.

[125] C. Ware, H. C. Purchase, L. Colpoys, and M. McGill. Cognitive measurements of graph aesthetics. *Information Visualization*, 1(2):103–110, 2002.

[126] H. Whitney. Congruent graphs and the connectivity of graphs. *American Journal of Mathematics*, 54(1):150–168, 1932.

[127] A. Wigderson. The complexity of the Hamiltonian circuit problem for maximal planar graphs. Technical Report 298, Princeton University, 1982.

[128] D. R. Wood. Bounded degree book embeddings and three-dimensional orthogonal graph drawing. In P. Mutzel, M. Jünger, and S. Leipert, editors, *Graph Drawing, 9th International Symposium, GD 2001 Vienna, Austria, September 23–26, 2001, Revised Papers*, volume 2265 of *Lecture Notes in Computer Science*, pages 312–327. Springer, 2001.

[129] A. Wotzlaw, E. Speckenmeyer, and S. Porschen. Generalized k-ary tanglegrams on level graphs: A satisfiability-based approach and its evaluation. *Discrete Applied Mathematics*, 160(16-17):2349–2363, 2012.

[130] M. Yannakakis. Four pages are necessary and sufficient for planar graphs (extended abstract). In *Proc. 18th Annu. ACM Sympos. Theory Comput. (STOC 1986)*, pages 104–108, 1986.

[131] M. Yannakakis. Embedding planar graphs in four pages. *J. Comput. Syst. Sci.*, 38(1):36–67, 1989.

Zusammenfassung

Diese Arbeit behandelt drei unterschiedliche Problemstellungen aus der Disziplin des Graphenzeichnens (Graph Drawing). Bei jedem der behandelten Probleme ist die gesuchte Darstellung planar.

Ordered Level Planarity. Wir führen das Problem Ordered Level Planarity ein, bei dem es darum geht, einen Graph so zu zeichnen, dass jeder Knoten an einer vorgegebenen Position der Ebene platziert wird und die Kanten als y-monotone Kurven dargestellt werden. Dies kann als eine Variante von Level Planarity interpretiert werden, bei der die Knoten jedes Levels in einer vorgeschriebenen Reihenfolge platziert werden müssen. Wir klassifizieren die Eingaben bezüglich ihrer Komplexität in Abhängigkeit von sowohl dem Maximalgrad, als auch der maximalen Anzahl von Knoten, die demselben Level zugeordnet sind. Wir motivieren die Ergebnisse, indem wir Verbindungen zu einigen anderen Graph Drawing Problemen herleiten: Mittels Reduktionen von Ordered Level Planarity zeigen wir die $\mathcal{NP}$-Schwere einiger Probleme, deren Komplexität bislang offen war. Insbesondere wird gezeigt, dass Clustered Level Planarity bereits für Instanzen mit zwei nichttrivialen Clustern $\mathcal{NP}$-schwer ist, was eine Frage von Angelini, Da Lozzo, Di Battista, Frati und Roselli [2015] beantwortet. Wir zeigen die $\mathcal{NP}$-Schwere des Bi-Monotonicity Problems und beantworten damit eine Frage von Fulek, Pelsmajer, Schaefer und Štefankovič [2013]. Außerdem wird eine Reduktion zu Manhattan Geodesic Planarity angegeben. Dies zeigt, dass ein bestehender [2009] Polynomialzeitalgorithmus für dieses Problem inkorrekt ist, es sei denn, dass $\mathcal{P} = \mathcal{NP}$ ist.

Bucheinbettungen von dreifach zusammenhängenden planaren Graphen mit zwei Seiten. Wir zeigen, dass jeder dreifach zusammenhängende planare Graph mit Maximalgrad 5 Teilgraph eines Hamiltonischen planaren Graphen ist. Dies ist äquivalent dazu, dass ein solcher Graph eine Bucheinbettung auf zwei Seiten hat. Der Beweis ist konstruktiv und zeigt in der Tat sogar, dass es für die Realisierbarkeit nur notwendig ist, den Grad von Knoten separierender 3-Kreise zu beschränken—die übrigen Knoten können beliebig hohe Grade aufweisen. Dieses Ergebnis ist bestmöglich: Wenn die Gradschranke auf 6 abgeschwächt wird, gibt es Gegenbeispiele. Diese Ergebnisse verbessern Resultate von Heath [1995] und von Bauernöppel [1987] und, unabhängig davon, Bekos, Gronemann und Raftopoulou [2016], die gezeigt haben, dass planare Graphen mit Maximalgrad 3 beziehungsweise 4 auf zwei Seiten realisiert werden können.

Konvexitätssteigernde Deformationen. Wir zeigen, dass jede planare geradlinige Zeichnung eines intern dreifach zusammenhängenden planaren Graphen stetig zu einer solchen deformiert werden kann, in der jede Fläche ein konvexes Polygon ist. Dabei erhält die Deformation die Planarität und ist konvexitätssteigernd—sobald ein Winkel konvex ist, bleibt er konvex. Wir geben einen effizienten Algorithmus an, der eine solche Deformation berechnet, die aus einer asymptotisch optimalen Anzahl von Schritten besteht. In jedem Schritt bewegen sich entweder alle Knoten entlang horizontaler oder entlang vertikaler Geraden.